



HAL
open science

On geometric models of epistemic logic

Roman Kniazev

► **To cite this version:**

Roman Kniazev. On geometric models of epistemic logic. Logic in Computer Science [cs.LO]. Institut Polytechnique de Paris, 2023. English. NNT : 2023IPPAX132 . tel-04534589v2

HAL Id: tel-04534589

<https://hal.science/tel-04534589v2>

Submitted on 5 Jul 2024

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

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



INSTITUT
POLYTECHNIQUE
DE PARIS

NNT : 2023IPPAX132

Thèse de doctorat



On Geometric Models of Epistemic Logic

Thèse de doctorat de l'Institut Polytechnique de Paris
préparée à l'École polytechnique

École doctorale n°626 École Doctorale de l'Institut Polytechnique de Paris (EDIPP)
Spécialité de doctorat : Informatique

Thèse présentée et soutenue à Palaiseau, le 16 Novembre 2023, par

ROMAN KNIAZEV

Composition du Jury :

Hans van Ditmarsch Directeur de recherche, IRIT, CNRS, Université de Toulouse	Président
Yoram Moses Full professor, Technion—Israel Institute of Technology	Rapporteur
Sonja Smets Full Professor, University of Amsterdam	Rapporteuse
Thomas Ågotnes Professor, University of Bergen	Examineur
Paul-André Melliès Directeur de recherche, IRIF, CNRS, Université Paris Cité	Examineur
Sergio Rajsbaum Full professor, Instituto de Matematicas, UNAM	Examineur
Eric Goubault Professeur, LIX, École Polytechnique	Directeur de thèse
Jean Goubault-Larrecq Professeur, LMF, CNRS, ENS Paris-Saclay	Co-directeur de thèse

**On Geometric Models
of
Epistemic Logic**

Roman Kniazev

RÉSUMÉ (EN FRANÇAIS)

Il y a environ 40 ans, deux approches majeures pour l'étude des systèmes distribués ont été développées. Une des approches, récompensée par le prix Gödel, a établi une perspective topologique sur le calcul distribué, exprimant la solvabilité des tâches distribuées à travers des notions standards de topologie algébrique, en modélisant les espaces d'états sous forme de complexes simpliciaux. Un complexe simplicial se compose d'un ensemble de sommets et d'une famille de simplexes – intervalles, triangles, tétraèdres, etc. Dans le contexte des systèmes distribués, les sommets représentent les états locaux possibles des processus dans le système, tandis que les simplexes sont des combinaisons globales compatibles de ces états locaux. L'existence d'un protocole, c'est-à-dire d'un programme qui résout une tâche, est alors considérée comme l'existence d'une fonction entre tels complexes, qui préserve de manière appropriée la structure.

En parallèle, une application importante de la logique épistémique, un type de logique modale, a apporté un point de vue alternatif sur la structure des systèmes distribués, en utilisant la notion de connaissance pour décrire le comportement des processus dans un système. Cette application a également remporté le prix Gödel. Dans ce contexte, la solvabilité des tâches s'exprime en termes de formules logiques. Telles formules permettent de formaliser dans un langage formel des expressions telles que "le processus A sait que le processus B stocke la valeur 1".

On s'est récemment rendu compte que ces deux approches sont étroitement liées et que les modèles topologiques peuvent, en fait, servir de modèles de logique épistémique. Cette thèse poursuit un programme de recherche visant à unifier ces deux approches en étendant et en enrichissant l'équivalence.

Le premier objectif est de généraliser la sémantique actuelle de la logique épistémique basée sur les complexes simpliciaux au cas des ensembles simpliciaux. La principale différence entre les complexes simpliciaux et les ensembles simpliciaux réside dans le fait que dans ces derniers, on peut avoir des simplexes qui partagent des sommets, mais ne partagent pas des faces de dimension supérieure. Nous montrons qu'avec ces modèles, on peut exprimer des connaissances de groupe non standard, c'est-à-dire des situations où la connaissance d'un groupe vu dans son ensemble est strictement supérieure à l'union des connaissances de ses membres. Le principal exemple vient de la cryptographie : une

clé secrète peut être distribuée en plusieurs parties entre un groupe d'agents, de sorte que personne ne connaisse par lui-même le message secret. Cependant, lorsque le groupe combine ses connaissances, il peut déchiffrer le message.

On considère ensuite une variante multi-sortes de la logique épistémique, dans laquelle les propriétés de l'environnement et les propriétés locales des agents sont exprimées séparément. Cette logique est interprétée dans les hypergraphes chromatiques, qui sont une généralisation supplémentaire des complexes simpliciaux, nous permettant de souligner le rôle des points de vue locaux des agents dans les systèmes distribués. Un avantage de cette approche est qu'elle permet d'exprimer succinctement des modèles qui ne sont pas nécessairement totalement tolérants aux pannes, mais juste résilients.

On étudie la dynamique de la connaissance dans les systèmes distribués en introduisant une variante dynamique des hypergraphes chromatiques. Dans ces modèles, les points de vue locaux des agents sont dotés d'une structure temporelle, ce qui permet de modéliser l'évolution des connaissances dans le temps. Cela présente une version duale des systèmes interprétés, étendant l'équivalence originale au cas dynamique. De plus, on discute de la relation entre la connaissance et la concurrence dans ce contexte.

Enfin, on présente des résultats sur l'utilisation de catégories supérieures dans l'étude des logiques modales multi-sortes. On montre que certaines logiques peuvent être représentées comme des catégories enrichies sur des ordres partiels, et qu'il y a une classe naturelle de modèles algébriques, par rapport à laquelle la logique est complète.

ACKNOWLEDGEMENTS

First and foremost, I would like to thank my *directeurs de thèse*, Éric and Jean for providing all the necessary guidance during this journey. I’ve learned a lot from you, sometimes even things you may not have intended. You bore with my endless pessimism and taught me how to channel this sentiment in more productive ways, for which I am deeply grateful.

I am thankful to all the members of the committee: to Sonja and Yoram, who kindly agreed to review this manuscript, for their remarks and advices; and to Thomas, Hans, Sergio, and Paul-André for agreeing to be part of the jury and their excellent questions.

During the three years, I have met many wonderful people who shared some pieces of wisdom with me, and many of these pieces of advice proved to be invaluable. The list is surely not complete, but the most striking realizations were catalyzed by the following individuals: Sam van Gool, Jérémy Dubut, Samuel Mimram, Nicolas Behr, Georg Struth, and Uli Fahrenberg. I would also like to thank Jérémy Ledent, who played a crucial role in making this thesis real.

I must mention my academic brother-in-arms, Cameron. Even though we didn’t end up with any results together, which is a pity, I greatly enjoyed all the time spent at La Montagne and other places. I do hope that everything will be better than it is now.

Participants and organizers of the post-pandemic conferences I could take part in also played an important role in this work. This includes the “Logic and Higher Structures” week at CIRM, GETCO 2022, TACL 2022, and Dagstuhl seminar 23272.

The friendly *ambiance* at LIX helped a lot to keep the spirits high, so I would like to thank all the members of the team for that. I would also like to express my gratitude to the members of IRIF, where the very last stretch of writing took place, for their hospitality.

Je remercie les citoyens et citoyennes de la France pour avoir soutenu mes études de master et doctorat dans ce pays, ainsi que pour leur accueil chaleureux.

To all friends that I had and still have, who are far and not so much, who were there when I needed it the most, I am grateful for your support.

All this would not be possible without Anya. Even if I knew all languages of the world, I would not be able to express my love and gratitude for your presence in my life. I am grateful for your support, and I am excited to face the future together.

Why am I writing this book?
Because I realize it's imperfect.
Dreamed, it would be perfection;
written, it becomes imperfect;
that's why I'm writing it.

The Book of Disquiet
Fernando Pessoa

CONTENTS

Résumé (en français)	iii
Acknowledgements	v
1 Preliminaries	1
1.1 Introduction	1
1.2 Vector-based description of tasks	3
1.3 Simplicial models of distributed computing	4
1.4 Distributed computing and epistemic logic	10
1.5 Simplicial models of epistemic logic	16
2 Epistemic logic and semi-simplicial models	21
2.1 Multi-agent epistemic logic with absent agents	22
2.2 Relational Kripke models	24
2.3 Semi-simplicial models	39
2.4 Conclusion and future work	50
3 Hypergraphs for knowledge	55
3.1 Two-level logic	57
3.2 Axiomatization	62
3.3 Hypergraphs and epistemic frames	66
3.4 Hypergraphs, neighborhood frames and Chu spaces	69
3.5 Conclusion and future work	75
4 Incorporating time	79
4.1 Epistemic-temporal frames	80
4.2 Dynamic chromatic hypergraphs	88
4.3 Knowledge and concurrency	94
4.4 Conclusion and future work	100
5 Conclusion	103
A Towards enriched categories for modal logic	105
A.1 Many-sorted algebras	106
A.2 Enriched categories	112
A.3 Final remarks	119
Bibliography	121

PRELIMINARIES

1.1	Introduction	1
1.2	Vector-based description of tasks	3
1.3	Simplicial models of distributed computing	4
1.4	Distributed computing and epistemic logic	10
1.5	Simplicial models of epistemic logic	16

1.1 Introduction

Distributed computer systems are omnipresent. In various kinds: from global networks of servers that ensure constant access to the internet to multicore processors found in smartphones that almost everyone carries in their pocket — distributed systems are the backbone of the modern world. But do we understand them well? Is there a theoretical justification why they work and, sometimes, do not?

Interestingly, until mid-80’s, we did not even know whether certain, fairly basic problems in distributed computing could be solved at all [FLP85]. Fortunately, nowadays, we have a plethora of theoretical tools that can be used to study distributed computing. In particular, about 40 years ago, two groundbreaking approaches emerged.

On the one hand, in [HM90], Halpern and Moses provided a formal framework for reasoning about knowledge in distributed systems. Behavior of agents that participate in a system can often be informally specified using epistemic notions: “a process sends a message when it knows that the previous message has been received”, “a server saves the data when it does not know whether other servers have it”, and so on. The paper, now considered a classic in computer science, uses epistemic logic to formalize such statements and to reason about distributed systems. In particular, it explores how information and common knowledge are needed for achieving coordinated action of agents and that certain problems are, in fact, unsolvable.

On the other hand, Herlihy and Shavit developed a topological approach to the analysis of distributed tasks in [HS99]. The paper shows that the set of possible states of processes, as well as the description of tasks and protocols, can be represented using

combinatorial topological structures called simplicial complexes. It turns out that the solvability of a task can be characterized by the topological properties of the corresponding simplicial complex.

Naturally, one might wonder about the relationship between these two approaches, as both can yield similar results. A research direction aimed at unification of the knowledge-based reasoning with topological models was initiated in [GLR18]. In this paper, Goubault, Ledent, and Rajsbaum have shown that simplicial complexes can serve as models for epistemic logic, and that unsolvability of certain tasks, which is usually shown by topological means, can be expressed in epistemic logic. This thesis expands these connections further.

Plan of the thesis and contributions. In the present chapter we recall the basics of the topological and epistemic approaches to distributed computing, and discuss the relationship between them. Nothing is novel here, except several remarks that generalize some definitions.

In Chapter 2, we consider semi-simplicial sets, a type of structures that generalize simplicial complexes, as models for epistemic logic. These models naturally allow interpreting non-standard, or implicit, group knowledge, that is, situations when a group of agents knows together more than the sum of their individual knowledge. We show that these models are equivalent to a generalized kind of Kripke frames, which subsumes equivalence results from [GLR21; GLR22; vDK] (Theorem 2.3.11). This allows us to show completeness of the logic with respect to these simplicial models. We also identify pertinent subclasses of simplicial models, and provide a complete axiomatization (Corollary 2.3.13). This chapter is an extended version of [Gou+23].

Next, in Chapter 3, we introduce a many-sorted epistemic logic that allows us to reason about properties of agents and the environment separately. We show that this logic is a conservative extension of the epistemic logic of [GLR21] (Theorem 3.3.7). The semantics of this logic is given using hypergraphs, geometrical structures that allow finer control over the interaction of agents and the environment. We show that the logic is complete with respect to these models (Theorem 3.2.11). This chapter extends the preprint [goubaultManysortedEpistemicLogic2023].

In Chapter 4, we discuss dynamic aspects of models of knowledge. In particular, we show how several types of temporal-epistemic structures that appear in the literature are related to each other (Theorem 4.1.20). We also extend the definition of models of true concurrency, called higher-dimensional automata, in order to accommodate knowledge and solve issues raised by Ledent in his thesis (Definition 4.3.6). In a similar vein, we

extend the definition of hypergraphs from Chapter 3 with dynamics, and show how they can be induced from temporal epistemic structures (Proposition 4.2.6).

The Appendix differs from the context of the main body of the thesis. There, we exemplify how categories can be used for completeness theorems in modal logic. We show algebraic completeness theorems for multi-sorted modal logic and its generalization for enriched categories (Theorem A.1.13, Theorem A.2.15).

1.2 Vector-based description of tasks

Before diving into geometric and epistemic models of distributed systems, let us first discuss the operational model we have in mind. For this exposition, we mainly follow [HS99].

A distributed system consists of a set of processes that communicate with each other through various means such as message passing or shared memory. The processes, often referred to as agents, are trying to achieve some goal or to solve a certain problem. The problem is usually formalized as a *decision task*. In such a task, each process has an input value, and the goal is to decide on an output value. More abstractly, a process starts in some initial state, and it should finish in some admissible final state, according to the task specification. States of agents are called *local*, and the state of the system as a whole is called *global*.

Fault-tolerant distributed computing is concerned with situations where some processes may crash during the execution. The reason for crashes may be various, from hardware failures to malicious attacks. We will use an abstract representation of a crash, where a process simply stops participating in the computation.

Let S_I and S_O be the set of possible initial and final states of a process, respectively.

Definition 1.2.1. An $(n + 1)$ -process *input vector* V^I is an $(n + 1)$ -component vector whose i -th component is an element of S_I or \perp . We additionally assume that there is at least one non- \perp component.

Input vectors formalize possible *global* initial states of the system: each component corresponds to the initial state of a process, and \perp means that the corresponding process crashes at the beginning of the computation. The definition is similar for output vectors. In general, a state of the system, not necessarily initial or final, is considered to be a combination of local states of processes and, possibly, the state of the environment.

A process i is called *participating* in V if $V_i \neq \perp$. A pair of vectors V, V' is called *matching* if for all $0 \leq i \leq n$, $V_i = \perp$ if and only if $V'_i = \perp$, where V_i denotes the i -th

component of V . That is, the vectors match if they have the same set of participating processes. A vector V is called a *prefix* of a vector V' if V is obtained from V' by replacing some non- \perp components with \perp , namely, for all $0 \leq i \leq n$, either $V_i = V'_i$ or $V_i = \perp$. A set of vectors is called *prefix-closed* if it contains all prefixes of its elements (except the vector consisting of \perp 's only). Notice that a set of vectors can be seen as a partially ordered set with respect to the prefix relation.

In fault-tolerant distributed computing, we are interested in the case where the set of possible vectors is prefix-closed. It corresponds to the fact that any process can crash at any time, and the system should still be able to decide.

We can now formally define a task.

Definition 1.2.2. A *task specification* for a pair of prefix-closed sets of vectors I and O is a relation $\Delta \subseteq I \times O$ that carries every input vector to a *non-empty* set of matching output vectors. A *decision task* is a tuple (I, O, Δ) consisting of a pair of prefix-closed sets of vectors I and O and a task specification Δ for I and O .

A well-known example of a decision task is the *binary consensus problem*. We define it by giving only the maximal vectors: I and O are then the closure by taking prefixes.

Example 1.2.3. The *binary consensus problem* for $n + 1$ agents is a decision task (I, O, Δ) , where the maximal vectors in I are of form $[x_0, \dots, x_n]$, where $x_i \in \{0, 1\}$, and there are just two maximal vectors in O : $[0, \dots, 0]$ and $[1, \dots, 1]$. If the values of participating processes in an input vector are the same, then the task specification Δ assigns to it the corresponding matching output vector, otherwise an input vector is assigned to both matching output vectors. The description formalizes the following condition: if all processes start with the same value, then they must agree on this value. In the case when processes start with different values, they may decide on either one of them.

1.3 Simplicial models of distributed computing

As we saw in the previous section, the set of possible global states represented by vectors is naturally a partially ordered set. This paves the way to the use of algebraic topology in the study of distributed computing: the poset formed by the global states of a distributed system can be seen as a simplicial complex. In this section, we discuss the basics of these structures.

Simplicial complexes

Definition 1.3.1. A *simplicial complex* K is a pair (V, S) that consists of a set V of vertices and a set $S \subseteq \mathcal{P}_\bullet(V)$ of non-empty finite subsets of V , called *simplices*, such that

- S is closed under inclusion, that is, if $X \in S$ and $Y \subseteq X$ with $Y \neq \emptyset$, then $Y \in S$;
- for every $v \in V$, the set $\{v\}$ is in S .

A simplex is called a *facet* if it is not contained in any other simplex, that is, it is maximal in S with respect to the inclusion order. The *dimension* (or *degree*) of a simplex X is defined as $|X| - 1$: for example, a simplex consisting of a single vertex is of dimension 0, a simplex consisting of two vertices is of dimension 1, and so on. A simplicial complex is called *pure* if all its maximal simplices have the same dimension.

A simplicial complex $K = (V, S)$ is a *subcomplex* of a simplicial complex $K' = (V', S')$ if $V \subseteq V'$ and $S \subseteq S'$, and this is denoted $K \subseteq K'$.

Example 1.3.2. An example of a simplicial complex is depicted in Figure 1.1. In this complex, the set of vertices V has six elements, they are depicted as white vertices. The set of simplices S has 21 elements: 1 simplex of size 4 which is represented as the interior of the tetrahedron, 5 simplices of size 3 that are depicted as green triangles, 9 simplices of size 2 that are depicted as edges between vertices, and the six simplices of size 1 which are the vertices.




Figure 1.1: A simplicial complex.

Definition 1.3.3. A *simplicial map* $f : K \rightarrow K'$ between two simplicial complexes $K = (V, S)$ and $K' = (V', S')$ is a map $f : V \rightarrow V'$ such that for every simplex $X \in S$, $f(X) = \{f(v) \mid v \in X\} \in S'$, that is, the image of a simplex is a simplex.

A simplicial map is just a function on vertices, but it should preserve the structure of simplices. From the perspective of distributed computing it can be seen as a function

that associates another local state with each *local* state of an agent, but in a way that respects the structure of *global* states.

In order to define an analogue of a task specification in the geometric setting, we introduce carrier maps.

Definition 1.3.4. Given two simplicial complexes $K = (V, S)$ and $K' = (V', S')$, a *carrier map* from K to K' , denoted $\Phi : K \rightarrow 2^{K'}$, is a function that assigns a subcomplex $\Phi(X) \subseteq K'$ to each simplex $X \in K$, such that if $X \subseteq Y$ then $\Phi(X) \subseteq \Phi(Y)$.

Intuitively, a carrier map is a relation between simplicial complexes: a simplex X is related to a simplex Y if $Y \in \Phi(X)$. The condition means that the carrier map is monotone with respect to the inclusion order on simplices, that is, if X is a subsimplex of Y , then $\Phi(X)$ is a subcomplex of $\Phi(Y)$. There are various constructions and properties of carrier maps. Given two carrier maps $\Phi, \Psi : K \rightarrow 2^{K'}$, Φ is said to be *carried* by Ψ if for every simplex $X \in K$, $\Phi(X) \subseteq \Psi(X)$. A carrier map $\Phi : K \rightarrow 2^{K'}$ is called *rigid* if for every simplex $X \in K$, $\Phi(X)$ is pure of the same dimension as X . If we consider a subcomplex $L \subseteq K$, then we write $\Phi(L) = \bigcup_{X \in L} \Phi(X)$. Using this notation, we can define the composition of carrier maps: given two carrier maps $\Phi : K \rightarrow 2^{K'}$ and $\Psi : K' \rightarrow 2^{K''}$, the composition $\Psi \circ \Phi : K \rightarrow 2^{K''}$ is defined as $(\Psi \circ \Phi)(X) = \Psi(\Phi(X)) = \bigcup_{Y \in \Phi(X)} \Psi(Y)$. A simplicial map defines a carrier map in a natural way: given a simplicial map $f : K \rightarrow K'$, we define a carrier map $\Phi_f : K \rightarrow 2^{K'}$ as the subcomplex $\Phi_f(X)$ induced by the simplex $f(X)$. As such, a carrier map can be post-composed and pre-composed with a simplicial map.

In fact, the notion of a carrier map is exactly the right notion of a relation between simplicial complexes from the point of view of category theory. In the standard presentation, the definition is asymmetric: we associate a subcomplex with each simplex. We can turn it into a symmetric version, with a carrier map defined as a set of pairs of simplices that satisfy certain conditions.

The key here is to see that the set of simplices is itself a partially ordered set (which in distributed computing is exactly the set of possible global configurations). Then, one can consider a generalization of a relation between sets to a relation between posets, called a *monotone relation*.

Recall that for a poset P , its *opposite* (or *dual*) P^{op} is the poset with the same underlying set, with $x \leq y$ in P^{op} if and only if $y \leq x$ in P .

Definition 1.3.5 ([Bil+11]). Given two posets X and Y , a *monotone relation* is an order-preserving map $R : Y^{\text{op}} \times X \rightarrow \mathbf{2}$, where $\mathbf{2}$ is the poset $\{\perp < \top\}$.

For a pair of elements $(x, y) \in X \times Y^{\text{op}}$ we write xRy if $R(y, x) = \top$, and we say that the pair is related by R . Then, the condition of monotonicity can be reformulated as follows: if $x \leq x'$, xRy and $y' \leq y$, then $x'Ry'$.

Proposition 1.3.6. Let $\Phi : K \rightarrow 2^{K'}$ be a carrier map between two simplicial complexes. It defines a monotone relation $\text{Rel}(\Phi) : (S')^{\text{op}} \times S \rightarrow \mathbf{2}$. Conversely, every monotone relation between sets of simplices $R : (S')^{\text{op}} \times S \rightarrow \mathbf{2}$ defines a carrier map $\Phi_R : K \rightarrow 2^{K'}$. Moreover, the two constructions are inverse to each other.

Proof. Given a carrier map $\Phi : K \rightarrow 2^{K'}$, we define a monotone relation $\text{Rel}(\Phi) : (S')^{\text{op}} \times S \rightarrow \mathbf{2}$ as follows: for every pair of simplices $X \in S$ and $Y \in S'$, XRY if and only if $Y \in \Phi(X)$. The relation is indeed monotone: suppose that $X \subseteq X'$, XRY and $Y' \subseteq Y$. By monotonicity of Φ , $X \subseteq X'$ implies $\Phi(X) \subseteq \Phi(X')$, so every $Y \in \Phi(X)$ is also in $\Phi(X')$. Thus, $X'RY$. As $\Phi(X')$ is a subcomplex, $Y \in \Phi(X')$ implies $Y' \in \Phi(X')$. So, $X'RY'$.

Conversely, given a monotone relation $R : S \times (S')^{\text{op}} \rightarrow \mathbf{2}$, we define a carrier map $\Phi_R : K \rightarrow 2^{K'}$ as follows: for every simplex $X \in S$, $\Phi_R(X) = \{Y \in S' \mid XRY\}$. It is straightforward to check that Φ_R is a carrier map.

The two constructions are inverse to each other by definition. \square

The proposition shows that carrier maps and monotone relations are essentially the same thing when restricted to posets that are simplicial complexes. For carrier maps, the definition is more geometric, and is given in terms of simplicial complexes and their properties, while in for posets it is given in terms of order and monotonicity.

The composition of two monotone relations is defined as the usual composition of relations: $R \circ S = \{(x, z) \mid \exists y, (x, y) \in S \wedge (y, z) \in R\}$. It is easy to see that this composition coincides with the composition of carrier maps.

Chromatic simplicial complexes

In distributed systems, processes often have unique identifiers that are used to distinguish between them. On the side of simplicial complexes, this can be modeled by assigning a color to each vertex. We fix now a set of colors \mathcal{A} .

Definition 1.3.7. A *chromatic simplicial complex* $K = (V, S, \chi)$ consists of a simplicial complex (V, S) together with a *coloring map* $\chi : V \rightarrow \mathcal{A}$, such that no two vertices of the same simplex have the same color, that is, for every $X \in S$, for any $v, v' \in X$, if $\chi(v) = \chi(v')$, then $v = v'$.




Figure 1.2: A chromatic simplicial complex.

Alternatively, a chromatic simplicial complex may be seen as a simplicial complex K together with a simplicial map $\chi : K \rightarrow A$ that preserves the dimension of simplices, where A is a simplicial complex with the set of vertices \mathcal{A} and the set of simplices consisting of all finite subsets of \mathcal{A} .

Given a simplex X in a chromatic simplicial complex K , the set of colors of its vertices is defined to be $\chi(X) = \{\chi(v) \mid v \in X\}$.

Definition 1.3.8. A *chromatic simplicial map* between two chromatic simplicial complexes $K = (V, S, \chi)$ and $K' = (V', S', \chi')$ is a simplicial map $f : K \rightarrow K'$ such that for every vertex $v \in V$, $\chi(v) = \chi'(f(v))$.

It is straightforward to see that every chromatic simplicial map is rigid.

The definition of a carrier map also extends to the chromatic case.

Definition 1.3.9 (Chromatic carrier map). A chromatic carrier map between two chromatic simplicial complexes $K = (V, S)$ and $K' = (V', S')$, written $\Phi : K \rightarrow 2^{K'}$, assigns a subcomplex of K' to every simplex $X \in K$, such that

- it is *monotone*: if $X \subseteq Y$ then $\Phi(X) \subseteq \Phi(Y)$;
- it is *rigid*: for every simplex X of dimension d , $\Phi(X)$ is pure of dimension d ;
- it is *color-preserving*: for every simplex X , $\chi(X) = \chi(\Phi(X))$.

Simplicial complexes in distributed computing

We now define the simplicial analogues of the notions of a task and a protocol.

Definition 1.3.10. A *simplicial task specification* is a triple (I, O, Δ) , where I and O are pure chromatic simplicial complexes, and $\Delta : I \rightarrow 2^O$ is a chromatic carrier map.

In the original formulation of the task specification in [HS99] (Definition 1.2.2), the input and output complexes are always based on some sets of vectors. As such, vertices of the complexes can always be labeled by the local values of processes.

The simplicial reformulation of the binary consensus problem is the following.

Example 1.3.11. For clarity, let us consider the simplicial formulation of binary consensus for two processes. The input is a chromatic simplicial complex I consisting of a pair of vertices for both colors, corresponding to the two possible starting values of each process, together with all possible four simplices, call them $(0, 0)$, $(0, 1)$, $(1, 0)$ and $(1, 1)$. The output complex O consists of two 1-dimensional simplices, one for each possible decision value 0 and 1, call them 0 and 1. The carrier map Δ assigns to input simplex $(0, 0)$ a single simplex 0, similarly to $(1, 1)$ only 1, and both 0 and 1 to $(0, 1)$ and $(1, 0)$.

So far, we have only discussed the way the problems are described. But how can the behavior of a system be modeled using simplicial complexes? In the operational approach, processes are represented as automata interacting with a shared object. Because of that, the properties of distributed systems become very implementation-dependent. Using simplicial complexes one can model the behavior of a system in a purely geometric way, without referring to any particular implementation. The idea behind the next definition is that if a simplex in a complex represents a global state of the system, then its image under a carrier map to another simplicial complex can describe a (possibly non-deterministic) evolution of the system to the end of an execution.

Definition 1.3.12. A *simplicial protocol* is a triple (I, P, Ψ) , where I and P are pure chromatic simplicial complexes, and $\Psi : I \rightarrow 2^P$ is a chromatic carrier map.

The simplicial complex P is called *the protocol complex* and contains the information about the possible global states of the system after the execution of the protocol.

Example 1.3.13. A simple example of a protocol is the following. A *test-and-set object* is a shared object that can be in two states: *free* and *taken*. A process can *test* the object, which will return the current state of the object. If the object is free, after being tested for the first time, it becomes taken. Thus, if there are several processes trying to access the object concurrently, only one of them can succeed: the *winner* will receive the value *free*, and all other processes will receive the value *taken*. For three processes, the initial complex is a simplex with three vertices, and the protocol complex consists of three simplices of dimension 2, glued pairwise by vertices. This is depicted on Figure 1.2. The carrier map Ψ sends the initial simplex to *all* simplices in the protocol complex.

In the simplicial setting, there is no formal difference between tasks and protocols: both just relate two simplicial complexes. The difference lies in the asymmetry that arises from the formalization of solvability.

Definition 1.3.14. A protocol (I, P, Ψ) solves a task (I, O, Δ) if there exists a chromatic simplicial map $\delta : P \rightarrow O$, such that for all $X \in I$, $\delta(\Psi(X)) \subseteq \Delta(X)$.

The idea is that a process must decide on its final value deterministically based on its local state after the execution of the protocol. This is formalized by the existence of a chromatic simplicial map δ . A simplicial map sends a vertex in P that corresponds to this local view to a vertex in O that corresponds to the decision value. Since it is a simplicial map, every global configuration is also sent to a global configuration. The condition $\delta(\Psi(X)) \subseteq \Delta(X)$ means that the decision value is in the set of possible decision values for the given input value, according to the task specification.

We can also reformulate the definition in more abstract categorical setting. Recall that a 2-category consists of a collection of objects, a set of 1-morphisms between objects, and a set of 2-morphisms between 1-morphisms. An example of a 2-category is the category of relations, whose objects are sets, 1-morphisms are relations, and 2-morphisms are inclusions of relations. It is also known [KV16] that monotone relations on posets form a 2-category, where 1-morphisms are monotone relations and 2-morphisms are inclusions of relations. The 1-morphisms in this context are usually denoted as $f : a \dashrightarrow b$.

By Proposition 1.3.6, both protocols and tasks can be seen as monotone relations. Thus, we can reformulate the solvability definition as follows:

Definition 1.3.15. A protocol $\Psi : I \dashrightarrow P$ solves a task $\Delta : I \dashrightarrow O$ if there exists a functional monotone relation $\delta : P \rightarrow O$ such that $\delta \circ \Psi \subseteq \Delta$.

Remark 1.3.16. In fault-tolerant distributed computing, the set of vectors is prefix-closed, and this is exactly why we can build the corresponding simplicial complex. When this is not the case, the geometric representation changes: this is the subject of Chapter 3.

1.4 Distributed computing and epistemic logic

Distributed protocols are often informally described in terms of knowledge of participants: “Once Alice knows that Bob sent the message, she should call Eve”, “after I know availabilities of others for the meeting, I can suggest a date” and so on. Such reasoning was indeed informal until Halpern and Moses provided a thorough account on applicability of a formal logical approach to analysis of protocols. The formal approach in question is called *epistemic logic*. It is a kind of modal logic: its syntax is a (classical) propositional logic with additional operators on formulas that models some idealized notion of knowledge.

In this section, we provide a brief introduction on the syntactic and semantics aspects of epistemic logic. The main reference for this section is [Fag+95].

Syntax of epistemic logic

We fix a set of agents \mathcal{A} and the set of atomic propositions \mathbf{Ap} that describe some basic facts. Then, the syntax of multi-agent epistemic logic is given by the following grammar.

$$\varphi, \psi ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_a \varphi \quad p \in \mathbf{Ap}, a \in \mathcal{A}$$

The additional connectives are encoded as usual in propositional logic. The operator K_a is the knowledge modality, and $K_a\varphi$ reads as “agent a knows that φ ”. The knowledge modality has a dual defined as $\widehat{K}_a\varphi := \neg K_a\neg\varphi$, which reads as “agent a considers possible that φ ”. The system includes rules standard for normal modal logic: modus ponens rule and the necessitation rule.

What are the axioms of knowledge? It is still a subject of discussion in the community as to what are the right axioms. The standard approach, however, is well-established: knowledge is modeled using axiom system **S5**. Let us go through the list of axioms.

The first axiom is axiom (**K**). This is the distributivity axiom: $K_a(\varphi \rightarrow \psi) \rightarrow (K_a\varphi \rightarrow K_a\psi)$. Together with modus ponens rule the axiom says that knowledge is closed under implication, that is if I know that φ implies ψ and I know φ , then I know ψ . In a way, it partially encodes the deductive capabilities modulo knowledge.

The next axiom is axiom (**T**). Formally, it is stated like this: $K_a\varphi \rightarrow \varphi$. In the context of epistemic logic this axiom is known as the axiom of veracity, or axiom of truth: if I know something, then it is true. Taking the contrapositive, it means that something which is false cannot be known. This axiom is precisely what distinguishes knowledge from belief.

The penultimate axiom uses the dual of knowledge operator. The axiom called axiom (**B**) is the following: $\varphi \rightarrow K_a\widehat{K}_a\varphi$. It means that if something is true, then I know that it is possible.

The last axiom in our list is axiom (**4**): $K_a\varphi \rightarrow K_aK_a\varphi$. It is known as the positive introspection axiom. If the previous axioms seemed almost self-evident, this axiom can be considered a deliberate choice. It says that if I know something, then I know that I know it. Even more interestingly, one can show that this set of axiom implies the axiom of negative introspection, called (**5**): $\neg K_a\varphi \rightarrow K_a\neg K_a\varphi$.

Relational semantics of epistemic logic

The semantics of epistemic logic is usually given by relational Kripke models. In such models one has a set of possible worlds that represent epistemic alternatives, and knowledge is interpreted via an equivalence relation on these alternatives.

Definition 1.4.1. An *epistemic frame* $F = (W, \sim)$ consists of a set of worlds W and an equivalence relation \sim on W . A *multi-agent epistemic frame* $F = (W, \{\sim_a\}_{a \in \mathcal{A}})$ for a set of agents \mathcal{A} consists of a set of worlds W and a family of equivalence relations $\{\sim_a\}_{a \in \mathcal{A}}$ on W .

Equivalence relations are usually called *indistinguishability* relations: if for an agent a , $w \sim_a w'$, from the point of view of a they are the same, that is they cannot be distinguished.

Definition 1.4.2. A *multi-agent epistemic model* $M = (W, \{\sim_a\}_{a \in \mathcal{A}}, V)$ consists of an epistemic frame (W, \sim) and a valuation function $V : W \rightarrow 2^{\text{AP}}$.

