Synergistic Knowledge*,**

Christian Cachin^{a}, David Lehnherr^{a}, Thomas Studer^{a}

^{a}*University of Bern, Neubrückstrasse 10, 3012 CH, Bern, Switzerland*

Abstract

In formal epistemology, group knowledge is often modelled as the knowledge that the group would have, if the agents shared 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. We interpret synergistic knowledge on simplicial models that are based on semi-simplicial sets. As examples, we investigate the use of consensus objects and the problem of the dining cryptographers. Furthermore, we introduce the axiom system Syn for synergistic knowledge and show that it is sound and complete with respect to the presented simplicial models.

Keywords: Distributed Knowledge, Synergy, Modal Logic

1. Introduction

The successful application of combinatorial topology to distributed systems has sparked the interest of modal logicians. A topological model interprets the different configurations of a distributed system as an abstract structure called a simplicial complex. Within a simplicial complex, one can express which global states are indistinguishable to processes. If a process cannot tell two configurations apart, then, for any distributed protocol, it must behave the same in both states. Naturally, this approach, exemplified by Herlihy, Kozlov, and Rajsbaum [15], has parallels with formal epistemology. Hence, studying simplicial interpretations of modal logic is a novel and promising subject.

Due to its application to well known problems such as agreement (cf. Halpern and Moses [13]), modal logic has been proven to be vital for the formal analysis of distributed Algorithms. Models for modal logic are usually based on a possible worlds approach, where the modal operator $\Box$ is evaluated on Kripke frames. A

Kripke frame is a graph whose vertices represent worlds, i.e., global states, and whose edges depict an indistinguishability relation. 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 [13], if and only if ϕ is true in all worlds that cannot be distinguished by any member of the group.

Some recent publications in logic have explored simplicial complexes as models for modal logics with multiple agents. In its simplest form, pure simplicial complexes describe settings in which agents never die and impure complexes are capable of capturing an agent's death. Goubault, Ledent, and Rajsbaum [11] pioneered the analysis of simplicial interpretations and laid the foundation for pure complexes as epistemic models. Work regarding impure complexes was conducted independently by van Ditmarsch, Kuznets, and Randrianomentsoa [7] and by Goubault, Ledent, and Rajsbaum [12]. The latter work shows the equivalence between their simplicial models and Kripke models for the logic $\mathsf{KB4}_n$, whereas van Ditmarsch, Kuznets, and Randrianomentsoa [7] propose a three-valued logic in which formulas that reference dead agents might be undefined. Van Ditmarsch, Kuzents, and Randrianomentsoa [6] further compare two-valued with three-valued semantics and categorise simplicial models accordingly. Lastly, Goubault et al. [10] generalise impure complexes based on their containment of sub-worlds, i.e., worlds whose set of alive agents is contained in another indistinguishable world's set of alive agents.

As pointed out by van Ditmarsch et al. [5], improper[1] Kripke frames cannot be depicted by simplicial complexes. However, relaxing the properties of simplicial complexes might enable us to do so and lead to novel notions of group knowledge. One example of a slightly relaxed version of a simplicial complex is a semi-simplicial set. Simply put, semi-simplicial sets allow us to model global states that consist of the same local states, but differ in their meaning. These kind of scenarios arise in distributed systems when accessing concurrent objects. For example, consider two processes P and Q that concurrently enqueue their values v_P and v_Q to a shared queue. Upon completion, each process locally knows that the queue contains its submitted value. However, they do not know whose value was enqueued first, i.e., P and Q cannot distinguish the worlds where the state of the queue is $v_P v_Q$ or $v_Q v_P$. The complex below in Figure 1 illustrates this scenario. The local state of a process is represented by a vertex that is labelled accordingly and global states correspond to edges. The processes cannot distinguish between the global states because their local states can be part of both. That is, v_Q and v_P belong to both edges. However, we must not restrict the notion of indistinguishability to containing vertices alone. Instead, we can look at higher-order connectivity as well. By doing so, we introduce a new kind of knowledge that has not yet been explored.

---

[1]A frame is proper if any two worlds with the same set of alive agents can be distinguished by at least one of them.
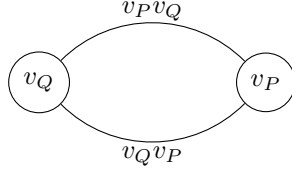
Figure 1: The two processes P and Q cannot distinguish between worlds $v_P v_Q$ and $v_Q v_P$.

