Synergistic Knowledge

Christian Cachin[0000−0001−8967−9213], David Lehnherr[0000−0002−4956−4064], and Thomas Studer[0000−0002−0949−3302]

University of Bern, Bern, Switzerland
{christan.cachin, david.lehnherr, thomas.studer}@unibe.ch

Abstract. In formal epistemology, group knowledge is often modelled as the knowledge that the group would have if the agents share all their individual knowledge. However, this interpretation does not account for relations between agents. In this work, we propose the notion of synergistic knowledge, which makes it possible to model those relationships. As examples, we investigate the use of consensus objects and the problem of the dining cryptographers. Moreover, we show that our logic can also be used to model certain aspects of information flow in networks.

Keywords: Distributed Knowledge · Synergy · Modal Logic

1 Introduction

A simplicial interpretation of the semantics of modal logic has gained recent interest, due to the success of applying topological methods to problems occurring in distributed systems. The topological approach to distributed computing, exemplified by Herlihy, Kozlov and Rajsbaum [10], interprets the configurations of a distributed system as a simplicial complex. The vertices of a simplicial complex represent local states of different agents and an edge between two vertices means that the two local states can occur together.

Modal logic has various applications to problems in distributed computing, such as agreement (c.f. Halpern and Moses [8]). Models for modal logic are usually based on a possible worlds approach where the operator $\Box$ is evaluated on Kripke frames. In a world w, a formula ϕ is known, denoted by $\Box\phi$, if and only if ϕ is true in each world indistinguishable from w. These frames can be extended to multi-agent systems by introducing an indistinguishability relation for each agent. A formula ϕ is distributed knowledge of a group, first introduced by Halpern and Moses [8], if and only if ϕ is true in all worlds that cannot be distinguished by any member of the group.

Given a set of agents, van Ditmarsch, Goubault, Ledent and Rajsbaum [2] define a simplicial model for settings in which all agents are present at any point in time. This semantics is shown to be equivalent to the modal logic $\mathsf{S5}_n$. In the same setting, Goubault, Ledent and Rajsbaum [5] look at distributed task computability through the lens of dynamic epistemic logic (c.f. van Ditmarsch, van der Hoek and Kooi [3]). Using dynamic epistemic logic makes it possible to model the relationship between input and output configurations of tasks, which is one of the core objectives of the classical topological approach to distributed systems (c.f. Herlihy and Shavit [11]). Work regarding models where some agents might not be present in a configuration was conducted independently by van Ditmarsch and Kuznets [4] and Goubault, Ledent and Rajsbaum [6]. The latter work shows the equivalence between their simplicial models and Kripke models for the logic $\mathsf{KB4}_n$, whereas van Ditmarsch and Kuznets [4] deal with crashed agents by letting formulas be undefined and show that their logic is sound. Randrianomentsoa, van Ditmarsch and Kuznets [12] show in a follow up work that the route taken by van Ditmarsch and Kuznets [4] leads to a sound and complete semantic for their axiom system. Both works remark that impure complexes cannot capture the information of improper Kripke frames, i.e. models in which some worlds cannot be distinguished from others by all agents. They point out the need for extending the interpretation of simplicial complexes to simplicial sets, i.e. simplicial complexes that may contain the same simplex arbitrarily often. Furthermore, the latter work also shed light on a new notion of group knowledge which differs from the usual definition of distributed knowledge. Their example is depicted in Figure 1 in which the agents a and b individually cannot distinguish the worlds X and Y since the vertices labelled with a and b belong to both X and Y. However, a and b together can distinguish between X and Y since the worlds do not share an edge between vertices a and b. In a follow up work, Goubault, Kniazev, Ledent and Rajsbaum [7] provide, among various results, a semantics for such simplicial sets and provide a higher order interpretation of distributed knowledge for a set of agents.

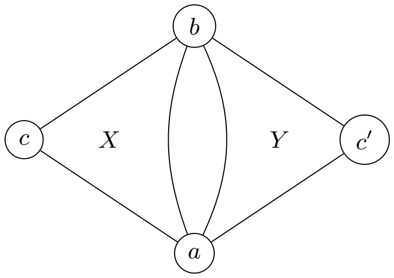

Fig. 1. A model in which two agents a and b can together distinguish between the worlds X and Y. However, they cannot do so individually.

In this paper, we propose the notion of synergistic knowledge, which allows a group of agents to know more than just the consequences of their pooled knowledge. That is, our newly introduced epistemic operator $[G]$ supports a principle that could be paraphrased as *the sum is greater than its parts*, hence the name synergistic knowledge. Different to the higher order interpretation of distributed knowledge by Goubault, Kniazev, Ledent and Rajsbaum [7], which analyses the knowledge of a set of agents, we interpret G as a simplicial complex over the set of agents. Hence, we refer to G as an agent pattern instead of a group. The operator $[G]$ allows us to model relations between subgroups of agents and how they interact with each other. Hence, two agent patterns G and H may contain the same agents, but differ in the relations among them. For example, in Figure 1, we can distinguish the pattern $\{\{a\},\{b\}\}$, which cannot distinguish between X and Y because the two worlds share vertices labelled with a and with b, from the pattern $\{\{a,b\}\}$, which can distinguish X and Y due to X and Y not sharing an edge. As applications for synergistic knowledge, we investigate the use of consensus objects (c.f. Herlihy [9]) and the problem of the dining cryptographers (c.f. Chaum [1]).

Our main contribution consists in i) providing a novel simplicial set semantics for modal logic that is simpler than previous approaches as it does not refer to category theory or make use of chromatic maps, and ii) the introduction of a new knowledge operator $[G]$ that allows us to express distributed knowledge in a more fine grained way, as well as iii) presenting a new notion of indistinguishability, called componentwise indistinguishability, which models the flow of information in networks. All points mentioned are accompanied by examples.