There are several possible definitions of morphisms between epistemic models. One of them is called *simulation*: for a pair of multi-agent models $M = (W, \{\sim_a\}_{a \in \mathcal{A}}, V)$, $M' = (W', \{\sim'_a\}_{a \in \mathcal{A}}, V')$, a simulation is a function $f : W \rightarrow W'$ such that for all $w, w' \in W$, if $w \sim_a w'$ then $f(w) \sim'_a f(w')$ and if $p \in V(w)$, then $p \in V'(f(w))$. A slightly stricter version of a morphism is considered in [GLR18]. There, for the needs of applications, the preservation of relations between agents is kept, but the condition on valuations is replaced with a stronger one: $p \in V(w)$ if and only if $p \in V'(f(w))$. These functions compose, so models form a category.

Let now $M = (W, \sim, V)$ be an epistemic model. We define the *satisfaction relation* \models_M as follows:

- $M, w \models p$ if and only if $p \in V(w)$;
- $M, w \models \neg\varphi$ if and only if $M, w \not\models \varphi$;
- $M, w \models \varphi \wedge \psi$ if and only if $M, w \models \varphi$ and $M, w \models \psi$;
- $M, w \models K_a\varphi$ if and only if for all $w' \in W$, if $w \sim_a w'$ then $M, w' \models \varphi$.

The idea is that an agent knows φ in a world w if and only if φ is true for all other worlds indistinguishable for the agent. For example, in a distributed system, a process knows a fact if and only if this fact holds in all global states compatible with its local state.

Remark 1.4.3. In possible world semantics, an agent doesn't know something only if he is presented with epistemic alternatives. In a way, we can see it as a quantitative ignorance: an agent is aware of all these alternatives as he considers each one of them possible, but he just doesn't know which one is "the real one". There are certain inconveniences with this approach. Say, we want to model a situation where Bob does not know a password. In the possible worlds approach, we have to consider as worlds all possible passwords, and that Bob cannot distinguish between them. It becomes even more complicated if we consider the multi-agent case: if an agent a does not know the state of another agent b , we have to consider all possible states of the other agent *from the point of view of a* . But some of these states of b may be in fact not permitted by the design of the system, that is b may not be actually be in these states from its own point of view. So, it means that agent a is always supposed to know the intrinsic rules of the system, for example all the states other agents can be in.

On the other hand, these effects may be mitigated by a coarser modelling of systems. In particular, if for analysis of the system we may care only about the fact that an agent does not know a password. In this case, it is enough to have two worlds where the passwords are different, thus avoiding the need for a possibly infinite set of indistinguishable worlds.

Knowledge in distributed computing

So far, we have only talked about situations where knowledge is static. In distributed computing, however, knowledge is always changing: agents communicate with each other, so new facts can be learned (or even forgotten) at any moment. One of the ways to capture the dynamics of knowledge is via *interpreted systems* [HM90; Fag+95].

To start, we have a set W of global states of the system. Usually, one considers it to be a subset of a hypercube: all possible combinations of local states of processes, together with the states of the environment. Formally, if we denote local states of agents as L_a of $a \in \mathcal{A}$, and L_e the set of possible states of the environment, then $W \subseteq L_e \times \prod_{a \in \mathcal{A}} L_a$. A *run* over W is a function $r : \mathbb{N} \rightarrow W$, that is, an infinite sequence of global states. A run models a possible execution of a system as a sequential change of global states. For a run r , and a time m , a pair (r, m) is called a *point* of the run. The global state in point (r, m) is thus $r(m)$, and local states in a point are defined by projections $\pi_a : W \rightarrow L_a$. A *system* \mathcal{R} over W consists of a set R of runs over W . The set of all points $R \times \mathbb{N}$ is denoted $\text{Pts}(R)$.

Definition 1.4.4. An *interpreted system* $\mathcal{I} = (\mathcal{R}, \ell)$ consists of a system \mathcal{R} of runs over

a set W of global states, together with a valuation function $\ell : \text{Pts}(R) \rightarrow 2^{\text{Ap}}$, where Ap is a set of atomic propositions.

The set of runs in an interpreted system describes the behavior of the system. Points are equipped with a valuation function which assigns some atomic properties to them. These properties may describe properties of agents or of the environment.

Local states of an agent naturally induce an equivalence relation on the set of global states. Two global states w, w' are indistinguishable by agent a , denoted $w \sim_a w'$ if and only if the local states of the agent a coincide in w and w' , that is, the projections $\pi_a(w)$ and $\pi_a(w')$ are equal. This allows us to see that every interpreted system has an underlying (static) multi-agent epistemic model. As such, interpreted systems can interpret epistemic logic. For an interpreted system I and a point (r, m) , the satisfaction relation is defined as follows:

- $\mathcal{I}, r, m \models p$ if and only if $p \in \ell((r, m))$;
- $\mathcal{I}, r, m \models K_a \varphi$ if and only if for all (r', m') such that $\pi_a(r(m)) = \pi_a(r'(m'))$
 $\mathcal{I}, r', m' \models \varphi$,

and the semantics of boolean connectives is defined as usual.

Surely, the point of supplying an epistemic frame (which underlies an interpreted system) with a temporal structure is to be able to give semantics of a temporal-epistemic logic, that is, a logic that has both epistemic and some temporal modalities. In the case of interpreted systems, the language is essentially a union of the epistemic logic with LTL [Pnu77]. Its syntax is defined in the following way:

$$\varphi, \psi ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_a \varphi \mid \bigcirc\varphi \mid \square\varphi \mid \varphi U \psi \quad p \in \text{Ap}, a \in \mathcal{A}$$

The \bigcirc modality means “on the next step of execution”, \square means “always”, and $\varphi U \psi$ means “ φ holds until ψ holds”.

The semantics for temporal modalities extends the epistemic part as follows:

- $\mathcal{I}, r, m \models \bigcirc\varphi$ if and only if $\mathcal{I}, r, m + 1 \models \varphi$;
- $\mathcal{I}, r, m \models \square\varphi$ if and only if for all $m' \geq m$, $\mathcal{I}, r, m' \models \varphi$;
- $\mathcal{I}, r, m \models \varphi U \psi$ if and only if there exists $m' \geq m$ such that $\mathcal{I}, r, m' \models \psi$ and for all m'' such that $m \leq m'' < m'$, $\mathcal{I}, r, m'' \models \varphi$.

Notice that the temporal aspect of this logic is *linear* and not branching. However, when epistemic modalities are interpreted, we can “jump” from one run to another via agents’ indistinguishability relations. Moreover, as we can now mix epistemic and temporal modalities, we can describe possible interactions of knowledge with time. For example, one can express “perfect recall”, called otherwise “no forgetting”, a property that informally means that an agent does not forget any information, as follows: $K_a \bigcirc \varphi \rightarrow \bigcirc K_a \varphi$.

For an extended review of the effects of different axiomatizations of knowledge-time interaction we refer to [BvdH14].

A small example of the proof-theoretic reformulation

Epistemic analysis of distributed protocols is often done semantically: one usually considers a particular model of a distributed system, and then shows that certain logical properties hold in this model. Some theorems about knowledge are shown to hold in an arbitrary model, but still using the semantic approach. An example of such a theorem is the *knowledge of preconditions principle* [Mos16]. Let us recall it, in a reduced form, here.

First, suppose there is a set of actions Act that can be performed by agents. The set of atomic propositions then contains propositions $\text{does}_a(\alpha)$, where $a \in \mathcal{A}$ is an agent and $\alpha \in Act$ is an action. When evaluated in an interpreted system, $\mathcal{I}, r, m \models \text{does}_a(\alpha)$ means that in an execution r and time m , agent a does α . A formula ψ is called a *necessary* condition for $\text{does}_a(\alpha)$ in \mathcal{I} if for all points (r, m) , if $\mathcal{I}, r, t \models \text{does}_a(\alpha)$ holds, then $\mathcal{I}, r, m \models \psi$ holds. Then, an action α is called *conscious* for agent a in \mathcal{I} if agent’s local state completely determines whether the action is performed: if $(r, m) \sim_a (r', m')$, then $\mathcal{I}, r, m \models \text{does}_a(\alpha)$ if and only if $\mathcal{I}, r', m' \models \text{does}_a(\alpha)$. Put differently, if action is performed, then the agent knows it is performed, that is, $\mathcal{I}, r, m \models \text{does}_a(\alpha) \rightarrow K_a \text{does}_a(\alpha)$.

The principle then consists in the following.

Theorem 1.4.5 ([Mos16, Theorem 3.1]). Let α be a conscious action for a in \mathcal{I} . If ψ is a necessary condition for $\text{does}_a(\alpha)$ in \mathcal{I} , then $K_a \psi$ is also a necessary condition for $\text{does}_a(\alpha)$ in \mathcal{I} .

Instead of the semantics-based approach, one can derive general properties using purely proof-theoretic methods. Which one to choose always depends on the particular situation, but in some cases it is interesting to see how the same statements can be

shown, or interpreted, in different ways. For example, the knowledge of preconditions principle can be reformulated as a meta-theorem of modal logic.

Proposition 1.4.6. If $\vdash \varphi \rightarrow \psi$ and $\vdash \varphi \rightarrow K\varphi$, then $\vdash \varphi \rightarrow K\psi$.

Proof. First, by necessitation, axiom **(K)** and modus ponens, from $\vdash \varphi \rightarrow \psi$, we have $\vdash K\varphi \rightarrow K\psi$. From this and $\vdash \varphi \rightarrow K\varphi$ by cut rule, we have $\vdash \varphi \rightarrow K\psi$. \square

In order to get back the knowledge of precondition principle, we substitute φ with $\text{does}_i(\alpha)$. Then the first premise means that ψ is a necessary condition for α , and the second premise means that the action is conscious. Using completeness of the logic with respect to the Kripke models, the proposition implies the following. If in a model $\text{does}_i(\alpha) \rightarrow \psi$ is valid (ψ is a necessary condition for α), and $\text{does}_i(\alpha) \rightarrow K\text{does}_i(\alpha)$ is valid (α is conscious), then $\text{does}_i(\alpha) \rightarrow K\psi$ is valid in the model too (knowledge of ψ is a necessary condition for α).

1.5 Simplicial models of epistemic logic

A line of research that aims to unify simplicial and epistemic approaches was initiated in [GLR18] by Goubault, Ledent, and Rajsbaum. The main observation is the following: the categories of multi-agent epistemic frames and chromatic simplicial complexes are isomorphic. Using this, one can transfer the notion of a valuation from frames to simplicial complexes, thus obtaining a simplicial model of epistemic logic.

In this section, we briefly describe the content of [GLR18] and [GLR22]. The theorems and definitions considered in this section are taken from there, and thus constitute the core of the approach, which is generalized in [Gou+23], the results of which are presented in Chapter 2.

Initially, it was shown that the category of *proper* multi-agent frames is isomorphic to the category of *pure* simplicial complexes. Recall that a frame $(W, \{\sim_a\}_{a \in \mathcal{A}})$ is called proper, if the intersection $\bigcap_{a \in \mathcal{A}} \sim_a$ is the identity on W , in other words, any pair of worlds is distinguishable by somebody.

Theorem 1.5.1 ([GLR18]). The category of proper multi-agent epistemic frames for the set of agents \mathcal{A} is isomorphic to the category of pure of chromatic simplicial complexes with the set of colors \mathcal{A} .

For the details of the proof, we refer to the original paper. The most important idea is the following. In a partial epistemic frame, we have a partial equivalence relation

associated with each agent. In a chromatic simplicial complex, the relation “to share a vertex of color a ” is a partial equivalence relation *too*. The isomorphism works by associating to every simplicial complex K a frame $\kappa(K)$, where the set of worlds is the set of facets of K , and $w \sim_a w'$ if and only if w and w' share a vertex of color a . The inverse of this construction is given by associating to every epistemic frame the simplicial complex $\sigma(F)$ with the set of facets being the set of worlds of F , and two facets share a vertex of color a if and only if the corresponding worlds are related by \sim_a .

By transferring the definition of a valuation from frames to simplicial complexes, we obtain the following.

Definition 1.5.2. A *simplicial model* $M = (K, \ell)$ consists of a pure chromatic simplicial complex K , and a labeling ℓ that associates a set of atomic propositions with each facet of K . A morphism of simplicial models $f : M \rightarrow M'$ is a chromatic simplicial map $f : K \rightarrow K'$ such that for every facet X of K , if $f(X) \subseteq Y$, then $p \in \ell(X)$ if and only if $p \in \ell(Y)$.

The semantics of epistemic logic in simplicial models is defined in as follows. Fix a simplicial model M and its facet X . Then:

- $M, X \models p$ if and only if $p \in \ell(X)$;
- $M, X \models \neg\varphi$ if and only if $M, X \not\models \varphi$;
- $M, X \models \varphi \wedge \psi$ if and only if $M, X \models \varphi$ and $M, X \models \psi$;
- $M, X \models K_a\varphi$ if and only if for all facets $Y \in K$ such that $a \in \chi(X \cap Y)$, $M, Y \models \varphi$,

where $X \cap Y$ is an intersection of two facets, and $a \in \chi(X \cap Y)$ means that the intersection has a vertex of color a .

Then, we have the following proposition.

Proposition 1.5.3. Given a simplicial model M with its facet X , $M, X \models \varphi$ if and only if $\kappa(M), X \models \varphi$. Conversely, given a proper epistemic model N together with a world w , $N, w \models \varphi$ if and only if $\sigma(N), w \models \varphi$.

The main application of the equivalence between the categories is the *logical obstruction method*. Recall that solvability of distributed tasks is defined as existence of a certain map between chromatic simplicial sets. In some cases, one can slightly change the setting for solvability by reformulating everything in terms of simplicial maps only.

One of the conditions for such cases is that every local view after execution of the protocol has a unique starting point, which can be interpreted as that processes do not forget the history of their actions. More formally, now instead of the output complex and the protocol complex together with a pair of carrier maps between them and the initial complex, we consider an *update* complex $I[A]$ and *task* complex $I[T]$ together with simplicial maps $I[A] \rightarrow I$ and $I[T] \rightarrow I$. The solvability is then defined as existence of a map $\delta : I[A] \rightarrow I[T]$.

The logical obstruction method then helps with establishing that a given task is not solvable by a given protocol by showing that no such map exists. The crucial component is the *knowledge gain lemma*.

Theorem 1.5.4 ([GLR18]). Given a morphism of simplicial models $f : M \rightarrow M'$, and φ a positive formula, that is, one that does not contain negations, except, possibly, in front of atomic propositions, and contains only K modalities. Then, $M', f(X) \models \varphi$ implies $M, X \models \varphi$.

The method then consists in finding a formula that holds in the task model but not in the update model. It has been used to reinterpret the impossibility of binary consensus in the original papers, and k -set agreement in [YN21] with further refinement using fixed point logic in [Nis22].

The next logical step is to consider different variants of simplicial models and see what their logical properties are. For example, what does it mean, logically, to have as models simplicial complexes that are not pure, that is with facets of different dimensions? This option was explored in [GLR22] by taking a similar route: a category of epistemic frames that is isomorphic to the whole category of chromatic simplicial complexes was found.

A binary relation is called a *partial equivalence* relation if it is symmetric and transitive.

Definition 1.5.5. A multi-agent *partial epistemic frame* $F = (W, \{\sim_a\}_{a \in \mathcal{A}})$ consists of a set of worlds W and a family of partial equivalence relations $\{\sim_a\}_{a \in \mathcal{A}}$ on W .

The idea behind the definition is that in certain worlds an agent can be just absent, which is represented exactly by such worlds where $w \not\sim_a w$. On the simplicial side that corresponds to facets without vertices of color a .

Theorem 1.5.6 ([GLR22]). The category of proper partial epistemic frames is isomorphic to the category of chromatic simplicial complexes with the set of colors \mathcal{A} .

Can other variants of combinatorial topological structures that appear in mainstream mathematics, such as simplicial *sets*, be seen as models for epistemic logic? And if so, what kind of epistemic content do they have? With these questions, the main part of the thesis starts.

EPISTEMIC LOGIC AND SEMI-SIMPLICIAL MODELS

2.1	Multi-agent epistemic logic with absent agents	22
2.2	Relational Kripke models	24
2.3	Semi-simplicial models	39
2.4	Conclusion and future work	50

As we have seen in Section 1.5, in a simplicial complex, “to share a vertex” is a partial equivalence relation on the set of simplices. It is a simple idea that links simplicial complexes with relational models of epistemic logic. A natural question to ask is how this idea can be generalized: what conditions on the Kripke frames can be dropped, while still allowing to have a corresponding simplicial model?

In this chapter we consider two generalizations of the idea. First, we consider a relation “to share a face”. Second, instead of simplicial complexes we consider semi-simplicial sets. The first generalization corresponds to going from individual knowledge of agents to knowledge of groups in epistemic logic called *distributed knowledge*. The second generalization allows us to express situations where knowledge of a group is greater than the sum of individual knowledge of agents.

Combined, the two generalizations provide a unified framework for the use of simplicial structures as epistemic models: Theorem 2.3.11 subsumes equivalence theorems shown for particular subcases in [GLR21; GLR22; vDK]. We identify these subcases and provide a complete system of axioms for each one of them (Theorem 2.2.14, Theorem 2.2.15). Surprisingly, models based on semi-simplicial sets happen to be equivalent to the models based on complexes, from the logical point of view, as we show in Theorem 2.2.24.

Plan of the chapter. We introduce syntax and axioms of the logic in Section 2.1. In Section 2.2, we define the relational Kripke models for the logic, and show that the logic is complete with respect to these models. We also consider subclasses of models and corresponding axioms in the same section. In Section 2.3 we introduce simplicial models and show that they are equivalent to the relational models. Similarly, we consider subclasses of simplicial models and their topological properties. Finally, in Section 2.4 we discuss possible directions for further research.

This chapter is based on [Gou+23], joint work with Éric Goubault, Jérémy Ledent and Sergio Rajsbaum.

2.1 Multi-agent epistemic logic with absent agents

Syntax

Let A be a finite set of agents, and \mathbf{Ap} a countable set of atomic propositions. We consider the language \mathcal{L}_D of epistemic logic with *distributed knowledge* [Fag+95], generated by the following BNF grammar:

$$\varphi, \psi ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid D_U \varphi \quad p \in \mathbf{Ap}, U \subseteq A$$

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

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

The operator $D_U \varphi$ is read “the group of agents U (collectively) knows φ ”, while its dual $\widehat{D}_U \varphi$ means that the group U considers the formula φ possible. The operator $K_a \varphi$ is the standard knowledge operator of epistemic logic that stands for “agent a knows φ ”. We also consider a sublanguage \mathcal{L}_K that consists of formulas that contain only K operator, that is, talk about individual knowledge only.

Axioms

For every modality D_U , we use the usual axiomatization of normal modal logics: we have all propositional tautologies, axiom **(K)**, closure by modus ponens, and the necessitation rule. Together with **(K)**, we have several additional axioms:

- **(K)** $D_U(\varphi \rightarrow \psi) \rightarrow (D_U \varphi \rightarrow D_U \psi)$
- **(4)** $D_U \varphi \rightarrow D_U D_U \varphi$
- **(B)** $\varphi \rightarrow D_U \widehat{D}_U \varphi$
- **(Mono)** for $U \subseteq U'$, $D_U \varphi \rightarrow D_{U'} \varphi$
- **(Union)** for U, U' , $\widehat{D}_U \text{true} \wedge \widehat{D}_{U'} \text{true} \rightarrow \widehat{D}_{U \cup U'} \text{true}$
- **(Empty)** $D_\emptyset \varphi \rightarrow \varphi$

We abbreviate the resulting system as **EC_n**, which stands for the logic of epistemic coverings (as we will see). The first three axioms are standard for the logic of knowledge.

However, in contrast with the more standard multi-agent epistemic logic $\mathbf{S5}_n$, we do not have the truth axiom **(T)**: $D_U\varphi \rightarrow \varphi$. It is also in contrast with the logic of belief, where the veracity of knowledge is substituted with a weaker axiom **(D)**: $B\varphi \rightarrow \widehat{B}\varphi$. By doing so, we model absence of agents: we think of an agent as being absent, or dead, if it “knows” false, which does not necessarily allow deriving false, as it would be the case with **(T)**.

The axiom **(Mono)** is a standard axiom for distributed knowledge: it says that if a group knows something, then any larger group also knows it. As we deal with situations where agents may be present or absent in certain worlds, we need to have a special axiom **(Union)** that says that if two groups exist, then their union also exists. The last axiom is needed to model the fact that an empty group always exists.

Remark 2.1.1. In [vdHM92], Van der Hoek and Meyer have shown that if we consider the axiom $D_{U \cup U'}\varphi \rightarrow D_U\varphi \vee D_{U'}\varphi$ for distributed knowledge, then $D_{U \cup U'}$ coincides with D_U or $D_{U'}$. The axiom **(Union)** can be seen as a weaker version of this axiom, as it can be rewritten as $D_{U \cup U'}\text{false} \rightarrow D_U\text{false} \vee D_{U'}\text{false}$.

Using the same approach as in [GLR22], we can express within the language \mathcal{L}_D the fact that agents can be dead or alive. For any agent $a \in A$ and group of agents $U \subseteq A$, we define the following formulas:

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

Note that $\text{dead}(a)$ is equivalent to $\neg\text{alive}(a)$, but $\text{dead}(U)$ is not equivalent to $\neg\text{alive}(U)$: the formula $\neg\text{alive}(U)$ is true when at least one agent in U is dead, whereas $\text{dead}(U)$ is true when all agents in U are dead.

Here are a few examples of valid formulas in \mathbf{EC}_n related to the life and death of agents.

- $\mathbf{EC}_n \vdash \text{dead}(a) \rightarrow K_a\varphi$: dead agents know everything, including false facts. More generally, whenever $a \in U$, $\mathbf{EC}_n \vdash \text{dead}(a) \rightarrow D_U\varphi$.
- $\mathbf{EC}_n \vdash \text{alive}(a) \rightarrow K_a\text{alive}(a)$: alive agents know that they are alive. The same holds for a group U of agents.
- $\mathbf{EC}_n \vdash \text{alive}(U) \rightarrow (D_U\varphi \rightarrow \varphi)$: axiom **(T)** is verified when restricted to groups of agents that are alive.

2.2 Relational Kripke models

As we have seen in Section 1.4, the usual semantics of epistemic logic is given by Kripke models, with a relation \sim_a for each agent a . In the case of distributed knowledge, it is also standard ([Fag+95]) to give the semantics of the group modality D_U by considering a relation \sim_U generated by the relations of each agent $a \in U$. More specifically, \sim_U is defined to be the intersection $\bigcap_{a \in U} \sim_a$. This allows to interpret distributed knowledge as the union of the knowledge of each agent in the group.

However, when one wants to show completeness of the logic with respect to these models by the canonicity argument, it can be quickly seen that the canonical model does not have this property. That is, in the canonical model, the relation \sim_U is not the intersection of the relations \sim_a . Because of this, it is usually called a “pseudo-model”. Thus, in order to get completeness, this pseudo-model needs to be “repaired”, or “unraveled” (the technique that first appeared in [Sah75]), in a way that the resulting model is (modally) equivalent to the original one, but has the desired property.

The original canonical pseudo-model is seen as an obstacle, but in fact it shows an important fact: the logic is also complete with respect to the models where the relation \sim_U is not the intersection of the relations \sim_a . This was first shown (implicitly) by Van der Hoek and Meyer in [vdHM92]. An even more curious fact was shown there: it is not possible to give an axiom that would characterize the models where \sim_U is exactly the intersection of the relations \sim_a .

In this section, instead of considering pseudo-models as a mere technical tool, we take them as the main semantic target of our logic. The main contribution is the adaptation to the setting where the group of agents is not constant, but can vary from world to world. We show the completeness result for this setting and consider various subclasses of models.

Generalized epistemic frames

As in [GLR22], we consider Kripke frames where the indistinguishability relation \sim is not an equivalence relation but just a partial equivalence relation. But we generalize one step further by associating an indistinguishability relation not only for each agent, but also for any set of agents, making it possible to interpret distributed knowledge in a very general manner. A similar definition appears in [BS10], except that indistinguishability relations are equivalence relations rather than PERs, and as a consequence condition (b) below is not needed.

Definition 2.2.1. A *generalized epistemic frame* is a structure $\mathcal{M} = \langle M, \sim \rangle$, where:

- $M = \{w_0, w_1, \dots\}$ is a set of *possible worlds*,
- \sim is a function that assigns to every group of agents $U \subseteq A$ a PER $\sim_U \subseteq M \times M$ called the *U -indistinguishability relation*, with \sim_\emptyset being an equivalence relation,
- the PERs \sim_U satisfy the following conditions:
 - (a) compatibility:

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

- (b) closure under union of the groups of alive agents:

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

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

We write $[w]_U$ for the equivalence class of w with respect to \sim_U . Usually, we write $w \sim_a w'$ as shorthand for $w \sim_{\{a\}} w'$. Additionally, we say that a world w is a *sub-world* of w' when $|w| \subsetneq |w'|$ and $w \sim_{|w|} w'$.

Example 2.2.2. On Figure 2.1 is an example of a generalized epistemic frame. It has seven worlds $M = \{w_0, \dots, w_5\} \cup \{w'_1\}$. Not all relations are shown, but only the generating ones. Sets of alive agents can be read off directly on the reflexive loops above a world; in w_5 , no agent is alive, that is $w_5 \not\sim_a w_5$ for all $a \in A$. The empty group can distinguish w_5 from other worlds: $w_5 \not\sim_\emptyset w_i, i \neq 5$. In world w_4 , agents a and c are alive, and they can distinguish it from all other worlds, but the empty group cannot. In w_2 and w_3 all three agents are alive. Agents b and c cannot individually distinguish w_2 from w_3 , however, together they can: we have $w_2 \sim_b w_3$ and $w_2 \sim_c w_3$, but $w_2 \not\sim_{\{b,c\}} w_3$. In w_1 and w'_1 , agents a and b are alive, but even together they cannot distinguish them, as we have $w_1 \sim_{\{a,b\}} w'_1$. World w_0 is a sub-world of both w_1 and w'_1 : only b is alive in w_0 and it cannot distinguish w_0 from w_1 or w'_1 .

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




Figure 2.1: A generalized epistemic frame

Definition 2.2.3. Let $\mathcal{M} = \langle M, \sim \rangle$ and $\mathcal{N} = \langle N, \sim' \rangle$ be two generalized epistemic frames. A morphism from \mathcal{M} to \mathcal{N} is a function $f : M \rightarrow N$ such that for all $U \subseteq A$, for all $u, v \in M$, $u \sim_U v$ implies $f(u) \sim'_U f(v)$,

We write \mathbf{GEF}_A for the category of generalized epistemic frames with agents A .

Models and semantics of epistemic logic

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

Definition 2.2.4. A (*generalized*) *epistemic model* $\mathcal{M} = \langle M, \sim, \ell \rangle$ over the set of agents A consists of a generalized frame $\langle M, \sim \rangle$ together with *valuation* function $\ell : M \rightarrow \mathcal{P}(\text{Ap})$. A morphism of epistemic models $f : \mathcal{M} \rightarrow \mathcal{N}$ is a morphism of underlying frames that preserves valuations, that is, if $p \in \ell_M(w)$, then $p \in \ell_N(f(w))$. The category of generalized epistemic models is denoted \mathcal{KM}_A .

Remark 2.2.5. Morphisms in \mathcal{KM}_A are in fact just functional simulations [BRV01]. They are different from the morphisms used in [GLR21]: there, the valuations of atomic propositions were preserved and reflected, that is $\ell_M(s) = \ell_N(f(s))$, whereas in our definition $\ell_M(s)$ is just a subset of $\ell_N(f(s))$. They are also different from morphisms in [GLR22]: there, morphisms can be seen as relations, as they are maps $f : M \rightarrow \mathcal{P}(N)$. This type of morphisms can be recovered by considering spans of functions.

Given a generalized epistemic model, we can define the satisfaction relation.

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

We are going to show soundness and completeness of \mathbf{EC}_n with respect to generalized epistemic models. In order to show completeness, we use the standard approach: by providing a canonical model, see for example [BRV01].

Definition 2.2.6. The *canonical generalized epistemic model* $\mathcal{M}^c = \langle M^c, \sim, L \rangle$ is defined as follows:

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

Lemma 2.2.7 (Truth Lemma). For any formula φ and any maximal consistent set of formulas $\Gamma \in M^c$, we have $\varphi \in \Gamma$ iff $\mathcal{M}^c, \Gamma \models \varphi$.

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

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

Lemma 2.2.8. In the canonical generalized epistemic model \mathcal{M}^c , for any $\Gamma \in M^c$, $U \subseteq |\Gamma|$ iff $\text{alive}(U) \in \Gamma$.

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

Theorem 2.2.9. The logic \mathbf{EC}_n is sound and complete with respect to the family of generalized epistemic models.

Proof. Soundness is straightforward to check: \mathbf{K} is valid in any relational frame, the axioms \mathbf{B} and $\mathbf{4}$ correspond exactly to the fact that the indistinguishability relations \sim_U are PERs. Axiom \mathbf{Mono} corresponds to condition (a) and axiom \mathbf{Union} corresponds to condition (b) in the definition of generalized epistemic frames.

For completeness, consider the canonical model \mathcal{M}^c from Definition 2.2.6. Axioms \mathbf{K}, \mathbf{B} and $\mathbf{4}$ ensure that \sim_U is a PER, as in the standard treatment of completeness. Similarly, axiom \mathbf{Mono} ensures that the generated family of PERs is monotone: assume $U \subseteq U'$, $\Gamma \sim_{U'} \Delta$ and $D_U \varphi \in \Gamma$, then by \mathbf{Mono} , $D_{U'}$ is also in Γ . As we assumed $\Gamma \sim_{U'}$, $\varphi \in \Delta$, so $\Gamma \sim_U \Delta$. Axiom \mathbf{Union} ensures that the condition (b) is satisfied. Assume that $\Gamma \sim_U \Gamma$, $\Gamma \sim_{U'} \Gamma$ and $D_{U \cup U'} \varphi \in \Gamma$. First, we have that if $\Gamma \sim_U \Gamma$, then $\widehat{D}_U \text{true} \in \Gamma$. Indeed, by contraposition we have that for any ψ , $\psi \notin \Gamma$ entails $D_U \psi \notin \Gamma$. In particular, as Γ is consistent, $\text{false} \notin \Gamma$, so $D_U \text{false} \notin \Gamma$. By standard reasoning, $\neg D_U \text{false} \in \Gamma$, as we intended. Hence, $\text{alive}(U) \wedge \text{alive}(U') \in \Gamma$. By \mathbf{Union} and modus ponens, $\text{alive}(U \cup U') \in \Gamma$. As the formula $\text{alive}(V) \rightarrow (D_V \varphi \rightarrow \varphi)$ is deducible when $V = U \cup U'$, we have $D_{U \cup U'} \varphi \rightarrow \varphi \in \Gamma$. As we assumed $D_{U \cup U'} \varphi \in \Gamma$, by modus ponens $\varphi \in \Gamma$, which shows that the canonical model is indeed a generalized epistemic frame.

Applying the Lindenbaum Lemma and the Truth Lemma, any consistent formula φ holds in some state of \mathcal{M}^c , and thus φ is satisfiable. \square

Subclasses of epistemic frames and axiomatization

Generalized epistemic models are very versatile. Depending on what kind of applications we have in mind, we might want to impose some extra properties on the structure of our models. For instance, in distributed computing, the model is usually assumed to have group knowledge to be exactly the sum of knowledge of individual agents, because a global state of the system is merely the sum of the local states of the agents, without any extra information. Similarly, in some applications, we might want to assume that agents never know whether other agents still participate or not; or, in opposite, they always know the exact set of agents that are still participating.

We now define a number of interesting properties of epistemic frames.

Definition 2.2.10. An epistemic frame $\mathcal{M} = \langle M, \sim \rangle$ is said to


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

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

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

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

Example 2.2.11. In the figure below, the leftmost frame has trivial empty-group knowledge, $(w_0 \sim_{\emptyset} w_1 \sim_{\emptyset} w_2)$, but also has an empty world w_2 , and is not proper $(w_0 \sim_{\{a,b\}} w_1)$. The middle frame has non-trivial empty-group knowledge $(w_1 \not\sim_{\emptyset} w_2)$, has no empty worlds, and is not proper. The rightmost frame has non-trivial empty-group knowledge, has an empty world, but is proper $(w_0 \not\sim_b w_1)$.




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

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

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

By “standard group knowledge”, we mean that group indistinguishability relations are generated by individual agents, that is, knowledge of the group is exactly the sum of individual agents’ knowledge. Dually, in terms of relations, it means that two worlds are indistinguishable precisely when everybody in the group cannot distinguish them. More formally, for every group U , $\sim_U = \bigcap_{a \in U} \sim_a$.

Example 2.2.12. In the figure below, the upper left frame is maximal, as it has all sub-worlds. The upper right frame is minimal as it does not have sub-worlds at all. The bottom left frame is pure as all its worlds have the same set of alive agents. The bottom right frame has non-standard group knowledge, contrary to all previous examples, since $w_0 \sim_a w_1$, $w_0 \sim_b w_1$, but $w_0 \not\sim_{a,b} w_1$.



We now show how to axiomatize the various sub-classes of epistemic models. As in the previous section, we establish soundness and completeness for each type of epistemic frames. The proof is modular: each property of the frames (Definition 2.2.10) corresponds to an extra axiom. We denote by U^c the complement of the set of agents U , that is $A \setminus U$.

- **(NE)** $\bigvee_{a \in A} \text{alive}(a)$;
- **(P)** $\text{alive}(U) \wedge \text{dead}(U^c) \wedge \varphi \rightarrow D_U(\text{dead}(U^c) \rightarrow \varphi)$;

- (**Max**) for $U \neq \emptyset$, $\text{alive}(U) \rightarrow \widehat{D}_U \text{dead}(U^c)$;
- (**Min**) $\text{alive}(U) \wedge \text{dead}(U^c) \rightarrow D_U \text{dead}(U^c)$;
- (**Pure**) $\text{alive}(A)$.

However, there are no axioms related to trivial empty-group knowledge and standard group knowledge. This is because those two properties cannot be expressed in the language \mathcal{L}_D : we will show that the logic is complete with respect to these models. As a consequence, they can be assumed “for free”, without additional axiom. Indeed, to show that the logic \mathbf{EC}_n is sound and complete with respect to frames with trivial empty-group knowledge, we can use the strategy from [BS10]. The proof that standard group knowledge requires no extra axiom is the subject of Theorem 2.2.24.

Remark 2.2.13. There are several interesting relationships between those axioms. Axiom **Pure**, which says that all agents are alive in all worlds, has many consequences. It entails the axioms **NE**, **Min**, and **Union**. Axiom **P** is greatly simplified and becomes $\varphi \rightarrow D_A \varphi$, where A is the set of all agents. Furthermore, $\mathbf{EC}_n + \mathbf{Pure}$ together entail Axiom **T**, so that the logic $\mathbf{KB4}_n$ becomes $\mathbf{S5}_n$ when **Pure** is assumed. Another possible interaction is $\mathbf{P} + \mathbf{Min}$, which can be reformulated together as $\text{alive}(U) \wedge \text{dead}(U^c) \wedge \varphi \rightarrow D_U \varphi$. This axiom appears in [GLR22], in the particular case of $U = \{a\}$.