In this paper, we present a semantics for epistemic reasoning on simplicial models that are based on semi-simplicial sets and introduce the synergistic knowledge operator $[G]$. Our notion of synergistic knowledge is capable of describing scenarios in which a group of agents can know more than just the consequences of their pooled knowledge. Different from the higher-order interpretation of distributed knowledge by Goubault et al. [9], which analyses the knowledge of a set of agents, we unfold relations within a group and analyse its knowledge with respect to the internal structure of the group. With this interpretation, the knowledge of two seemingly equal groups can differ due to internal relations between agents. Limiting oneself to sets of agents cannot model this without introducing hybrid agents that represent relations between agents.

Given a finite set of agents Ag, we introduce a new modality $[G]$, where the set $G \subseteq \mathsf{Pow}(\mathsf{Ag}) \setminus \{\emptyset\}$ is called an agent pattern. Agent patterns are best understood representing relations within a group. For example, if $G = \{\{a, b, c\}\}$, then, the agents a, b, and c mutually benefit from each other and $[G]$ reasons about the synergy of the three agents; whereas if $G = \{\{a\}, \{b\}, \{c\}\}$, the modality $[G]$ describes traditional distributed knowledge. When interpreting the language of synergistic knowledge on simplicial models based on semi-simplicial sets, $[G]\phi$ is true in a simplex S if and only if ϕ is true in all simplices that intersect with S in G. For example, in Figure 1, the agent pattern $\{\{P, Q\}\}$ can distinguish between $v_P v_Q$ and $v_Q v_P$ because the two edges are different.

Another contribution of this work is the axiom system Syn. We show that Syn is sound and complete with respect to our simplicial models. Our proof employs a variant of the unravelling method used to establish completeness for traditional distributed knowledge (cf. Fagin et al. [8] and Goubault et al. [9]). However, our logic has new axioms that close knowledge under sub-simplices. For example, we cannot only rely on the usual monotonicity axiom that states $[G]\varphi \to [H]\varphi$, whenever $G \subseteq H$. Instead, we need that

$$[\{B\}]\phi \to [G]\phi \quad \text{if there exists } A \in G \text{ and } B \subseteq A. \tag{1}$$

Most notably, (1) can be proved by adding an axiom enforcing that adding agent subpatterns does not create new insights, i.e.,

$$[G \cup \{B\}]\phi \to [G]\phi \quad \text{if there exists } A \in G \text{ and } B \subseteq A.$$

3

After introducing the syntax and semantics of Syn, we proceed by proving sound-ness and completeness of Syn with respect to simplicial models. Soundness can be proven right away, but completeness is more involved. To prove complete-ness, we first introduce different types of Kripke models and show that a weaker version of Syn, namely Syn⁻, is sound and complete with respect to those mod-els. After that, we show that our proofs can be extended to Kripke models that are equivalent to simplicial models which are based on semi-simplicial sets. The difference between Syn and Syn⁻ is that Syn characterises proper models.

We first introduce κ-models, which are similar to the standard $\mathsf{KB4}_n$ Kripke models. It is worth noting that, κ-models are pseudo-models because they do not satisfy standard group knowledge. That is, given two agent patterns G and H, if both cannot distinguish two worlds, neither can their union. However, not satisfying standard group knowledge lets us prove soundness and complete-ness of Syn⁻ with respect to κ-models using a canonical model construction. Subsequently, we present δ-models, which are κ-models that satisfy standard group knowledge. We use unravelling of κ-models to show that Syn⁻ is sound and complete with respect to δ-models. In a last step, we extend the previous proofs to proper δ-models which shows soundness and completeness of Syn with respect to them. In order to show completeness of Syn with respect to simplicial models, we develop a novel construction that takes a proper δ-frame as input and transforms it into an equivalent complex. This implies completeness of Syn with respect to our simplicial models. The constructive nature of the translation intuitively demonstrates the relation between the two semantic approaches.

Section 2 presents our notion of a simplicial complex together with an in-distinguishability relation for synergistic knowledge. In Section 3 we introduce the axiom system Syn and show its soundness with respect to the earlier defined simplicial models. Section 4 accompanies the presented theory with examples from distributed computing. In Section 5 we prove that Syn is complete with respect to simplicial models. Lastly, in Section 6 and Section 7 we conclude our work and discuss future research directions.