In Section 2, we introduce our new simplicial set model together with a corresponding indistinguishability relation. Section 3 studies the logic induced by our model. In Section 4, we present examples that illustrate the use of our logic. In Section 5, we adapt our notion of indistinguishability and show how it can be used in order to model the information flow in a network. Lastly, we draw a conclusion of our work in Section 6.

2 Indistinguishability

In this section, we introduce the indistinguishability relation that is used to model synergistic knowledge. Let Ag denote a set of finitely many agents and let

$$\mathsf{Agsi} = \{(A,i) \mid A \subseteq \mathsf{Ag} \text{ and } i \in \mathbb{N}\}.$$

We may think of a pair $(A,i) \in \mathsf{Agsi}$ as representing a set of agents A in local state i. Further, let $S \subseteq \mathsf{Agsi}$. An element $(A,i) \in S$ is *maximal in S* if and only if

$$\forall (B,j) \in S. |A| \geq |B|, \text{ where } |X| \text{ denotes the cardinality of the set } X.$$

Definition 1 (Simplex). *Let $\emptyset \neq S \subseteq \mathsf{Agsi}$. S is a* simplex *if and only if*

S1*: The maximal element is unique, i.e.*

$$\textit{if } (A,i) \in S \textit{ and } (B,j) \in S \textit{ are maximal in } S \textit{ then, } A = B \textit{ and } i = j.$$

The maximal element of S is denoted as $\max(S)$.

S2*: S is uniquely downwards closed, i.e. for all $(B,i) \in S$ and $\emptyset \neq C \subseteq B$*

$$\exists! j \in \mathbb{N}.(C,j) \in S, \textit{ where } !\exists j \textit{ means that there exists exactly one } j.$$

$\mathsf{S3}$*: S contains nothing else, i.e.*

$$(B,i) \in S \textit{ and } (A,j) = \max(S) \textit{ implies } B \subseteq A.$$

Definition 2 (Complex). *Let $\mathbb{C}$ be a set of simplexes. $\mathbb{C}$ is a* complex *if and only if*

C*: For any $S, T \in \mathbb{C}$, if there exist A and i with $(A,i) \in S$ and $(A,i) \in T$, then*

$$\textit{for all } B \subseteq A \textit{ and all } j \quad (B,j) \in S \Longleftrightarrow (B,j) \in T.$$

Condition C guarantees that the maximal element of a simplex uniquely determines it within a given complex.

Lemma 1. *Let $\mathbb{C}$ be a complex and $S, T \in \mathbb{C}$. We find*

$$\max(S) = \max(T) \textit{ implies } S = T.$$

Proof. We show $S \subseteq T$. The other direction is symmetric. Let $(A,i) = \max(S)$. Assume $(B,j) \in S$. Because of $\mathsf{S3}$, we have $B \subseteq A$. By Condition C, we conclude $(B,j) \in T$. □

Whenever it is clear from the context, we abbreviate $(\{a_1, ..., a_n\}, i)$ as $a_1...a_n i$ in order to enhance readability. Furthermore, we may use a row (or a mixed row-column) notation to emphasize simplexes. For example,

$$\left\{ \begin{Bmatrix} ab0 \\ a0 \\ b0 \end{Bmatrix}, \begin{Bmatrix} ab1 \\ a0 \\ b1 \end{Bmatrix} \right\}$$

is a complex that contains 2 simplexes. Whenever we refer to a simplex within a complex, we write $\langle Ai \rangle$ for the simplex with maximal element (A,i). Condition C guarantees that this notation is well-defined.

Oberserve that Condition C ensures that neither

$$\left\{ \begin{Bmatrix} ab0 \\ a0, b0 \end{Bmatrix}, \begin{Bmatrix} ab0 \\ a1, b1 \end{Bmatrix} \right\} \quad \text{nor} \quad \left\{ \begin{Bmatrix} abc0 \\ ab0, ac0, bc0 \\ a0, b0, c0 \end{Bmatrix}, \begin{Bmatrix} ab0 \\ a1, b1 \end{Bmatrix} \right\}$$

is a complex, although each individual simplex is well-formed.

Definition 3 (Indistinguishability). *Let $S \subseteq \mathsf{Agsi}$, we define*

$$S^\circ = \{A \mid \exists i \in \mathbb{N} : (A,i) \in S\}.$$

An agent pattern *is a subset of* $\mathsf{Pow}(\mathsf{Ag}) \setminus \{\emptyset\}$*. An agent pattern* cannot distinguish *between two simplexes S and T, denoted by $S \sim_G T$, if and only if $G \subseteq (S \cap T)^\circ$.*

Definition 4 (Partial equivalence relation (PER)). *A relation $R \subseteq \mathcal{S} \times \mathcal{S}$ is a partial equivalence relation if and only if it is symmetric and transitive.*

Lemma 2 (PER). *$\sim_G$ is a PER.*

Proof. Symmetry immediately follows from the fact that set intersection is commutative. To show transitivity, let S, T, U be simplexes with $S \sim_G T$ and $T \sim_G U$, i.e.

$$G \subseteq (S \cap T)^\circ \tag{1}$$

$$G \subseteq (T \cap U)^\circ \tag{2}$$

Let $A \in G$. Because of (1), there exists i with

$$(A, i) \in S \quad \text{and} \quad (A, i) \in T. \tag{3}$$

Because of (2), there exists j with

$$(A, j) \in T \quad \text{and} \quad (A, j) \in U. \tag{4}$$

From (3), (4), and Condition $\mathsf{S2}$ we obtain $i = j$. Thus by (3) and (4), we get $A \in (S \cap U)^\circ$. Since A was arbitrary in G, we conclude $G \subseteq (S \cap U)^\circ$. □

Lemma 3. *Let $G \subseteq \mathsf{Pow}(\mathsf{Ag})$ be an agent pattern and*

$$\mathsf{noSym}(G) := \{\{a\} \mid \exists A \in G \textit{ and } a \in A\}.$$

Let $\mathcal{S}_G$ be a set of simplexes such that for any $S \in \mathcal{S}_G$ we have $\mathsf{noSym}(G) \subseteq S^\circ$. The indistinguishability relation $\sim_G$ is reflexive on $\mathcal{S}_G \times \mathcal{S}_G$ and empty otherwise.