To demonstrate how to show soundness and completeness with those axioms, we prove the following theorem.

Theorem 2.2.14. The logic $\mathbf{EC}_n + \mathbf{NE} + \mathbf{P} + \mathbf{Max}$ is sound and complete with respect to proper maximal epistemic models with trivial empty-group knowledge and no empty worlds.

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

Now we show completeness. We define the canonical model \mathcal{M}^c as in Definition 2.2.6, except that consistency now refers to the new logic. The proof of the Truth Lemma (Lemma 2.2.7) can easily be adapted. Let φ_0 be a consistent formula. We shall show that φ_0 is satisfiable in some good epistemic model. Let \mathcal{M}^c be the canonical model for the logic, as in Theorem 2.2.9. We need to show that the canonical model is a good model. Recall that in the canonical model $\Gamma \sim_U \Delta$ iff for any φ , $D_U\varphi \in \Gamma \Rightarrow \varphi \in \Delta$. The model \mathcal{M}^c is:

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

Maximal Let $U \subsetneq |\Gamma|$. We want to exhibit a sub-world Δ of Γ , such that $|\Delta| = U$. By Lemma 2.2.8, $\text{alive}(U) \in \Gamma$, so by axiom **Max**, maximality and consistency, $\neg D_U \neg \text{dead}(U^c)$ is in Γ . Then the set $\Delta^- = \{\text{dead}(U^c)\} \cup \{\psi \mid D_U\psi \in \Gamma\}$ is consistent, by the same reasoning as in the proof of the Truth Lemma. By Lindenbaum's Lemma, Δ^- can be extended to a maximal consistent set Δ . Moreover, $\Delta \sim_U \Gamma$ by construction, so $U \subseteq |\Delta|$. Also, as $\text{dead}(U^c) \in \Delta$, $|\Delta| \subseteq U$. Hence, $|\Delta| = U$ and Δ is a sub-world of Γ .

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

Trivial empty-group knowledge The canonical model \mathcal{M}^c , however, does not have trivial empty-group knowledge. Nevertheless, for every φ_0 , we can extract a sub-model, \mathcal{M}_0^c , which consists of the set $\{\Gamma \in M^c \mid \Gamma \sim_\emptyset \Gamma_0\}$, where Γ_0 is some maximal consistent theory that contains φ_0 . By monotonicity of \sim , this restriction preserves all worlds U -accessible from Γ_0 and properties of \sim . Moreover, \mathcal{M}_0^c has trivial empty-group knowledge. Thus, φ_0 holds in at Γ_0 in \mathcal{M}_0^c , which is proper, maximal, has no empty worlds and trivial empty-group knowledge. \square

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

Theorem 2.2.15. The logic $\mathbf{EC}_n + \mathbf{NE} + \mathbf{P} + \mathbf{Min}$ is sound and complete with respect to proper minimal epistemic models with trivial empty-group knowledge and no empty worlds.

Proof. Call an epistemic frame that satisfies properties from the statement of the theorem a good minimal model. The proof of all clauses, except the correspondence between axiom **Min** and minimality, is the same as in Theorem 2.2.14. So for soundness, let us show the validity of axiom **Min** in any good minimal model. Consider a good minimal model \mathcal{M} and a world $w \in \mathcal{M}$. Assume that $\text{alive}(U) \wedge \text{dead}(U^c)$ holds in w . It means that $|w|$ is exactly U . Then by minimality of \mathcal{M} , as w is not a subworld of any other world, in any w' , such that $w \sim_U w'$, we have $|w'| = U$. Hence, $\mathcal{M}, w' \models \text{dead}(U^c)$ and $\mathcal{M}, w \models D_U \text{dead}(U^c)$, so axiom **Min** is valid.

For completeness, it remains only to verify that the canonical model \mathcal{M}^c for this logic is minimal. Consider $\Gamma, \Delta \in \mathcal{M}^c$, such that $|\Gamma| \subsetneq |\Delta|$ and $|\Gamma| = U$. Suppose that $\Gamma \sim_U \Delta$. By Lemma 2.2.8, $\text{alive}(U) \wedge \text{dead}(U^c)$ belongs to Γ . By modus ponens and axiom **Min**, $D_U \text{dead}(U^c)$ is in Γ . By definition of \sim_U , $\text{dead}(U^c) \in \Delta$. But $|\Delta|$ is strictly bigger than U , thus there $a \in U^c$ such that $a \in |\Delta|$, so $\text{alive}(a) \in \Delta$ which contradicts with consistency of Δ . Thus, $\Gamma \not\sim_U \Delta$ as required. \square

For the case of pure models:

Theorem 2.2.16. The logic $\mathbf{EC}_n + \mathbf{NE} + \mathbf{P} + \mathbf{Pure}$ is sound and complete with respect to proper pure epistemic models with trivial empty-group knowledge and no empty worlds.

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

For completeness, we need to verify that the canonical model \mathcal{M}^c is pure. As $\text{alive}(A)$ is in any Γ in \mathcal{M}^c , by Lemma 2.2.8 we have $|\Gamma| = A$, so \mathcal{M}^c is pure. \square

The relationship between all the axiom systems that we consider is summarized in Figure 2.2.

Standard group knowledge and bisimulations

The last property that we did not axiomatize is the standard group knowledge, or, as we will see later from geometrical point of view, the distinction between simplicial




Figure 2.2: Variants of epistemic logic

sets and simplicial complexes. In this section, we show that restricting to standard group knowledge does not require any additional axiom. For that purpose, we use a construction called *unraveling*, which turns a generalized epistemic model \mathcal{M} into a new model $U(\mathcal{M})$ that is bisimilar to \mathcal{M} and has standard group knowledge.

Let \mathcal{M} be a generalized epistemic model. A *history* in \mathcal{M} is a finite sequence of the form $h = (w_0, U_1, w_1, \dots, U_k, w_k)$ for some $k \geq 0$, such that $w_{i-1} \sim_{U_i} w_i$ for all $1 \leq i \leq k$ and U_i is a maximal such subgroup of agents. Notice that since epistemic models do not necessarily have standard group knowledge, every pair of worlds w, w' can have several groups U maximal with respect to inclusion such that $w \sim_U w'$. For example, if $w \sim_U w'$, $w \sim_{U'} w'$ and both U, U' are maximal, then one has, in particular, two different histories (w, U, w') and (w, U', w') . We write $\text{last}(h) = w_k$ for the last element of a history, and we write $h \rightarrow_U h'$ if $h' = (h, U_{k+1}, w_{k+1})$ with $U \subseteq U'$.

Definition 2.2.17. The *unraveling* of \mathcal{M} is a generalized epistemic model $U(\mathcal{M}) = (H, \sim^u, \ell^u)$ defined as follows:

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

The unraveling construction is similar to the tree unraveling from [Sah75], adjusted to the case of multiple relations related by monotonicity.

Lemma 2.2.18. Let M be a generalized epistemic model. Then its unraveling $U(M)$ is a generalized epistemic model.

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

Lemma 2.2.19. For every model M , its unraveling $U(M)$ has standard group knowledge.

Proof. We need to show that for any two h, h' in $U(M)$, if $h \sim_U^u h'$ and $h \sim_{U'}^u h'$, then $h \sim_{U \cup U'}^u h'$. There are two cases: if $h = h'$ and $h \neq h'$. For the first case, the statement holds since $U(M)$ is a generalized epistemic model by Lemma 2.2.18.

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

Remark 2.2.20. In the proof we have also shown that if $h \sim_U^u h'$, then $\text{last}(h) \sim_U \text{last}(h')$.

Definition 2.2.21. Given two generalized epistemic models \mathcal{M} and \mathcal{N} , we say that morphism $f : \mathcal{M} \rightarrow \mathcal{N}$ is a *(collective) functional bisimulation* if the following conditions hold:

- (atoms) for any $w \in M$, $\ell_{\mathcal{N}}(f(w)) = \ell_{\mathcal{M}}(w)$;

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

The definition we give is an extension of the standard notion of bisimulation, which links structural similarity of models with validity of formulas. In particular, we have the following proposition by adapting the standard construction (see [BRV01]). A similar definition, together with the statement, appeared in [Roe07].

Proposition 2.2.22. If $f : M \rightarrow N$ is a functional bisimulation, then for any formula φ , we have that $M, w \models \varphi$ if and only if $N, f(w) \models \varphi$.

Proof. By induction on φ . □

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

Lemma 2.2.23. For every generalized epistemic model M , its unraveling $U(M)$ is bisimilar to M .

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

Suppose now that $h \sim_U^u h'$. By Remark 2.2.20, $\text{last}(h) \sim_U h'$, so (forth) condition is satisfied.

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

Theorem 2.2.24. The logic \mathbf{EC}_n is sound and complete with respect to models with standard group knowledge.

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

Consider the canonical model M^c for \mathbf{EC}_n from Theorem 2.2.9. It is shown that M^c is a generalized epistemic frame. By Lemma 2.2.18, $U(M^c)$ is a generalized epistemic frame, and by Lemma 2.2.19 it has standard group knowledge. For any formula φ , there is Γ in M^c , such that $M^c, \Gamma \models \varphi$. Since $\text{last} : U(M^c) \rightarrow M^c$ is a bisimulation by Lemma 2.2.23, for any $h \in U(M^c)$ such that $\text{last}(h) = \Gamma$, we have $U(M^c), h \models \varphi$. Thus, any formula φ is valid in a model with classic group knowledge, which concludes completeness. □

Principle of full communication

The principle of full communication was introduced in [vdHvLM99] as a measure of usefulness of an interpretation of a notion of group knowledge. Intuitively, group knowledge complies with the principle if whenever some fact is considered to be group knowledge, then it should be possible to establish this fact by means of communication between agents.

To make the notion of full communication precise, recall some definitions. For now, suppose that in the models we consider, the group knowledge is standard, that is $\sim_U = \bigcap_{a \in U} \sim_a$. This is an assumption in [Roe07; vdHvLM99].

The *knowledge set* of a group of agents U in a model is the set of formulas in \mathcal{L}_K (that is with individual knowledge modalities only) that at least one agent in U knows to be true in a given world:

$$\text{KS}_U(M, w) = \{\varphi \in \mathcal{L}_K \mid M, w \models K_a \varphi \text{ for some } a \in U\}.$$

Then, it is said that a formula can be shown through communication by a group of agents iff the formula is entailed by the knowledge set of the group. We say that distributed knowledge complies with the principle of full communication if for all $\varphi \in \mathcal{L}_K$:


$$M, w \models D_U \varphi \quad \Rightarrow \quad \text{KS}_U(M, w) \vdash \varphi.$$

Here is an example of a model that complies with the principle of full communication, and also describes a case of distributed group knowledge.

Example 2.2.25. Consider the following situation. Alice knows, that if it is going to rain, than her neighbor Eve takes an umbrella. However, Alice does not know if it is going to rain today or not. Bob knows that it is going to rain, but is not aware of Eve's rule. If we say that p means "it rains", and q means "Eve took her umbrella", then the epistemic situation is formalized by the following set of formulas: $K_a(p \rightarrow q), K_b p$. Neither Alice nor Bob knows that Alice took her umbrella: $\neg K_a q \wedge \neg K_b q$. At the same time, if Alice tells Bob about Eve's rule, and Bob tells Alice that it rains, then together they can conclude that q : $D_{\{a,b\}} q$. The corresponding model is depicted below. This model complies with the principle of full communication since q is entailed by $\text{KS}_{\{a,b\}}(M, w) = \{p, p \rightarrow q\}$.

$$\neg p, \neg q \xrightarrow{a} p, q \xrightarrow{b} \neg p, q$$

The usual example of a model for which the full communication principle fails is the following:



It has four worlds, in two proposition p holds, in two holds its negation $\neg p$. Moreover, we can see that this model has standard group knowledge: the group indistinguishability relation is the intersection of the individual ones, and the group can distinguish all four worlds. However, the full communication principle fails for this model: the knowledge set of the group consists only of tautologies in every world, while the group knows which proposition holds in every world.

Roelofsen in [Roe07] describes the exact class of models with standard group knowledge for which the full communication principle holds.

Definition 2.2.26 ([Roe07]). A model \mathcal{M} is called a *full communication model* if for all worlds $w \in \mathcal{M}$ and all $U \subseteq A$, every \mathcal{L}_K -formula that is consistent with $\text{KS}_U(\mathcal{M}, w)$ is satisfiable in any w' such that $w \sim_U w'$.

A stronger, just sufficient condition was identified by Van der Hoek, van Linder, and Meyer.

Definition 2.2.27 ([vdHvLM99]). A model \mathcal{M} is called *distinguishing* when for every pair of worlds w, w' there is a formula $\varphi_{w,w'}$ such that $\mathcal{M}, w \models \varphi_{w,w'}$, but $\mathcal{M}, w' \not\models \varphi_{w,w'}$.

The following proposition shows how these two notions are related.


Proposition 2.2.28 ([Roe07]). A finite distinguishing model is a full communication model.

Notice, however, that in this context the semantics of the distributed knowledge is assumed to be generated by the indistinguishability relations of individual agents, that is

$$M, w \models D_U \varphi \quad \text{iff} \quad \forall w' (\forall a \in U w \sim_a w' \Rightarrow M, w' \models \varphi).$$

But what happens if we don't have standard group knowledge? Then even to be finite and distinguishing is not enough to comply with the principle. Consider the following example.

Example 2.2.29. Here is a finite, distinguishing model with non-standard group knowledge where full communication fails.



The model is distinguishing as in the left world p holds, and in the right world $\neg p$ holds. The knowledge sets contain only tautologies, as agents know only tautologies, but the group knows whether p or $\neg p$ holds. Notice, this model is (collectively) bisimilar to the previous one, which means that distinguishability of models is not invariant under bisimulation.

We can even say that if group knowledge is non-standard, then the principle of full communication is not relevant: this type of group knowledge appears exactly when knowledge is not transferrable by communication, but appears when agents fuse into one entity. Such effects may appear in quantum systems, for further discussion we refer to [BS10].

At this point we might want to understand what kind of effects non-standard group knowledge can express. One of the examples is an environment interaction. Suppose agent a has a locked chest that either contains an orange or does not. The agent has no possibility of knowing what the situation really is, unless she opens the box with a key. Agent B has the key. Obviously, agent b also does not know whether the chest contains an orange or not. We can model the situation by an epistemic frame with two worlds, w_0 and w_1 and an atomic proposition p that means that the chest contains an orange. We let $p \in v(w_0)$ and $p \notin v(w_1)$, and $w_0 \sim_a w_1$, $w_0 \sim_b w_1$, thus modeling that neither a nor b knows the real situation: this is exactly the model from Example 2.2.29. However, if agent b brings the key to a , then they can unlock the chest and find out the real situation: $w_0 \not\sim_{\{a,b\}} w_1$. In this example we can see that non-standard group knowledge may arise in the situations that do not concern the information held by agents, but rather the tools they have, or levels of access to information.

2.3 Semi-simplicial models

In Section 1.5 we discussed simplicial models of epistemic logic. The correspondence worked by restricting the class of epistemic Kripke structures: this allowed to find a nice representation in terms of simplicial complexes. A natural question is to what extent the conditions on frames can be relaxed. In this section, we will show that

the correspondence can be extended to generalized epistemic frames from the previous section. On the side of simplicial models, we will consider *augmented semi-simplicial sets*, and *epistemic coverings* based on them.

Simplicial complexes and semi-simplicial sets

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

Less abstractly, an augmented semi-simplicial set F consists of a family of sets $F([n])$ for every $n \geq -1$, and a family of face maps $\partial_i : F([n]) \rightarrow F([n-1])$ for $i = 0, \dots, n$, satisfying the simplicial identities $\partial_i \partial_j = \partial_{j-1} \partial_i$ for $i < j$.

For example, the simplicial set depicted in Figure 2.3 has 1 simplex of dimension -1 , 4 vertices, 6 edges and 2 triangles. So, accordingly, the sets $F([-1])$, $F([0])$, $F([1])$ and $F([2])$ are chosen to have 1, 4, 6 and 2 elements, respectively. There are three face maps $\partial_i : F([2]) \rightarrow F([1])$ for $i = 0, 1, 2$ which assign to each triangle its three boundary edges. The maps $\partial_i : F([1]) \rightarrow F([0])$ for $i = 0, 1$ assign to each edge its source and target vertices.




Figure 2.3: A (semi-)simplicial set which is not a simplicial complex.

Notice the two differences with simplicial complexes: first, we have -1 -simplices, which can be thought as generalized connected components, and second, two simplices can have the same set of vertices, but still be different. That is, simplices are not defined by the set of their vertices.

Chromatic augmented semi-simplicial sets

As in the case of simplicial models [GLR21], our first step will be to decorate the vertices of a simplicial set with colors, representing the names of the agents in A . The resulting

structure is called a *chromatic* augmented semi-simplicial set, or *cset* for short. We identify A with the linear order $[n] = \{0 < \dots < n\}$.

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

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

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

For the rest of the section, we can either see this category as a slice category, or notice that, by the fundamental theorem of topos theory, the category of csets is once again a presheaf category on a site Γ made of simplices of the standard n -simplex. As with the site of semi-simplicial sets (see e.g., [Rie11]), Γ is the posetal category of subsets of A with the inclusion partial order, defined below.


Definition 2.3.1. The category Γ is such that:

- objects are (possibly empty) subsets of A ,
- there is a unique morphism $\delta_{U,V} : U \rightarrow V$ in Γ whenever $U \subseteq V$. Composition $\delta_{V,W} \circ \delta_{U,V} = \delta_{U,W}$ is given by the fact that $U \subseteq V \subseteq W$ implies $U \subseteq W$.


We write \mathbf{Cset}_A for the presheaf category on Γ . This category is equivalent to $\widehat{\Delta}_{\text{inj}}^+ / S_A$, hence a cset can equivalently be viewed as a functor $F : \Gamma^{\text{op}} \rightarrow \mathbf{Set}$. Given a cset $F \in \mathbf{Cset}_A$, and a group of agents $U \subseteq A$, the elements of $F(U)$ are called the *U -simplices*. When there is no ambiguity, we write $\partial_{U,V} : F(V) \rightarrow F(U)$ for the boundary operator $F(\delta_{U,V})$. For x a V -simplex, $\partial_{U,V}(x)$ is called the *U -face* of x . If it is clear which V is considered, we simply write $\partial_U(x)$.

Given $U \subseteq A$, the standard U -simplex $\Gamma[U]$, is defined as the representable presheaf $\Gamma(-, U)$, image of U by the Yoneda embedding $y : \Gamma \rightarrow \mathbf{Cset}_A$.

In less categorical terms, a cset F consists of a family of sets $F(U)$ for every $U \subseteq A$, called *U -simplices*, together with a family of boundary operators $\partial_{U,V} : F(V) \rightarrow F(U)$ for every $U \subseteq V \subseteq A$, such that $\partial_{U,V} \partial_{V,W} = \partial_{U,W}$ for every $U \subseteq V \subseteq W \subseteq A$.

Figure 2.4: The category Γ for three agents.

Example 2.3.2. A cset X is depicted in the figure below. We represent it as a simplicial set together with colors on the vertices. Elements of X_{-1} are depicted as dashed regions (and interpreted as generalized connected components). Elements of X_0 are depicted as vertices, X_1 as edges, X_2 as triangles, etc. The boundary operators ∂_i give the equations that permit to glue these simplices together, along lower dimensional simplices. The cset in the picture below is composed of seven vertices (colored with three agents), eight edges, two triangles, and two (-1) -simplices.



In the rest, by a (chromatic) simplicial *complex* we mean a cset in which if two simplices share the same set of vertices, then they coincide. More formally, if for two U -simplices s, s' $\partial_a(s) = \partial_a(s')$ for all $a \in U$, then $s = s'$.

Epistemic coverings

The notion of a cset already allows us to relax the conditions on the frames. In particular, we can consider frames with empty worlds, that is worlds where no agent is present: these are represented by the (-1) -simplices of the cset. Similarly, we can represent non-standard group knowledge by having simplices that share the same set of vertices, but

are not equal. In this situation, no individual agent can distinguish the two simplices, but a group of agents can.

However, we still can not represent a situation where no group of agents can distinguish a pair of worlds. In order to do so, we need to introduce the notion of an *epistemic covering*. The idea of an epistemic covering is to capture the geometrical structure of indistinguishability relations in a simplicial set, abstracting away from the notion of a world. However, in order to interpret the logic in a fashion similar to Kripke structures, we still need to keep a link with set-theoretical representation. That is, we want to obtain a category of structures more topological than Kripke structures, but that still would be isomorphic to the category of epistemic frames. We will do so by representing a frame by a pair of chromatic semi-simplicial sets, together with a morphism between them, where one cset represents worlds, thus preserving the necessary rigidity of a frame, whereas the other cset represents the indistinguishability relations.


In a cset X , given a V -simplex t and a subset $U \subseteq V$, we say that $s = \partial_U(t)$ is a *subsimplex* of t . We say that X is *projective* when for every simplex $s \in X$, there is a unique maximal simplex $\uparrow s$, such that s is a subsimplex of $\uparrow s$. Categorically, a projective cset is a disjoint union of representables.

Definition 2.3.3. An *epistemic covering* is a morphism $f : E \rightarrow B$ in \mathbf{Cset}_A such that E is projective and $f : E \rightarrow B$ is surjective. Equivalently, a morphism $f : E \rightarrow B$ is an epistemic covering if E is projective and every maximal simplex of B has a preimage. We refer to E as *total cset* and to B as *base cset*.

Epistemic coverings form a category where a morphism from $f : E \rightarrow B$ to $f' : E' \rightarrow B'$ is a pair of morphisms $\alpha_E : E \rightarrow E'$, $\alpha_B : B \rightarrow B'$, such that the following square commutes:

$$\begin{array}{ccc} E & \xrightarrow{\alpha_E} & E' \\ f \downarrow & & \downarrow f' \\ B & \xrightarrow{\alpha_B} & B' \end{array}$$

An epistemic covering $f : E \rightarrow B$ can be visually represented as an annotated cset by taking the base cset B and writing on a simplex s the set of maximal simplices from E that are mapped on s . In particular, as f is surjective, every maximal simplex in B must have an annotation. For instance, the annotation on the example of a cset from previous section represents an epistemic covering with the following maximal simplices in E : two 2-simplices w_2, w_3 ; three 1-simplices w_1, w'_1, w_4 with $f(w_1) = f(w'_1)$; one 0-simplex w_0 ; one (-1) -simplex w_5 .



Remark 2.3.4. There are two ways to produce an epistemic covering out of a given cset X . We can take X as the base of the covering, but we need to choose the space E . Choosing E amounts to deciding which simplices of X constitute the possible worlds. There are two natural choices: either take all of the simplices of X , or take only the maximal ones (a.k.a. the facets).

The first choice, that we call the “maximal” one, where E is the disjoint union of all the simplices of X , yields a faithful functor from \mathbf{Cset}_A to \mathbf{eCov}_A . A morphism $g : X \rightarrow Y$ is sent to the morphism of coverings $\langle \alpha_E, \alpha_B \rangle$, where $\alpha_B = g$ and α_E sends a simplex s to t if g sends s to t . This functor is injective on objects, so it makes \mathbf{Cset}_A a subcategory of \mathbf{eCov}_A , though not full. The maximal interpretation appears (implicitly) in [vDK] for example, where a formula can be evaluated in every simplex of a model.

The second choice, where E is the disjoint union of the maximal simplices of X , is the “minimal” one. It is the one that is studied in [GLR21; GLR22], where a formula is only evaluated in a facet of a simplicial complex. However, this construction is not functorial.

Isomorphism of categories

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

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

Now we define this functor on morphisms. Suppose we are given a morphism of epistemic coverings $\langle \alpha_E, \alpha_B \rangle$. Then, the morphism of epistemic frames $\kappa(\langle \alpha_E, \alpha_B \rangle) :$

$\kappa(f) \rightarrow \kappa(f')$ sends a world w_s to a world $w_{s'}$ if the image of the maximal simplex s in E is included in the maximal simplex s' in E' .

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

$$E_M = \bigcup_{w \in M} \Gamma[|w|].$$

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

Lemma 2.3.5. For any frame $\langle M, \sim \rangle$, B_M is a cset.

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

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

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

Proposition 2.3.6. The functors κ and σ define an isomorphism of categories: $\kappa \circ \sigma = \text{id}_{\mathbf{GEF}_A}$ and $\sigma \circ \kappa = \text{id}_{\mathbf{eCov}_A}$.

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

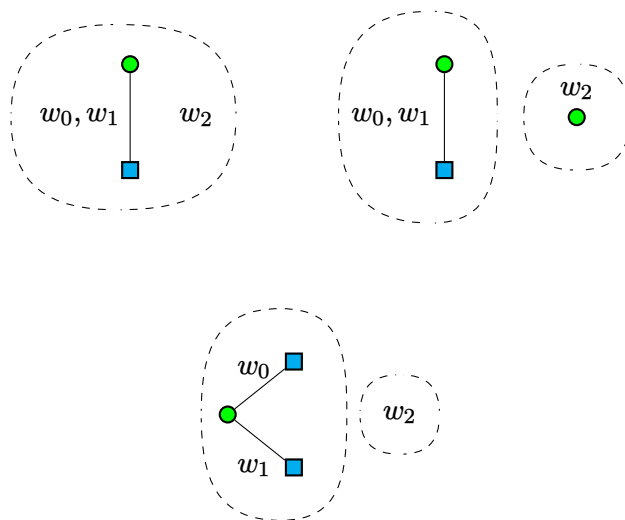
Remark 2.3.7. In previous articles, the equivalences that Proposition 2.3.6 generalizes have been called dualities. It is not a duality in the sense of category theory, as we have an equivalence, and not a dual equivalence of categories. The isomorphism theorem rather shows that Kripke frames already have an underlying topological structure which is being exposed in the corresponding simplicial set, thus making it closer to dual polyhedra, or as we will see in the next chapter, dual hypergraphs.

Subclasses of epistemic coverings


Definition 2.3.8. An epistemic covering $f : E \rightarrow B$ is said to

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

The intuition behind the definitions is exactly the same as in the case of frames. We give a few illustrative examples. In the picture on the next page, the top left covering has trivial empty-group knowledge as there is only one (-1) -simplex in the base space (one dashed region), but it has one maximal (-1) -simplex in the top space, so it has empty worlds. The top right covering has no empty worlds as it has two maximal simplices of dimension 1 and one 0-simplex in the top space. It, however, does not have trivial empty-group knowledge as it has two (-1) -simplices in the base space. It is not proper neither. The bottom covering is proper, as every simplex is annotated with at most one world.



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



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

Lemma 2.3.9. The properties of Definition 2.2.10 agree with the ones of Definition 2.3.8 up to the isomorphism in Proposition 2.3.6. Namely, if f is a covering of a certain type, then $\kappa(f)$ is of the same type, and conversely for σ .

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

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

Consider a covering $f : E \rightarrow B$ with standard group knowledge. Suppose that for a pair w_1, w_2 in the frame $\kappa(f)$, $w_1 \sim_a w_2$ for all a in some U . It means that $\partial_a(s_1) = \partial_a(s_2)$, where s_1 corresponds to w_1 and s_2 to w_2 . As B is a simplicial complex, it follows that $\partial_U(s_1) = \partial_U(s_2)$, so $\kappa(f)$ has standard group knowledge. Take now a frame M with standard group knowledge. Suppose there are two worlds w_1, w_2 in M such that $|w_1| = |w_2| = U$ and for all $a \in U$, $[w_1]_a = [w_2]_a$. Since M has standard group knowledge, $[w_1]_U = \bigcap_{a \in U} [w_1]_a$. Thus, $[w_1]_U = [w_2]_U$, which means precisely that if two simplices in $B_{\sigma(M)}$ have the same set of vertices, then they are equal. \square

Definition 2.3.10. An *epistemic covering model* $X = \langle f : E \rightarrow B, \ell \rangle$ consists of an epistemic covering $f : E \rightarrow B$, together with a labelling $\ell : \max(E) \rightarrow \mathcal{P}(\mathbf{Ap})$ that associates a set of atomic propositions $\ell(s)$ with each maximal simplex s of E . A morphism of epistemic covering models $\alpha : X \rightarrow Y$ is a morphism of epistemic coverings that preserves the labelling: $\ell(s) \subseteq \ell'(\uparrow \alpha(s))$, where $\uparrow \alpha(s)$ is the maximal simplex of E that contains $\alpha(s)$. We denote by \mathcal{EC}_A the category of epistemic covering models.

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

$$\begin{aligned}
X, s \models p & \quad \text{iff } p \in \ell(s) \\
X, s \models \neg \varphi & \quad \text{iff } X, s \not\models \varphi \\
X, s \models \varphi \wedge \psi & \quad \text{iff } X, s \models \varphi \text{ and } X, s \models \psi \\
X, s \models D_U \varphi & \quad \text{iff } X, s' \models \varphi \text{ for all maximal } s' \in E \\
& \quad \text{such that } \partial_U(f(s')) = \partial_U(f(s))
\end{aligned}$$

Proposition 2.3.6 can be readily extended to show that the categories of models are also isomorphic:

Theorem 2.3.11. The category of epistemic covering models \mathcal{EC}_A is isomorphic to the category of generalized epistemic models \mathcal{KM}_A .

Proof. We provide two functors κ, σ . On the underlying coverings and frames, they act as in Proposition 2.3.6. We just need to extend them to valuations. For $\kappa : \mathcal{EC}_A \rightarrow \mathcal{KM}_A$, for each maximal simplex s in E , we have an associated world w_s in $\kappa(f)$. We set $\ell_{\kappa(f)}(w_s) = \ell(s)$. Similarly, for $\sigma : \mathcal{KM}_A \rightarrow \mathcal{EC}_A$, there is an associated maximal simplex s_w for every world w . We set $\ell_{\sigma(M)}(s_w) = \ell(w)$. The rest of the proof is the same as in Proposition 2.3.6. \square

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


Lemma 2.3.12. Given a pointed epistemic covering model (X, s) , we have $X, s \models \varphi$ iff $\kappa(X), w_s \models \varphi$. Conversely, given a pointed generalized epistemic model (M, w) , we have $M, w \models \varphi$ iff $\sigma(M), s_w \models \varphi$.

Proof. We prove the first equivalence by induction on the structure of φ . The case of atomic propositions comes from the fact that we keep the labelling $\ell(w_s) = \ell(s)$. The case of boolean connectives is straightforward. For a formula of the form $D_U\varphi$ one can notice that we defined $w_s \sim_U w_{s'}$ iff $\partial_U(f(s)) = \partial_U(f(s'))$, which coincides with the semantics of D_U . The second equivalence follows from the first one, together with Theorem 2.3.11. \square

The corollary follows from Lemma 2.3.12.

Corollary 2.3.13. The logic \mathbf{EC}_n is sound and complete with respect to epistemic covering models.

Corollary 2.3.13 can be specialized to subclasses of models in the same way as Theorems 2.2.14 to 2.2.16. More interestingly, we also have completeness with respect to simplicial models with standard group knowledge, analogously to Theorem 2.2.24. It means that there is no formula of \mathbf{EC}_n that is valid in models where the base space is a simplicial complex, but not valid in models where the base space is a simplicial set. We get a geometric interpretation of the result from [vdHM92] that says that there is no formula that can distinguish models with standard and non-standard group knowledge: in the simplicial setting it means that epistemic logic cannot distinguish between simplicial complexes and simplicial sets. The unraveling construction from the proof of Theorem 2.2.24 thus has a topological meaning, as it can be seen as taking a universal cover of a simplicial complex, akin to the universal cover of a topological space, as shown in the figure on the next page.



2.4 Conclusion and future work

Distributed knowledge as update

As we could see in this chapter, the notion of distributed group knowledge allows multiple interpretations. In particular, it is not clear what we count as a group: should it be just an arbitrary group of agents, or a group of agents that can communicate everything they know to each other, or something else? Perhaps, we should consider a group of agents as a separate entity? If so, how do we model this fusion?

One possible line of research is to consider distributed knowledge as a form of update. We call this approach *dynamic distributed knowledge**, opposed to static distributed knowledge.

Let us briefly recall the setting of dynamic epistemic logic. In DEL, the syntax of epistemic logic is extended with an additional modal operator $[\alpha]$, which expresses that after the action α is performed, the formula φ holds. On the semantics side, in order to interpret $[\alpha]\varphi$, in a world w of a model M , we need to define *an update model* $M[\alpha]$, and a world $w[\alpha]$ that represent the knowledge of agents after the action was performed. The satisfaction relation is then defined as $M, w \models [\alpha]\varphi$ iff $M[\alpha], w[\alpha] \models \varphi$. For more details, we refer to [vDvdHK07].

We could try to apply the same idea to distributed knowledge. In this case, we would have a modality $[U]$ of distributed knowledge, which we interpret as the knowledge after the fusion, or full exchange of information, of agents in U . The semantics for this modality is then defined as

$$M, w \models [U]\varphi \Leftrightarrow M[U], v \models \varphi,$$

for every $v \sim_U w$, where $M[U]$ is the update model that has the same set of worlds as M , but the accessibility relation of each agent $a \in U$ is substituted with $\sim_U = \bigcap_{b \in U} \sim_b$. We can consider an even more general approach, where we allow the accessibility relation of each group to be explicitly defined as part of the model. This modality has several

*the name is suggested by Hans van Ditmarsch

interesting properties. For example, some axioms that we would expect to hold for distributed knowledge, such as $[U]\varphi \rightarrow \varphi$, are not valid.

It is also interesting to study how this modality is related to *resolution of distributed knowledge* [ÅW17], along with the topological content of this operation.

Belief in simplicial models

A question we have not addressed in this chapter is how to model belief in simplicial sets. There are various approaches to belief in epistemic logic. If we do not want to consider knowledge and belief at the same time, that is we want to change the modality K to another modality B , one way is to relax the veracity axiom $K\varphi \rightarrow \varphi$ to $B\varphi \rightarrow \widehat{B}\varphi$. This corresponds to the position that knowledge is “justified true belief”, thus if we substitute the truth axiom with a weaker option that allows beliefs be false, we get belief.

If we want to consider knowledge and belief at the same time, we can use *plausibility frames* [BS08]. Let us recall the definitions.

Definition 2.4.1. A *local plausibility frame* is a tuple $\langle S, \{\preceq_a\}_{a \in \mathcal{A}} \rangle$, where S is a set of worlds, and \preceq_a is a preorder on S for every agent a .