2. Simplicial structures

In this section, we introduce simplicial structures based on semi-simplicial sets. Let Ag denote the set of finitely many agents and let

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

Let $S \subseteq \mathsf{Agsi}$. An element $(A, i) \in S$ is *maximal in S* if and only if

$\forall (B, j) \in S.|A| \geq |B|$, where $|X|$ 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*

S*1: The maximal element is unique, i.e.,*

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

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

S*2: S is uniquely downwards closed, i.e., for $(B, i) \in S$ and $C \subseteq B$*

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

S*3: S contains nothing else, i.e.,*

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

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

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

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

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

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

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

Whenever it is clear from the context, we abbreviate $(\{a_1, ..., a_n\}, i)$ as $a_1...a_n i$. Furthermore, we omit elements of the form $(\emptyset, i)$ and may use a row (or a mixed row-column) notation to emphasize simplices. For example,

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

is a complex that contains 2 simplices. Whenever we refer to a simplex within a complex, we write $\langle a_1, .., a_n i \rangle$ for the simplex with maximal element $(\{a_1, ..., a_n\}, i)$. Condition C guarantees that this notation is well-defined.

Remark 1. *Complexes can contain sub-simplices. That is, if $S \in \mathbb{C}$, then there may exist $T \in \mathbb{C}$ with $S \subseteq T$. An example of such a complex is*

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

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

We will now prove some properties of $\sim_G$ that motivate it as an indistinguishability relation between simplices.

Lemma 2. *$\sim_G$ is transitive and symmetric.*

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

$$G \subseteq (S \cap T)^\circ \tag{2}$$
$$G \subseteq (T \cap U)^\circ \tag{3}$$

Let $A \in G$. Because of (2), there exists i with

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

Because of (3), there exists j with

$$(A, j) \in T \quad \text{and} \quad (A, j) \in U. \tag{5}$$

From (4), (5), and Condition S2 we obtain $i = j$. Thus by (4) and (5), we get $A \in (S \cap U)^\circ$. Since A was arbitrary in G, we conclude $G \subseteq (S \cap U)^\circ$. $\qquad \square$

Lemma 3. *Let G be an agent pattern and*

$$G^\star := \{\{a\} \mid \exists A \in G \text{ and } a \in A\}.$$

Let $\mathcal{S}_G$ be a maximal set of simplices such that for any $S \in \mathcal{S}_G$ we have $G^\star \subseteq S^\circ$. 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

$$\text{there exists } i \text{ with } (B, i) \in S. \tag{6}$$

Let $(A, i) := \max(S)$. Let $b \in B$. Because of $G^\star \subseteq S^\circ$, there exists l such that $(\{b\}, l) \in S$. By S3 we get $b \in A$. Since b was arbitrary in B, we get $B \subseteq A$. By S2 we conclude that (6) holds and symmetry is established.

We now show that $\sim_G$ is empty otherwise. Let S be a simplex such that $G^\star \not\subseteq S^\circ$ 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.,

$$\text{for all } i, (\{a\}, i) \notin S. \tag{7}$$

Suppose towards a contradiction that

$$G \subseteq (S \cap T)^\circ \tag{8}$$

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 $(\{a\}, j) \in S$. This is a contradiction to (7). Thus (8) cannot hold. $\qquad \square$

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

Lemma 4. *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 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 $B \in (S \cap T)^{\circ}$. $\qquad\square$

Corollary 2. *Let G be an agent pattern and let $A, B \subseteq \mathsf{Ag}$ such that $B \subseteq A \in G$. It holds that*

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

Lemma 5. *Let G be an agent pattern. It holds that*

$$\bigcap_{B \in G} \sim_{\{B\}} = \sim_G . \tag{D}$$

Proof. $(S, T) \in \bigcap_{B \in G} \sim_{\{B\}}$ iff for each $B \in G$, we have $B \in (S \cap T)^{\circ}$ iff $G \subseteq (S \cap T)^{\circ}$ iff $S \sim_G T$. $\qquad\square$

Remark 2 (Standard group knowledge). *The property* (D) *is also referred to as standard group knowledge.*

Lemma 6 (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 3 and hence $S \sim_G T$. $\qquad\square$

The next lemma states that adding synergy to an agent pattern makes it stronger in the sense that it can distinguish more simplices.