Proof. We first show reflexivity. If $G = \emptyset$, then trivially $G \subseteq (S \cap S)^\circ$ for any S. Assume $G \neq \emptyset$. Let $S \in \mathcal{S}_G$. For each $B \in G$, we have to show that $B \in (S \cap S)^\circ$, i.e. that

$$\text{there exists } i \text{ with } (B, i) \in S. \tag{5}$$

Let $(A, i) := \max(S)$. Let $b \in B$. Because of $\mathsf{noSym}(G) \subseteq S^\circ$, there exists l such that $(\{b\}, l) \in S$. By $\mathsf{S3}$ we get $b \in A$. Since b was arbitrary in B, we get $B \subseteq A$. By $\mathsf{S2}$ we conclude that (5) holds and symmetry is established.

We now show that $\sim_G$ is empty otherwise. Let S be a simplex such that $\mathsf{noSym}(G) \nsubseteq S^\circ$ and let T be an arbitrary simplex. Then there exists a, A with $a \in A \in G$ and $\{a\} \notin S^\circ$, i.e.

$$\text{for all } i,\ (\{a\}, i) \notin S. \tag{6}$$

Suppose towards a contradiction that

$$G \subseteq (S \cap T)^\circ \tag{7}$$

Because of $A \in G$ we get $A \in (S \cap T)^\circ$. Hence $A \in S^\circ$, i.e. there exists l with $(A, l) \in S$. With S2 and $\{a\} \subseteq A$ we find that there exists j with $(\{a\}, j) \in S$. Contradiction to (6). Thus (7) cannot hold. □

Corollary 1. *$\sim_G$ is an equivalence relation on $\mathcal{S}_G \times \mathcal{S}_G$.*

The following two lemmas establish basic properties of the indistinguishability relation.

Lemma 4 (Anti-Monotonicity). *$G \subseteq H$ implies $\sim_H \subseteq \sim_G$.*

Proof. Assume $G \subseteq H$. For any two simplices S and T with $S \sim_H T$, we have $G \subseteq H \subseteq (S \cap T)^\circ$ by definition and hence $S \sim_G T$, which concludes the proof. □

Lemma 5 (Downward closure). *Let $\mathbb{C}$ be a complex and $S, T \in \mathbb{C}$. Further, let $A \in (S \cap T)^\circ$ and $B \subseteq A$. We find $B \in (S \cap T)^\circ$.*

Proof. From $A \in (S \cap T)^\circ$, we obtain that there exists i such that $(A, i) \in S$ and $(A, i) \in T$. From $\mathsf{S2}$ we find that there exists j such that $(B, j) \in S$. Thus by C, we get $(B, j) \in T$ and we conclude $B \in (S \cap T)^\circ$. □

From the previous two lemmas we immediately obtain the following:

Corollary 2. *Let G be an agent pattern. Let $A, B \subseteq \mathsf{Ag}$ such that $A \subseteq B \in G$. We have*

$$\sim_{G \cup \{A\}} = \sim_G .$$

The next lemma states that adding synergy to an agent pattern makes it stronger in the sense that it can distinguish more simplices. This is shown in Example 1 where the pattern $\{\{a\}, \{b\}\}$ cannot distinguish $\langle abc0 \rangle$ and $\langle abc1 \rangle$ but $\{\{a, b\}\}$ can distinguish these two simplices.

Lemma 6. *Let $H_1, H_2, \ldots, H_n \subseteq \mathsf{Ag}$ with $n \geq 2$ We have*

$$\sim_{\{H_1 \cup H_2, \ldots, H_n\}} \subseteq \sim_{\{H_1, H_2, \ldots, H_n\}} .$$

Proof. From the Lemma 5 and Lemma 4 we find that

$$\sim_{\{H_1 \cup H_2, \ldots, H_n\}} = \sim_{\{H_1 \cup H_2, H_1, H_2, \ldots, H_n\}} \subseteq \sim_{\{H_1, H_2, \ldots, H_n\}} .$$

□

In traditional Kripke semantics, distributed knowledge of a set of agents is modeled by considering the accessibility relation that is given by the intersection of the accessibility relations of the individual agents. The following lemma states that in our framework, this intersection corresponds to the agent pattern consisting of singleton sets for each agent.

Lemma 7. *Let $G \subseteq \mathsf{Ag}$ and $H = \bigcup_{a \in G} \{\{a\}\}$. We have*

$$\bigcap_{a \in G} \sim_{\{\{a\}\}} = \sim_H .$$

Proof. $(S, T) \in \bigcap_{a \in G} \sim_{\{\{a\}\}}$ iff for each $a \in G$, we have $\{a\} \in (S \cap T)^\circ$ iff (by the definition of H) $H \subseteq (S \cap T)^\circ$ iff $S \sim_H T$. □

3 Logic

The logic of synergistic knowledge is a normal modal logic that includes a modality $[G]$ for each agent pattern G. It is closely related to the logic of distributed knowledge but has some additional validities concerning the pattern-based modalities, see, e.g., (Sub) and (Clo) below.