From the preorder \preceq_a we can recover the equivalence relation \sim_a as the union $\preceq_a \cup \preceq_a^{-1}$. Thus, our interpretation is that if for two worlds there is a chain of plausibility comparisons, then the two worlds are in fact indistinguishable for the agent. We can see local plausibility frames as usual epistemic frames, where each equivalence class with respect to the indistinguishability relation is equipped with a plausibility preorder, and worlds from different equivalence classes are incomparable.

This allows us to extend the isomorphisms of categories of epistemic frames and simplicial models we presented in this chapter to the case of plausibility frames and simplicial belief models. For simplicity, again, consider the case of pure models where the set of present agents is the same in every world (equivalently, the indistinguishability relations are equivalence relations).

Definition 2.4.2. A simplicial belief model is a pure simplicial epistemic model, such that every vertex is equipped with a preorder on the set of adjacent simplices.

Example 2.4.3. Consider the model in Figure 2.5. It consists of six 2-simplices, and every vertex is equipped with a preorder on the set of adjacent simplices. In fact, for every agent, the local orders are total orders. For agent \bullet in the center, the worlds are ordered clockwise, from one to six. Then, local orders do not necessarily coincide, for

example, for agent \bullet the north-east world is more plausible than the south-east world, while for agent \blacktriangle it is the other way around.




Figure 2.5: An example of a simplicial belief model.

The study of the categorical structure of such models, in particular what is the suitable notion of morphism should be, is left for future work. It is easy to see, however, that without morphisms, simplicial belief models and local plausibility frames are equivalent: given a local plausibility frame, we can construct a simplicial belief model by taking the set of worlds to be the set of equivalence classes with respect to the indistinguishability relation, and the local orders are the plausibility preorders. Conversely, given a simplicial belief model, we can construct a local plausibility frame by taking the set of worlds to be the set of simplices, and the plausibility preorders are the local orders.

Nominal extension

Hybrid logic [AtC07] is an extension of modal logic with *nominals*. Nominals is an additional sort of atomic propositions i, j, k, \dots that serve as names of worlds. In addition to nominals, hybrid logic has a *satisfaction operator* $@_i$ that gives access to the world with corresponding name.

Formal semantics of hybrid logic is defined as follows. Suppose we are given a set of atomic propositions \mathbf{Ap} and a set of nominals \mathbf{Nom} . Then, given a Kripke model \mathcal{M} , an assignment is a function g that assigns to every nominal a world. The semantics is then given for models and assignments as follows:

$$\begin{aligned}
 M, g, w \models p & \quad \text{iff } w \in \ell(p) \\
 M, g, w \models i & \quad \text{iff } w = g(i) \\
 M, g, w \models \varphi \wedge \psi & \quad \text{iff } M, g, w \models \varphi \text{ and } M, g, w \models \psi \\
 M, g, w \models \Box \varphi & \quad \text{iff } M, g, w' \models \varphi \text{ for all } w' \text{ such that } w \sim w' \\
 M, g, w \models @_i \varphi & \quad \text{iff } M, g, g(i) \models \varphi
 \end{aligned}$$

The models of hybrid logic, when they are epistemic models, can be seen as simplicial models with additional structure: we just equip a simplicial model with a function that assigns a simplex to each nominal. By doing so, we break the modal equivalence of simplicial models with standard and non-standard group knowledge. For example, in hybrid logic, it is possible to express the fact that a model has at most two worlds [BRV01]: the formula $@_i(\neg j \wedge \neg k) \rightarrow @_j k$. Consider the following model with non-standard group knowledge: it has two worlds w, w' where $w \sim_a w'$, $w \sim_b w'$ and $w \sim_{a,b} w'$. As it has just two worlds, the formula is true in the model. However, its unraveled version with standard group knowledge has an infinite number of worlds, so the formula is not true in the model.

Relation with three-valued logic

In our approach to knowledge of dead agents we chose to model it by making dead agents know everything: in particular, they know false. There is an alternative approach to knowledge of dead agents, which is to add a relation between models and formulas, which is called *definability*. This approach is taken in [RvDK23] and [vDK], where it was called *three-valued logic*. Let us recall the setting.

The language is defined by the grammar $\varphi ::= p_a \mid \neg\varphi \mid \varphi \wedge \psi \mid K_a\varphi$. Models are simplicial models, which reformulated in terms of this chapter, are maximal: a formula can be interpreted in every simplex. The big difference is in valuations: the set atomic proposition \mathbf{Ap} is assumed to be partitioned by agents, that is $\mathbf{Ap} = \bigcup_{a \in A} \mathbf{Ap}_a$. Then, a valuation is a function ℓ_a from vertices of color a to $\mathcal{P}(\mathbf{Ap}_a)$, and atomic propositions of other agents are assumed to be false. If a simplex s has a vertex of color a in which a proposition p_a holds, we say that p_a holds in the simplex s , denoted $p \in \ell(s)$.

Then, for a model X and a simplex s , the *definability*, denoted \bowtie is defined as follows:

$$\begin{aligned} X, s \bowtie p_a & \quad \text{iff } s \text{ has a vertex of color } a \\ X, s \bowtie \neg\varphi & \quad \text{iff } X, s \not\bowtie \varphi \\ X, s \bowtie \varphi \wedge \psi & \quad \text{iff } X, s \bowtie \varphi \text{ and } X, s \bowtie \psi \\ X, s \bowtie K_a\varphi & \quad \text{iff } X, s' \bowtie \varphi \text{ for some } s' \text{ which is } a\text{-adjacent to } s \end{aligned}$$

The satisfaction relation is then defined as follows:

$$\begin{aligned} X, s \models p_a & \quad \text{iff } p_a \in \ell(s) \\ X, s \models \neg\varphi & \quad \text{iff } X, s \not\models \varphi \\ X, s \models \varphi \wedge \psi & \quad \text{iff } X, s \models \varphi \text{ and } X, s \models \psi \\ X, s \models K_a\varphi & \quad \text{iff } X, s \bowtie K_a\varphi \text{ and for all } a\text{-adjacent } s', X, s' \models \varphi \end{aligned}$$

In order to distinguish this satisfiability from the one defined in this chapter, we will denote it by \models_3 .

We can define a translation of formulas from [RvDK23] to formulas of our two-valued logic. First, we define a translation $\text{def}^?(\varphi)$ such that $X, s \models \text{def}^?(\varphi)$ iff $X, s \bowtie \varphi$. It is done by induction on φ :

$$\begin{aligned} \text{def}^?(p_a) &= \text{alive}(a) \\ \text{def}^?(\neg \varphi) &= \text{def}^?(\varphi) \\ \text{def}^?(\varphi \wedge \psi) &= \text{def}^?(\varphi) \wedge \text{def}^?(\psi) \\ \text{def}^?(K_a \varphi) &= \widehat{K}_a \text{def}^?(\varphi) \end{aligned}$$

We then can see that $X, s \models_3 \varphi$ if and only if $X, s \models \text{def}^?(\varphi) \wedge \varphi$.

The question is studied from another angle, more similar to hybrid logic, in the upcoming “On Two- and Three-valued Semantics for Impure Simplicial Complexes” by Randrianomentsoa, van Ditmarsch, and Kuznets.

Conclusion

In this chapter we considered a multi-agent epistemic logic with distributed knowledge. We have shown completeness of the logic with respect to a class of generalized epistemic Kripke models. This approach has two main differences from the standard case: agents may be absent from some worlds, and the indistinguishability relation of groups is not necessarily the intersection of the indistinguishability relations of its members. We also considered simplicial models corresponding to the Kripke models, and showed that the two categories are isomorphic.

However, at the end, this correspondence is somewhat unsatisfactory. If two approaches are equivalent, why prefer one over the other? What do simplicial structures bring that Kripke models do not? Moreover, even though the epistemic coverings have certain topological content, the construction is a bit convoluted, so it is not clear how to use this structure for geometric reasoning. We address these issues in the next chapter.

HYPERGRAPHS FOR KNOWLEDGE

3.1	Two-level logic	57
3.2	Axiomatization	62
3.3	Hypergraphs and epistemic frames	66
3.4	Hypergraphs, neighborhood frames and Chu spaces	69
3.5	Conclusion and future work	75

As we have seen in the previous chapter, simplicial models provide a semantics for multi-agent epistemic logic. However, the fact that the category of epistemic coverings is isomorphic to the category of epistemic frames is somewhat ambiguous. If the models are entirely the same, why do we prefer one over the other? In this chapter, we will see that simplicial models pave the way for a more general framework that exhibits a richer structure. Namely, we will consider *hypergraphs* as a generalization of simplicial complexes, and show that they provide a semantics for a many-sorted modal logic that allows us to reason about local properties of agents and global properties of the environment separately.

A *hypergraph* is a generalization of a graph, where instead of just edges between pairs of vertices, one has *hyperedges* that can connect multiple vertices at once.

Definition 3.0.1. A (simple) hypergraph H is a pair (V, E) , where V is the set of vertices, and E is the set of hyperedges, which are subsets of V .

Recall that in Chapter 2, we described which simplices of a simplicial set were actually worlds by having an annotation in the form of a covering. Hypergraphs allow to express the same idea in a more direct way: we just do not have hyperedges that are not worlds. This was not possible with simplicial complexes which are hypergraphs that are downwards closed under inclusion, that is, if $e \in E$ and $e' \subseteq e$, then $e' \in E$.

Hypergraphs also can be not extensional, that is, there can be multiple hyperedges with the same set of vertices. In this case, we define a hypergraph as a pair of sets

This chapter is based on the preprint [goubaultManysortedEpistemicLogic2023], joint work with Éric Goubault and Jérémy Ledent.

(V, E) , and a binary relation \in between V and E . In epistemic systems, this allows us to represent models that are not proper, that is when there are several worlds that are not distinguishable by anybody.

We will consider a generalization of hypergraphs, where each vertex is assigned a color, which corresponds to an agent. So, these structures are a generalization of chromatic simplicial complexes. Notice that in this case we consider a generalization for complexes, and not semi-simplicial sets. We will then consider only individual knowledge, and not group knowledge.

Now let us discuss the ideas behind the logic. A familiar example of a many-sorted modal logic is CTL* [EH83] that has two sorts of formulas: path formulas and state formulas. It is natural for such a logic to have two sorts as it talks about two types of objects: possible states of the system and executions of the system. In the same way, we consider a two-sorted logic that talks about two types of objects that exist in epistemic multi-agent systems: agents and the environment. Corresponding sorts in the logic are *agent formulas* and *world formulas*. The former are interpreted in the points of view of agents (that is, vertices of hypergraphs), and the latter are interpreted in the worlds, that are represented in the models as hyperedges. As agents are different and may play different roles in the system, we have a separate sort for every agent. So, in total our logic has $|\mathcal{A}| + 1$ sorts, one for each agent and an extra one for the environment.

The main observation is the following. Recall the semantics of the knowledge modality. It says that an agent a knows that φ in a world w precisely when φ holds in every w' which is indistinguishable from w by a . In a way, the only important thing is the equivalence class of w with respect to \sim_a , and not the actual world w . This suggests that we can consider a logic where the knowledge modality is interpreted in the equivalence classes of worlds, that is points of view, and not in the worlds. In addition, we can also consider local atomic propositions that are relevant to the points of view of the agents. So, we want to be able to have agent formulas that look like this: $p_a \wedge K_a \Phi$, where p_a is an atomic proposition about agent a , and Φ is a world formula that describes some global property of the environment.

We are dealing with situations where agent may be present or absent in a world. Thus, we cannot talk about the knowledge of a specific agent in world formulas, or we will run into the issue of how to define the knowledge of a dead agent. Instead, we will have two modal operators that can test whether an agent exists in the world. This way, we are forced to explicitly quantify over points of view on an agent in a given world, if we want to talk about the knowledge of that agent. One modality is existential, denoted E_a , and the other is universal, denoted A_a . The intuitive meaning of those operators is

the following:

- $E_a\varphi_a$: “there exists a point of view for agent a such that φ_a holds”.
- $A_a\varphi_a$: “for every point of view of agent a , φ_a holds”.

Such an interpretation naturally suggests that we can have cases when an agent has multiple points of view on the same world. This provides a connection with neighborhood frames.

Related work. Recently, a single-sorted epistemic logic on hypergraphs was considered in [DLW23] to study weakly aggregative logics. There, vertices of hypergraphs are not colored as they are interpreted as worlds, so these models lie in-between epistemic frames and neighborhood frames. A framework that uses adjoint modalities was studied in [SD09] in the context of epistemic modalities “agent is uncertain about” and “agent has information that”. In interpreted systems [Fag+95], epistemic frames are generated by explicitly modeling the local states of agents and global states of the environment. However, at the level of syntax, no difference is made between local properties of the agents, and global properties of the environment.

Plan of the chapter. In Section 3.1 we introduce the logic **2CH** together with its semantics in chromatic hypergraphs. In Section 3.2 we show that the logic is complete with respect to these structures. The relation with the logic **KB4** and partial epistemic frames is shown in Section 3.3. In Section 3.4 we discuss further connections with Chu spaces and neighborhood frames.

3.1 Two-level logic

Syntax

Fix a finite set \mathcal{A} of agents. For each $a \in \mathcal{A}$, we have a set Ap_a of atomic propositions about agent a . We also have a set Ap_e of atomic propositions for the environment. We use lowercase letters with subscripts φ_a, ψ_a, \dots to denote agent formulas, and uppercase letters Φ, Ψ, \dots for world formulas.

Definition 3.1.1. The language of the logic **2CH** is defined as follows. For each agent $a \in \mathcal{A}$, there is a sort of agent formulas generated by the following grammar:

$$\varphi_a ::= p_a \mid \neg\varphi \mid \varphi \wedge \psi \mid \widehat{K}_a\Phi \quad \text{where } p_a \in \text{Ap}_a$$

The sort of world formulas is generated by the following grammar:

$$\Phi ::= p_e \mid \neg\Phi \mid \Phi \wedge \Psi \mid E_a\varphi_a \quad \text{where } a \in \mathcal{A} \text{ and } p_e \in \text{Ap}_e$$

We will use standard propositional connectives like **true**, \vee , \rightarrow , defined as usual. There are also dual modalities: $K_a\Phi := \neg\widehat{K}_a\neg\Phi$, and $A_a\varphi_a := \neg E_a\neg\varphi_a$. Modalities are read as follows: \widehat{K}_a means “agent a considers possible that”, K_a means “agent a knows that”, E_a means “there exists a point of view of agent a such that”, and A_a means “for all points of view of agent a ”. We call \widehat{K}_a and E_a *existential* modalities and K_a and A_a *universal* modalities.

Notice that for world formulas, there are $|\mathcal{A}|$ modal operators to choose from, whereas for agent formulas for a given agent, there is only one. Moreover, notice that the modalities in world formulas are not binders, so E_a should not be read as “there exists an agent a such that”. Here is an example of a syntactically well-formed world formula: $E_a K_a \Phi \wedge A_b K_b \Psi$. This formula says that there exists a point of view of agent a such that a knows Φ and for every point of view of b , b knows Ψ . Here is an example of an agent formula which is not syntactically well-formed: $K_a \Phi \wedge K_b \Psi$.

Semantics

A *hypergraph* is a generalization of a graph, where instead of just edges between pairs of vertices, one has *hyperedges* that can connect multiple vertices at once. In the introduction, we defined *simple* hypergraphs (Definition 3.0.1), to explain the proximity with simplicial complexes. In fact, we will be slightly more general than that, and allow multiple hyperedges to have the same set of vertices; this will allow us to model non-proper behavior. Namely, a (non-simple) hypergraph H is a triple (V, E, P) , where V is the set of vertices, E is the set of hyperedges, and $P : E \rightarrow 2^V$ assigns to each hyperedge a set of vertices.

In the context of multi-agent systems, we need to consider *chromatic* hypergraphs, where each vertex is assigned an agent name. This could be done by adding an extra piece of data $\chi : V \rightarrow \mathcal{A}$, as in chromatic simplicial complexes from [GLR21], described in Section 1.3, and the slice category for semi-simplicial sets from Chapter 2. Similarly to the way it is done in Chapter 2, we can instead tweak the original definition of a hypergraph a bit in order to make explicit the set of vertices assigned to each individual agent.

Definition 3.1.2. A chromatic hypergraph H is a tuple $(E, \{V_a, \rho_a\}_{a \in \mathcal{A}})$, where:

- for all $a \in \mathcal{A}$, V_a is the set of *views* of agent a ;

- E is a set of hyperedges;
- for each agent $a \in \mathcal{A}$, $p_a : E \rightarrow V_a$ is a surjective partial function. Additionally, we require that for each $e \in E$, $p_a(e)$ is defined for at least one $a \in \mathcal{A}$.

We will use the words *view* and *vertex* interchangeably, as well as *world* and *hyperedge*. Given a world $e \in E$, when $p_a(e)$ is undefined, we say that agent a is *dead* in e . Otherwise, a is *alive* in e , and we call $p_a(e)$ the view of a in e . Moreover, when $p_a(e) = v$, we say that v belongs to e , or that e contains v , and occasionally write $v \in_a e$. If a hyperedge consists of views v_0, \dots, v_{n-1} then we say that these views are *compatible*. Let us explain each condition imposed on p_a .

- p_a is surjective: every view belongs to at least one world.
- p_a is partial: not all agents are required to be alive in a world.
- p_a is functional: every world contains at most one view of each agent.
- $p_a(e)$ is defined for at least one $a \in \mathcal{A}$: every world contains at least one live agent.

Remark 3.1.3. The main difference between chromatic hypergraphs and relational models for epistemic logic is, as we will see, that worlds and points of view are of equal importance. In multiple papers on simplicial semantics ([GLR18], [GLR22], [vDK]), atomic propositions were associated either with vertices, or with facets, depending on the context. When associated with vertices, properties are local to agents, which is natural for distributed computing. On the other hand, the semantics is always defined via Kripke frames, which forces one to (re)define valuations on the level of worlds. We combine the two possibilities in one by capturing properties of both environments and agents in the two-level syntax.

Example 3.1.4. Three examples of chromatic hypergraphs are depicted in Figure 3.1. The three agents a, b, c are represented as colored shapes \bullet , \blacksquare , \blacktriangle , respectively. In all three hypergraphs, there is only one view per agent.

On the leftmost figure, all possible combinations of views are compatible, giving 7 hyperedges. Intuitively, in this situation, the agents do not know whether other agents exist. A scenario like this is standard in distributed computing, where agents are processes, and they do not know whether other processes are concurrently running.

The hypergraph depicted in the middle has a single hyperedge containing the three views. In this situation, all agents know that everyone is alive, as it is the only possible

world. This represents a scenario where every agent has guarantees that the other two agents are running.

The rightmost figure is a hypergraph with three hyperedges, each of which contains two views. It represents a situation where the points of views are pairwise compatible, but not all three of them are compatible. That is, there is no possible world that realizes all three of them at once. This could model a scenario where each agent receives a message from one of the other two agents, but they do not know who sent the message. Another interpretation might be in a quantum setting, as an example of *contextuality* [Abr+15].




Figure 3.1: Three examples of chromatic hypergraphs

Definition 3.1.5. A chromatic hypergraph model is a tuple $(H, \{\ell_a\}_{a \in \mathcal{A}}, \ell_e)$, where H is a chromatic hypergraph, and $\ell_a : \text{Ap}_a \rightarrow P(V_a)$, $\ell_e : \text{Ap}_e \rightarrow P(E)$ are valuation functions.

From now on, we sometimes omit the adjective “chromatic” when clear from context.

We can now define the semantics of **2CH** formulas with respect to chromatic hypergraph models. Given a hypergraph model H , the satisfaction relations are defined for every sort by mutual induction. As expected, world formulas are interpreted in a world $e \in E$, and agent formulas are interpreted in a point of view $v \in V_a$ of that agent.

$H, v \models_a p_a$	iff	$v \in \ell_a(p_a)$	$H, e \models_e p_e$	iff	$e \in \ell_e(p_e)$
$H, v \models_a \neg \varphi$	iff	$H, v \not\models_a \varphi$	$H, e \models_e \neg \Phi$	iff	$H, e \not\models_e \Phi$
$H, v \models_a \varphi \wedge \psi$	iff	$H, v \models_a \varphi$ and $H, v \models_a \psi$	$H, e \models_e \Phi \wedge \Psi$	iff	$H, e \models_e \Phi$ and $H, e \models_e \Psi$
$H, v \models_a \widehat{K}_a \Phi$	iff	$H, e \models_e \Phi$ for some $e \in E$ such that $p_a(e) = v$	$H, e \models_e E_a \varphi$	iff	$H, v \models_a \varphi$ for some $v \in V_a$ such that $p_a(e) = v$

Examples


Let us illustrate the semantics of **2CH** on a few examples.

Example 3.1.6. Consider again the three hypergraphs from Example 3.1.4, which we denote by H_1, H_2, H_3 , from left to right. For now, we do not worry about atomic propositions: we just need to assume that the sets Ap_* are non-empty, in order to get the constant $\text{true} = p \vee \neg p$. Also, recall that the agents are depicted as $a = \bullet$, $b = \blacksquare$, and $c = \blacktriangle$.

In the model H_1 , let us consider the hyperedge $e = \{\bullet, \blacksquare\}$. Then we have $H_1, e \models E_a \text{true}$ and $H_1, e \models E_b \text{true}$, but $H_1, e \not\models E_c \text{true}$. Indeed, the edge e does not contain a point of view of agent c , i.e., $p_c(e)$ is undefined. In fact, the world formula “ $E_a \text{true}$ ” is satisfied exactly in the worlds where agent a is alive. So let us write $\text{alive}(a) := E_a \text{true}$, so that $H, e \models \text{alive}(a)$ iff e contains a point of view of a , that is, iff $p_a(e)$ is defined.

We can now talk about whether agents know that other agents are alive. Let $v = \bullet$ be the (unique) point of view of a in the model. In model H_1 , all combinations of alive and dead agents are possible. So as expected, $H_1, v \models \neg K_a \text{alive}(b)$. In model H_2 however, a knows that all agents are alive, since this is the only possible world: $H_2, v \models K_a(\text{alive}(b) \wedge \text{alive}(c))$. In model H_3 the situation is more complicated: a knows that another agent is alive, but does not know which one. Thus, $H_3, v \models K_a(\text{alive}(b) \vee \text{alive}(c)) \wedge \neg K_a \text{alive}(b) \wedge \neg K_a \text{alive}(c)$.

Example 3.1.7 (2-agent binary input model with solo executions). As an example where atomic propositions play a role, we consider the situation where two agents are given a binary input value, either 0 or 1. Moreover, there can be solo executions (a.k.a. initial crash failures, in distributed computing), so that agents do not know if they are running alone or not. As before, the agents are depicted as $a = \bullet$ and $b = \blacksquare$. The sets of atomic propositions are $\text{Ap}_a = \{0_a, 1_a\}$, $\text{Ap}_b = \{0_b, 1_b\}$, and $\text{Ap}_e = \{\text{solo}\}$. The agent atomic propositions hold in the points of view indicated on the picture below. The environment atomic proposition solo holds in the four singleton hyperedges where only one agent is alive.



Let v be the top-left vertex, where agent a has input value 0. Then by definition $H, v \models 0_a$. Moreover, a does not know whether agent b is alive: $H, v \models \neg K_a E_b \text{true}$, which we could also reformulate as $H, v \models \neg K_a \text{solo}$, that is, a does not know whether this is a solo execution. We can also say that a considers possible that b is alive with

value 1: $H, v \models \widehat{K}_a E_b 1_b$. As a last example, we can express that a knows that if b is alive, its value is either 0 or 1: $H, v \models K_a(\neg \text{solo} \rightarrow E_b(0_b \vee 1_b))$. Alternatively, we could have used the operator A_b to express the same fact without relying on a property solo of worlds: $H, v \models K_a A_b(0_b \vee 1_b)$.

3.2 Axiomatization

We have all standard axioms and rules of classical propositional logic in every sort, as well as the following rules for modalities:

$$\frac{\vdash_E \Phi}{\vdash_a K_a \Phi} \text{ NEC-A} \qquad \frac{\vdash_a \varphi}{\vdash_E A_a \varphi} \text{ NEC-E}$$

$$\frac{\vdash \Phi \rightarrow \Psi}{\vdash \heartsuit \Phi \rightarrow \heartsuit \Psi} \text{ RM}$$

$$\frac{\vdash_E \Phi \rightarrow A_a \psi}{\vdash_a \widehat{K}_a \Phi \rightarrow \psi} \text{ ADJ-1} \qquad \frac{\vdash_a \varphi \rightarrow K_a \Psi}{\vdash_E E_a \varphi \rightarrow \Psi} \text{ ADJ-2}$$

where \heartsuit is any modality. The first two rules are the usual necessitation rules. The third rule is monotonicity of modalities. The last two rules are called *adjunction* rules, and they describe the interaction between the two pairs of modalities.

We also have the following axioms for modalities:

- $\vdash_a \varphi \rightarrow \widehat{K}_a E_a \varphi$ (surjectivity) every point of view belongs to some world;
- $\vdash_a \widehat{K}_a E_a \varphi \rightarrow \varphi$ (functionality) every world has at most one point of view of a given agent;
- $\vdash_E \bigvee_a E_a \text{true}_a$ (non-emptiness) every world contains at least one point of view.

These axioms correspond exactly to the properties of the function that relates hyperedges and vertices in a chromatic hypergraph.

Proposition 3.2.1. The logic **2CH** is sound with respect to chromatic hypergraph models.

In this logic, we can show that the universal modalities satisfy axiom **K**:

Proposition 3.2.2. For $\heartsuit \in \{A_a, K_a\}$, the axiom K_{\heartsuit} holds: $\heartsuit(\varphi \rightarrow \psi) \rightarrow (\heartsuit\varphi \rightarrow \heartsuit\psi)$.

Proof. First, we show that universal modalities distribute over conjunction, that is $\heartsuit(\xi \wedge \eta) \leftrightarrow (\heartsuit\xi \wedge \heartsuit\eta)$. Left-to-right direction: we have that $(\xi \wedge \eta) \rightarrow \xi$ and $(\xi \wedge \eta) \rightarrow \eta$. Applying the RM rule, we get that $\heartsuit(\xi \wedge \eta) \rightarrow \heartsuit\xi$ and $\heartsuit(\xi \wedge \eta) \rightarrow \heartsuit\eta$. From this, the left-to-right direction follows. Right-to-left direction: denote the modality adjoint to \heartsuit by \spadesuit , that is, if $\heartsuit = A_a$ then $\spadesuit = \widehat{K}_a$ and if $\heartsuit = K_a$ then $\spadesuit = E_a$. We have that $\spadesuit\heartsuit\xi \rightarrow \xi$ from $\heartsuit\xi \rightarrow \heartsuit\xi$ and the corresponding adjunction rule, similarly for η . From this we have that $\spadesuit\heartsuit\xi \wedge \spadesuit\heartsuit\eta \rightarrow \xi \wedge \eta$. By the same proof as in the left-to-right direction, we have that $\spadesuit(\heartsuit\xi \wedge \heartsuit\eta) \rightarrow \spadesuit\heartsuit\xi \wedge \spadesuit\heartsuit\eta$. Thus, we have that $\spadesuit(\heartsuit\xi \wedge \heartsuit\eta) \rightarrow \xi \wedge \eta$. Applying the adjunction rule, we get that $\heartsuit(\xi \wedge \eta) \rightarrow \heartsuit\xi \wedge \heartsuit\eta$, which is the right-to-left direction.

The fact that K follows from the distribution of universal modalities over conjunction is a standard proof: From $((\xi \rightarrow \eta) \wedge \xi) \rightarrow \eta$ by RM we have $\heartsuit((\xi \rightarrow \eta) \wedge \xi) \rightarrow \heartsuit\eta$. By distribution, we have $(\heartsuit(\xi \rightarrow \eta) \wedge \heartsuit\xi) \rightarrow \heartsuit((\xi \rightarrow \eta) \wedge \xi)$. Combining these two, we get $(\heartsuit(\xi \rightarrow \eta) \wedge \heartsuit\xi) \rightarrow \heartsuit\eta$, and thus $\heartsuit(\xi \rightarrow \eta) \rightarrow (\heartsuit\xi \rightarrow \heartsuit\eta)$. \square

We give a list of useful formulas that are derivable in the logic **2CH**.

Proposition 3.2.3. The following statements are derivable in the two-level logic **2CH**:

- | | |
|--|--|
| 1. $E_a\varphi \rightarrow A_a\varphi$ | 4. $\Phi \rightarrow A_a\widehat{K}_a\Phi$ |
| 2. $K_a\Phi \rightarrow K_aE_aK_a\Phi$ | 5. $\varphi \rightarrow K_aE_a\varphi$ |
| 3. $E_aK_a\Phi \rightarrow \Phi$ | 6. $K_a\Phi \rightarrow \widehat{K}_a\Phi$ |

Proof. For the first formula, we just apply the adjunction rule to $\widehat{K}_aE_a\varphi \rightarrow \varphi$, which is an axiom. In order to show the second formula, just apply the adjunction rule to $E_aK_a\Phi \rightarrow E_aK_a\Phi$, which is a tautology. Formulas 3, 4 and 5 are derived from $K_a\Phi \rightarrow K_a\Phi$, $\widehat{K}_a\Phi \rightarrow \widehat{K}_a\Phi$, and $E_a\varphi \rightarrow E_a\varphi$ respectively by applying the adjunction rule. The last formula is shown as follows: from formulas 3 and 4 we have $E_aK_a\Phi \rightarrow A_a\widehat{K}_a\Phi$. Applying the adjunction rule, we get $\widehat{K}_aE_aK_a\Phi \rightarrow \widehat{K}_a\Phi$. We also have $K_a\Phi \rightarrow \widehat{K}_aE_aK_a\Phi$, which is an axiom. From the last two formulas, we get $K_a\Phi \rightarrow \widehat{K}_a\Phi$. \square

Here is an intuitive explanation of these formulas:

1. If φ holds in a point of view of a , then it holds in all points of view of a . That is, there can be at most one point of view per world.

2. This is a form of positive introspection: if an agent knows something, then he knows that he knows it. However, we can put stress on the last “he”, that is, “if he knows something, then he knows that *he* knows it”.
3. This is a form of veracity: if a fact about a world is known by someone, then it is true.
4. If a certain fact holds in a world, then the agents in this world consider this fact possible. This is related to negative introspection.
5. Agents know the local facts about themselves.
6. This is the usual modal axiom D. It reflects the fact that every point of view belongs to at least one world.

In [GLR21], the use of local atomic propositions leads to a so-called assumption of locality, $K_a(p_{a,x}) \vee K_a(\neg p_{a,x})$: agents know the value of every local atomic proposition. In hypergraph models, valuations are local by construction:

Proposition 3.2.4. For any $p_a \in \text{Ap}_a$, the formula $K_a E_a p_a \vee K_a E_a \neg p_a$ is derivable in **2CH**.

Proof. By adjunction rules, we have $\widehat{K}_a A_a p_a \rightarrow p_a$ and $p_a \rightarrow K_a E_a p_a$. By cut rule, we have $\widehat{K}_a A_a p_a \rightarrow K_a E_a p_a$. This is equivalent to $K_a E_a p_a \vee K_a E_a \neg p_a$ by propositional logic. \square

Note that the above proof does not use the fact that p_a is atomic, so in fact, an agent knows *any* formula about itself. That is, for any φ , $K_a E_a \varphi \vee K_a E_a \neg \varphi$ is derivable.

Completeness

The proof of completeness uses the standard canonicity argument, extended to the many-sorted case. There is nothing surprising: the canonical model consists of the maximal consistent sets of formulas, now of several sorts. These sets satisfy standard properties, and together form a chromatic hypergraph.

Definition 3.2.5. A set of formulas S_* of sort $*$ is *inconsistent* if false can be derived from it. Otherwise, it is called consistent. A consistent set of formulas is *maximal* if it is not a proper subset of any other consistent set of formulas.

Proposition 3.2.6. Maximal consistent sets (MCS) of formulas of sort $*$ are closed under modus ponens: for any formula ξ of sort $*$, either $\xi \in S_*$ or $\neg\xi \in S_*$, and every (non-maximal) consistent set of formulas of sort $*$ is contained in a maximal consistent set of formulas of sort $*$.

Definition 3.2.7. The canonical hypergraph model consists of the following:

- the set of hyperedges is $E = \{S_e \mid S_e \text{ is an MCS of sort } e\}$;
- for each agent a , the set of vertices is $V_a = \{S_a \mid S_a \text{ is an MCS of sort } a\}$;

together with relations $R_a \subseteq E \times V_a$ for every agent a , which are defined as follows: $S_E R_a S_a$ iff for all formulas Φ , if $\Phi \in S_E$, then $\widehat{K}_a \Phi \in S_a$. The valuation function is defined by: $\ell_*(p_*) = \{S_* \mid p_* \in S_*\}$.

Proposition 3.2.8. In the canonical model, if $S_E R_a S_a$, if $K_a \Phi \in S_a$, then $\Phi \in S_E$.

Lemma 3.2.9. The canonical model is a chromatic hypergraph, that is, R_a is surjective, functional, and for any S_E there is S_a such that $S_E R_a S_a$ for at least one a .

Proof. First, suppose that $S_E R_a S_a$, $S_E R_a S'_a$, and $S_a \neq S'_a$. It means that there is a formula φ which is in S_a , but is not in S'_a . By adjoint axiom and modus ponens, $K_a E_a \varphi$ is in S_a . By Proposition 3.2.8, $E_a \varphi$ is in S_E . Using the definition of the canonical model, $\widehat{K}_a E_a \varphi$ belongs to S'_a . From there, by functionality axiom and modus ponens, $\varphi \in S'_a$, which is a contradiction. Thus, $S_a = S'_a$.

Second, we need to show that in the canonical model every vertex belongs to a hyperedge. Assume this is not the case, that is, there is a vertex S_a that does not belong to any hyperedge. It means that there is no maximal consistent set of formulas that contains $S = \{\Phi \mid K_a \Phi \in S_a\}$. In particular, it means that S is itself not consistent, that is, there is a finite set of formulas $\{\Phi_i\}$ such that $\bigwedge_i \Phi_i \rightarrow \text{false}$ is derivable. By applying necessitation, we get that $K_a(\bigwedge_i \Phi_i \rightarrow \text{false})$ is derivable, thus belongs to S_a . As K_a distributes over conjunction, and every $A_a \Phi_i$ is in S_a , we get that $K_a \bigwedge_i \Phi_i$ is in S_a . Applying modus ponens, we get that $K_a \text{false}$ is in S_a . Using surjectivity axiom, we get that false is in S_a , that is S_a is not consistent, which is a contradiction.