Lemma 7. *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 Lemma 4 and Lemma 6 we find that

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

$\square$

In traditional Kripke semantics, distributed knowledge of a set of agents is modelled by considering the indistinguishability relation that is given by the intersection of the indistinguishability 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 8. *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 $S \sim_H T$. $\qquad\square$

3. Logic

Let Prop be a countable set of atomic propositions. The logic of synergistic knowledge is a normal modal logic that includes a modality $[G]$ for each agent pattern G. Formulas of the language of synergistic knowledge $\mathcal{L}$ are inductively defined by the following grammar:

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

where $p \in$ 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 $p \in$ Prop, and we write alive(G) for $\neg[G]\bot$. If G is an agent pattern, then G^C denotes its complement. That is,

$$G^C := \{H \in \mathsf{Pow}(\mathsf{Ag}) \setminus \{\emptyset\} \mid \nexists B \in G.H \subseteq B\}.$$

Moreover, we define

$$\mathsf{dead}(G) := \bigwedge_{B \in G} \neg\mathsf{alive}(\{B\}).$$

Notice that $\mathsf{dead}(G) \not\equiv \neg\mathsf{alive}(G)$. Indeed, $\mathsf{dead}(G)$ expresses that for each $B \in G$, the pattern $\{B\}$ is dead, whereas $\neg\mathsf{alive}(G)$ is true if some $\{B\} \subseteq G$ is dead. The axiom system Syn consists of the axioms:

$$
\begin{array}{ll}
\text{all propositional tautologies} & \text{(Taut)} \\
[G](\phi \rightarrow \psi) \rightarrow ([G]\phi \rightarrow [G]\psi) & \text{(K)} \\
\phi \rightarrow [G]\neg[G]\neg\phi & \text{(B)} \\
[G]\phi \rightarrow [G][G]\phi & \text{(4)} \\
\mathsf{alive}(G) \rightarrow ([G]\phi \rightarrow \phi) & \text{(T)} \\
\mathsf{alive}(G) \wedge \mathsf{dead}(G^C) \wedge \phi \rightarrow [G](\mathsf{dead}(G^C) \rightarrow \phi) & \text{(P)} \\
\displaystyle\bigvee_{G \subseteq \mathsf{Pow}(\mathsf{Ag})\setminus\{\emptyset\}} \mathsf{alive}(G) & \text{(NE)} \\
[G]\phi \rightarrow [H]\phi \quad \text{if } G \subseteq H & \text{(Mono)} \\
[G \cup \{B\}]\phi \rightarrow [G]\phi \quad \text{if there exists } A \in G \text{ and } B \subseteq A & \text{(Equiv)} \\
\mathsf{alive}(G) \wedge \mathsf{alive}(H) \rightarrow \mathsf{alive}(G \cup H) & \text{(Union)} \\
\mathsf{alive}(G) \rightarrow \mathsf{alive}(\{A \cup B\}) \quad \text{if } A, B \in G & \text{(Clo)}
\end{array}
$$

and the inference rules modus ponens (MP) and $[G]$-necessitation ($[G]$-Nec). We write $\vdash \varphi$ to denote that $\varphi \in \mathcal{L}$ can be deduced in the system Syn.

$$\frac{A \quad A \rightarrow B}{B} \quad \text{(MP)} \qquad \frac{A}{[G]A} \quad ([G]\text{-Nec})$$

Notice that

$$\mathsf{alive}(G) \wedge \mathsf{dead}(G^C) \wedge \phi \rightarrow [G]\phi, \tag{9}$$

is not derivable in Syn, because, as mentioned in Remark 1, complexes can contain sub-worlds. In order to derive Equation (9) we would have to add

$$\mathsf{alive}(G) \wedge \mathsf{dead}(G^C) \rightarrow [G]\mathsf{dead}(G^C)$$

to Syn (cf. Goubault et al. [10]). However, if $\mathsf{Ag} \in G$, we have $G^C = \emptyset$, i.e.,

$$\mathsf{alive}(G) \wedge \phi \rightarrow [G]\phi$$

because the empty conjunction evaluates to $\top$.