Let Prop be a countable set of atomic propositions. Formulas of the language of synergistic knowledge $\mathcal{L}_{\mathsf{Syn}}$ are inductively defined by the following grammar:

$$\phi ::= p \mid \neg\phi \mid \phi \wedge \phi \mid [G]\phi$$

where $p \in \mathsf{Prop}$ and G is an agent pattern. The remaining Boolean connectives are defined as usual. In particular, we set $\bot := p \wedge \neg p$ for some fixed $p \in \mathsf{Prop}$.

Definition 5 (Model). *A model $\mathcal{M} = (\mathbb{C}, V)$ is a pair where*

1. *$\mathbb{C}$ is a complex and*
2. *$V : \mathbb{C} \to \mathsf{Pow}(\mathsf{Prop})$ is a valuation.*

Definition 6 (Truth). *Let $\mathcal{M} = (\mathbb{C}, V)$ be a model, $w \in \mathbb{C}$, and $\phi \in \mathcal{L}_{\mathsf{Syn}}$. We define $\mathcal{M}, w \Vdash \phi$ inductively by*

$$\begin{array}{lll}
\mathcal{M}, w \Vdash p & \textit{iff} & p \in V(w) \\
\mathcal{M}, w \Vdash \neg\phi & \textit{iff} & \mathcal{M}, w \nVdash \phi \\
\mathcal{M}, w \Vdash \phi \wedge \psi & \textit{iff} & \mathcal{M}, w \Vdash \phi \textit{ and } \mathcal{M}, w \Vdash \psi \\
\mathcal{M}, w \Vdash [G]\phi & \textit{iff} & w \sim_G v \textit{ implies } \mathcal{M}, v \Vdash \phi \quad \textit{for all } v \in \mathbb{C}.
\end{array}$$

We write $\mathcal{M} \Vdash \phi$ if $\mathcal{M}, w \Vdash \phi$ for all $w \in \mathbb{C}$. A formula ϕ is valid *if $\mathcal{M} \Vdash \phi$ for all models $\mathcal{M}$.*

The following formulas are valid:

$$[G](\phi \to \psi) \to ([G]\phi \to [G]\psi) \tag{K}$$

$$[G]\phi \to [G][G]\phi \tag{4}$$

$$\phi \to [G]\neg[G]\neg\phi \tag{B}$$

$$[G]\phi \to [H]\phi \quad \text{if } G \subseteq H \tag{Mono}$$

Let G be an agent pattern. We define, as usual, the formula $\mathsf{alive}(G)$ to be $\neg[G]\bot$. For a single agent a we write $\mathsf{alive}(a)$ instead of $\mathsf{alive}(\{a\})$. The expected equivalences hold:

$$\mathcal{M}, w \Vdash \mathsf{alive}(G) \quad \text{iff} \quad G \subseteq w^\circ \quad \text{iff} \quad w \sim_G w. \tag{8}$$

Indeed, we have $\mathcal{M}, w \Vdash \neg[G]\bot$ iff it is not the case that for all v with $w \sim_G v$, it holds that $\mathcal{M}, v \Vdash \bot$. This is equivalent to there exists v with $w \sim_G v$, which is equivalent to there exitsts v with $G \subseteq (w \cap v)^\circ$. This is equivalent to $w \sim_G w$ and also to $G \subseteq w^\circ$.

Related to $\mathsf{alive}(\cdot)$, the following formulas are valid:

$$\mathsf{alive}(G) \land \mathsf{alive}(H) \to \mathsf{alive}(G \cup H) \tag{Union}$$

$$\mathsf{alive}(G) \to \mathsf{alive}(\{B\}) \quad \text{if there is } A \text{ with } A \in G \text{ and } B \subseteq A \tag{Sub}$$

$$\mathsf{alive}(G) \to \mathsf{alive}(\{A \cup B\}) \quad \text{if } A, B \in G \tag{Clo}$$

(Union) is an immediate consequence of (8). For (Sub), assume $w \sim_G w$, $A \in G$, and $B \subseteq A$. We have $A \in (w \cap w)^\circ$. By Lemma 5 we find $B \in (w \cap w)^\circ$. Hence $w \sim_{\{B\}} w$, which yields (Sub). To show Clo, assume $w \sim_G w$ and $A, B \in G$. Hence $A \in (w \cap w)^\circ$. That is $A \in w^\circ$, i.e. there exists i with $(A, i) \in w$. Let C, j be such that $(C, j) = \max(w)$. By S3, we get $A \subseteq C$. Similarly, we find $B \subseteq C$, and thus $A \cup B \subseteq C$. Using S2, we obtain $A \cup B \in w^\circ$. Therefore, $w \sim_{\{A\cup B\}} w$ and (Clo) is estabished.

Further, note that axiom (T) holds when restricted to groups of agents that are alive:

$$\mathsf{alive}(G) \to ([G]\phi \to \phi) \tag{T}$$

Question 1. Do the axioms (K), (4), (B), (Mono), (Union), (Sub), and (Clo) together with all propositional tautologies and the rules of modus ponens and [G]-necessitation provide a complete axiom system for our notion of validity?

Lemma 7 motivates the following abbreviation. Let $G \subseteq \mathsf{Ag}$ be a set of agents and set $H := \bigcup_{a \in G}\{\{a\}\}$. Then we let D_G be the modality $[H]$. We call this the distributed knowledge modality and let $\mathcal{L}_\mathsf{D}$ be the restriction of $\mathcal{L}_\mathsf{Syn}$ that contains distributed knowledge D_G as the only modality. Note that the usual axioms for the logic of distributed knowledge, formulated in $\mathcal{L}_\mathsf{D}$, hold with respect to synergistic models.