Lastly, we need to show that every hyperedge contains some vertex. Suppose it is not the case, that is there is a hyperedge S_E that does not contain any vertex. It means that for every a , the set $\widehat{K}_a S_E$ is not consistent. Thus, for all a , there is a finite set of formulas $\{\Phi_i^a\} \subset S_E$, such that $\bigwedge_i \widehat{K}_a \Phi_i^a \rightarrow \text{false}$ is derivable. We now show that this implies that S_E is not consistent. By applying necessitation, we get that $A_a(\bigwedge_i \widehat{K}_a \Phi_i^a \rightarrow \text{false})$ is

derivable for every a . By (K) and modus ponens, we derive $A_a(\bigwedge \widehat{K}_a \Phi_i^a) \rightarrow A_a \text{false}$. Combining them all together, we have that $\bigwedge_a \bigwedge_i (A_a \widehat{K}_a \Phi_i^a) \rightarrow \bigwedge_a A_a \text{false}$ is derivable too. The antecedent is in S_E because every Φ_i^a is in S_E and $\Phi \rightarrow A_a \widehat{K}_a \Phi$ is an axiom. Thus, $\bigwedge_a A_a \text{false}$ is in S_E , which means that S_E is not consistent since $\bigvee_a E_a \text{true}_a$ is an axiom, which is its negation. We have a contradiction, which means that every hyperedge contains some vertex. \square

Lemma 3.2.10. In the canonical model, $S_* \models_* \xi$ iff $\xi \in S_*$.

Theorem 3.2.11. The logic **2CH** is complete with respect to chromatic hypergraph models.

3.3 Hypergraphs and epistemic frames

Equivalence with partial epistemic frames

Let us recall first the definition of a partial epistemic frame, which has been one of the main models used in the study of epistemic logics such as **KB4** in [GLR22]:

Definition 3.3.1. Given the set of agents \mathcal{A} , a partial epistemic frame \mathcal{M} consists of a set of worlds M together with a family of partial equivalence relations $\{\sim_a\}_{a \in \mathcal{A}}$, such that for every $w \in M$, $w \sim_a w$ for at least one $a \in \mathcal{A}$. A morphism of partial epistemic frames is a function $f : M \rightarrow M'$ such that for every $a \in \mathcal{A}$ and $w, w' \in M$, $w \sim_a w'$ implies $f(w) \sim_a f(w')$.

We can transform a partial epistemic frame into a chromatic hypergraph, and vice versa, using the following construction. Suppose we are given a partial epistemic frame \mathcal{M} . We construct a chromatic hypergraph $\eta(\mathcal{M})$ by setting $E = M$ and $V_a = M/\sim_a$, that is the set of hyperedges is exactly the set of worlds, and the set of vertices of color a is the set of equivalence classes of \sim_a . We then set $p_a(w)$ to be $[w]_a$, that is the equivalence class of w under \sim_a . It is easy to check that this indeed defines a chromatic hypergraph.

Conversely, given a chromatic hypergraph H , we can construct a partial epistemic frame $\kappa(H)$. We set the set of worlds M to be equal to the set of hyperedges of H , and $e \sim_a e'$ if and only if e and e' share an a -colored vertex. This yields a partial equivalence relation.

These maps can be seen as the dual hypergraph construction: if $H = (V, E)$ is a (non-chromatic) hypergraph, then H^* is the hypergraph (E, V) where the hyperedges are the

vertices of H and the vertices are the hyperedges of H . Partial epistemic frames can be seen as hypergraphs that have colored hyperedges which are defined by equivalence classes. The correspondence is exemplified in Figure 3.2.




Figure 3.2: Example of correspondence between chromatic hypergraphs and frames.

In fact, η and κ can be extended to morphisms of partial epistemic frames and chromatic hypergraphs, giving an equivalence of categories. First, we need to define the corresponding morphisms of chromatic hypergraphs:

Definition 3.3.2. A morphism of hypergraphs $f : H \rightarrow H'$ is a family of functions $f_a : V_a \rightarrow V'_a$ for each agent a , together with a function $f_e : E \rightarrow E'$ such that for all $a \in \mathcal{A}$, if $p_a(e) = v$ then $p'_a(f_e(e)) = f_a(v)$. A morphism of hypergraph models $f : H \rightarrow H'$ is a morphism of underlying hypergraphs, such that if $e \in \ell_e(p_e)$, then $f(e) \in \ell_e(p_e)$, as well as if $v \in \ell_a(p_a)$, then $f(v) \in \ell_a(p_a)$.

Theorem 3.3.3. The category of partial epistemic frames is isomorphic to the category of chromatic hypergraphs. In particular, for any chromatic hypergraph H , $\eta(\kappa(H))$ is isomorphic to H , and for any partial epistemic frame \mathcal{M} , $\kappa(\eta(\mathcal{M}))$ is isomorphic to \mathcal{M} .

Proof. We just need to show how η and κ are extended to morphisms. Functoriality is then straightforward, and checking that $\eta(\kappa(H)) \simeq H$ and $\kappa(\eta(\mathcal{M})) \simeq \mathcal{M}$ is also straightforward. Let $f : H \rightarrow H'$ be a morphism of chromatic hypergraphs. Then $\eta(f) : \eta(H) \rightarrow \eta(H')$ just sends a world e to $f_E(e)$. This is indeed a morphism of partial epistemic frames: suppose two worlds e and e' in $\eta(H)$ are \sim_a -equivalent. It means that in H these two hyperedges share a vertex, and thus in H' the two hyperedges $f_E(e)$ and $f_E(e')$ share a vertex, and thus $f_E(e) \sim_a f_E(e')$.

Now let $g : \mathcal{M} \rightarrow \mathcal{M}'$ be a morphism of partial epistemic frames. Then $\kappa(g) : \kappa(\mathcal{M}) \rightarrow \kappa(\mathcal{M}')$ sends a hyperedge w to $g(w)$, thus $\kappa(g)_E$ is defined. We need to show that it induces a map on vertices, and that the condition for morphisms is satisfied. As g preserves \sim_a , it induces a map on equivalence classes, which is exactly $\kappa(g)_a$. Let w be a hyperedge and v be a vertex in $\kappa(\mathcal{M})$, such that $p_a(w) = v$. It means that w belongs

to the equivalence class corresponding to v in \mathcal{M} . Thus, $g(w)$ belongs to the equivalence class corresponding to $g(v)$ in \mathcal{M}' , and thus $p_a(g(w)) = g(v)$. \square

Theorem 3.3.3 shows that chromatic hypergraphs and partial epistemic frames contain exactly the same information. In particular, it means that, in fact, we can define the semantics of **2CH** in partial epistemic frames. It is just a bit inconvenient to do, since we would have to attach agent atomic propositions and interpret agent formulas in the equivalence classes of \sim_a .

Instead, we can still embed partial epistemic models (with only world atomic propositions) into a subclass of chromatic hypergraphs models, such that $\mathbf{Ap}_a = \emptyset$ for every agent. From this, we can extend Theorem 3.3.3 to work at the level of models:

Corollary 3.3.4. The category of partial epistemic models is isomorphic to the category of chromatic hypergraph models with empty sets of atomic propositions for agents.

Translation from KB4 to 2CH

We can use the equivalence of epistemic frames and hypergraphs for showing how the logics **KB4** and **2CH** are related: we will show that **2CH** is a conservative extension of **KB4 + NE**, where axiom **NE** ensures that there is an alive agent in each world (see [GLR22] for further details). From the semantic side, the worlds of epistemic frames are the hyperedges of hypergraphs, thus the formulas of **KB4** are to be translated to the *world* formulas of **2CH**. In particular, when defining the translation, we set the set of world atomic propositions to be the set of atomic propositions of **KB4**. The translation of formulas is defined recursively as follows:

$$\lceil p \rceil := p \quad \lceil \neg \Phi \rceil := \neg \lceil \Phi \rceil \quad \lceil \Phi \wedge \Psi \rceil := \lceil \Phi \rceil \wedge \lceil \Psi \rceil \quad \lceil K_a \Phi \rceil := A_a K_a \lceil \Phi \rceil$$

Proposition 3.3.5. For a partial epistemic frame \mathcal{M} and a formula Φ of **KB4**, $\mathcal{M}, w \models \Phi$ iff $\eta(\mathcal{M}), \eta(w) \models_e \lceil \Phi \rceil$.

Proof. We show the statement by induction on the structure of Φ . For atomic propositions, as well as boolean connectives, the proof is trivial. For the modality, we have: $\mathcal{M}, w \models K_a \Phi$ if and only if for all $w' \in M$ such that $w \sim_a w'$, $\mathcal{M}, w' \models \Phi$. By induction, this is equivalent to for all $w' \in M$ such that $w \sim_a w'$, $\eta(\mathcal{M}), \eta(w') \models \lceil \Phi \rceil$. By definition of η , it is the same as for all hyperedges $e \in \eta(\mathcal{M})_E$ that share an a vertex with $\eta(w)$, $\eta(\mathcal{M}), e \models \lceil \Phi \rceil$. This is equivalent to $\eta(\mathcal{M}), \eta(w) \models \lceil K_a \Phi \rceil$. \square

Corollary 3.3.6. Φ is valid in a partial epistemic frame \mathcal{M} iff $\lceil \Phi \rceil$ is valid in $\eta(\mathcal{M})$.

Using this and Theorem 3.2.11, we can show that **2CH** is a conservative extension of **KB4**:

Theorem 3.3.7. For every **KB4**-formula Φ , $\vdash_{\mathbf{KB4+NE}} \Phi$ if and only if $\vdash_e \ulcorner \Phi \urcorner$.

Proof. By completeness, for **KB4** we have $\vdash_{\mathbf{KB4+NE}} \Phi \Leftrightarrow \models \Phi$. By Corollary 3.3.6, we have $\models \Phi \Leftrightarrow \models_e \ulcorner \Phi \urcorner$. And by completeness for **2CH**, we have $\models_e \ulcorner \Phi \urcorner \Leftrightarrow \vdash_e \ulcorner \Phi \urcorner$. \square

As a corollary, we get that the combination of modalities $A_a K_a$ satisfies axioms **K**, **B**, and **4**.

The crucial part of the proof of Theorem 3.3.3 hides in the way morphisms of chromatic hypergraphs are organized. In fact, in order to define a morphism, it is enough to define a function between the sets of hyperedges, and check whether a condition is verified. This condition is that if a hyperedge e has a vertex of color a , then its image must have a vertex of the same color. In this way, the functions between vertices are automatically induced, since there is at most one vertex of each color in a hyperedge.

For usual Kripke models of modal logic, the morphisms have a preservation property for a particular class of formulas, namely, positive existential formulas (see [BRV01]). We can adapt the lemma to chromatic hypergraphs. We call a formula *positive existential* if it contains only atomic propositions, \wedge , \vee and existential modalities, that is **E** and \widehat{K} . The lemma is then the following:

Lemma 3.3.8. If φ is a positive existential formula, then for any morphism of hypergraph models $f : H \rightarrow H'$, if $H, w \models_* \varphi$, then $H', f(w) \models_* \varphi$, where $*$ is either e or $a \in \mathcal{A}$, and w is either a hyperedge or a vertex, respectively.

Proof. By induction on φ . The base case of atomic propositions and boolean connectives is straightforward. Suppose now that $\varphi = E\psi$. Then $H, e \models_a E\psi$ if and only if there is $v_a \in_a e$, such that $H, v_a \models_a \psi$. By induction, we have that $H', f(v_a) \models_a \psi$. By definition of morphism, we have that $f(v_a) \in_a f(e)$. Thus, $H', f(e) \models_a E\psi$. Similarly for \widehat{K} . \square

In the same way, we can see that formulas that contain only universal modalities and where every atomic proposition is preceded by a negation are *reflected* along the morphisms, that is, if $H', f(w) \models_* \varphi$, then $H, w \models_* \varphi$.

3.4 Hypergraphs, neighborhood frames and Chu spaces

In the definition of chromatic hypergraphs, we required that every hyperedge contains at most one vertex of each color. Thus, in every world, an agent can have either 0 or 1 point

of view. But what happens if we drop this condition and allow agents to have multiple points of view about a given world? Technically, this can be achieved by replacing, the partial function $p_a : E \rightarrow V_a$ by a relation $p_a \subseteq E \times V_a$, i.e., get rid of the functionality requirement.

This leads to a connection with neighborhood frames and Chu spaces. Let us first recall the definitions.

Neighborhood frames

In the context of epistemic logic, neighborhood frames generalize epistemic frames by allowing agents to have multiple points of view (or “frames of mind” [FH87]) on the same world. The general definition of a multi-agent neighborhood frame is as follows.

Definition 3.4.1. A neighborhood frame is a pair $\mathbf{M} = (S, \{N_a\}_{a \in \mathcal{A}})$, where


- S is a set of states;
- for every agent $a \in \mathcal{A}$, N_a is a function that assigns to every state $s \in S$, a set $N_a(s) \subseteq \mathcal{P}(S)$ called the a -neighborhoods of s .

The function N_a associates a set of subsets of S with every world. These subsets are analogous to equivalence classes in relational frames, but now one world can belong to multiple classes at the same time. Thus, we can interpret the sets in $N_a(s)$ as the possible points of view of agent a on state s : in one point of view $U \in N_a(s)$, the agent thinks that the set of possible worlds is U , in another point of view $V \in N_a(s)$, the agent thinks that the set of possible worlds is V , and so on.

Remark 3.4.2. In the case of a single agent, neighborhood frames generalize topological spaces, and do so in two ways. First, the neighborhoods are not required to satisfy any topological axioms, such as being closed under intersection. Second, the neighborhoods are associated with each point of the space, and not the whole space: this way a point does not have to belong to its own neighborhood.

Example 3.4.3. Suppose we have two processes, communicating through shared memory. The memory has two cells, and each cell stores a bit of information: 0 or 1. Processes are given access to memory cells arbitrarily, and both can be assigned the same cell. They know which cell is assigned to them, and they know the value that is stored in this cell, that is, they read the value of the cell. Therefore, a process can have two points of view on the same situation, depending on which cell it is given access to. Assume for

example that the shared memory stores values $(0, 1)$. Process a , when assigned the first cell, knows that the memory stores 0, and when assigned the second cell, knows that the memory stores 1. So the set of possible states is $S = \{(0, 0), (0, 1), (1, 0), (1, 1)\}$. In state $(0, 1)$, the two possible points of view of process a are described by neighborhoods: $N_a((0, 1)) = \{\{(0, 0), (0, 1)\}, \{(0, 1), (1, 1)\}\}$. The four states and the neighborhoods for one agent are depicted in the figure below. The neighborhoods for the other agent are the same.



Neighborhood semantics

Neighborhood frames are models of modal logic. As usual, we fix a set of atomic propositions Ap , and define valuation functions on frames.

Definition 3.4.4. A neighborhood model is a tuple $\mathbf{M} = (S, \{N_a\}_{a \in \mathcal{A}}, \ell)$, where $\mathbf{M} = (S, \{N_a\}_{a \in \mathcal{A}})$ is a neighborhood frame and $\ell : \text{Ap} \rightarrow \mathcal{P}(S)$ is a valuation function.

These models are quite general. For example, axiom **(K)** is valid in every Kripke frame, but this is not the case for neighborhood frames, similarly for the necessitation rule. More formally, the minimal modal logic, which is not normal, is sound and complete with respect to the class of all neighborhood frames. For a more detailed discussion of the correspondence between various logical systems and neighborhood frames, see [Pac17].

In epistemic logic, we do not need the full generality, since we still want to capture some properties of knowledge. The use of neighborhood frames, however, allows us to deal with the problem of logical omniscience: we may have frames that do not validate axiom **(K)**, thus even when an agent knows φ and $\varphi \rightarrow \psi$, she does not necessarily know ψ .

The way it works can be seen from the definition of the semantics.

Definition 3.4.5. Given a neighborhood model $\mathbf{M} = (S, \{N_a\}_{a \in \mathcal{A}}, \ell)$, and a state $s \in S$, the satisfaction relation \models is defined as follows:

- $\mathbf{M}, s \models p$ iff $s \in \ell(p)$;

- $M, s \models \neg\varphi$ iff $M, s \not\models \varphi$;
- $M, s \models \varphi \wedge \psi$ iff $M, s \models \varphi$ and $M, s \models \psi$;
- $M, s \models K_a\varphi$ iff there is $U \in N_a(s)$, for all $s' \in U$, $M, s' \models \varphi$.

The knowledge modality is then interpreted not just as “in world s , agent knows that”, but as “there is a point of view of agent on s such that agent knows that”. From this we can see how axiom **(K)** can be invalidated: agent can have a point of view in which she knows φ , and another point of view in which she knows $\varphi \rightarrow \psi$, but she does not have a point of view in which she knows ψ .

Remark 3.4.6. The semantics above is not the only possible semantics for neighborhood frames. There is another one, which is usually called *strict* when compared with the semantics above. There, for a $K\varphi$ to be true in a state w , it is required that w has a neighborhood which is exactly the set of states satisfying φ . That is, $M, w \models K\varphi$ iff there is $U \in N_w$ such that $U = \{w' \mid w' \models \varphi\}$. See [AF09] for a detailed discussion.

For the rest of the chapter, we consider only neighborhood frames for which the following condition holds: for every $a \in \mathcal{A}$ and for every $s \in S$, if $N_a(s)$ is not empty then $s \in \bigcap N_a(s)$. It is slightly different from another condition which is usually imposed on neighborhood frames in epistemic context: for every $s \in S$, $s \in \bigcap N_a(s)$. The difference is that we add a conditional “if $N_a(s)$ is not empty”. This corresponds to the idea that an agent may be absent from a world, and thus not have any point of view on it. In the standard setting it is usually assumed that every agent is present in every world. Notice that the function $N_a : S \rightarrow \mathcal{P}\mathcal{P}(S)$ induces a relation $\nu_a \subseteq S \times \mathcal{P}(S)$, such that $(s, U) \in \nu_a$ if and only if $U \in N_a(s)$. This relation does not necessarily coincide with the membership relation $\in \subseteq S \times \mathcal{P}(S)$. The condition above can be then reformulated as follows: for every $a \in \mathcal{A}$ and for every $s \in S$, if $(s, U) \in \nu_a$ then $s \in U$.

We additionally ask that every point belongs to a neighborhood of some agent, and that an empty set is never a neighborhood. These conditions correspond to the conditions we imposed on chromatic hypergraphs.

Chu spaces

Another kind of structures we consider is Chu spaces. Similarly to neighborhood frames, Chu spaces can be seen as a generalization of topological spaces. For an introduction, see [Pra99], discussion of categorical Chu construction, see [Bar96], and for applications in concurrency theory, see [Pra03].

We concentrate on the simplest version, adapted to our multi-agent setting. Fix a finite set of agent \mathcal{A} .

Definition 3.4.7. A *chromatic Chu space* is a tuple $M = (X, \{S_a, \in_a\}_{a \in \mathcal{A}})$, where

- X is a set of points, and
- for every agent $a \in \mathcal{A}$, S_a is a set of states of agent a together with a relation $\in_a \subseteq X \times S_a$.

One can see that neighborhood frames are exactly *extensional* Chu spaces, that is, Chu spaces where if two states v and v' contain the same set of points, then $v = v'$.

Remark 3.4.8. It is known ([Bar98]), that the category of extensional Chu spaces is a coreflective subcategory of the category of Chu spaces. We expect the proof to extend to the case of chromatic Chu spaces.

Morphisms of Chu spaces are defined as follows.

Definition 3.4.9. A *chromatic Chu transformation* for a pair of chromatic Chu spaces $M = (X, \{S_a, \in_a\}_{a \in \mathcal{A}})$ and $M' = (X', \{S'_a, \in'_a\}_{a \in \mathcal{A}})$ is a family of functions $f : X \rightarrow X'$ and $g_a : S'_a \rightarrow S_a$ for each agent a , such that for all $a \in \mathcal{A}$, $x \in X$, and $s'_a \in S'_a$, $x \in_a g_a(s'_a)$ iff $f(x) \in'_a s'_a$.

Mildly generalized hypergraphs

We can now relax the condition from the definition of chromatic hypergraphs, and allow hyperedges to contain multiple vertices of the same color. This will lead us to a connection with Chu spaces and neighborhood frames.

Definition 3.4.10. A chromatic hypergraph H is a tuple $(E, \{V_a, \in_a\}_{a \in \mathcal{A}})$, where:

- for all $a \in \mathcal{A}$, V_a is the set of *views* of agent a ;
- E is a set of hyperedges;
- for each agent $a \in \mathcal{A}$, $\in_a \subseteq V_a \times E$ is a total relation. Additionally, we require that for each $e \in E$, for at least one $a \in \mathcal{A}$ there is $v_a \in V_a$ such that $v_a \in_a e$.

Though it is a bit confusing, we call these structures chromatic hypergraphs too. Notice that we reversed the order of the relation: now the relation $\in_a \subseteq V_a \times E$, and not $\rho_a \subseteq E \times V_a$.

The proof of the completeness theorem 3.2.11 steadily adapts to this generalized setting: we just need to omit the axioms $\widehat{K}_a E_a \varphi \rightarrow \varphi$ that corresponded exactly to the at most uniqueness of the vertex of each color in a hyperedge.

Given a chromatic hypergraph H , we can construct a Chu space $\chi(H)$. We set $X = E$, and for each agent $a \in \mathcal{A}$, we set $S_a = V_a$. The relation $\chi(\in_a)$ is the opposite of \in_a , that is $(e, v_a) \in \chi(\in_a)$ iff $(v_a, e) \in \in_a$. Of course, the Chu space we get satisfies some additional conditions: for each $x \in X$, there is $a \in \mathcal{A}$ and $s_a \in S_a$ such that $(x, s_a) \in \chi(\in_a)$, and for each $a \in \mathcal{A}$ and $s_a \in S_a$, there is $x \in X$ such that $(x, s_a) \in \chi(\in_a)$. Let us call Chu spaces that satisfy these conditions *epistemic* Chu spaces.

Similarly, if we have an epistemic Chu space $M = (X, \{S_a, \in_a\}_{a \in \mathcal{A}})$, we can construct a chromatic hypergraph $\eta(M)$, by setting $E = X$, $V_a = S_a$, and $\eta(\in_a) = \in_a^{\text{op}}$. It is easy to see that these constructions are inverse to each other, so the semantics of **2CH** can be equivalently defined in terms of chromatic Chu spaces.

Proposition 3.4.11. For every chromatic hypergraph H , $\chi(\eta(H))$ is isomorphic to H , and for every epistemic Chu space M , $\eta(\chi(M))$ is isomorphic to M .

Let us call a chromatic hypergraph *separated* if every two vertices v_a and v'_a of color a are contained in the same set of hyperedges are equal. If we consider separated chromatic hypergraphs, then the Chu space $\chi(H)$ is extensional, and thus we get a correspondence between separated chromatic hypergraphs and neighborhood frames.

Is this construction functorial? Unfortunately, for morphisms of hypergraphs defined as in Definition 3.3.2, it is not. Notice the direction of the functions $g_a : S'_a \rightarrow S_a$ in the definition of a morphism of chromatic Chu spaces. They go from the target to the source, whereas in the definition of a morphism of chromatic hypergraphs, the functions $f_a : V_a \rightarrow V'_a$ go from the source to the target.

What about neighborhood frames? We may consider morphisms that differ from that of Chu spaces. Moreover, for the neighborhood frames, there is no one standard notion of a morphism. Usually, especially in the context of coalgebras, morphisms are taken to be bisimulations, and even then, there are several kinds of them, see [HKP07].

As morphisms of hypergraphs are not bisimulations (as they do not preserve validity of all formulas), we may try to consider a notion of a morphism which should just resemble the notion of a simulation. For the definition, we just keep the “forth” direction of the definition of a monotonic bisimulation.

Definition 3.4.12. Given two neighborhood models M, M' , a (*functional*) *simulation* is a function $f : S \rightarrow S'$ such that if $f(w) = w'$ then:

- (atoms) for all $p \in \mathbf{Ap}$, if $w \in \ell(p)$, then $w' \in \ell(p)$;
- (forth) if $U \in N_w$, then there is $U' \in N_{w'}$ such that for all $x' \in U'$ there is $x \in U$ such that $f(x) = x'$.

For this choice of morphisms, the construction is not functorial either. One way to see this is to look what kind of formulas are preserved along the morphisms when models are considered. In the case of chromatic hypergraphs, formulas that contain existential modalities are preserved, as we have seen in Lemma 3.3.8. In the case of neighborhood frames, formulas that contain the K modality are preserved (see [GS13]). However, as we will see in a moment, the K modality can be seen as a combination of an existential and a universal modality of **2CH**. Thus, the morphisms of neighborhood frames preserve a different class of formulas of **2CH**, so they are not morphisms of chromatic hypergraphs.

Similarly to the case of relational frames, we can relate the semantics of epistemic logic in neighborhood frames and in chromatic hypergraphs. We define a translation of formulas in the following way:

$$\lceil p \rceil := p \quad \lceil \neg \Phi \rceil := \neg \lceil \Phi \rceil \quad \lceil \Phi \wedge \Psi \rceil := \lceil \Phi \rceil \wedge \lceil \Psi \rceil \quad \lceil K_a \Phi \rceil := E_a K_a \lceil \Phi \rceil$$

The main difference with the translation we had for relational frames, is that now the modality K_a is translated to $E_a K_a$ instead of $A_a K_a$. It captures the semantics in neighborhood frames: for a formula $K_a \Phi$, we require that there is a point of view of agent a such that for all worlds in this point of view, Φ holds. This translation is similar to the one used in [HKP07], but instead of going to first-order logic we stay in the modal case, and also can consider atomic propositions on neighborhoods.

The proof of the following proposition is similar to the proof of Proposition 3.3.5.

Proposition 3.4.13. For a neighborhood frame \mathbf{M} and a formula Φ , $\mathbf{M}, w \models \Phi$ iff $\eta(\mathbf{M}), \eta(w) \models_e \lceil \Phi \rceil$.

3.5 Conclusion and future work

In this chapter, we introduced a new many-sorted modal logic for epistemic reasoning in multi-agent systems. We gave its semantics in terms of chromatic hypergraphs, and showed that it is sound and complete with respect to this semantics. We also showed that this system is a conservative extension of **KB4**, and explored some connections with neighborhood frames and Chu spaces.

The perspective on epistemic logic uncovered by simplicial semantics does not appear to be fully developed on the syntactical side. The main contribution of the chapter is the further development of logical tools for local reasoning of agents through the use of a two-level syntax.

Hypergraph invariants in distributed computing

Every chromatic simplicial complexes can be represented as a chromatic hypergraph. The leftmost hypergraph from Figure 3.1 is an example: for every hyperedge, which in the context of simplicial complexes we would call a simplex, all (non-empty) combinations of its vertices are also hyperedges. This is exactly the definition of a simplicial complex. Moreover, simplicial maps can be translated to morphisms of hypergraphs. Thus, chromatic simplicial complexes form a subcategory of chromatic hypergraphs. However, this subcategory is not full, as the definition of a morphism of chromatic hypergraphs is slightly weaker and allows for more morphisms between hypergraphs. For example, a hyperedge can be sent to a hyperedge that contains more vertices, whereas simplices must be sent to simplices of the same dimension. However, if we consider only morphisms that preserve the size of hyperedges, then the subcategory is reflective, as the inclusion functor has a left adjoint.

Regardless of categorical relation of these two categories, it is interesting to consider chromatic hypergraphs as a model of distributed computing, as they generalize simplicial complexes and allow for a more flexible modeling of systems. There is also a non-trivial theory of algebraic invariants of hypergraphs (see, for example, [Gri+19]), which can be employed in the context of distributed computing similarly to the way it is used for simplicial complexes.

Additional modalities

In this chapter, we considered only the modalities of individual knowledge. In distributed computing, modalities of group knowledge are also important. Hence, it is interesting to consider extensions of the logic with modalities such as common knowledge, everybody knows, and so on. Since groups exist only in worlds, group modalities can be seen as operators on world formulas, that themselves form world formulas. Once we decided how we formalize knowledge of an agent, that is whether it is $E_a K_a$ or $A_a K_a$, we can define the modalities of group knowledge in the same way it is done in [HM90].

For example, “everybody knows” can be defined in the language in the following way:

$$E_{\mathcal{A}}\Phi = \bigwedge_{a \in \mathcal{A}} E_a K_a \Phi.$$

In hypergraphs, this corresponds to looking at the set of hyperedges that are adjacent to a given hyperedge, and verifying that Φ holds in all of them.

Modality “somebody knows” is defined in the same way, but with a disjunction instead of a conjunction:

$$S_{\mathcal{A}}\Phi = \bigvee_{a \in \mathcal{A}} E_a K_a \Phi.$$

Semantically, it corresponds to the fact that given a hyperedge e , Φ holds in at least one adjacent hyperedge.

The modality of common knowledge cannot be defined in the language, as it is the fixed point of the operator $E_{\mathcal{A}}$. Interpretation of the modality in hypergraphs, however, is straightforward: it corresponds to considering a connected component of a hypergraph.

 INCORPORATING TIME

4.1	Epistemic-temporal frames	80
4.2	Dynamic chromatic hypergraphs	88
4.3	Knowledge and concurrency	94
4.4	Conclusion and future work	100

So far in this thesis, we have considered only static models of knowledge. The main point of the study of distributed systems, however, is to understand a behavior of a system over time. One usually wants to show whether certain state of one agent, or of the whole system, is achievable by performing actions or communicating messages. In this context, knowledge that participants have helps to guide the system towards the desired state, or to prove that it is impossible to reach.

At the same time, since agents in the system often act asynchronously, distributed systems are naturally concurrent. Knowledge and concurrency, however, are formally independent: one can have a system where all actions are fully concurrent, but participants have perfect knowledge of the execution, or a system where all actions are sequential, but agents do not know what actions were performed. In some situations, however, knowledge and concurrency are intertwined. For example, the impossibility of consensus in asynchronous systems arises exactly because agents do not know whether other agents have sent their messages or not.

In this chapter, we provide a discussion, together with preliminary results, on various interactions of knowledge, concurrency, and time.

Related work. A prominent approach to the dynamics of knowledge in distributed systems is the framework of interpreted systems [Fag+95]. In this chapter we consider an alternative approach based on histories of actions, akin to the one used in [Pac07]. Another way to describe the interaction of actions and knowledge is dynamic epistemic logic [vDvdHK07]. In the same vein, the dynamics is studied using epistemic programs in [Bal03]. To our knowledge, for the first time an intentional mix of the standard models of concurrency with knowledge was considered in [KMP12].

Plan of the chapter. In Section 4.1, we recall the definitions of existing models that combine knowledge and time, and show how they are related. In particular, we show a connection between epistemic temporal frames and epistemic transition systems. We then extend the framework of chromatic hypergraphs in Section 4.2 to the dynamic case by introducing a new model, and show how to build one from a given epistemic temporal frame. Finally, we discuss how knowledge dynamics can be seen in the true concurrency setting, and how true concurrency models can be enriched with knowledge in Section 4.3.

4.1 Epistemic-temporal frames

There are several ways to combine time and knowledge in a single framework. For a thorough overview of the literature, see [BvdH14]. In this section, we will concentrate on history-based structures, alternately known as epistemic-temporal frames, since they are the most suitable for our purposes. For the moment, we consider its simplest version.

Let Σ be a set of actions. A *history* h is a sequence of actions from Σ . We say that a set of histories H is a *protocol* if it is closed under prefixes, that is, if $h \in H$ and h' is a prefix of h , then $h' \in H$.

Definition 4.1.1. An *epistemic temporal frame* is a tuple $(H, \Sigma, \{\sim_i\}_{i \in \mathcal{A}})$, where H is a protocol with the action set Σ , and for every agent i an equivalence relation \sim_i on H .

The equivalence relations model the knowledge of agents in the system, the dynamics of which is given by the set of histories. For example, if for agent i history ab is indistinguishable from ba , then the agent does not know what was the order of actions. At the same time, an agent may not distinguish between histories that consist of different actions: $a \sim_i b$, means that the agent does not know which one of the actions actually happened.

Remark 4.1.2. This is not the only way to define knowledge in protocols based on histories of actions. Another way, presented, for example, in [Pac07], is to consider alphabets of *observable* actions O_i for every agent, together with maps $\lambda_i : H \rightarrow O_i^+$. A history is mapped into a sequence of actions that an agent can observe, that is, these functions describe the perception of “real” executions by agents. In this case, two histories h, h' are indistinguishable for an agent i if $\lambda_i(h) = \lambda_i(h')$.

Remark 4.1.3. Another well-established approach to the dynamics of knowledge in distributed systems is *interpreted systems*, which is discussed in the introduction. The

(modal) equivalence of history-based structures and interpreted systems is studied in [Pac07].

We now give several examples of epistemic temporal frames. These examples describe some well-known distributed objects (or instructions): *test-and-set*, *compare-and-swap* and *read-write registers*.

Example 4.1.4 (Test-and-set with single bit.). We start with three actions $\{a_1, a_2, a_3\}$, one for each agent. The histories we consider are *one-shot*, that is, every agent performs its action only once. The set of histories H consists of six maximal histories together with all their prefixes. The indistinguishability relation of every agent divides the set of histories H into three equivalence classes: one where the agent does not yet know the outcome, one where he knows that he was first to access the object, and one where he knows that he was not the first. More formally, for agent i , the first class contains all histories that do not contain a_i , the second contains all histories that have a_i on the first place, and the last class has all histories that have a_i on the second or the third place. The frame is depicted in Figure 4.1.




Figure 4.1: ETF for test-and-set

Example 4.1.5 (Compare-and-swap.). In this case, the corresponding frame is similar to the previous one, except that now, when an agent knows it didn't win, it also knows which of two other agents won. That is, for every agent, the equivalence classes are defined by the first action of the history (and coincide for everybody). The frame is depicted in Figure 4.2.




Figure 4.2: ETF for compare-and-swap

Example 4.1.6 (Read-write registers). Suppose now that there are two agents that can write to an assigned register and read from whole memory. That is, the set of actions is $\{w_1, w_2, r_1, r_2\}$. An agent knows whether it wrote something to the register or not, and also, after reading the memory, knows whether the other agent already performed its writing action or not. The frame is depicted in Figure 4.3.




Figure 4.3: ETF for read-write registers

We will now define morphisms of epistemic temporal frames. First, a partial function between sets of actions $\lambda : \Sigma \rightarrow \Sigma'$ extends to histories in the following way:

$$\widehat{\lambda}(he) = \begin{cases} \widehat{\lambda}(h)\lambda(e), & \text{if } \lambda(e) \text{ is defined,} \\ \widehat{\lambda}(h), & \text{otherwise.} \end{cases}$$

Definition 4.1.7. A *morphism* $\lambda : F \rightarrow F'$ of epistemic temporal frames is a partial map between sets of actions $\lambda : \Sigma \rightarrow \Sigma'$, such that:

- for any $h \in F$, $\widehat{\lambda}(h) \in F'$, and
- if $h \sim_a h'$, then $\widehat{\lambda}(h) \sim'_a \widehat{\lambda}(h')$.