Definition 4 (Simplicial model). *A simplicial model $\mathcal{C} = (\mathbb{C}, L)$ is a pair such that*

1. *$\mathbb{C}$ is a complex;*

2. *$L : \mathbb{C} \rightarrow \mathsf{Pow}(\mathsf{Prop})$ is a valuation.*

Definition 5 (Truth). *Let $\mathcal{C} = (\mathbb{C}, L)$ be a simplicial model, $S \in \mathbb{C}$, and $\phi \in \mathcal{L}$. We define $\mathcal{C}, S \Vdash_\sigma \phi$ inductively by*

$$
\begin{array}{llll}
\mathcal{C}, S \Vdash_\sigma p & \textit{iff} & p \in L(S) \\
\mathcal{C}, S \Vdash_\sigma \neg\phi & \textit{iff} & \mathcal{C}, S \nVdash_\sigma \phi \\
\mathcal{C}, S \Vdash_\sigma \phi \wedge \psi & \textit{iff} & \mathcal{C}, S \Vdash_\sigma \phi \textit{ and } \mathcal{C}, S \Vdash_\sigma \psi \\
\mathcal{C}, S \Vdash_\sigma [G]\phi & \textit{iff} & S \sim_G T \textit{ implies } \mathcal{C}, T \Vdash_\sigma \phi \quad \textit{for all } T \in \mathbb{C}.
\end{array}
$$

We write $\mathcal{C} \Vdash_\sigma \phi$, if $\mathcal{C}, w \Vdash_\sigma \phi$ for all $S \in \mathbb{C}$. A formula ϕ is σ-valid, denoted by $\Vdash_\sigma \phi$, if $\mathcal{C} \Vdash_\sigma \phi$ for all models $\mathcal{C}$. Whenever it is clear from the context we omit the subscript σ.

Corollary 3 is an immediate consequence of previously introduced Corollary 1. It relates the formula $\mathsf{alive}(G)$ to the structure of the underlying complex as expected.

Corollary 3. *Let $\mathcal{C}$ be a simplicial model. For any agent pattern G, we find*

$$\mathcal{C}, S \Vdash_\sigma \mathsf{alive}(G) \quad \textit{iff} \quad S \sim_G S.$$

Soundness of Syn with respect to simplicial models follows as usual. We present the proof of completeness in Section 5.

Theorem 1 (Soundness). *$\vdash \varphi$ implies $\Vdash_\sigma \varphi$.*

Proof. We only show (T), (P), (NE), (Union), (Clo), (Mono), (Equiv), and ([G]-Nec). Let $\mathcal{C} = (\mathbb{C}, L)$ be an arbitrary model.

1. (T): Consider a world $S \in \mathbb{C}$ and assume that $\mathcal{C}, S \Vdash_\sigma \mathsf{alive}(G)$ and $\mathcal{C}, S \Vdash_\sigma [G]\varphi$. By Corollary 3, we find $S \sim_G S$ and thus $\mathcal{C}, S \Vdash_\sigma \varphi$ because $\mathcal{C}, S \Vdash_\sigma [G]\varphi$.

9

2. (P): Consider a world $S \in \mathbb{C}$ and assume $\mathcal{C}, S \Vdash \mathsf{alive}(G) \wedge \mathsf{dead}(G^C) \wedge \phi$ for some $\phi \in \mathcal{L}$. By assumption, G must have an unique maximal element. Indeed, towards a contradiction, assume that it is not the case and G has the maximal elements $B_1, B_2, ..., B_n$. Consider now

$$B = \bigcup_{i=1}^{n} B_i.$$

It is straightforward to verify that $\mathcal{C}, S \Vdash \mathsf{alive}(\{B\})$. However, since $B \in G^C$, by assumption, it holds that $\mathcal{C}, S \Vdash \mathsf{dead}(\{B\})$ which is a contradiction. Additionally, by S3, the maximal element of G, say A, is the set of all agents alive in S. Let $T \in \mathbb{C}$ be such that $S \sim_G T$ and assume $\mathcal{C}, T \Vdash \mathsf{dead}(G^C)$. By transitivity of $\sim_G$ and Corollary 3, we have that $\mathcal{C}, T \Vdash \mathsf{alive}(G)$. Thus, by the same reasoning as before, A is the set of all alive agents for T as well, i.e., $\max(S) = (A, i)$ and $\max(T) = (A, i)$ for some $i \in \mathbb{N}$. Finally, by Lemma 1, we find that $S = T$, and thus $\mathcal{C}, T \Vdash \phi$.

3. (NE): Follows because simplices are not empty.