Question 2. Is the logic of synergistic knowledge a conservative extension (with respect to $\mathcal{L}_\mathsf{D}$) of the logic of distributed knowledge?

4 Examples

In this section, we present some examples that illustrate possible applications of our logic to distributed systems. Example 1 highlights one of the main characteristics of synergetic knowledge. That is, the agents a and b can together distinguish between the worlds $\langle abc0\rangle$ and $\langle abc1\rangle$ although they cannot do so individually. Hence, our logic can express the difference between the patterns $\{\{a\}, \{b\}\}$ and $\{\{a, b\}\}$.

Regarding the notation, we will omit the set parentheses for agent patterns whenever it is clear from the context and write for example $[abc, ab, ac]$ instead of $[\{\{a, b, c\}, \{a, b\}, \{a, c\}\}]$.

Example 1 (Two triangles). Let $\mathsf{Ag} = \{a, b, c\}$, $p \in \mathsf{Prop}$, and consider the model $\mathcal{M} = (\mathbb{C}, V)$ in Figure 2 which is given by the complex

$$\mathbb{C} = \left\{ \left\{ \begin{matrix} abc0 \\ ab0, bc0, ac0 \\ a0, b0, c0 \end{matrix} \right\}, \left\{ \begin{matrix} abc1 \\ ab1, bc1, ac1 \\ a0, b0, c1 \end{matrix} \right\} \right\}$$

and by a valuation V such that $p \in V(\langle abc0 \rangle)$ and $p \notin V(\langle abc1 \rangle)$. We find

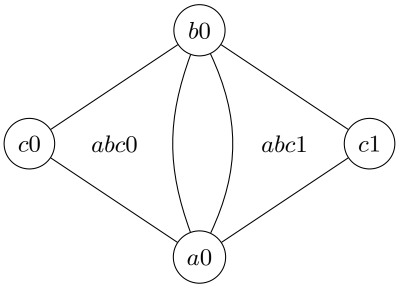

Fig. 2. A model in which two agents a and b can together distinguish between the worlds X and Y. However, they cannot do so individually.

$$\mathcal{M}, \langle abc0 \rangle \Vdash [ab]p \text{ and } \mathcal{M}, \langle abc1 \rangle \Vdash [ab]\neg p,$$

because the worlds $\langle abc0 \rangle$ and $\langle abc1 \rangle$ can be distinguished due to

$$\{\{a, b\}\} \not\subseteq (\langle abc0 \rangle \cap \langle abc1 \rangle)^\circ = \{\{a\}, \{b\}\}.$$

However, for the pattern $H = \{\{a\}, \{b\}\}$ it holds that

$$H = \{\{a\}, \{b\}\} \subseteq (\langle abc0 \rangle \cap \langle abc1 \rangle)^\circ,$$

and hence $\mathcal{M}, \langle abc0 \rangle \nVdash [H]p$. Lastly, if we add c to H, the agents know p:

$$\mathcal{M}, \langle abc0 \rangle \Vdash [a, b, c]p.$$

Another motivation for simplicial sets is that we can model how agents can reason about each others death. As remarked by van Ditmarsch and Kuznets [4] as well as by Goubault, Ledent and Rajsbaum [6], simplicial complexes are not enough to model a setting where an agent considers it possible to be the only one alive. Such scenarios are important, because they arise in failure detection protocols. Example 2 shows such a model.

Example 2 (Two-agents). Let $\mathsf{Ag} = \{a, b\}$, and consider the model $\mathcal{M} = (\mathbb{C}, V)$ in Figure 3 which is given by an arbitrary valuation and the complex

$$\mathbb{C} = \left\{ \left\{ \begin{matrix} ab0 \\ a0, b0 \end{matrix} \right\}, \{a0\} \right\}$$

It is straightforward to verify that $\mathcal{M}, \langle ab0 \rangle \Vdash \mathsf{alive}(a)$ and $\mathcal{M}, \langle ab0 \rangle \Vdash \mathsf{alive}(b)$. However, $\mathcal{M}, \langle a0 \rangle, \nVdash \mathsf{alive}(b)$ because $\{b\} \not\subseteq \langle a0 \rangle^\circ$ and hence $\mathcal{M}, \langle a0 \rangle \Vdash [b]\bot$. Moreover, a alone does not know whether $\mathsf{alive}(b)$ because a cannot distinguish $\langle a0 \rangle$ from $\langle ab0 \rangle$ due to $\{\{a\}\} \subseteq (\langle a0 \rangle \cap \langle ab0 \rangle)^\circ$.

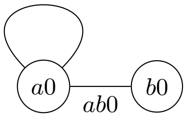

Fig. 3. A model in which a considers it possible that it is the only agent alive.

In Examples 3 and 4 we interpret synergy as having access to some shared primitives. Given three agents, Example 3 captures the idea that for some applications, the agent pattern must include the area of the triangle and not just its edges. Example 4 demonstrates that the patterns $\{\{a, b\}, \{a, c\}\}$, $\{\{a, b\}, \{b, c\}\}$, and $\{\{b, c\}, \{a, c\}\}$ are weaker than the pattern $\{\{a, b\}, \{a, c\}, \{b, c\}\}$.

Example 3 (Consensus number). An n-consensus protocol is implemented by n processes that communicate through shared objects. The processes each start with an input of either 1 or 0 and must decide a common value. A consensus protocol must ensure that

1. Consistency: all processes must decide on the same value.
2. Wait-freedom: each process must decide after a finite number of steps.
3. Validity: the common decided value was proposed by some process.

Herlihy [9] defines the consensus number of an object O as the largest n for which there is a consensus protocol for n processes that only uses finitely many instances of O and any number of atomic registers. It follows from the definition that no combination of objects with a consensus number of $k < n$ can implement an object with a consensus number of n.