Morphisms describe knowledge-aware simulations, or refinements of systems. If there is a morphism from F to F' we can say that F refines F' : if the image of an action in Σ is not defined under λ , then in F' “nothing happens” for this action, so it can be seen as a hidden transition. Additionally, morphisms need to preserve the indistinguishability relations. Epistemic temporal frames and morphisms between them form a category, denoted by **ETF**.

Example 4.1.8. There is a morphism of frames from test-and-set to compare-and-swap, given by $\lambda(a_1) = a_1$, $\lambda(a_2) = a_2$, $\lambda(a_3) = a_3$. In this case, for every agent, if two histories are not distinguishable in the first frame, then they are also not distinguishable in the second frame.

Labeled transition systems

Now, let us take a step back, and forget for a moment about knowledge. In this subsection, we recall some definitions and theorems from the theory of labeled transition systems: structures that are commonly used to model behavior of concurrent systems. The definitions and proofs of theorems can be found in [NSW94] [WN93].

Definition 4.1.9. A labeled transition system (LTS) is a quadruple $(S, \Sigma, \rightarrow, s_0)$, where S is a set of states, Σ is a set of labels, $\rightarrow \subseteq S \times \Sigma \times S$ is a transition relation and $s_0 \in S$ is the initial state.

A morphism of LTSs is a pair (σ, λ) , where $\sigma : S \rightarrow S'$ and $\lambda : \Sigma \rightarrow \Sigma'$, such that:

- $\sigma(s_0) = s'_0$ and
- if $(s, e, s') \in \rightarrow$ and $\lambda(e)$ is defined, then $(\sigma(s), \lambda(e), \sigma(s')) \in \rightarrow'$. Otherwise, if $\lambda(e)$ is undefined, then $\sigma(s) = \sigma(s')$.

The category of LTSs is denoted by **TS**.

Definition 4.1.10. An LTS is called *deterministic* if for every $s \in S$ and $e \in \Sigma$, there is at most one $s' \in S$ such that $(s, e, s') \in \rightarrow$.

Deterministic labeled transition systems form a full subcategory of the category of LTSs. This category is denoted by **dTS**.

Definition 4.1.11. A *synchronization tree* is an LTS $T = (S, \Sigma, \rightarrow, s_0)$ such that

- every state is reachable from the initial state,
- it is acyclic, and
- if $(s', e', s) \in \rightarrow$ and $(s'', e'', s) \in \rightarrow$, then $e' = e''$ and $s' = s''$.

Synchronization trees form a full subcategory of the category of LTSs, and this category is denoted by **Sync**.

Definition 4.1.12. A *language* over set Σ consists of a pair (H, Σ) , where H is a nonempty subset of strings Σ^* closed under prefixes.

Partial function $\lambda : \Sigma \rightarrow \Sigma'$ extends to strings in the following way:


$$\widehat{\lambda}(he) = \begin{cases} \widehat{\lambda}(h)\lambda(e), & \text{if } \lambda(e) \text{ is defined,} \\ \widehat{\lambda}(h), & \text{otherwise.} \end{cases}$$

A *morphism* of languages $\lambda : (H, \Sigma) \rightarrow (H', \Sigma')$ consists of a partial function $\lambda : \Sigma \rightarrow \Sigma'$ such that $\widehat{\lambda}(H) \subseteq H'$, where $\widehat{\lambda}(H) = \{\widehat{\lambda}(h) \mid h \in H\}$.

Theorem 4.1.13 ([NSW94]). There are the following pairs of adjoint functors:

- the inclusion functor **L** \hookrightarrow **dTS** has a right adjoint;
- the inclusion functor **L** \hookrightarrow **Sync** has a left adjoint;
- the inclusion functor **dTS** \hookrightarrow **TS** has a left adjoint;
- the inclusion functor **Sync** \hookrightarrow **TS** has a right adjoint.

The relation between the four categories is summarized in the following diagram.



What is the relation between ETFs and various kinds of transition systems? Here is a simple observation: if we forget the equivalence relations in an ETF, we obtain a language in the sense of Definition 4.1.12. That is, given an ETF $F = (H, \Sigma, \{\sim_i\}_{i \in \mathcal{A}})$, we build a language $L(F) = (H, \Sigma)$ by forgetting the equivalence relations. This construction extends to morphisms of ETFs: by definition, a morphism of ETFs is a morphism of languages if we do not take into account the equivalence relations, so we obtain a functor $L : \mathbf{ETF} \rightarrow \mathbf{L}$.

Going the opposite direction, if we are given a language (H, Σ) , we can build an ETF $E(H, \Sigma) = (H, \Sigma, \{=\}_{i \in \mathcal{A}})$. That is, the set of histories is H , the set of actions is Σ , and for every agent i , the equivalence relation \sim_i is the equality relation. It is also easy to see that E extends to morphisms of languages: as the equality is the minimal equivalence relation, it is automatically preserved by any morphism of languages, thus we have a functor $E : \mathbf{L} \rightarrow \mathbf{ETF}$.

Proposition 4.1.14. The forgetful functor $L : \mathbf{ETF} \rightarrow \mathbf{L}$ is right adjoint to the free functor $E : \mathbf{L} \rightarrow \mathbf{ETF}$. Moreover, E is fully faithful, so languages form a coreflective subcategory of ETFs.

Proof. We show that there is a natural bijection between morphisms $E(H, \Sigma) \rightarrow F$ and morphisms $(H, \Sigma) \rightarrow L(F)$. For every morphism $f : E(H, \Sigma) \rightarrow F$, we associate a morphism $g : (H, \Sigma) \rightarrow L(F)$ defined by $g(e) = f(e)$. It is a morphism of languages by definition. Conversely, given a morphism $g : (H, \Sigma) \rightarrow L(F)$, we define a morphism $f : E(H, \Sigma) \rightarrow F$ also by setting $f(e) = g(e)$. It is straightforward that this is a morphism of ETFs: since the relations on $E(H, \Sigma)$ are equalities, they are automatically preserved by f . It is also easy to see that these two constructions are inverse to each other and natural.

The unit of adjunction $\eta : \text{id}_{\mathbf{L}} \Rightarrow L \circ E$ is a natural isomorphism since every component $\eta_{(H, \Sigma)} : (H, \Sigma) \rightarrow L(E(H, \Sigma))$ is the identity morphism. Thus, E is fully faithful. \square

Remark 4.1.15. The forgetful functor L has a further right adjoint that sends a language to a frame where the indistinguishability relations are *total* relations, that is, every two histories are indistinguishable for all agents. As such, languages also form a reflective subcategory of ETFs. This is similar to the notion of cohesion in topos theory [Law07]: epistemic relations in ETFs are additional structure over languages, akin to topology on a set. The forgetful functor L forgets this structure, and the two adjoints provide free and co-free constructions, like discrete and codiscrete topologies on a set.

It is clear that a language itself carries no information about knowledge of the agents that are participating in the system described by it. Hence, knowledge has to be modeled by additional structure, such as the indistinguishability relations in ETFs. So why should we stop at languages? We now define epistemic variants of structures used in concurrency theory.

Definition 4.1.16. An *epistemic transition system* is a tuple $(S, \Sigma, \rightarrow, s_0, \{\sim_i\}_{i \in \mathcal{A}})$, where $(S, \Sigma, \rightarrow, s_0)$ is an LTS, and \sim_i is an equivalence relation on S for every agent i .

Similarly to how transition systems generalize languages, epistemic transition systems generalize ETFs. We can also define a notion of morphism of epistemic transition systems, and obtain a category \mathbf{TS}_{\sim} . In the same way it is done for transition systems, we define *epistemic trees* and *deterministic epistemic transition systems* as full subcategories of \mathbf{TS}_{\sim} consisting of epistemic transition systems that are trees and deterministic, respectively.

Definition 4.1.17. An epistemic transition system is called an *epistemic tree* if its underlying transition system is a tree. The category of epistemic trees is denoted by \mathbf{Sync}_{\sim} . An epistemic transition system is called *deterministic* if its underlying transition system is deterministic. The category of deterministic epistemic transition systems is denoted by \mathbf{dTS}_{\sim} .

Remark 4.1.18. Epistemic transition systems and epistemic trees were introduced in [KMP12]. There, the main object of study were the trees, as they served as models for a mix of Hennesy-Milner logic and epistemic logic. We will not concentrate on the logical aspects of these systems.

These concurrent epistemic structures allow us to describe knowledge in concurrent systems. Similarly to the non-epistemic case, the choice of the model depends on the application: if one is concerned with how actions affect knowledge of agents, then languages or trees should be used; if the main goal is to understand the states of the system and knowledge of agents about them, then epistemic transition systems will be more useful.

Example 4.1.19. The example of the ETF for the read-write registers can be into an epistemic transition system in the following way. In this case, we consider states of the system, that is possible combinations of local memories of agents and the global memory which is shared with them. In total, there are 13 possible states (including the initial one). This is because certain combinations of actions lead to the same global state: for




Figure 4.4: Epistemic transition system of read-write registers.

example, w_1 and w_2 result in the same state where the global state contains both writes of two agents. The resulting epistemic transition system is depicted in Figure 4.4.

The categories of epistemic structures are related in a way similar to their non-epistemic counterparts.

Theorem 4.1.20. There are the following pairs of adjoint functors:

- the inclusion functor $\mathbf{ETF} \hookrightarrow \mathbf{dTS}_{\sim}$ has a right adjoint;
- the inclusion functor $\mathbf{ETF} \hookrightarrow \mathbf{Sync}_{\sim}$ has a left adjoint;
- the inclusion functor $\mathbf{dTS}_{\sim} \hookrightarrow \mathbf{TS}_{\sim}$ has a left adjoint;
- the inclusion functor $\mathbf{Sync}_{\sim} \hookrightarrow \mathbf{TS}_{\sim}$ has a right adjoint.

Proof. We show just one of the adjunctions as the proofs are similar for the other three, and all four are done essentially the same way as for systems without equivalence relations.

Consider first the inclusion functor $i : \mathbf{Sync}_{\sim} \hookrightarrow \mathbf{TS}_{\sim}$. Its right adjoint is the unraveling functor $un : \mathbf{TS}_{\sim} \rightarrow \mathbf{Sync}_{\sim}$ which is defined as follows. Given an epistemic transition system $T = (S, \Sigma, \rightarrow, s_0, \{\sim_i\}_{i \in \mathcal{A}})$, we construct a synchronization tree $un(T) = (Path(T), \Sigma, un(\rightarrow), s_0, \{\sim_i\}_{i \in \mathcal{A}})$, where $Path(T)$ is the set of paths in T that start in s_0 . The transition relation $un(\rightarrow)$ is such that $(p, e, p') \in un(\rightarrow)$ if there is a path $p = s_0 \rightarrow \dots \rightarrow s$ and transition e with $s \xrightarrow{e} s'$, and $p' = s_0 \rightarrow \dots \rightarrow s \xrightarrow{e} s'$. The equivalence relations \sim_i are inherited from T by $p \sim_i p'$ if $s \sim_i s'$ for the last states of p and p' .

We shall show that there is a natural bijection between hom-sets $\text{Hom}_{\text{TS}_{\sim}}(i(-), -)$ and $\text{Hom}_{\text{Syn}_{\sim}}(-, \text{un}(-))$. Suppose we are given a morphism $f : i(S) \rightarrow T$. Notice that every state s in $i(S)$ defines a unique path in T that goes from the initial state to $f(s)$. By definition of the unraveling, it defines a morphism $g : S \rightarrow \text{un}(T)$ that sends every s from S to the corresponding path in $\text{un}(T)$. It is straightforward to check that g is a morphism of epistemic trees. In the opposite direction, given a morphism $g : S \rightarrow \text{un}(T)$, we define a morphism $f : i(S) \rightarrow T$ by sending every state s to the last state of the path $g(s)$. It is easy (though tedious) to check that these two constructions are inverse to each other and natural. \square

4.2 Dynamic chromatic hypergraphs

In this section, we introduce a new dynamic model that we call *dynamic hypergraphs*. It extends static chromatic hypergraphs by adding local transition systems for agents, similarly to how interpreted systems extend Kripke models with runs. In this framework, however, the main focus is on the evolution of local view of agents, rather than on the dynamics of the global state of the system.

Suppose again that we are given a set of agents \mathcal{A} , and for each agent $i \in \mathcal{A}$, a set of actions Σ_i . The local sets of actions are not necessarily disjoint, that is, agents can have shared actions which serve as synchronization points.

We can now define the dynamic extension of chromatic hypergraphs.

Definition 4.2.1. A *dynamic hypergraph* is a triple $(E, \{V_i, \rightarrow_i\}_{i \in \mathcal{A}})$, where $(E, \{V_i\}_{i \in \mathcal{A}})$ is a chromatic hypergraph and $\rightarrow_i \subseteq V_i \times \Sigma_i \times V_i$ is a transition relation for every agent.

In other words, we have a local transition system for every agent, and hyperedges bind together local states of agents into global configurations. We assume that actions are performed by agents, that is there are transitions between vertices, but no transitions between hyperedges.

The following examples illustrate the definition. In figures, for simplicity, we will denote hyperedges as simple edges of a graph. Transitions are depicted as arrows between vertices.

Example 4.2.2 (Coarse-grained read-write registers). In this case, there are only two actions, one for each agent: $\Sigma_1 = \{a_1\}$, $\Sigma_2 = \{a_2\}$. Action a_i is performed by agent i and stands for immediate snapshot: the agent writes its value to a register, and then reads the value of the whole memory. From the point of view of the agent, the action is

atomic and non-deterministic: after execution it will either know the value of the other agent or not. This corresponds to having one initial vertex for both agents, and two final vertices for both agents: one where the agent knows the value of the other agent, and one where it does not. There are four hyperedges: one that contains the initial vertices and three that contain the final vertices. The transitions are depicted in Figure 4.5.




Figure 4.5: Dynamic hypergraph of coarse-grained specification of read-write registers.

Example 4.2.3 (Fine-grained read-write registers). In this example two agents have two actions: $\Sigma_1 = \{w_1, r_1\}$, $\Sigma_2 = \{w_2, r_2\}$. Action w_i is performed by agent i and stands for writing its value to a register, and r_i stands for reading the memory. In total, there are four points of view for an agent. The first one, the initial, where he has not done anything yet. The second one where he has written his value, that is performed w_i . The last two are the following: one where he has learned the value of the other agent after reading the memory, and one where he has not learned the value of the other agent after reading the memory, both after performing r_i . Again, from the point of view of the agent the execution is non-deterministic: after performing r_i he will either know the value of the other agent or not. The hyperedges and transitions are depicted in Figure 4.6.




Figure 4.6: Fine-grained model of read-write registers.

Notice that here, even after executing r_i , the agent does not know whether the other agent has performed w_j or r_j . Thus, in order to introduce rounds, we need to have a synchronization action. For that, we add a common action s to both local alphabets. The action s is performed after reading the memory, and allows agents to know that other agents completed their round. Corresponding hyperedges and transitions are depicted in Figure 4.7.




Figure 4.7: Fine-grained model of read-write registers with synchronization.

Example 4.2.4 (Test-and-set, compare-and-swap). Suppose there are three agents. For both examples, every agent has one action, a_i . In the case of test-and-set, every agent has three possible states: the initial state s_i^0 , the state s_i^w where he knows that he was the first to access the object, and the state s_i^f where he knows that he was not the first. There is one initial hyperedge that contains s_1^0, s_2^0, s_3^0 , and three final hyperedges $\{s_1^w, s_2^f, s_3^f\}$, $\{s_1^f, s_2^w, s_3^f\}$, and $\{s_1^f, s_2^f, s_3^w\}$. Every agent has two transitions from the initial state to the final states, both labeled by a_i . Notice that in this case, the hyperedges can be seen as triangles that are pairwise connected by vertices. These vertices are exactly the states of agents that do not know who was the first to access the object.

In the case of compare-and-swap, every agent has four possible states: the initial state s_i^0 , the state s_i^w where he knows that he was the first to access the object, and two states $s_i^{j,k}$ where he knows that he was not the first, and that the other agent j or k was the first. The hyperedges and transitions are similar to the case of test-and-set, the only difference is that now triangles are disconnected.

The reader might have guessed that the dynamic hypergraphs are related to epistemic transition systems. Indeed, given an epistemic transition system M , we can construct a dynamic hypergraph $\chi(M)$, similarly to how given an epistemic Kripke structure one can build a chromatic hypergraph. Vertices V_i of $\chi(M)$ of color i are the equivalence classes M/\sim_i . The hyperedges are states of M , and $v_i \in s$ iff the world corresponding to s is in the equivalence class corresponding to v_a . Every local set of action Σ_i is equal to the set of actions of M . We put an arrow $v_i \rightarrow_e v'_i$ between two vertices of color i if in M there is a transition e from a world in the equivalence class v_a to a world in the equivalence class v'_a .

In other words, the local agent transition systems are just quotients of the base LTS by the corresponding equivalence relations.

Definition 4.2.5 (Quotient of LTS). Given an LTS $L = (S, \Sigma, \rightarrow, s_0)$ and an equivalence relation \sim on S , we define the quotient LTS L/\sim as follows. The set of states is the

set of equivalence classes S/\sim , with the initial state being the equivalence class of the initial state of L . The transition relation of the quotient \rightarrow/\sim_i is the following: there is a transition labeled by $e \in \Sigma$ from s to s' if there is $s_1 \in s$ and $s_2 \in s'$ such that $(s_1, e, s_2) \in \rightarrow$.

Basically, for an epistemic LTS $M = (S, \Sigma, \rightarrow, s_0, \{\sim_i\}_{i \in \mathcal{A}})$, in order to construct the corresponding dynamic hypergraph, we are taking a family of quotient LTSs $\{M/\sim_i\}_{i \in \mathcal{A}}$, and then connect vertices together using the hyperedges.

Proposition 4.2.6. For every epistemic transition system $M = (S, \Sigma, \rightarrow, s_0, \{\sim_i\}_{i \in \mathcal{A}})$, $\chi(M) = (S, \{S/\sim_i, \rightarrow/\sim_i\}_{i \in \mathcal{A}})$ is a dynamic hypergraph.

In this construction, all possible actions of the system are in local alphabets of agents. We want, however, actions to be specific to agents, that is, an action which can be performed only by agent i should not be in the alphabet of agent j . For example, in the case of read-write registers, an agent cannot directly observe actions of other agents, so it does not make sense to include them in the local set of actions. We can redefine χ by instead including only the actions that change the epistemic state of agents. Given an epistemic transition system, the local sets of actions in χ_ℓ are then defined as follows: $\Sigma_i = \{e \in \Sigma \mid \exists s \in S, \exists s' \in S, s \rightarrow_e s' \text{ and } s \not\sim_i s'\}$. That is, an action e is local to agent i if it changes its point of view. This way, on the one hand, we do not have transitions that go from a vertex to itself, and on the other, the alphabet contains only actions that are relevant.

Definition 4.2.7. We call a dynamic hypergraph *action-aware* if for every local transition system

- there are no transitions from a vertex to itself, and
- all actions from the local alphabet are actually used.

Epistemic transition systems can also be similarly refined. Instead of considering just a full set of possible transitions Σ , we additionally consider a set of *local* transitions Σ_i for every agent i , such that $\Sigma_i \subseteq \Sigma$ and $\Sigma = \bigcup_{i \in \mathcal{A}} \Sigma_i$. For convenience, we add two conditions that describe the interaction of agents knowledge with their own actions. Informally, the conditions go like this:

1. agents are always aware of their own actions, but
2. agents cannot directly observe actions of other agents.

For the first, the intuition is that an agent always knows its very last action, and can distinguish its state right before the action from the state right after the action. An agent does not necessarily remember the whole history of its actions. For example, suppose there are five states $\{s_0, s_1, s_2, s_3, s_4\}$ and four transitions, all labeled by local actions, $s_0 \rightarrow s_1$, $s_1 \rightarrow s_2$, $s_0 \rightarrow s_3$, $s_3 \rightarrow s_4$. Then, the agent must distinguish s_0 from s_1 , s_1 from s_2 , s_0 from s_3 and s_3 from s_4 , but not necessarily s_2 from s_4 . In other words, this condition is not equivalent to having perfect recall.

For the second condition, the idea is that in order to observe actions of others, an agent needs to perform an action itself. This is a reasonable assumption, as in distributed computing, an agent can only observe the state of the system after it performed some kind of “reading” or “receiving” action. The condition also does not mean that actions of others cannot be observed at all. For example, in the case of read-write registers, an agent can observe whether the other agent has performed a write action or not, but actions with index j do not change the epistemic state of agent i .

Formally, the conditions are:

$$\forall i \in \mathcal{A}, \forall s \in S, \forall e \in \Sigma_i, \forall s' \text{ such that } s \rightarrow_e s' \Rightarrow s \not\sim_i s' \quad (*)$$

$$\forall i \in \mathcal{A}, \forall s \in S, \forall e \in \Sigma \setminus \Sigma_i, \forall s' \text{ such that } s \rightarrow_e s' \Rightarrow s \sim_i s' \quad (**)$$

On the side of dynamic hypergraphs, these conditions translate into the following properties. The condition (*) means that there are no transitions that go from a vertex to itself: every local action changes the epistemic state of an agent. The condition (**) means that only local actions can change the epistemic state, and all effects of external actions are considered as non-determinism from the point of view of an agent.

Example 4.2.8. The dynamic hypergraph of read-write registers in Figure 4.7 is obtained by applying χ_ℓ to the epistemic transition system from Example 4.1.19.

It is clear that one should have an adjunction between the categories of action-aware dynamic hypergraphs and epistemic transition systems that satisfy conditions (*) and (**). We leave this for future work.

The main observation about dynamic hypergraphs is the following: the epistemic transition systems already contain the geometric information usually considered in the topological approach to distributed computing, which is exposed in the corresponding dynamic hypergraph. It can be formalized as follows.

Definition 4.2.9. In a dynamic hypergraph, two views v_i, v'_i of an agent i are said to be *related*, if there is a path from v_i to v'_i in the local transition system of i .

Definition 4.2.10. A (sub)set of views V is said to be *achronal* if for every $v_i, v'_i \in V$, v_i and v'_i are not related. A set of views V is said to be *maximal achronal* if it is a maximal achronal subset of V_i with respect to inclusion. A set of hyperedges is said to be *achronal* if the set of views consisting of its vertices for every color is achronal.

In the examples we considered, the *terminal* achronal sets, that is, such sets of hyperedges that every vertex in it does not have any outgoing transitions, are exactly the protocol complexes of the corresponding protocols. For example, in the case of read-write registers from Figure 4.7, the bottom line of vertices and edges forms the chromatic subdivision of an interval. In the case of “test-and-set”, we have the three connected triangles, and in the case of “compare-and-swap”, we have the three disconnected triangles.

Logic

The obvious question is what kind of logic can be used to reason about dynamic hypergraphs. As a tentative answer, we suggest a temporal extension of the many-sorted epistemic logic from Chapter 3.

Definition 4.2.11. The syntax of the language is given by the following grammar:

$$\begin{aligned} \varphi &::= p_a \mid \neg\varphi \mid \varphi \wedge \psi \mid \widehat{K}_a\Psi \mid \overrightarrow{EG}\varphi \mid \overrightarrow{EF}\varphi \mid \overrightarrow{EX}\varphi \\ \Psi &::= \Psi \vee \Psi \mid \neg\Psi \mid E_a\varphi \end{aligned}$$

In this logic, we have a copy of CTL per agent. It allows to reason about the dynamics from the perspective of an agent, and actions it may perform. As the views are connected by hyperedges, there are modalities that allow to move between the views.

The semantics is given in the following way:

- $H, v^a \models p$ iff $p \in L(v^a)$;
- $H, v^a \models \neg\varphi$ iff $H, v^a \not\models \varphi$;
- $H, v^a \models \varphi \wedge \psi$ iff $H, v^a \models \varphi$ and $H, v^a \models \psi$;
- $H, v^a \models \widehat{K}_a\Psi$ iff there is a hyperedge e such that $v^a \in e$ and $H, e \models \Psi$;
- $H, e \models E_a\varphi$ iff there is a vertex v^a in e such that $H, v^a \models \varphi$;
- $H, v^a \models \overrightarrow{EG}\varphi$ iff there is a path $v_0^a \rightarrow_a v_1^a \rightarrow_a \dots$ such that $H, v_i^a \models \varphi$ for all i ;

- $H, v^a \models \overrightarrow{EF}\varphi$ iff there is a path $v_0^a \rightarrow_a v_1^a \rightarrow_a \dots$ such that $H, v_i^a \models \varphi$ for some i ;
- $H, v^a \models \overrightarrow{EX}\varphi$ iff $H, v_0^a \models \varphi$ for some v_0^a such that $v^a \rightarrow_a v_0^a$.

4.3 Knowledge and concurrency

The models of concurrency discussed in the first section of this chapter are not the only ones that exist in the literature. Another approach to concurrency is provided by *true concurrency* models, such as events structures [Win86] or higher-dimensional automata [Pra91]. These models are based on the idea that actions can happen simultaneously, and not as a non-deterministic choice of order of actions.

In this section we provide a preliminary discussion on how knowledge and concurrency are related in distributed systems, and how geometric models of true concurrency can be redefined in order to capture epistemic information.

Concurrency from distributed knowledge

It is interesting whether temporal epistemic models, such as ETFs, possess any information about concurrency. In interleaving models of concurrency, such as transition systems, a pair of actions a, b is considered to be happening in parallel if two possible orderings of these actions ab and ba lead to the same state.

In an epistemic temporal frame, however, all possible histories are considered to be different states, so we cannot say if two actions may happen simultaneously or not. At the same time, there are indistinguishable histories for agents, and we can use this information to define a notion of concurrency. We observe that sometimes concurrency arises from ignorance: two actions may happen in parallel if nobody knows what was the order of these actions.

Formally, given an ETF $F = (H, \Sigma, \{\sim_i\}_{i \in \mathcal{A}})$, we consider distributed knowledge $\sim = \bigcap_{i \in \mathcal{A}} \sim_i$. The quotient H/\sim is then can be seen as a transition system, where the equivalence classes are states, and the transitions are defined by the transitions in H .

Example 4.3.1. The observation is illustrated by the following example: the underlying transition system of the ETS from Example 4.1.19 is exactly the quotient by distributed knowledge of the ETF from Example 4.1.6.

Higher-dimensional automata

Higher-dimensional automata constitute a geometric model of true concurrency. Contrary to interleaving models, in true concurrency models it is possible to capture actions happening simultaneously, and not as a non-deterministic choice of order of actions. An example of an interleaving model is a labeled transition systems, considered in Section 4.1. Let us recall the definitions.

Definition 4.3.2. A *precubical set* consist of a graded set $X = \cup_{n \leq 0} X_n$ of n -cubes, together with face maps $\delta_{i,n}^{\{-,+\}} : X_n \rightarrow X_{n-1}$, $i \in \{0, \dots, n\}$, such that $\delta_{i,n-1}^\alpha \delta_{j,n}^\beta = \delta_{j-1,n-1}^\beta \delta_{i,n}^\alpha$ for all $i \leq j$ and $\alpha, \beta \in \{-, +\}$.

Intuitively, a precubical set consists of a set of n -cubes, together with face maps that specify how to glue them together, similarly how semi-simplicial sets consist of simplices glued together along faces.

Suppose now that we have a set of actions Σ .

Definition 4.3.3. A *higher-dimensional automaton* is a tuple (X, s_0, λ) where X is a precubical set, $s_0 \in X_0$ is the initial state, and $\lambda : X_1 \rightarrow \Sigma$ is a labeling function, such that for every cube $x \in X_n$, $n > 1$, $\lambda(\delta_{1,n}^-(x)) = \lambda(\delta_{1,n}^+(x))$.

The idea behind the definition is to have a set of states X_0 modelled by 0-cubes in a precubical set, together with transitions that are represented by 1-cubes, that is intervals. The transitions that can happen in parallel are edges of a cube: for example, if action a can happen simultaneously with action b , then they are faces of a square. The condition on the labelling of transitions is such that the opposite edges of a cube are labeled by the same action. In particular, the HDAs that do not have cubes of dimension higher than 2 are exactly the labeled transition systems.

Historically, HDA models are used for static analysis of concurrent systems. In other words, we have a situation when only the source code of the system is given and the set of possible real states is not known. Thus, we deal only with the *stateless* syntactic structure of the system, its flow-control graph. In such models, every vertex corresponds to a point of advancement in the (concurrent) program, and every edge corresponds to an instruction in the code. We give several examples for systems that we considered before, now in the form of HDAs. These examples were first introduced in [Led19] and [GMT18].

Example 4.3.4 (Read-write registers). Suppose we have two processes that share a read-write register. Each process i has just a simple program to execute: $w_i; r_i$, that is,

it writes to the register, and then reads from it. The writing commands can happen in parallel since every process is attributed an isolated register. Similarly, the reading commands can happen in parallel. However, the reading and writing are in conflict and cannot happen simultaneously.

The corresponding HDA is depicted in Figure 4.8.




Figure 4.8: HDA for read-write registers.

Example 4.3.5 (Test-and-set and compare-and-swap). Recall the description of test-and-set and compare-and-swap from Example 4.1.4 and Example 4.1.5. In both cases, there are three processes, each can perform a single action. The first access to the object, in both cases, cannot be done in parallel with other actions. Because of that, the cube on Figure 4.9 is empty inside and the faces that contain the first action in the sequence are empty too. However, it does not matter what is the order of the access of the last two processes, so the squares on the outer side of the cube are filled. As it was noticed in [Led19], this construction provides the same model for both test-and-set and compare-and-swap.




Figure 4.9: HDA for test-and-set *and* compare-and-swap

It is interesting to compare the expressivity of HDAs in the examples above with epistemic models from Section 4.1. On one hand, the epistemic models express only interleaving semantics, and thus cannot capture true concurrency. On the other hand, as we can see in the examples, the HDAs do not have any epistemic structure, thus

they cannot capture any knowledge of agents. Moreover, as these particular examples do not take into account the full state of the system, but only how far in the execution every process is, we cannot introduce any epistemic structure in them: there are just not enough states to distinguish. Thus, as the first step, we need to introduce a *stateful* counterpart.*

The stateful HDA for read-write registers is depicted in Figure 4.10. Notice that this is just the transition system from Example 4.1.19 with two squares filled in.




Figure 4.10: Stateful HDA for read-write registers.

The stateful HDA for test-and-set is depicted in Figure 4.11. Nevertheless, it coincides with the HDA for compare-and-swap.




Figure 4.11: Stateful HDA for test-and-set.

In fact, one of the questions that Ledent asks in his thesis [Led19] is whether it is possible to reconstruct the protocol complex from the HDA models of the objects that we use. This question is directly related to the problem of having the same HDA for test-and-set and compare-and-swap. As one of the reasons why the answer is negative, he suggests that an HDA models does not have an appropriate notion of a *view*. We can fix it using a brute-force approach, simply by adding indistinguishability relations on the set of states for every agent, similarly to epistemic transition systems.

*The stateless/stateful dichotomy is inspired by [MS18].

Definition 4.3.6. An *epistemic HDA* is a tuple $(X, s_0, \lambda, \{\sim_i\}_{i \in \mathcal{A}})$ where (X, s_0, λ) is an HDA and \sim_i is an equivalence relation on X_0 for every agent i . An epistemic HDA is called *agentive* if the every agent is supplied with a set of actions $\Sigma_i \subseteq \Sigma$ and the labeling function λ is such that for every n -cube, $n > 1$, no two elements of $\{\lambda(\delta_{k,n}^-(x)) \mid k \in \{1, \dots, n\}\}$ are in the same Σ_i .

Alternatively, the labelling for agentive HDAs can be defined as a pre-cubical map $\lambda : X \rightarrow C(\Sigma_0, \dots, \Sigma_n)$, where pre-cubical C has one vertex, edges are elements of $\Sigma = \bigcup_{i \in \mathcal{A}} \Sigma_i$, and higher cubes S are such that the opposite edges are always equal, that is $\delta_{j,n}^-(S) = \delta_{j,n}^+(S)$, and any pair $\delta_{j,n}^-(S), \delta_{k,n}^-(S)$ belongs to different Σ_i .

The condition for an agentive epistemic HDA means that there is no autoconcurrency of actions, that is, all actions that happen simultaneously are performed by different agents.

These models allow us to fix the problem of having the same HDA for different primitives: now we just need to specify what states are distinguishable. For the case of compare-and-swap, every agent distinguishes all three final states, whereas for test-and-set, two of them are indistinguishable.

It is also interesting to notice that now knowledge and concurrency may interact in several ways. Suppose there are two agents i, j , and each can do only one action a_i and a_j , respectively. Then there are four different possible situations, depicted in Figure 4.12. In the first, the actions are not concurrent, and lead to two different states, depending on the order. Both agents can distinguish the final states, so they know what was the order of actions after the execution. In the second situation, the actions are not concurrent, but the agents do not distinguish final states (depicted as green and blue lines), so they do not know the order of actions. From their point of view, actions may be concurrent. In the third case, both orders lead to the same global state, so agents have no means of distinguishing what was the real sequence of actions. In the last case, actions are truly concurrent, which leads to the same epistemic situation as in the previous two cases.




Figure 4.12: Interaction of knowledge and true concurrency

A morphism of epistemic HDAs $(k, l) : E \rightarrow E'$ consists of a pre-cubical map

$k : E \rightarrow E'$ of underlying pre-cubical sets, that preserves the initial state and the indistinguishability relations, together with a function between alphabets $l : \Sigma \rightarrow \Sigma'$, such that the labeling is respected by k , that is, for every edge $e \in E_1$, $l(\lambda(e)) = \lambda'(k_1(e))$. We denote the category of epistemic HDAs by \mathbf{HDA}_\sim .

It is known [GM12] that if we restrict morphisms of LTS to *total* functions between alphabets, then there is an adjunction between the category of higher dimensional automata and the category of labelled transition systems. The adjunction is based on the fact that LTSs are just one-dimensional HDAs. Doing the same for epistemic versions, we extend it to the category of epistemic HDAs and the category of epistemic transition systems.

Theorem 4.3.7. There is an adjunction between the categories of epistemic HDAs and epistemic transition systems.