4. (Union): Consider a world $S \in \mathbb{C}$ and assume that $\mathcal{C}, S \Vdash_\sigma \mathsf{alive}(G)$ and $\mathcal{C}, S \Vdash_\sigma \mathsf{alive}(H)$. By Corollary 3 it holds that $S \sim_G S$ and $S \sim_H S$. By Lemma 5 we have

$$\sim_{G \cup H} = \bigcap_{B \in G \cup H} \sim_{\{B\}} = \left(\bigcap_{B \in G} \sim_{\{B\}} \right) \cap \left(\bigcap_{B \in H} \sim_{\{B\}} \right),$$

and thus $S \sim_{G \cup H} S$ and $\mathcal{C}, S \Vdash_\sigma \mathsf{alive}(G \cup H)$.

5. (Clo): Consider a world $S \in \mathbb{C}$ and assume that $\mathcal{C}, S \Vdash_\sigma \mathsf{alive}(G)$, and let $A, B \in G$. By Lemma 4 we find $S \sim_{\{A\}} S$ and $S \sim_{\{B\}} S$, i.e., there exist $i, j \in \mathbb{N}$ such that $(A, i) \in S$ and $(B, j) \in S$. Furthermore, let $(C, k) = \max(S)$. By S3, we find $A \subseteq C$ as well as $B \subseteq C$, and thus $A \cup B \subseteq C$. Since S is downwards closed by S2, there exists $k \in \mathbb{N}$ such that $(\{A \cup B\}, k) \in S$. Hence, $S \sim_{\{A \cup B\}} S$ and $\mathcal{C}, S \Vdash_\sigma \mathsf{alive}(\{A \cup B\})$ by Lemma 11.

6. (Mono): Follows from Lemma 6.

7. (Equiv): Follows from Corollary 2.

Lastly, we show ([G]-Nec). Let $A \in \mathcal{L}$ and assume that A is σ-valid, i.e., for any simplicial model $\mathcal{C} = (\mathbb{C}, L)$ and $S \in \mathbb{C}$, we have $\mathcal{C}, S \Vdash_\sigma A$. Let $S \in \mathbb{C}$ be arbitrary. By assumption, it holds that for all T with $S \sim_G T$, we have $\mathcal{C}, T \Vdash_\sigma A$. Thus, $\mathcal{C}, S \Vdash_\sigma [G]A$ by the definition of truth. Since S was arbitrary, we have $\mathcal{C} \Vdash_\sigma [G]A$. Lastly, since $\mathcal{C}$ was arbitrary, $[G]A$ is σ-valid, i.e., $\Vdash_\sigma [G]A$. $\square$

4. Examples

This section illustrates the application of our logic to distributed systems with the help of two examples. We interpret synergy as having access to shared primitives in both examples. In Example 1, the shared primitive is a consensus object and in Example 2 it is a shared coin. Given three agents, Example 1 captures the idea that for some applications, the agent pattern must include the area of the triangle and not just its edges. Thus, Example 1 shows the difference between mutual and pairwise synergy. A similar example can be found in the extended abstract of this work [1]. Example 2 demonstrates that the patterns $\{\{a,b\},\{a,c\}\}$, $\{\{a,b\},\{b,c\}\}$, and $\{\{b,c\},\{a,c\}\}$ are weaker than the pattern $\{\{a,b\},\{a,c\},\{b,c\}\}$. In other words, pairwise synergy between all agents is stronger than pairwise synergy among some agents.

Regarding notation, from now on 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 (Consensus number). *A 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 [14] 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. Such a state is called a critical state.

In order to show that an object has a consensus number strictly 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

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

i.e., if φ_i is true, then the i-th process moved first. Let $\mathcal{C} = (\mathbb{C}, L)$ be a model, if $\mathcal{C} \Vdash_\sigma [G]\varphi_1 \vee [G]\varphi_2 \vee \cdots \vee [G]\varphi_n$ 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

$$\mathcal{C} \not\Vdash_\sigma [G']\varphi_1 \vee [G']\varphi_2 \vee \cdots \vee [G']\varphi_n,$$

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

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

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

with a valuation L representing that someone moved first, i.e.,

$$\mathcal{C}, \langle abc0 \rangle \Vdash_\sigma \varphi_a \qquad \mathcal{C}, \langle abc1 \rangle \Vdash_\sigma \varphi_b \qquad \mathcal{C}, \langle abc2 \rangle \Vdash_\sigma \varphi_c.$$