We can represent the executions of a n-consensus protocol as a tree in which one process moves at a time. By validity and wait-freedom, the initial state of the protocol must be bivalent (i.e. it is possible that 0 or 1 are decided), and there must exist a state from which on all successor states are univalent. Hence, the process that moves first in such a state decides the outcome of the protocol. This state is called the critical state.

In order to show that an object has a consensus number lower than k, we derive a contradiction by assuming that there is a valid implementation of a k-consensus protocol. Next, we maneuver the protocol into a critical state and show that the processes will not be able to determine which process moved first. Therefore, for some process P, there exist two indistinguishable executions in which P decides differently. However, if the object has a consensus number of k, the processes will be able to tell who moved first.

Synergetic knowledge is able to describe the situation from the critical state onwards. We interpret an element $\{p_1, ..., p_k\}$ of a synergy pattern G as the processes p_1 up to p_k having access to objects with a consensus number of k. For each process p_i, we define a propositional variable move_i that is true if p_i

moved first at the critical state. Furthermore, we define

$$\varphi_i := \mathsf{move}_i \wedge \bigwedge_{1 \leq j \leq n \text{ and } j \neq i} \neg \mathsf{move}_j,$$

i.e., if φ_i is true, then the i-th process moved first. Let $\mathcal{M} = (\mathbb{C}, V)$ be a model, if $\mathcal{M} \Vdash [G]\varphi_1 \vee [G]\varphi_2 \vee \cdots \vee [G]\varphi_n$ holds in the model, then it is always possible for the processes in G to tell who moved first. Lastly, if G has n agents, we have for any G' with less than n agents

$$\mathcal{M} \nVdash [G']\varphi_1 \vee [G']\varphi_2 \vee \cdots \vee [G']\varphi_n,$$

which means that the access to objects with a consensus number of n is required.

For three agents a, b and c, the model $\mathcal{M} = (\mathbb{C}, V)$ is given by

$$\mathbb{C} = \left\{ \left\{ \begin{matrix} abc0 \\ ab0 \\ bc0 \\ ac0 \\ a0, b0, c0 \end{matrix} \right\}, \left\{ \begin{matrix} abc1 \\ ab0 \\ bc0 \\ ac0 \\ a0, b0, c0 \end{matrix} \right\}, \left\{ \begin{matrix} abc2 \\ ab0 \\ bc0 \\ ac0 \\ a0, b0, c0 \end{matrix} \right\} \right\}$$

with a valuation V that represents that someone moved first, i.e.

$$\mathcal{M}, \langle abc0 \rangle \Vdash \varphi_a \qquad \mathcal{M}, \langle abc1 \rangle \Vdash \varphi_b \qquad \mathcal{M}, \langle abc2 \rangle \Vdash \varphi_c.$$

It is easy to check that $\langle abc0 \rangle \sim_{ab,ac,bc} \langle abc1 \rangle$ and hence, having access to an object with consensus number 2 is not enough in order to distinguish those worlds. However,

$$\mathcal{M} \Vdash [abc]\varphi_a \vee [abc]\varphi_b \vee [abc]\varphi_c$$

is true and shows that access to objects with consensus number 3 suffices.

Example 4 (Dining cryptographers). The dining cryptographers problem, proposed by Chaum [1], illustrates how a shared-coin primitive can be used by three cryptographers (i.e. agents) to find out whether their employer or one of their peers paid for the dinner. However, if their employer did not pay, the payer wishes to remain anonymous.

For lack of space, we do not give a full formalisation of the dining cryptographers problem. Instead, we solely focus on the ability of agreeing on a coin-flip and the resulting knowledge. In what follows, we will provide a model in which the agents a, b and c can determine whether or not their employer paid if and only if they have pairwise access to a shared coin.

Let the propositional variable p denote that their employer paid. We interpret an agent pattern $G = \{\{a, b\}\}$ as a and b, having access to a shared coin. Our model $\mathcal{M} = (\mathbb{C}, V)$, depicted in Figure 4, is given by the complex

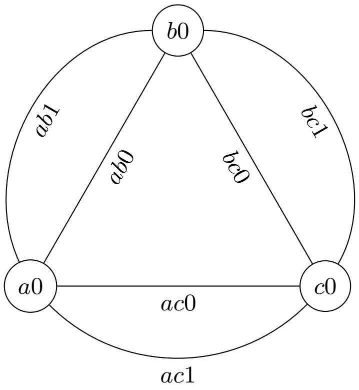

Fig. 4. Dining cryptographers model.

$$\mathbb{C} = \left\{ \begin{matrix} \left\{ \begin{matrix} abc0 \\ ab0 \\ bc0 \\ ac0 \\ a0, b0, c0 \end{matrix} \right\}, \left\{ \begin{matrix} abc1 \\ ab1 \\ bc0 \\ ac0 \\ a0, b0, c0 \end{matrix} \right\}, \left\{ \begin{matrix} abc2 \\ ab0 \\ bc1 \\ ac0 \\ a0, b0, c0 \end{matrix} \right\}, \left\{ \begin{matrix} abc3 \\ ab0 \\ bc0 \\ ac1 \\ a0, b0, c0 \end{matrix} \right\}, \\ \left\{ \begin{matrix} abc4 \\ ab1 \\ bc1 \\ ac0 \\ a0, b0, c0 \end{matrix} \right\}, \left\{ \begin{matrix} abc5 \\ ab1 \\ bc0 \\ ac1 \\ a0, b0, c0 \end{matrix} \right\}, \left\{ \begin{matrix} abc6 \\ ab0 \\ bc1 \\ ac1 \\ a0, b0, c0 \end{matrix} \right\}, \left\{ \begin{matrix} abc7 \\ ab1 \\ bc1 \\ ac1 \\ a0, b0, c0 \end{matrix} \right\} \end{matrix} \right\}$$

and the valuation V is chosen such that

$$\begin{array}{llll} p \in V(\langle abc0 \rangle), & p \notin V(\langle abc1 \rangle), & p \notin V(\langle abc2 \rangle), & p \notin V(\langle abc3 \rangle), \\ p \in V(\langle abc4 \rangle), & p \in V(\langle abc5 \rangle), & p \in V(\langle abc6 \rangle), & p \notin V(\langle abc7 \rangle). \end{array}$$