Proof. First we define a functor $i : \mathbf{TS}_\sim \rightarrow \mathbf{EHDA}_\sim$. Given an epistemic transition system $T = (S, \Sigma, \rightarrow, s_0, \{\sim_i\}_{i \in \mathcal{A}})$, we associate a one-dimensional epistemic HDA $i(T)$ with it by setting the set of vertices to be S , the set of edges to be \rightarrow , with the face morphisms $\delta_{1,1}^-((s, e, s')) = s$ and $\delta_{1,1}^+(s, e, s') = s'$, the initial state to be s_0 and the labeling function $\lambda(s, e, s') = e$. The equivalence relations are the same as in T . The definition of the functor extends to morphism in the obvious way.

The functor $j : \mathbf{EHDA}_\sim \rightarrow \mathbf{TS}_\sim$ is defined in the following way. Given an epistemic HDA $X = (X, s_0, \lambda, \{\sim_i\}_{i \in \mathcal{A}})$ with the set of labels Σ , we define an epistemic transition system $j(X)$ by setting the set of states to be X_0 , the set of actions to be Σ , the initial state to be s_0 , the transition relation to be $\rightarrow = \{(s, \lambda(x), s') \mid x \in X_1, \delta_{1,1}^-(x) = s, \delta_{1,1}^+(x) = s'\}$, and the equivalence relations to be \sim_i . Similarly, the functor is defined on morphisms in the obvious way.

The functor j is right adjoint to i . In order to show this, we construct a natural bijection between morphisms $i(T) \rightarrow X$ in \mathbf{EHDA}_\sim and morphisms $T \rightarrow j(X)$ in \mathbf{TS}_\sim . To every morphism $(k, l) : i(T) \rightarrow X$ we associate a morphism $(\sigma, l) : T \rightarrow j(X)$ with it by setting $\sigma(s) = f(s)$. It is clear that this is indeed a morphism of epistemic transition systems: l is a total function between alphabets that respects transitions, the indistinguishability relations are preserved too since they are preserved by k . In the opposite direction, we associate a morphism $(k, l) : i(T) \rightarrow X$ with every morphism $(\sigma, l) : T \rightarrow j(X)$ by setting $k(s) = \sigma(s)$. Similarly, it is a morphism of epistemic HDAs, since σ preserves the indistinguishability relations and respects the labelling. Two constructions are clearly mutually inverse, and naturality can be easily checked. \square

4.4 Conclusion and future work

In this chapter, we suggested several ways of generalizing models of concurrency to capture epistemic information. The most straightforward way is to just supplement languages, transition systems or higher-dimensional automata with epistemic indistinguishability relations on the set of states. We have shown the relation between various models arising in this way. We have also introduced dynamic chromatic hypergraphs. These structures capture the knowledge of agents using the view-based approach, similarly to chromatic simplicial complexes. We demonstrated that they can be naturally constructed from epistemic transition systems.

Rounds, endofunctors and induced models

In many cases, distributed systems are *round-based*, especially in message-passing models. Executions in such systems are divided into rounds inside which agents asynchronously perform actions. Iterated snapshot [BG97] is an example.

One of the way to model this situation from a categorical point of view it by using a notion of an *algebra*. For a category C and an endofunctor $F : C \rightarrow C$ on it, an algebra is an object $c \in C$ together with a morphism $\alpha : F(c) \rightarrow c$. In the topological approach to distributed computing, the main category one works with is that of chromatic simplicial complexes. Chromatic subdivision, which models wait-free computation using read-write registers [HKR14], is, in fact, an endofunctor: it sends a complex to its subdivision and a simplicial map of complexes to an induced simplicial map. Interestingly, for every simplicial complex S , there is a map $\alpha : Sub(S) \rightarrow S$ from its subdivision $Sub(S)$, which sends a simplex in the subdivision to the simplex it was divided from: from computational perspective, this map sends a history to its previous step. Thus, there is an algebra structure induced by chromatic subdivision on every simplicial complex. This construction can be generalized to chromatic hypergraphs.

Notice that every morphism of chromatic hypergraphs induces a dynamic hypergraph. For $f : H \rightarrow H'$, we take the set of vertices of the corresponding dynamic hypergraph $D(f)$ to be the union of sets vertices of H and H' , and the set of hyperedges to be the union of sets of hyperedges of H and H' . There is a transition between a vertex v' from H' to a vertex v from H iff $f(v) = v'$. By construction, every vertex has at most one predecessor. Thus, every algebra $\alpha : F(H) \rightarrow H$ on chromatic hypergraphs induces a dynamic hypergraph. For the case of subdivision, the induced dynamic hypergraph models *one* round of communication.

Similarly, every *endomorphism* of hypergraphs generates a dynamic hypergraph. In this case, for a morphism $f : H \rightarrow H$, $D(f)$ has the same set of vertices and hyperedges as H , and there is a transition from v to v' if $f(v') = v$. This construction is needed for modeling possibly infinite executions. One can consider an infinite union $I(H) = H \sqcup \text{Sub}(H) \sqcup \text{Sub}(\text{Sub}(H)) \sqcup \dots$ of iterated subdivisions. There is a map $H \sqcup \text{Sub}(I(H)) \rightarrow I(H)$ that sends every iterated subdivision to its previous iteration (and the starting hypergraph to itself). It is easy to see that in fact $H \sqcup \text{Sub}(I(H))$ is isomorphic to $I(H)$, so the map induces a dynamic hypergraph. This hypergraph models iterated communication via read-write registers.

One question is whether other models can be represented as endofunctors on the category of chromatic simplicial complexes and chromatic hypergraphs. It is also interesting to see what kind of logical axioms these systems satisfy, especially in the language suggested in Section 4.2. It is reasonable to suppose that the property of perfect recall will appear in some form, since every point of view of an agent comes from a unique predecessor. Possible links with dynamic epistemic logic are also interesting to explore: steps between subdivisions can be seen as epistemic actions, so a relation between two-level temporal epistemic logic and DEL is not too far-fetched, akin to the relation with usual temporal epistemic logic [vBen+09; DLW11].

The standard questions of completeness, decidability and complexity of the two-level temporal epistemic logic are also interesting to explore.

CONCLUSION

This thesis started with the aim to deepen the ties between the topological and epistemic frameworks in distributed computing, first formalized in [GLR18]. The pursuit was driven by the core question: if we can interpret knowledge in simplicial complexes, what is the epistemic meaning of semi-simplicial sets?


The answer to this question is the main focus of Chapter 2, where we have shown that semi-simplicial sets can be used to model non-standard (or implicit) group knowledge. A surprising result was that epistemic logic cannot differentiate between these two classes of models. In parallel, the question of how to accommodate varying number of agents in the simplicial setting arose. We dealt with this by introducing epistemic coverings, which are annotated semi-simplicial sets.

It quickly became clear that epistemic coverings are somewhat clumsy: for example, if we are not interested in wait-free, but just t -resilience, then simplices of dimension lower than t are redundant. We wanted to have a simple combinatorial structure that would not have this redundancy. For this end, we introduced chromatic hypergraphs, the central object of Chapter 3. Another inconvenience of simplicial models of epistemic logic was related to atomic propositions: one always had to deliberately choose whether to annotate local states of agents or global states. Hypergraphs provide an elegant solution: we explicitly separate the local points of view of processes from states of the environment on the level of logic by having a many-sorted language.

Upon examining the logical properties of chromatic hypergraphs, our focus shifted to dynamics. We aimed to extend the simplicial or hypergraph approach with a temporal dimension, akin to how interpreted systems extend static Kripke frames. For this, in Chapter 4, we introduced hypergraphs equipped with temporal structure. This approach is orthogonal to what has been done so far in the simplicial framework: there, one models a step of a protocol as a function between two models; in dynamic hypergraphs one has a transition system per process inside the models. Such structures arise naturally from models that combine knowledge and concurrency, like epistemic transition system. So it was natural to ask whether one can add knowledge to models of true concurrency. Accidentally, such an extension provided a solution to a problem posed by Ledent in his PhD thesis related to the modeling of various synchronization primitives using higher-dimensional automata.

There are several ways to continue this research program, many of which are listed at the end of each chapter. From a broader point of view, the most important direction is to understand the rules that govern the dynamic of knowledge in the simplicial or hypergraph settings. We now have a topological interpretation of local knowledge evolution inside a model. Is there a complete axiomatization of a corresponding temporal-epistemic logic? Can we logically characterize classes of tasks and protocols in this language using only local flow of information? These are questions we plan to address in subsequent work.

Further, is it possible to adapt modern mathematical theories to the study of distributed systems, similar to how simplicial complexes were brought to the field of distributed computing? For example, can we use cobordism theory for characterization of task solvability? Formally, a cobordism consists of a manifold with boundary W , a pair of two closed manifolds M, N together with embeddings $M \hookrightarrow W$ and $N \hookrightarrow W$, such that the boundary of W consists of images of M and N . In the picture below, a cobordism between a circle and a pair of circles is shown.



There is a certain resemblance between this cobordism and the dynamic hypergraph that describes binary consensus problem. It is interesting to see whether we can represent protocols and tasks in a distributed system as cobordisms, suitably adapted, such that their boundaries would constitute initial and final states. Moreover, is it possible to reformulate the solvability condition in these terms? We do not have answers to these questions yet, but we believe that they are worth exploring.

TOWARDS ENRICHED CATEGORIES FOR MODAL LOGIC

A.1 Many-sorted algebras	106
A.2 Enriched categories	112
A.3 Final remarks	119

The leitmotif of the core chapters was the interaction of logics and models for them via completeness theorems. It is natural to ask whether such results can be produced automatically: given a description of a modal logic, possibly multi-modal and many-sorted, can one identify a class of models this logic is complete with respect to? Additionally, as in practice one usually has a class of models in mind, can we canonically find a logic that captures all important features of this class?

In this chapter, we describe an attempt to deal with this problem. We will use categories to describe both languages and corresponding models at the same time. We do so by using a *mode theory*, a category that on the one hand, specifies the sorts and modalities and on the other hand serves as a mold for models. As models, however, we have lattices, instead of set-based frames. Because of this, the results can hardly be considered to be surprising: we simply extend the usual algebraic completeness to the many-sorted case.

We also make use of *enriched categories*, which allow us to accommodate additional axioms. Informally, the enrichment makes it possible to encode an ordering between modalities in the mode theory. An example of such a relation is one of the axioms of the logic of belief $B\varphi \rightarrow \widehat{B}\varphi$, that is, if I believe in φ , then I consider φ possible. In this case, the modality B can be seen as included into \widehat{B} .

Related work. The logic of this chapter consists of modal sequents, instead of formulas, which were first considered in [Dun95] for positive modal logic, in [GNV05] for distributive modal logic, and more recently in [Bez+23] for non-distributive positive modal logic. The approach of this chapter is also remotely related to *modal type theory* [KG22; Shu23]. There, models for modal type theory are functors from a mode theory to the category of categories. Since we are interested in propositional logic, we consider instead functors to the category of distributive lattices.

Plan of the chapter. In Section 1, we introduce mode theories, together with logic and algebraic models induced by them. We show completeness of the logic with respect to many-sorted algebras. In Section 2, we extend the setting by considered enriched mode theories. This allows us to enlarge the kind of axioms we can express. Finally, in Section 3, we discuss perspectives for further research.

A.1 Many-sorted algebras

In this section, we discuss how the standard notion of a category can manifest itself in the study of many-sorted modal logic.

Syntax

Suppose we are given a set of sorts M , a set of atomic propositions Ap_a for every sort $a \in M$, and a (possibly empty) set of modalities $M(a, b)$ that change a sort from a to b for every pair of sorts. Every sort has a special identity modality $\{\text{id}_a\}$, which does not change the sort. We ask for modalities to be composable: if there is a modality $\{f\}$ that changes the sort from a to b , and a modality $\{g\}$ that changes the sort from b to c , then there is a composed modality $\{g \circ f\}$ that changes the sort from a to c . Additionally, composition of modalities is associative: $\{(h \circ g) \circ f\} = \{h \circ (g \circ f)\}$. The identity modality is asked to be neutral with respect to composition: $\{\text{id}_a \circ f\} = \{f\} = \{f \circ \text{id}_a\}$. Finally, there are equalities of modalities, that is, we can declare a pair of (non-empty) sequences of composable modalities $\{f_1\}, \dots, \{f_n\}$ and $\{g_1\}, \dots, \{g_k\}$ to be equal: $\{f_1\} \circ \dots \circ \{f_n\} = \{g_1\} \circ \dots \circ \{g_k\}$.

We call this data a *mode theory* and denote it \mathbf{M} .

Remark A.1.1. It is clear that we just defined a small category whose objects are sorts, and morphisms are modalities, together with a set of atomic propositions associated to each object.

We now define a many-sorted language $\text{Lan} = \{\text{Lan}_a \mid a \in M\}$ that consists of a parametrized family of languages, where Lan_a is the language of formulas of sort a , defined by the following grammar:

$$\begin{aligned} \varphi_a ::= \mathbf{p} \mid \text{true}_a \mid \text{false}_a \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \{f\} \psi_b \\ a, b \in M, f \in M(b, a), \mathbf{p} \in \text{Ap}_a, \psi_b \in \text{Lan}_b \end{aligned}$$

Formulas are defined by mutual recursion over sorts: if a formula φ belongs to $\text{Lan}_a(\mathbf{M})$ and $\{f\}$ is a modality from a to c , then $\{f\}\varphi$ is of sort c and belongs to $\text{Lan}_c(\mathbf{M})$.

Example A.1.2. The language of the logic \mathbf{K} with a single box modality can be seen as generated by a mode theory with one sort $*$, and a unique modality $\mathbf{K}(*, *) = \{\Box\}$.

Example A.1.3. Usually, \mathbf{CTL}^* is defined as having two sorts: state and path formulas, and one says that state formulas are in particular path formulas. In fact, one can say that there is a hidden modality \downarrow that turns every state formula into a path formula. Thus, a (very restricted) fragment of \mathbf{CTL}^* that contains only “for all paths” and “there is a path” modalities A, E , can be presented as a mode theory that has two sorts s, p , and three modalities $A, E \in \mathbf{CTL}^*(p, s)$ and $\downarrow \in \mathbf{CTL}^*(s, p)$.

We define the logic based on a mode theory to be a collection of *modal sequents* $\varphi \vdash_a \psi$, which are just pairs of modal formulas of the same sort.

For every sort $a \in \mathbf{M}$, we have the axioms and rules in Figure A.1. They talk about the logic internally to a sort.

$\varphi \vdash_a \text{true}_a$	$\text{false}_a \vdash_a \varphi$	<i>top and bottom</i>	
$\varphi \vdash_a \varphi$	$\frac{\varphi \vdash_a \psi \quad \psi \vdash_a \xi}{\varphi \vdash_a \xi}$	<i>identity and cut</i>	
$\varphi \wedge \psi \vdash_a \varphi$	$\varphi \wedge \psi \vdash_a \psi$	$\frac{\xi \vdash_a \varphi \quad \xi \vdash_a \psi}{\xi \vdash_a \varphi \wedge \psi}$	<i>conjunction</i>
$\varphi \vdash_a \varphi \vee \psi$	$\psi \vdash_a \varphi \vee \psi$	$\frac{\varphi \vdash_a \xi \quad \psi \vdash_a \xi}{\varphi \vee \psi \vdash_a \xi}$	<i>disjunction</i>
$\varphi \wedge (\psi \vee \xi) \vdash_a (\varphi \wedge \psi) \vee (\varphi \wedge \xi)$		<i>distributivity</i>	

Figure A.1: Internal mode axioms and rules

The rules and axioms that govern the modalities are in Figure A.2. The monotonicity rule says that all modalities are monotone, if φ implied ψ in sort a , then for any modality $\{f\}$ from a to b , $\{f\}\varphi$ implies $\{f\}\psi$. This rule is also known as Becker rule. The axiom for identity modality says that it does not affect the formula (we use $\varphi \dashv\vdash_a \psi$ as a shorthand for $\varphi \vdash_a \psi$ and $\psi \vdash_a \varphi$). The axiom for equivalence of modalities concerns the sequences that are defined to be equal in the mode theory. Finally, the composition axiom says that there is no difference between a successive application of two composable modalities and an application of the composition.

$$\begin{array}{c}
\frac{\varphi \vdash_b \psi}{\{f\}\varphi \vdash_a \{f\}\psi} \quad \textit{monotonicity} \\
\{id_a\}\varphi \dashv\vdash_a \varphi \quad \textit{identity modality} \\
\{f_1\} \dots \{f_n\}\varphi \dashv\vdash_a \{g_1\} \dots \{g_k\}\varphi \quad \textit{equivalence of modalities} \\
\{f \circ g\}\varphi \dashv\vdash_a \{f\}\{g\}\varphi \quad \textit{composition of modalities}
\end{array}$$

Figure A.2: Modal axioms and rules

Definition A.1.4. Let $\text{ML}(\mathbf{M})$ be the smallest parametrized set of modal sequents closed under the axioms and rules from Figure A.1 and Figure A.2.

When a modal sequent $\varphi \vdash_a \psi$ belongs to $\text{ML}(\mathbf{M})$, we say that it is *derivable*. Even in this comparatively poor language, modalities have some properties. For example, because of monotonicity and local rules for conjunction and disjunction, we have the following proposition.

Proposition A.1.5. The following sequents are derivable.

1. $\{f\}(\varphi \wedge \psi) \vdash \{f\}\varphi \wedge \{f\}\psi$;
2. $\{f\}\varphi \vee \{f\}\psi \vdash \{f\}(\varphi \vee \psi)$.

Proof. 1.

$$\frac{\frac{\varphi \wedge \psi \vdash \varphi}{\{f\}(\varphi \wedge \psi) \vdash \{f\}\varphi} \quad \frac{\varphi \wedge \psi \vdash \psi}{\{f\}(\varphi \wedge \psi) \vdash \{f\}\psi}}{\{f\}(\varphi \wedge \psi) \vdash \{f\}\varphi \wedge \{f\}\psi}$$

2.

$$\frac{\frac{\varphi \vdash \varphi \vee \psi}{\{f\}\varphi \vdash \{f\}(\varphi \vee \psi)} \quad \frac{\psi \vdash \varphi \vee \psi}{\{f\}\psi \vdash \{f\}(\varphi \vee \psi)}}{\{f\}\varphi \vee \{f\}\psi \vdash \{f\}(\varphi \vee \psi)}$$

□

Algebraic models

We will define the semantics of the logic in many-sorted distributive lattices. A mode theory can be seen as a category, so the models are functors from it to the category of bounded distributive lattices. A functor \mathcal{M} associates a lattice to every sort. The lattice is then the domain of the interpretation of formulas of this sort. The only unexpected thing is that the morphisms in the target category are not homomorphisms, but rather *monotone functions*, which allows us to interpret arbitrary modalities. We denote the category of bounded distributive lattices and monotone functions by \mathbf{DL}_m .

Definition A.1.6. A many-sorted algebra of a mode theory \mathbf{M} is a functor $\mathcal{M} : \mathbf{M} \rightarrow \mathbf{DL}_m$.

In order to define the semantics of atomic propositions, we will use *natural transformations*.

Definition A.1.7. Given a pair of categories C and D , together with a pair of functors $F, G : C \rightarrow D$, a *natural transformation* $\alpha : F \Rightarrow G$ is a family of morphisms $\alpha_c : F(c) \rightarrow G(c)$, such that for every morphism $f : c \rightarrow c'$ in C , the following diagram commutes.

$$\begin{array}{ccc} F(c) & \xrightarrow{F(f)} & F(c') \\ \alpha_c \downarrow & & \downarrow \alpha_{c'} \\ G(c) & \xrightarrow{G(f)} & G(c') \end{array}$$

Now consider a discrete category \mathbf{AP} , that is a category that has only identity morphisms, objects of which are pairs (p, a) , where a is a sort and $p \in \mathbf{Ap}_a$. There is a projection functor $\pi : \mathbf{AP} \rightarrow \mathbf{M}$, which maps (p, a) to a . We also have a functor $\mathbf{1} : \mathbf{AP} \rightarrow \mathbf{DL}_m$, which maps (p, a) to the trivial distributive lattice $\{*\}$. We then define an assignment ℓ to be a natural transformation $\ell : \mathbf{1} \Rightarrow \mathcal{M} \circ \pi$.

$$\begin{array}{ccc} \mathbf{AP} & & \\ \pi \downarrow & \searrow \mathbf{1} & \\ \mathbf{M} & \xrightarrow{\mathcal{M}} & \mathbf{DL}_m \end{array}$$

$\swarrow \ell$

Explicitly, an assignment ℓ is a family of morphisms $\ell_{(p,a)} : \{*\} \rightarrow \mathcal{M}(a)$. Since \mathbf{AP} has only identity morphisms, the naturality condition is trivially satisfied.

Definition A.1.8. A *model* is a many-sorted algebra \mathcal{M} equipped with an assignment $\ell : 1 \Rightarrow \mathbf{M} \circ \pi$.

Morphisms of models are natural transformations of functors, such that the components are homomorphisms of distributive lattices.

Definition A.1.9. A *morphism of many-sorted algebras* $F : \mathcal{M} \rightarrow \mathcal{N}$ is a natural transformation $F : \mathcal{M} \Rightarrow \mathcal{N}$, such that for every $a \in \mathbf{M}$, the component $F_a : \mathcal{M}(a) \rightarrow \mathcal{N}(a)$ is a homomorphism of distributive lattices. A *morphism of models* $F : \mathcal{M} \rightarrow \mathcal{N}$ is a morphism of underlying algebras that preserves the interpretation of atomic propositions, that is, for every $a \in \mathbf{M}$ and $\mathfrak{p} \in \text{Ap}_a$, $F_a(\ell_{\mathfrak{p},a}^{\mathcal{M}}(*)) = \ell_{\mathfrak{p},a}^{\mathcal{N}}(*)$.

An assignment ℓ induces an interpretation of all formulas $\varphi \in \text{Lan}(\mathbf{M})$ in a model \mathcal{M} . The interpretation $[-]_{\ell}^{\mathcal{M}}$, which assigns an element of the algebra to every formula, is defined inductively as follows.

- $[\mathfrak{p}]_{\ell}^{\mathcal{M}} = \ell_{\mathfrak{p},a}(*)$ for every atomic proposition $\mathfrak{p} \in \text{Ap}_a$;
- $[\text{true}_a]_{\ell}^{\mathcal{M}} = \top_{\mathcal{M}(a)}$;
- $[\text{false}_a]_{\ell}^{\mathcal{M}} = \perp_{\mathcal{M}(a)}$;
- $[\varphi \wedge \psi]_{\ell}^{\mathcal{M}} = [\varphi]_{\ell}^{\mathcal{M}} \wedge [\psi]_{\ell}^{\mathcal{M}}$;
- $[\varphi \vee \psi]_{\ell}^{\mathcal{M}} = [\varphi]_{\ell}^{\mathcal{M}} \vee [\psi]_{\ell}^{\mathcal{M}}$;
- $[\{f\}\varphi]_{\ell}^{\mathcal{M}} = \mathcal{M}(f)([\varphi]_{\ell}^{\mathcal{M}})$.

We then say that a modal sequent $\varphi \vdash \psi$ is valid in an algebra \mathcal{M} , if for all assignments ℓ on \mathcal{M} , $[\varphi]_{\ell}^{\mathcal{M}} \leq [\psi]_{\ell}^{\mathcal{M}}$.

It is easy to show that the logic is sound and complete with respect to these models. Recall that we write $\varphi \dashv\vdash_a \psi$ if $\varphi \vdash_a \psi$ and $\psi \vdash_a \varphi$. Notice that $\dashv\vdash_a$ is an equivalence relation on $\text{Lan}_a(\mathbf{M})$: it is symmetric by definitions, reflexive by the identity rule, and transitive by the cut rule. We write $\mathcal{L}_a(\mathbf{M})$ for the set of equivalence classes of $\dashv\vdash_a$, and $[\varphi]$ for the equivalence class of φ . Every $\mathcal{L}_a(\mathbf{M})$ is a distributive lattice by the axioms and rules on Figure A.1. In particular, $[\varphi] \wedge_{\mathcal{L}_a(\mathbf{M})} [\psi] = [\varphi \wedge \psi]$ and $[\varphi] \vee_{\mathcal{L}_a(\mathbf{M})} [\psi] = [\varphi \vee \psi]$, $\top_{\mathcal{L}_a(\mathbf{M})} = [\text{true}_a]$ and $\perp_{\mathcal{L}_a(\mathbf{M})} = [\text{false}_a]$. Moreover, for every modality $\{f\}$, a function $\mathcal{L}(\mathbf{M})(f) : \mathcal{L}(\mathbf{M})(a) \rightarrow \mathcal{L}(\mathbf{M})(b)$ that sends $[\varphi]$ to $[\{f\}\varphi]$ is monotone by the monotonicity rule. Finally, there is an assignment $\ell_{\mathcal{L}}$ that sends every atomic proposition $\mathfrak{p} \in \text{Ap}_a$ to $[\mathfrak{p}]$. Altogether, it constitutes the Lindenbaum algebra of the logic.

Lemma A.1.10. For every pair of formulas $\varphi, \psi \in \text{Lan}_a$, $\varphi \vdash_a \psi$ if and only if $[\varphi] \leq [\psi]$ in $\mathcal{L}_a(\mathbf{M})$.

Proof. Let us first notice that $\varphi \vdash_a \psi$ if and only if $\varphi \wedge \psi \dashv\vdash_a \varphi$: for right-to-left direction just use the cut rule and one of conjunction axioms, for left-to-right direction the cut rule and the conjunction rule, together with the identity axiom. By construction, this is equivalent to $[\varphi \wedge \psi] = [\varphi]$. By the definition of the order in lattices, this is equivalent to $[\varphi] \leq [\psi]$. \square

Lemma A.1.11. A sequent $\varphi \vdash \psi$ is valid in every algebra if and only if $[\varphi] \leq [\psi]$ in the Lindenbaum algebra.

Proof. The proof amounts to showing that the Lindenbaum algebra is the initial object in the category of models, that is, for any model there is a unique morphism of models from the Lindenbaum algebra to it.

Suppose we are given an algebra \mathcal{M} with an assignment ℓ . Then there is a morphism of models $e : \mathcal{L}(\mathbf{M}) \rightarrow \mathcal{M}$, that for every atomic proposition $\mathfrak{p} \in \text{Ap}_a$ sends $[\mathfrak{p}]$ to $\ell_{\mathfrak{p},a}(\ast)$, $[\varphi \wedge \psi]$ to $[\varphi]_{\ell}^{\mathcal{M}} \wedge [\psi]_{\ell}^{\mathcal{M}}$, and so on. It is easy to see that this is indeed a morphism of models. Moreover, it is unique: suppose there is another morphism of models $e' : \mathcal{L}(\mathbf{M}) \rightarrow \mathcal{M}$. Then for every atomic proposition $\mathfrak{p} \in \text{Ap}_a$, $e'([\mathfrak{p}]) = \ell_{\mathfrak{p},a}(\ast) = e([\mathfrak{p}])$. Then, by induction on formulas, $e'([\varphi]) = e([\varphi])$ for every formula φ . Also, $e([\varphi]) = [\varphi]_{\ell}^{\mathcal{M}}$.

By monotonicity of lattice homomorphisms, we have that $[\varphi] \leq [\psi]$ implies $e([\varphi]) \leq e([\psi])$, thus $[\varphi]_{\ell}^{\mathcal{M}} \leq [\psi]_{\ell}^{\mathcal{M}}$. This is the “if” direction of the statement.

As the Lindenbaum algebra is a model itself, we have that the validity of a sequent $\varphi \vdash \psi$ in it is equivalent to $[\varphi] \leq [\psi]$. This concludes the proof. \square

Remark A.1.12. The proof is essentially a restatement of the fact that the Lindenbaum algebra represents the syntax in the category of models, and that every assignment (or valuation) is the unique morphism from it to the model.

Combined, these two lemmas give us the following theorem.

Theorem A.1.13 (Soundness and completeness). A sequent $\varphi \vdash \psi$ is derivable if and only if it is valid in every algebra.

A.2 Enriched categories

In this section, we accommodate more axioms in the setting by allowing axioms of type $\{f_1\} \dots \{f_n\} \varphi \vdash_a \{g_1\} \dots \{g_k\} \varphi$, instead of equivalences. On the side of categories, this corresponds to the notion of a *poset-enriched* category, where instead of a set of morphisms between two objects, we have a poset.

Syntax

As in the previous section, suppose we are given a set of sorts M and a set of atomic propositions Ap_a for every sort, the set of modalities for every pair of sorts, identity modalities and the composition. However, now, instead of an equality of modalities, we allow *inequalities*: we can say that a modality $\{f\}$ is smaller than $\{g\}$, and write $\{f\} \leq \{g\}$. We ask the set of modalities that turn formulas of sort a to formulas of sort b to be a *poset*. That is, we have the following three conditions: (i) $\{f\} \leq \{f\}$; (ii) $\{f\} \leq \{g\}$ and $\{g\} \leq \{h\}$ implies $\{f\} \leq \{h\}$; (iii) $\{f\} \leq \{g\}$ and $\{g\} \leq \{f\}$ implies $\{f\} = \{g\}$. We also call this data a mode theory, but in order to distinguish it from the previous case, we will denote it by \mathbb{M} .

Definition A.2.1. Suppose we are given a mode theory \mathbb{M} . For every $a \in \mathbb{M}$, let $\text{Lan}_a(\mathbb{M})$ denote the language generated by the following grammar:

$$\begin{aligned} \varphi_a ::= & \mathbf{p} \mid \text{true}_a \mid \text{false}_a \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \{f\} \psi_b \\ & a, b \in \mathbb{M}, \{f\} \in \mathbb{M}(b, a), \mathbf{p} \in \text{Ap}_a, \psi_b \in \text{Lan}_b(\mathbb{M}) \end{aligned}$$

That is, the language is the same as in the previous section. The only difference lies in the axioms allowed.

Definition A.2.2. Let $\text{ML}(\mathbb{M})$ be the smallest parametrized set of modal sequents closed under the axioms and rules from Figure A.1 and Figure A.2, where the equivalence of modalities is replaced with the axiom $\{f\} \varphi \vdash_a \{g\} \varphi$ for every pair of modalities $\{f\}, \{g\}$ such that $\{f\} \leq \{g\}$.

Adjoint modalities

Using the inequality axiom one can express properties of modalities that are often encountered in practice.

Definition A.2.3. A pair of modalities $\{f\}, \{g\}$ is called *adjoint* if and only if $\{\text{id}_a\} \leq \{g \circ f\}$ and $\{f \circ g\} \leq \{\text{id}_b\}$. In this case, we say that $\{f\}$ is left adjoint to $\{g\}$, and $\{g\}$ is right adjoint to $\{f\}$, and write $\{f\} \dashv \{g\}$.

Example A.2.4. In the language of normal modal logic \mathbf{K} one can add a dual modality \diamond . This modality is not adjoint to \square , however, the language can be extended with *backward-modalities* \blacklozenge and \blacksquare , such that \blacklozenge is left adjoint to \square and \diamond is left adjoint to \blacksquare .

Example A.2.5. In the restricted CTL^* from Example A.1.3, the modalities are pairwise adjoint: $\mathbf{E} \dashv \downarrow$ and $\downarrow \dashv \mathbf{A}$.

Suppose that $\{f\} \dashv \{g\}$ in \mathbb{M} . Then, in the corresponding logic we have the following pair of axioms:

- $\varphi \vdash_a \{g\}\{f\}\varphi$;
- $\{f\}\{g\}\varphi \vdash_b \varphi$.

Proposition A.2.6. If $f \dashv g$, then the following rule is admissible.

$$\frac{\frac{\{f\}\varphi \vdash \psi}{\varphi \vdash \{g\}\psi}}{\varphi \vdash \{g\}\psi}$$

Proof.

$$\frac{\varphi \vdash \{g\}\{f\}\varphi \quad \frac{\{f\}\varphi \vdash \psi}{\{g\}\{f\}\varphi \vdash \{g\}\psi}}{\varphi \vdash \{g\}\psi}$$

Similarly for the other direction. □

Remark A.2.7. Looking at the proof, we can notice that if instead of having both inequalities from the definition of adjointness, we had only one of them, then the rule from the statement would go only in one direction.

The following proposition show that adjoint modalities behave like box and diamond modality.

Proposition A.2.8. If $\{f\} \dashv \{g\}$, then the following sequents are derivable.

1. $\{g\}\varphi \wedge \{g\}\psi \vdash \{g\}(\varphi \wedge \psi)$;

2. $\{f\}(\varphi \vee \psi) \vdash \{f\}\varphi \vee \{f\}\psi$.

Proof. 1. We subdivide the proof into several steps.

a)

$$\frac{\{f\}\{g\}\varphi \wedge \{f\}\{g\}\psi \vdash \{f\}\{g\}\psi \quad \{f\}\{g\}\psi \vdash \psi}{\{f\}\{g\}\varphi \wedge \{f\}\{g\}\psi \vdash \psi}$$

b)

$$\frac{\text{Proposition A.1.5} \quad \text{a)}}{\frac{\{f\}(\{g\}\varphi \wedge \{g\}\psi) \vdash \{f\}\{g\}\varphi \wedge \{f\}\{g\}\psi \quad \{f\}\{g\}\varphi \wedge \{f\}\{g\}\psi \vdash \psi}{\{f\}(\{g\}\varphi \wedge \{g\}\psi) \vdash \psi}}$$

c)

$$\frac{\text{as in b)} \quad \text{b)}}{\frac{\{f\}(\{g\}\varphi \wedge \{g\}\psi) \vdash \varphi \quad \{f\}(\{g\}\varphi \wedge \{g\}\psi) \vdash \psi}{\{f\}(\{g\}\varphi \wedge \{g\}\psi) \vdash \varphi \wedge \psi}}{\{g\}\varphi \wedge \{g\}\psi \vdash \{g\}(\varphi \wedge \psi)}$$

2. Similar. □

Proposition A.2.9. If $f \dashv g$, then the following sequents are derivable.

1. $\text{true} \vdash \{g\}\text{true}$;

2. $\{f\}\text{false} \vdash \text{false}$.

Proof. 1.

$$\frac{\{f\}\text{true} \vdash \text{true}}{\text{true} \vdash \{g\}\text{true}}$$

2. Similar. □

The sequents from Proposition A.2.8 and Proposition A.2.9 can be understood as the condition for the right adjoint modality to be *normal*, that is $\{g\}(\varphi \wedge \psi) \dashv\vdash \{g\}\varphi \wedge \{g\}\psi$ and $\{g\}\text{true} \dashv\vdash \text{true}$. In this case, $\{g\}$ behaves like a box modality, so we can write $[g]$. Similarly, the left adjoint modality can be thought of as a diamond-like modality, and we write $\langle f \rangle$.