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,

$$\mathcal{C} \Vdash_\sigma [abc]\varphi_a \vee [abc]\varphi_b \vee [abc]\varphi_c$$

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

Example 2 (Dining cryptographers). *The dining cryptographers problem, proposed by Chaum [4], 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 the sake 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 $\mathcal{C} = (\mathbb{C}, L)$, depicted in Figure 2, is given by the complex
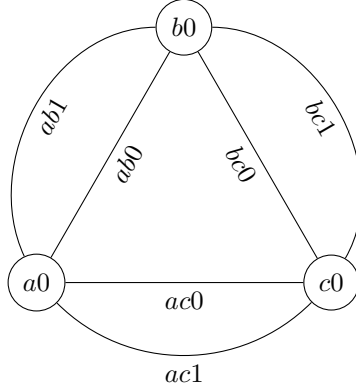
Figure 2: Dining cryptographers model.

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

and the valuation L is chosen such that

$$p \in L(\langle abc0 \rangle), \qquad p \notin L(\langle abc1 \rangle), \qquad p \notin L(\langle abc2 \rangle), \qquad p \notin L(\langle abc3 \rangle),$$
$$p \in L(\langle abc4 \rangle), \qquad p \in L(\langle abc5 \rangle), \qquad p \in L(\langle abc6 \rangle), \qquad p \notin L(\langle abc7 \rangle).$$

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

$$\mathcal{C} \Vdash_\sigma [G]p \vee [G]\neg p, \tag{10}$$

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 $S \in \mathbb{C}$

$$\mathcal{C}, S \not\Vdash_\sigma [H]p \vee [H]\neg p. \tag{11}$$

Notice that (11) states that there is no world, where an agent pattern H can know whether p or $\neg p$, and hence, it is stronger than $\mathcal{C} \not\Vdash_\sigma [H]p \vee [H]\neg p$.

13

5. Completeness of Syn

In order to show completeness of Syn with respect to simplicial models, we take a detour via Kripke-models and focus on Syn without the axiom (P), which is hereafter denoted by Syn⁻. We start by introducing κ-models which represent common multi-agent models in which (D) is not necessarily satisfied. Thus, we can employ the standard techniques to show that Syn⁻ is sound and complete with respect to κ-models. Next, we present δ-models, which are κ-models that satisfy (D). The proof that Syn⁻ is complete with respect to δ-models is more challenging and requires the so-called unravelling method. Once we showed soundness and completeness of Syn⁻ with respect to those models, we show that Syn is sound and complete with respect to their proper[2] versions. Lastly, we proceed with proving completeness with respect to simplicial models by relating them to proper δ-models.

5.1. κ-models

In this section, we introduce κ-models and show that Syn⁻ is sound and complete with respect to κ-models. Definition 6 introduces pre-models, which are our most simple models. All subsequent models are pre-models.

Definition 6 (Pre-model). *A pre-model $\mathcal{M} = (W, \sim, V)$ is a tuple where*

1. *W is a set of possible worlds;*

2. *$\sim$ is a function that assigns to each agent pattern G a symmetric and transitive relation $\sim_G$ on W;*

3. *$V : W \to \mathsf{Pow}(\mathsf{Prop})$ is a valuation.*

Remark 3 (Notation). *In order to avoid confusion due to the overloading of $\sim_G$, we use lower-case letters for worlds of a pre-model, i.e., $w \sim_G v$, and capital letters for simplices, i.e., $S \sim_G T$.*

Definition 7 (Truth). *Let $\mathcal{M} = (W, \sim, V)$ be a pre-model, $w \in W$, and $\phi \in \mathcal{L}$. We define $\mathcal{M}, w \Vdash \phi$ inductively by*

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

We write $\mathcal{M} \Vdash \phi$, if $\mathcal{M}, w \Vdash \phi$ for all $w \in W$.

Definition 8. *Let $\mathcal{M} = (W, \sim, V)$ be a pre-model, we define*

$$
\mathsf{Alive}(G)_{\mathcal{M}} = \{w \mid w \sim_G w\}.
$$

---

[2]Properness will be defined later.

If the pre-model is clear from the context, we omit the subscript $\mathcal{M}$ and write $\mathsf{Alive}(G)$ instead of $\mathsf{Alive}(G)_{\mathcal{M}}$.