Consider the agent pattern $G = \{\{a, b\}, \{a, c\}, \{b, c\}\}$, then

$$\mathcal{M} \Vdash [G]p \vee [G]\neg p, \tag{9}$$

i.e. in any world, if all agents have pairwise access to shared coins, they can know the value of p. Furthermore, for each $H \subsetneq G$ and each $w \in \mathbb{C}$

$$\mathcal{M}, w \nVdash [H]p \vee [H]\neg p. \tag{10}$$

Notice that (10) states that there is no world, where an agent pattern H can know whether p or $\neg p$, and hence, it is stronger than $\mathcal{M} \nVdash [H]p \vee [H]\neg p$.

5 Communication

In this section, we will explore a different reading of agent patterns, namely as a description of the communication happening between the agents. Let G be the pattern $\{\{a\},\ \{b,c\}\}$. We interpret this as *b and c communicate with each other* but *there is no communication between a and b or c*. A formula $[G]\phi$ will thus be interpreted as *a knows ϕ* ***and*** *the group b,c has distributed knowledge of ϕ*. We can also distinguish the patterns $\{\{a,b\},\ \{b,c\}\}$ and $\{\{a,b\},\ \{b,c\},\ \{a,c\}\}$. In the first one, a and c can only communicate via b whereas in the second one, a and c have a direct communication channel.

Definition 7 (Connected). *Let $C \subseteq \mathsf{Pow}(\mathsf{Ag})$, we call two elements $X, Y \in C$* connected in C *if and only if there exist $Z_0, ..., Z_k \in C$ with $Z_i \cap Z_{i+1} \neq \emptyset$ for $0 \leq i < k$ and $Z_0 = X$ and $Z_k = Y$.*

Definition 8 (Connected Component). *Let $C \subseteq \mathsf{Pow}(\mathsf{Ag})$, we call C a* connected component *if and only if for any $X, Y \in C$ with $X \neq Y$ it holds that X and Y are connected in C.*

Let $G \subseteq \mathsf{Pow}(\mathsf{Ag})$. We call H a maximal connected component of G *if and only if $H \subseteq G$ and there is no connected component $H' \subseteq G$ such that H is a proper subset of H'.*

We can represent an agent pattern G as the union of its maximal connected components. Let $C_1, \ldots, C_k$ be the maximal connected components of G. We have $G = \bigcup_{i=1}^{k} C_i$ and if $X \in C_i$ and $Y \in C_j$ with $i \neq j$, then $X \cap Y = \emptyset$.

Definition 9 (Componentwise indistinguishability). *Let $G = \bigcup_{i=1}^{k} C_i$ be an agent pattern with k maximal connected components C_i. We say that G cannot distinguish componentwise two simplices S and T, denoted by $S \mathsf{E}_G T$, if and only if*

$$\exists 1 \leq j \leq k.\ S \sim_{C_j} T,$$

i.e. there is some maximal component of G that cannot distinguish S and T.

We use the notation E_G since this relation is used to model something like *every component of G knows that*. For this section, we adapt the truth definition as follows:

$$\mathcal{M}, w \Vdash [G]\phi \qquad \text{iff} \qquad w \mathsf{E}_G v \text{ implies } \mathcal{M}, v \Vdash \phi \quad \text{for all } v \in \mathbb{C}.$$

Let $G := \{\{a\} \mid a \in \mathsf{Ag}\}$. Then we can read $[G]\phi$ as everybody knows that ϕ.

We immediately obtain the following properties:

Lemma 8. *Let $G = \bigcup_{i=1}^{k} C_i$ be an agent pattern with k maximal connected components C_i. Then E_G is symmetric. Moreover, let $\mathcal{S}_G$ be a set of simplexes such that for any $S \in \mathcal{S}_G$ we have $\mathsf{noSym}(C_i) \subseteq S^\circ$ for some $1 \leq i \leq n$. Then the indistinguishability relation E_G is reflexive on $\mathcal{S}_G \times \mathcal{S}_G$.*

Note that E_G is not transitive. Also, anti-monotonicity does not hold in general. It does, however, hold componentwise.

Lemma 9 (Anti-monotonicity). *Let $G = \bigcup_{i=1}^{k} C_i$ be an agent pattern with k maximal connected components C_i. Let C be a connected component with $C \supseteq C_i$ for some $1 \leq i \leq k$ and let $H := G \cup C$. We find that $\mathsf{E}_H \subseteq \mathsf{E}_G$.*

Lemma 10 (Link). *Let $F, G, H \subseteq \mathsf{Pow}(\mathsf{Ag})$ be connected components such that $F \cup G$ is connected and $F \cup H$ is connected. The following formula is valid:*

$$[G]A \wedge [H]B \to [F \cup G \cup H](A \wedge B).$$

Proof. First, observe that $F \cup G \cup H$ is connected. Thus, by Lemma 9, $[G]A$ implies $[F \cup G \cup H]A$ and $[H]B$ implies $[F \cup G \cup H]B$. Since $[F \cup G \cup H]$ is a normal modality, we conclude $[F \cup G \cup H](A \wedge B)$. □