Algebraic models

Similarly to how in the previous section a mode theory was a category, in this section, a mode theory with inequalities of modalities can be represented as a poset-enriched category. We recall the basics of the theory of enriched categories instantiated for the case of posets. For further details, see [Kel82] and [Rie14].

Definition A.2.10. A mode theory \mathbb{M} is a poset-enriched category. Explicitly, it consists of the following data:

- a set of objects $\text{Ob}(\mathbb{M})$;
- for every pair of objects $a, b \in \text{Ob}(\mathbb{M})$, a *poset* of morphisms $\text{Hom}(a, b)$;
- together with monotone functions between posets
 - composition $\circ : \text{Hom}(b, c) \times \text{Hom}(a, b) \rightarrow \text{Hom}(a, c)$ and
 - identity $\text{id}_a : \{*\} \rightarrow \text{Hom}(a, a)$

for all $a, b, c \in \text{Ob}(\mathbb{M})$, satisfying associativity and unitality axioms.

The category of bounded distributive lattices and monotone maps is naturally enriched over posets: for $A, B \in \mathbf{DL}_m$, and morphisms $f, g : A \rightarrow B$, $f \leq g$ iff $f(a) \leq g(a)$ for every $a \in A$. In order to distinguish the enriched version of the category, we will write \mathbb{DL}_m . Thus, the definitions of algebras, assignments, and models readily extend to the enriched setting.

We first give the definition of a poset-enriched functor.

Definition A.2.11. A *poset-enriched* functor $F : \mathbb{M} \rightarrow \mathbb{N}$ consists of a map $F : \text{Ob}(\mathbb{M}) \rightarrow \text{Ob}(\mathbb{N})$ and a family of order-preserving functions for every pair of objects $a, b \in \text{Ob}(\mathbb{M})$ $F_{a,b} : \text{Hom}_{\mathbb{M}}(a, b) \rightarrow \text{Hom}_{\mathbb{N}}(F(a), F(b))$, such that the following diagrams commute.

$$\begin{array}{ccc}
 \text{Hom}_{\mathbb{M}}(b, c) \times \text{Hom}_{\mathbb{M}}(a, b) & \xrightarrow{\circ} & \text{Hom}_{\mathbb{M}}(a, c) & & \{*\} & \xrightarrow{\text{id}_a} & \text{Hom}_{\mathbb{M}}(a, a) \\
 \downarrow F_{b,c} \times F_{a,b} & & \downarrow F_{a,c} & & \searrow \text{id}_{F(a)} & & \downarrow F_{a,a} \\
 \text{Hom}_{\mathbb{N}}(F(b), F(c)) \times \text{Hom}_{\mathbb{N}}(F(a), F(b)) & \xrightarrow{\circ} & \text{Hom}_{\mathbb{N}}(F(a), F(c)) & & & & \text{Hom}_{\mathbb{N}}(F(a), F(a))
 \end{array}$$

Basically, enriched functors should preserve the structure of the posets of morphisms, and respect composition and identity morphisms.

Definition A.2.12. An algebra of a mode theory \mathbb{M} is a poset-enriched functor $\mathcal{M} : \mathbb{M} \rightarrow \mathbb{DL}_m$.

Basically, a multi-sorted algebra is a diagram of shape \mathbb{M} in \mathbb{DL}_m . The enrichment allows us to interpret inequalities of modalities as inequalities on Hom-posets. For example, if we have $f \leq g$ in \mathbb{M} , then in the model $\mathcal{M}(f) \leq \mathcal{M}(g)$ as monotone functions between lattices.

The notion of a natural transformation between functors can be adapted for the enriched setting.

Definition A.2.13. A poset-enriched natural transformation $\alpha : F \Rightarrow G$ between a pair of poset-enriched functors $F, G : \mathbb{M} \rightarrow \mathbb{N}$ is a family of functions $\alpha_c : \{*\} \rightarrow \mathbb{N}(F(c), G(c))$ in the category of partial orders for every object $c \in \mathbb{M}$, such that for any pair of objects $c, c' \in \mathbb{M}$, the following diagram commutes in the category of partial orders:

$$\begin{array}{ccc}
 \mathbb{M}(c, c') \times \{*\} & \xrightarrow{G_{c,c'} \times \alpha_c} & \mathbb{N}(G(c), G(c')) \times \mathbb{M}(F(c), G(c)) \\
 \uparrow & & \downarrow \circ \\
 \mathbb{M}(c, c') & & \mathbb{M}(F(c), G(c')) \\
 \downarrow & & \uparrow \circ \\
 \{*\} \times \mathbb{M}(c, c') & \xrightarrow{\alpha_{c'} \times F_{c,c'}} & \mathbb{N}(F(c'), G(c')) \times \mathbb{M}(F(c), F(c))
 \end{array}$$

Similarly to the non-enriched case, there is a discrete category \mathbb{AP} , with objects being pairs (p, a) , where a is a sort and $p \in \mathbf{Ap}_a$; there is a projection functor $\pi : \mathbb{AP} \rightarrow \mathbf{M}$, which maps (p, a) to a , and a functor $\mathbf{1} : \mathbb{AP} \rightarrow \mathbb{DL}_m$, which maps (p, a) to the trivial distributive lattice $\{*\}$. We then define an assignment ℓ to be an enriched natural transformation $\ell : \mathbf{1} \Rightarrow \mathcal{M} \circ \pi$.

$$\begin{array}{ccc}
 \mathbb{AP} & & \\
 \pi \downarrow & \searrow \mathbf{1} & \\
 \mathbf{M} & \xrightarrow{\mathcal{M}} & \mathbb{DL}_m
 \end{array}$$

ℓ (indicated by a double arrow from $\mathbf{1}$ to \mathcal{M})

Definition A.2.14. A model is a multi-sorted algebra \mathcal{M} equipped with a valuation $\ell : \mathbf{1} \Rightarrow \mathcal{M} \circ \pi$.

The completeness theorem Theorem A.1.13 readily extends to the setting of poset-enriched models and mode theories with inequalities. The definition of the Lindenbaum

algebra stays the same, one just should notice that now it is an enriched functor, because of the axiom induced by inequality of modalities. The proof of completeness is then essentially the same, with relevant modifications for enrichment.

Theorem A.2.15. The logic $\text{ML}(\mathbb{M})$ is sound and complete with respect to models $\mathcal{M} : \mathbb{M} \rightarrow \text{DL}_m$.

Negation

So far we have considered only the positive fragment of modal logic, that is a logic without negation. Since negation is a unary operator on formulas, we can consider it as a modality. Seen as such, negation is not monotone: the usual property of negation is that it reverses the entailment, that is, if $\varphi \vdash \psi$, then $\neg\psi \vdash \neg\varphi$. Because of this, we need to consider a different kind of mode theories with both monotone and antitone modalities.

Definition A.2.16. A posetal category \mathbb{M} with contravariance consists of

- a set of objects $\text{Ob}(\mathbb{M})$;
- for every $a, b \in \mathbb{M}$, a pair of posets $\text{Hom}^+(a, b)$ and $\text{Hom}^-(a, b)$;
- for each $a \in \mathbb{M}$, a distinguished element $\text{id}_a \in \text{Hom}^+(a, a)$;
- for each $a, b, c \in \mathbb{M}$, monotone composition functions

$$\text{Hom}^+(b, c) \times \text{Hom}^+(a, b) \rightarrow \text{Hom}^+(a, c)$$

$$\text{Hom}^-(b, c) \times \text{Hom}^-(a, b)^{\text{op}} \rightarrow \text{Hom}^+(a, c)$$

$$\text{Hom}^+(b, c) \times \text{Hom}^-(a, b) \rightarrow \text{Hom}^-(a, c)$$

$$\text{Hom}^-(b, c) \times \text{Hom}^+(a, b)^{\text{op}} \rightarrow \text{Hom}^-(a, c)$$

such that unitality and associativity diagrams commute.

Example A.2.17. Posets with $\text{Hom}^+(a, b)$ being the poset of order-preserving functions between a and b , and $\text{Hom}^-(a, b)$ being the poset of order-reversing functions between a and b is an example of locally posetal category with contravariance. Composition of two monotone functions is monotone, as well as composition of two antitone functions. On the other hand, composition of a monotone and an antitone function is antitone, in any order.

Remark A.2.18. One can change the setting from poset-enriched to category-enriched categories. The definition is essentially the same. This situation is studied and generalized in [Shu18].

Definition A.2.19. A mode theory with contravariance induces the following syntax:

$$\varphi_a ::= \text{true}_a \mid \text{false}_a \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \{f\}(\psi_b) \mid \overline{\{g\}}(\psi_b) \mid$$

$$a, b \in \mathbb{M}, \{f\} \in \text{Hom}^+(b, a), \overline{\{g\}} \in \text{Hom}^-(b, a)$$

We add a rule for antitone modalities:

$$\frac{\varphi \vdash_a \psi}{\overline{\{g\}}\psi \vdash_a \overline{\{g\}}\varphi} \text{ANTI}$$

Proposition A.2.20. The antitone modality has the following properties:

1. $\overline{\{g\}}\varphi \vee \overline{\{g\}}\psi \vdash \overline{\{g\}}(\varphi \wedge \psi)$;
2. $\overline{\{g\}}(\varphi \vee \psi) \vdash \overline{\{g\}}\varphi \wedge \overline{\{g\}}\psi$.

Proof. 1.

$$\frac{\frac{\varphi \wedge \psi \vdash \varphi}{\overline{\{g\}}\varphi \vdash \overline{\{g\}}(\varphi \wedge \psi)} \quad \frac{\varphi \wedge \psi \vdash \psi}{\overline{\{g\}}\psi \vdash \overline{\{g\}}(\varphi \wedge \psi)}}{\overline{\{g\}}\varphi \vee \overline{\{g\}}\psi \vdash \overline{\{g\}}(\varphi \wedge \psi)}$$

2. Similar. □

In this setup we can capture back the definition of a *De Morgan negation*, that is an operator which reverses disjunction and conjunction ($\neg(\varphi \wedge \psi) \dashv\vdash \neg\varphi \vee \neg\psi$ and $\neg(\varphi \vee \psi) \dashv\vdash \neg\varphi \wedge \neg\psi$) and such that $\varphi \dashv\vdash \neg\neg\varphi$. In order to do that, we just add a single modality for a sort, together with equality $\neg\neg = \{\text{id}\}$.

A.3 Final remarks

In this chapter, we have shown how one can use enriched categories in order to specify certain class of modal logics, and how to obtain the corresponding algebraic semantics. However, the content presented in this chapter falls short of being satisfactory complete for two reasons. First, we provide only algebraic semantics. For many modal logicians, Kripke-style semantics is more appealing. Second, the scope of the axioms that can be expressed in the logics of this chapter is limited. We aim to address these issues in the future work.

Kripke-style completeness

Algebraic semantics is very convenient as in many cases the logic is almost automatically complete with respect to the models. As we could see in this chapter, the proof is fairly unsophisticated. Kripke semantics is usually of greater interest for modal logicians. The underlying set-theoretic structure of frames, be it relational Kripke frames, neighborhood structures or general frames, is appealing and easier to understand. It would be interesting to see whether duality theorems (like Jónsson-Tarski duality) can be extended to multi-sorted and enriched settings. This would be the main line of research in the direction of Kripke-style completeness.

Many-sorted enriched algebraic theories

The range of axioms that can be expressed in logics of this chapter is quite limited. It is enough to encode usual axioms such as **(T)** of **(4)**, but we cannot express, for example, the interaction axioms of positive modal logic, such as $\diamond\varphi \wedge \Box\psi \vdash \diamond(\varphi \wedge \psi)$ or $\Box(\varphi \vee \psi) \vdash \Box\varphi \vee \Box\psi$. One way to overcome this problem is to use *multisorted enriched Lawvere theories*.

Definition A.3.1. A *poset-enriched multisorted Lawvere theory* is a poset-enriched category \mathbb{L} with finite products and a set of distinguished objects S , such that every object of \mathbb{L} is (isomorphic to) a finite product of objects from S .

Using standard (not multisorted, not enriched) Lawvere theories one can describe algebraic theories, such as groups or distributive lattices categorically. Morphisms from X^n to X define the set of n -ary operations that can be defined in the theory. For example, for distributive lattices, there are two (except projections) morphisms $\wedge, \vee : X^2 \rightarrow X$ that represent the binary operation in lattices. The properties of operations are encoded

as equalities of morphisms. The set-theoretic models of the theory are product-preserving functors from \mathbb{L} to the category of sets. Moving to enriched multisorted setting, we can add additional operations between sorts, that represent modalities of the logic, and encode more axioms. Suppose we have two sorts a, b and two modalities \Box, \Diamond from b to a . In a Lawvere theory, the axiom $\Diamond\varphi \wedge \Box\psi \vdash \Diamond(\varphi \wedge \psi)$ would be represented by the following inequality of morphisms:

$$\begin{array}{ccccc}
 & & \wedge_b & \longrightarrow & b & & \Diamond & & \\
 & & \curvearrowright & & & & \curvearrowleft & & \\
 b \times b & & & & & & & & a \\
 & & & & \uparrow & & & & \\
 & & & & \uparrow & & & & \\
 & & & & & & & & \\
 & & \Diamond \times \Box & \longrightarrow & a \times a & & \wedge_a & &
 \end{array}$$

BIBLIOGRAPHY

- [Abr+15] Samson Abramsky, Rui Soares Barbosa, Kohei Kishida, Raymond Lal, and Shane Mansfield. “Contextuality, Cohomology and Paradox”. Comment: 18 pages, 4 figures. 2015. DOI: 10.4230/LIPIcs.CSL.2015.211. arXiv: 1502.03097 [quant-ph].
- [AF09] Carlos Areces and Diego Figueira. “Which Semantics for Neighbourhood Semantics?” In: *IJCAI*. 2009, pp. 671–676.
- [AtC07] Carlos Areces and Balder ten Cate. “Hybrid Logics”. In: *Studies in Logic and Practical Reasoning*. Ed. by Patrick Blackburn, Johan Van Benthem, and Frank Wolter. Vol. 3. Handbook of Modal Logic. Elsevier, 2007, pp. 821–868. DOI: 10.1016/S1570-2464(07)80017-6.
- [ÅW17] Thomas Ågotnes and Yi N. Wáng. “Resolving Distributed Knowledge”. In: *Artificial Intelligence* 252 (2017), pp. 1–21. ISSN: 0004-3702. DOI: 10.1016/j.artint.2017.07.002.
- [Bal03] Alexandru Baltag. “A Coalgebraic Semantics for Epistemic Programs”. In: *Electronic Notes in Theoretical Computer Science*. CMCS’03, Coalgebraic Methods in Computer Science (Satellite Event for ETAPS 2003) 82.1 (2003), pp. 17–38. ISSN: 1571-0661. DOI: 10.1016/S1571-0661(04)80630-3.
- [Bar96] Michael Barr. “The Chu Construction”. In: *Theory and Applications of Categories* 2.2 (1996), pp. 17–35.
- [Bar98] Michael Barr. “The Separated Extensional Chu Category”. In: *Theory and Applications of Categories* 4.6 (1998), pp. 137–147.
- [Bez+23] Nick Bezhanishvili, Anna Dmitrieva, Jim de Groot, and Tommaso Moraschini. *Positive Modal Logic Beyond Distributivity*. Version 3. 2023. DOI: 10.48550/arXiv.2204.13401. arXiv: 2204.13401 [cs, math]. preprint.
- [BG97] Elizabeth Borowsky and Eli Gafni. “A Simple Algorithmically Reasoned Characterization of Wait-Free Computation (Extended Abstract)”. In: *Proceedings of the Sixteenth Annual ACM Symposium on Principles of Distributed Computing*. PODC ’97. New York, NY, USA: Association for Computing Machinery, 1997, pp. 189–198. ISBN: 978-0-89791-952-4. DOI: 10.1145/259380.259439.

- [Bíl+11] Marta Bílková, Alexander Kurz, Daniela Petrişan, and Jiří Velebil. “Relation Liftings on Preorders and Posets”. In: *Algebra and Coalgebra in Computer Science: 4th International Conference, CALCO 2011, Winchester, UK, August 30–September 2, 2011. Proceedings 4*. Springer, 2011, pp. 115–129.
- [BRV01] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science. Cambridge: Cambridge University Press, 2001. ISBN: 978-0-521-52714-9. DOI: 10.1017/CB09781107050884.
- [BS08] Alexandru Baltag and Sonja Smets. “The Logic of Conditional Doxastic Actions”. In: *New Perspectives on Games and Interaction*. Ed. by Krzysztof R. Apt and Robert Van Rooij. Amsterdam University Press, 2008, pp. 9–32. ISBN: 978-90-8964-057-4. JSTOR: j.ctt46mwfz.4.
- [BS10] Alexandru Baltag and Sonja Smets. “Correlated Knowledge: An Epistemic-Logic View on Quantum Entanglement”. In: *International Journal of Theoretical Physics* 49.12 (2010), pp. 3005–3021. ISSN: 0020-7748, 1572-9575. DOI: 10.1007/s10773-010-0411-5.
- [BvdH14] Nick Bezhanishvili and Wiebe van der Hoek. “Structures for Epistemic Logic”. In: *Johan van Benthem on Logic and Information Dynamics*. Ed. by Alexandru Baltag and Sonja Smets. Outstanding Contributions to Logic. Cham: Springer International Publishing, 2014, pp. 339–380. ISBN: 978-3-319-06025-5. DOI: 10.1007/978-3-319-06025-5_12.
- [DLW11] Cédric Dégrement, Benedikt Löwe, and Andreas Witzel. “The Synchronicity of Dynamic Epistemic Logic”. In: *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge*. TARK XIII. New York, NY, USA: Association for Computing Machinery, 2011, pp. 145–152. ISBN: 978-1-4503-0707-9. DOI: 10.1145/2000378.2000395.
- [DLW23] Yifeng Ding, Jixin Liu, and Yanjing Wang. “Someone Knows That Local Reasoning on Hypergraphs Is a Weakly Aggregative Modal Logic”. In: *Synthese. An International Journal for Epistemology, Methodology and Philosophy of Science* 201.2 (2023), p. 46.
- [Dun95] J. Michael Dunn. “Positive Modal Logic”. In: *Studia Logica: An International Journal for Symbolic Logic* 55.2 (1995), pp. 301–317. ISSN: 0039-3215. JSTOR: 20015820.
- [EH83] E. Allen Emerson and Joseph Y. Halpern. ““Sometimes” and “Not Never” Revisited: On Branching versus Linear Time (Preliminary Report)”. In: *Proceedings of the 10th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*. POPL ’83. New York, NY, USA: Association for Computing Machinery, 1983, pp. 127–140. ISBN: 978-0-89791-090-3. DOI: 10.1145/567067.567081.

- [Fag+95] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Vardi. *Reasoning About Knowledge*. MIT Press, 1995.
- [FH87] Ronald Fagin and Joseph Y. Halpern. “Belief, Awareness, and Limited Reasoning”. In: *Artificial Intelligence* 34.1 (1987), pp. 39–76. ISSN: 0004-3702. DOI: 10.1016/0004-3702(87)90003-8.
- [FLP85] Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. “Impossibility of Distributed Consensus with One Faulty Process”. In: *Journal of the ACM* 32.2 (1985), pp. 374–382. ISSN: 0004-5411. DOI: 10.1145/3149.214121.
- [GLR18] Éric Goubault, Jérémy Ledent, and Sergio Rajsbaum. “A Simplicial Complex Model for Dynamic Epistemic Logic to Study Distributed Task Computability”. In: *Electronic Proceedings in Theoretical Computer Science* 277 (2018), pp. 73–87. ISSN: 2075-2180. DOI: 10.4204/EPTCS.277.6. arXiv: 1809.03095.
- [GLR21] Éric Goubault, Jérémy Ledent, and Sergio Rajsbaum. “A Simplicial Complex Model for Dynamic Epistemic Logic to Study Distributed Task Computability”. In: *Information and Computation*. Special Issue on Ninth International Symposium on Games, Automata, Logics and Formal Verification (GandALF 2018) 278 (2021), p. 104597. ISSN: 0890-5401. DOI: 10.1016/j.ic.2020.104597.
- [GLR22] Éric Goubault, Jérémy Ledent, and Sergio Rajsbaum. “A Simplicial Model for KB4n: Epistemic Logic with Agents That May Die”. In: *39th International Symposium on Theoretical Aspects of Computer Science (STACS 2022)*. Ed. by Petra Berenbrink and Benjamin Monmege. Vol. 219. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022, 33:1–33:20. ISBN: 978-3-95977-222-8. DOI: 10.4230/LIPIcs.STACS.2022.33.
- [GM12] Eric Goubault and Samuel Mimram. “Formal Relationships between Geometrical and Classical Models for Concurrency”. In: *Electronic Notes in Theoretical Computer Science* 283 (2012), pp. 77–109.
- [GMT18] Éric Goubault, Samuel Mimram, and Christine Tasson. “Geometric and Combinatorial Views on Asynchronous Computability”. In: *Distributed Computing* 31.4 (2018), pp. 289–316. ISSN: 1432-0452. DOI: 10.1007/s00446-018-0328-4.
- [GNV05] Mai Gehrke, Hideo Nagahashi, and Yde Venema. “A Sahlqvist Theorem for Distributive Modal Logic”. In: *Annals of Pure and Applied Logic* 131.1 (2005), pp. 65–102. ISSN: 0168-0072. DOI: 10.1016/j.apal.2004.04.007.

- [Gou+23] Éric Goubault, Roman Kniazev, Jérémy Ledent, and Sergio Rajsbaum. “Semi-Simplicial Set Models for Distributed Knowledge”. In: *2023 38th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. 2023 38th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS). 2023, pp. 1–13. DOI: 10.1109/LICS56636.2023.10175737.
- [Gri+19] Alexander Grigor’yan, Rolando Jimenez, Yuri Muranov, and Shing-Tung Yau. “Homology of Path Complexes and Hypergraphs”. In: *Topology and its Applications* 267 (2019), p. 106877. ISSN: 0166-8641. DOI: 10.1016/j.topol.2019.106877.
- [GS13] Daniel Gorín and Lutz Schröder. *Simulations and Bisimulations For Coalgebraic Modal Logics*. Version 2. 2013. DOI: 10.48550/arXiv.1303.2467. arXiv: 1303.2467 [cs, math]. preprint.
- [HKP07] Helle Hvid Hansen, Clemens Kupke, and Eric Pacuit. “Bisimulation for Neighbourhood Structures”. In: *Algebra and Coalgebra in Computer Science: Second International Conference, CALCO 2007, Bergen, Norway, August 20–24, 2007. Proceedings 2*. Springer, 2007, pp. 279–293.
- [HKR14] Maurice Herlihy, D. N Kozlov, and Sergio Rajsbaum. *Distributed Computing through Combinatorial Topology*. Waltham, Mass.: Morgan Kaufmann, 2014. ISBN: 978-0-12-404728-0.
- [HM90] Joseph Y. Halpern and Yoram Moses. “Knowledge and Common Knowledge in a Distributed Environment”. In: *Journal of the ACM* 37.3 (1990), pp. 549–587. ISSN: 0004-5411. DOI: 10.1145/79147.79161.
- [HS99] Maurice Herlihy and Nir Shavit. “The Topological Structure of Asynchronous Computability”. In: *Journal of the ACM* 46.6 (1999), pp. 858–923. ISSN: 0004-5411, 1557-735X. DOI: 10.1145/331524.331529.
- [Kel82] Gregory Maxwell Kelly. *Basic Concepts of Enriched Category Theory*. Vol. 64. CUP Archive, 1982.
- [KG22] G. A. Kavvos and Daniel Gratzer. *Under Lock and Key: A Proof System for a Multimodal Logic*. 2022. DOI: 10.48550/arXiv.2211.06217. arXiv: 2211.06217 [cs, math]. preprint.
- [KMP12] Sophia Knight, Radu Mardare, and Prakash Panangaden. “Combining Epistemic Logic and Hennessy-Milner Logic”. In: *Logic and Program Semantics: Essays Dedicated to Dexter Kozen on the Occasion of His 60th Birthday*. Springer, 2012, pp. 219–243.
- [KV16] Alexander Kurz and Jiří Velebil. “Relation Lifting, a Survey”. In: *Journal of Logical and Algebraic Methods in Programming* 85.4 (2016), pp. 475–499. ISSN: 23522208. DOI: 10.1016/j.jlamp.2015.08.002.
- [Law07] F William Lawvere. “Axiomatic Cohesion”. In: *Theory and Applications of Categories* 19 (2007), pp. 41–49.

- [Led19] Jérémy Ledent. “Geometric Semantics for Asynchronous Computability”. PhD thesis. Université Paris Saclay (COMUE), 2019.
- [Mos16] Yoram Moses. “Relating Knowledge and Coordinated Action: The Knowledge of Preconditions Principle”. In: *Theoretical Aspects of Rationality and Knowledge*. 2016.
- [MS18] Paul-André Melliès and Léo Stefanescu. “An Asynchronous Soundness Theorem for Concurrent Separation Logic”. In: *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science*. LICS ’18: 33rd Annual ACM/IEEE Symposium on Logic in Computer Science. Oxford United Kingdom: ACM, 2018, pp. 699–708. ISBN: 978-1-4503-5583-4. DOI: 10.1145/3209108.3209116.
- [Nis22] Susumu Nishimura. *Proving Unsolvability of Set Agreement Task with Epistemic Mu-Calculus*. 2022. DOI: 10.48550/arXiv.2205.06452. arXiv: 2205.06452 [cs]. preprint.
- [NSW94] Mogens Nielsen, Vladimiro Sassone, and Glynn Winskel. “Relationships between Models of Concurrency”. In: *A Decade of Concurrency Reflections and Perspectives*. Ed. by J. W. de Bakker, W. -P. de Roever, and G. Rozenberg. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 1994, pp. 425–476. ISBN: 978-3-540-48423-3. DOI: 10.1007/3-540-58043-3_25.
- [Pac07] Eric Pacuit. “Some Comments on History Based Structures”. In: *Journal of Applied Logic*. Selected Papers from the 4th International Workshop on Computational Models of Scientific Reasoning and Applications 5.4 (2007), pp. 613–624. ISSN: 1570-8683. DOI: 10.1016/j.jal.2006.03.012.
- [Pac17] Eric Pacuit. *Neighborhood Semantics for Modal Logic*. Springer, 2017. 165 pp. ISBN: 978-3-319-67149-9.
- [Pnu77] Amir Pnueli. “The Temporal Logic of Programs”. In: *18th Annual Symposium on Foundations of Computer Science (Sfcs 1977)*. 18th Annual Symposium on Foundations of Computer Science (Sfcs 1977). 1977, pp. 46–57. DOI: 10.1109/SFCS.1977.32.
- [Pra03] Vaughan R. Pratt. “Transition and Cancellation in Concurrency and Branching Time”. In: *Mathematical Structures in Computer Science* 13.4 (2003), pp. 485–529. ISSN: 09601295, 14698072. DOI: 10.1017/S0960129503004031.
- [Pra91] Vaughn Pratt. “Modeling Concurrency with Geometry”. In: *Proceedings of the 18th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. POPL ’91. New York, NY, USA: Association for Computing Machinery, 1991, pp. 311–322. ISBN: 978-0-89791-419-2. DOI: 10.1145/99583.99625.
- [Pra99] Vaughan Pratt. “Chu Spaces”. In: *School on Category Theory and Applications (Coimbra, 1999)* 21 (1999), pp. 39–100.

- [Rie11] Emily Riehl. “A Leisurely Introduction to Simplicial Sets”. In: *Unpublished expository article* (2011).
- [Rie14] Emily Riehl. *Categorical Homotopy Theory*. Vol. 24. Cambridge University Press, 2014.
- [Roe07] Floris Roelofsen. “Distributed Knowledge”. In: *Journal of Applied Non-Classical Logics* 17.2 (2007), pp. 255–273. DOI: 10.3166/janc1.17.255-273.
- [RvDK23] Rojo Randrianomentsoa, Hans van Ditmarsch, and Roman Kuznets. *Impure Simplicial Complexes: Complete Axiomatization*. 2023. DOI: 10.48550/arXiv.2211.13543. arXiv: 2211.13543 [cs]. preprint.
- [Sah75] Henrik Sahlqvist. “Completeness and Correspondence in the First and Second Order Semantics for Modal Logic”. In: *Studies in Logic and the Foundations of Mathematics*. Ed. by Stig Kanger. Vol. 82. Proceedings of the Third Scandinavian Logic Symposium. Elsevier, 1975, pp. 110–143. DOI: 10.1016/S0049-237X(08)70728-6.
- [SD09] Mehrnoosh Sadrzadeh and Roy Dyckhoff. “Positive Logic with Adjoint Modalities: Proof Theory, Semantics and Reasoning about Information”. In: *Electronic Notes in Theoretical Computer Science*. Proceedings of the 25th Conference on Mathematical Foundations of Programming Semantics (MFPS 2009) 249 (2009), pp. 451–470. ISSN: 1571-0661. DOI: 10.1016/j.entcs.2009.07.102.
- [Shu18] Michael Shulman. *Contravariance through Enrichment*. 2018. DOI: 10.48550/arXiv.1606.05058. arXiv: 1606.05058 [math]. preprint.
- [Shu23] Michael Shulman. *Semantics of Multimodal Adjoint Type Theory*. 2023. DOI: 10.48550/arXiv.2303.02572. arXiv: 2303.02572 [cs, math]. preprint.
- [vBen+09] Johan van Benthem, Jelle Gerbrandy, Tomohiro Hoshi, and Eric Pacuit. “Merging Frameworks for Interaction”. In: *Journal of Philosophical Logic* 38.5 (2009), pp. 491–526. ISSN: 1573-0433. DOI: 10.1007/s10992-008-9099-x.
- [vdHM92] W. van der Hoek and J.-J. Ch. Meyer. “Making Some Issues of Implicit Knowledge Explicit”. In: *International Journal of Foundations of Computer Science* 03.02 (1992), pp. 193–223. ISSN: 0129-0541. DOI: 10.1142/S0129054192000139.
- [vdHvLM99] Wiebe van der Hoek, Bernd van Linder, and John-Jules Meyer. “Group Knowledge Is Not Always Distributed (Neither Is It Always Implicit)”. In: *Mathematical Social Sciences* 38.2 (1999), pp. 215–240. ISSN: 0165-4896. DOI: 10.1016/S0165-4896(99)00013-X.

- [vDK] Hans van Ditmarsch and Roman Kuznets. “Wanted Dead or Alive : Epistemic Logic for Impure Simplicial Complexes”. to appear in *Journal of Logic and Computation*. DOI: 10.48550/arXiv.2103.03032. arXiv: 2103.03032 [cs].
- [vDvdHK07] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*. Springer Science & Business Media, 2007. 303 pp. ISBN: 978-1-4020-5839-4. Google Books: dKRQPHv0IGQC.
- [Win86] Glynn Winskel. “Event Structures”. In: *Advanced Course on Petri Nets*. Springer, 1986, pp. 325–392.
- [WN93] Glynn Winskel and Mogens Nielsen. “Models for Concurrency”. In: *DAIMI Report Series 463* (1993).
- [YN21] Koki Yagi and Susumu Nishimura. *Logical Obstruction to Set Agreement Tasks for Superset-Closed Adversaries*. 2021. DOI: 10.48550/arXiv.2011.13630. arXiv: 2011.13630 [cs]. preprint.

Titre : Modèles géométriques de la logique épistémique

Mots clés : logique modale, logique épistémique, calcul distribué

Résumé : Il y a environ 40 ans, deux approches majeures pour l'étude des systèmes distribués ont été développées. Une des approches a établi une perspective topologique sur le calcul distribué, exprimant la solvabilité des tâches distribuées à travers des notions standards de topologie algébrique, en modélisant les espaces d'états sous forme de complexes simpliciaux. En parallèle, une application importante de la logique épistémique, un type de logique modale, a apporté un point de vue alternatif sur la structure des systèmes distribués, en utilisant la notion de connaissance pour décrire le comportement des processus dans un système. On s'est récemment rendu compte que ces deux approches sont étroitement liées et que les modèles topologiques peuvent, en fait, servir de modèles de logique épistémique. Cette thèse poursuit un programme de recherche visant à unifier ces deux approches. Le premier objectif est de généraliser la sémantique actuelle de la logique épistémique basée sur les complexes simpliciaux au cas des ensembles simpliciaux. Nous montrons qu'avec ces modèles,

on peut exprimer des connaissances de groupe non standard, c'est-à-dire des situations où la connaissance d'un groupe vu dans son ensemble est strictement supérieure à l'union des connaissances de ses membres. On considère ensuite une variante multi-sortes de la logique épistémique, dans laquelle les propriétés de l'environnement et les propriétés locales des agents sont exprimées séparément. Cette logique est interprétée dans les hypergraphes chromatiques, qui sont une généralisation supplémentaire des complexes simpliciaux, nous permettant de souligner le rôle des points de vue locaux des agents dans les systèmes distribués. On étudie la dynamique de la connaissance dans les systèmes distribués en introduisant une variante dynamique des hypergraphes chromatiques. Dans ces modèles, les points de vue locaux des agents sont dotés d'une structure temporelle, ce qui permet de modéliser l'évolution des connaissances dans le temps. De plus, on discute de la relation entre la connaissance et la concurrence dans ce contexte.

Title : On Geometric Models of Epistemic Logic

Keywords : modal logic, epistemic logic, distributed computing

Abstract : About 40 years ago, two major approaches to the study of distributed systems were developed. One approach established a topological perspective on distributed computing, expressing solvability of distributed tasks through standard notions of algebraic topology, by modelling spaces of states as simplicial complexes. In parallel, a prominent application of epistemic logic, a type of modal logic, provided an alternative point of view on the structure of distributed systems, by using the notion of knowledge to describe the behavior of processes in a system. It has recently been realized that these two approaches are closely related, and that the topological models can, in fact, serve as models of epistemic logic. This thesis continues a research program aimed at unification of these two approaches. Our first goal is to generalize the existing semantics of epistemic logic based on simplicial complexes to the case of simplicial sets. We show that with these models one can ex-

press non-standard group knowledge, that is situations when knowledge of a group seen as a whole is strictly greater than the union of knowledge of its members. We then proceed to study a many-sorted variant of epistemic logic, where properties of the environment and local properties of agents are expressed separately. This logic is interpreted in chromatic hypergraphs, which are a further generalization of simplicial complexes, allowing us to emphasize the role of local points of view of agents in distributed systems. We also investigate the dynamics of knowledge in distributed systems by introducing a dynamic variant of chromatic hypergraphs. In these models, local views of agents are supplemented with a temporal structure, which allows us to model the evolution of knowledge in time. Additionally, we discuss the relationship between knowledge and concurrency in this setting.