Example 5 (Missing link). Two networks G and H, each modelled as a connected component, both know that if malicious activity is detected, certain services must be stopped. Let mact be a propositional variable that indicates whether an intruder has been spotted and let stop indicates that the services are disabled. Since the procedure is known to both networks, we have

$$[G](\mathsf{mact} \to \mathsf{stop}) \wedge [H](\mathsf{mact} \to \mathsf{stop}) \text{ as well as } [G \cup H](\mathsf{mact} \to \mathsf{stop}).$$

Suppose now that G detects malicious activity, i.e. $[G]\mathsf{mact}$. Thus, G will stop certain services, i.e. $[G]\mathsf{stop}$. If the networks cannot communicate with each other, i.e. $G \cup H$ is not connected, then H will not stop the services. Hence, G and H as a whole are not following the security protocol, i.e. $\neg[G \cup H]\mathsf{stop}$, and might leave the system in a vulnerable state. However, if a coordinating node relays messages from G to H, then H could shut down its services as well. By Lemma 10 we find that for some network F, such that $F \cup G$ as well as $F \cup H$ is connected, it holds that

$$([G \cup H](\mathsf{mact} \to \mathsf{stop}) \wedge [G]\mathsf{mact}) \to [F \cup G \cup H]\mathsf{stop}.$$

6 Conclusion

In this paper we present a semantics for epistemic reasoning on simplicial sets and introduce the synergistic knowledge operator $[G]$. Synergistic knowledge describes relations among agents of a group and enables us to reason about what the group can know beyond traditional distributed knowledge. For example, in Example 4, the pattern $\{\{a, b\}, \{a, c\}\}$ differs from $\{\{a, b\}, \{a, c\}, \{b, c\}\}$, although both contain the same agents.

Furthermore, we develop a logic based on our model and study some of its validities. We show that classical distributed knowledge, as introduced by Halpern and Moses [8], can be expressed with the operator $[G]$, if G is a set of singleton sets.

Moreover, we provide various examples of how our logic can be used to describe problems that arise in distributed computing. In Example 2 we illustrate how to model scenarios that arise in failure detection protocols, and in Examples 3 and 4 we showcase how synergistic knowledge may occur in distributed systems, if agents access shared primitives.

Lastly, we discussed a new notion of indistinguishability that accounts for the connectivity of the agent pattern G. Componentwise indistinguishability seems fruitful for analysing knowledge in networks with respect to their underlying topology.

Acknowledgments

This work has been funded by the Swiss National Science Foundation (SNSF) under grant agreement Nr. 200021_188443 (Advanced Consensus Protocols).

References

1. Chaum, D.: The dining cryptographers problem: Unconditional sender and recipient untraceability. J. Cryptol. **1**(1), 65–75 (1988). https://doi.org/10.1007/BF00206326, https://doi.org/10.1007/BF00206326
2. van Ditmarsch, H., Goubault, É., Ledent, J., Rajsbaum, S.: Knowledge and simplicial complexes. In: Lundgren, B., Nuñez Hernández, N.A. (eds.) Philosophy of Computing. pp. 1–50. Springer International Publishing, Cham (2022)
3. van Ditmarsch, H., van der Hoek, W., Kooi, B.: Dynamic Epistemic Logic. Springer Publishing Company, Incorporated, 1st edn. (2007)
4. van Ditmarsch, H., Kuznets, R.: Wanted dead or alive : Epistemic logic for impure simplicial complexes. Journal of Logic and Computation (to appear)
5. Goubault, É., Ledent, J., Rajsbaum, S.: A simplicial complex model for dynamic epistemic logic to study distributed task computability. Inf. Comput. **278**, 104597 (2021). https://doi.org/10.1016/j.ic.2020.104597, https://doi.org/10.1016/j.ic.2020.104597
6. Goubault, É., Ledent, J., Rajsbaum, S.: A simplicial model for KB4$_n$: Epistemic logic with agents that may die. In: Berenbrink, P., Monmege, B. (eds.) 39th International Symposium on Theoretical Aspects of Computer Science, STACS 2022, March 15-18, 2022, Marseille, France (Virtual Conference). LIPIcs, vol. 219, pp. 33:1–33:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2022). https://doi.org/10.4230/LIPIcs.STACS.2022.33, https://doi.org/10.4230/LIPIcs.STACS.2022.33
7. Goubault, É.G., Kniazev, R., Ledent, J., Rajsbaum, S.: Semi-simplicial set models for distributed knowledge (2023)
8. Halpern, J.Y., Moses, Y.: Knowledge and common knowledge in a distributed environment. In: Kameda, T., Misra, J., Peters, J.G., Santoro, N. (eds.) Proceedings of the Third Annual ACM Symposium on Principles of Distributed Computing, Vancouver, B. C., Canada, August 27-29, 1984. pp. 50–61. ACM (1984). https://doi.org/10.1145/800222.806735, https://doi.org/10.1145/800222.806735
9. Herlihy, M.: Wait-free synchronization. ACM Trans. Program. Lang. Syst. **13**(1), 124–149 (1991). https://doi.org/10.1145/114005.102808, https://doi.org/10.1145/114005.102808

10. Herlihy, M., Kozlov, D.N., Rajsbaum, S.: Distributed Computing Through Combinatorial Topology. Morgan Kaufmann (2013), https://store.elsevier.com/product.jsp?isbn=9780124045781
11. Herlihy, M., Shavit, N.: The topological structure of asynchronous computability. J. ACM **46**(6), 858–923 (1999). https://doi.org/10.1145/331524.331529, https://doi.org/10.1145/331524.331529
12. Randrianomentsoa, R.F., van Ditmarsch, H., Kuznets, R.: Impure simplicial complexes: Complete axiomatization. Logical Methods in Computer Science (to appear)