Lemma 9. *Let $\mathcal{M} = (W, \sim, V)$ be a pre-model. It holds that*

$$\mathcal{M}, w \Vdash \mathsf{alive}(G) \quad iff \quad w \in \mathsf{Alive}(G)_{\mathcal{M}}.$$

Proof. We first show that

$$\mathcal{M}, w \Vdash \mathsf{alive}(G) \quad \text{implies} \quad w \in \mathsf{Alive}(G)_{\mathcal{M}}.$$

Assume $\mathcal{M}, w \Vdash \mathsf{alive}(G)$, i.e., $\mathcal{M}, w \not\Vdash [G]\bot$. By the definition of truth, there must exist $v \in W$ with $w \sim_G v$ and $\mathcal{M}, v \Vdash \top$. By symmetry, we have $v \sim_G w$ and by transitivity we have $w \sim_G w$. Hence, $w \in \mathsf{Alive}(G)_{\mathcal{M}}$. We now show that

$$w \in \mathsf{Alive}(G)_{\mathcal{M}} \quad \text{implies} \quad \mathcal{M}, w \Vdash \mathsf{alive}(G).$$

Assume $w \in \mathsf{Alive}(G)_{\mathcal{M}}$. By definition of $\mathsf{Alive}(G)_{\mathcal{M}}$, we have that $w \sim_G w$. Therefore, $\mathcal{M}, w \not\Vdash [G]\bot$ by the definition of truth. $\square$

Definition 9 (κ-model)**.** *Let $\mathcal{M} = (W, \sim, V)$ be a pre-model. $\mathcal{M}$ is called a κ-model if and only if, for all agent patterns G and H:*

K1*: $\mathsf{Alive}(G)_{\mathcal{M}} \cap \mathsf{Alive}(H)_{\mathcal{M}} \subseteq \mathsf{Alive}(G \cup H)_{\mathcal{M}}$;*

K2*: $\mathsf{Alive}(G)_{\mathcal{M}} \subseteq \mathsf{Alive}(\{A \cup B\})_{\mathcal{M}}$ for $A, B \in G$;*

K3*: $\sim_H \subseteq \sim_G$, if $G \subseteq H$;*

K4*: $\sim_G \subseteq \sim_{G \cup \{B\}}$ if there exists $A \in G$ with $B \subseteq A$;*

NE*: for all $w \in W$, there exists an agent pattern G such that $w \sim_G w$.*

A formula ϕ is κ-valid, denoted by $\Vdash_\kappa \phi$, if $\mathcal{M} \Vdash \phi$ for all κ-models $\mathcal{M}$.

Lemma 10. *Let $\mathcal{M} = (W, \sim, V)$ be a κ-model. Furthermore, let G be an agent pattern such that there exists $A \in G$ with $B \subseteq A$. It holds that $\sim_G \subseteq \sim_{\{B\}}$.*

Proof. By K3 and K4 we have $\sim_G \subseteq \sim_{G \cup \{B\}} \subseteq \sim_{\{B\}}$. $\square$

Remark 4. *Let $\mathcal{M} = (W, \sim, V)$ be a κ-model and let G be an agent pattern as in Lemma 10. By K3, we have $\sim_G = \sim_{G \cup \{B\}}$.*

The property K1 ensures that for each world, there exists a maximal alive agent pattern and K2 forces $\mathsf{alive}(G)$ to be downwards closed. K3 guarantees that an agent pattern G cannot know more than its supersets, and K4 states that adding subpatterns to G does not strengthen its knowledge. Moreover, NE ensures that there are no empty-worlds, i.e., worlds in which no agent is alive. However as shown below in Example 3, κ-models do not necessarily satisfy (D), i.e.,
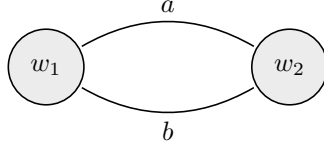
$$\bigcap_{B \in G} \sim_{\{B\}} = \sim_G .$$

Figure 3: The κ-model used in Example 3 does not satisfy standard group knowledge. Reflexive arrows are implicit.

Example 3. *Let* $\mathsf{Ag} = \{a, b\}$ *and consider the κ-model $\mathcal{M}$ shown in Figure 3, where reflexive arrows are implicit. The agents a and b cannot distinguish between the worlds w_1 and w_2, i.e., $(w_1, w_2) \in \sim_{\{a\}}$ and $(w_1, w_2) \in \sim_{\{b\}}$. However, together they are able to tell the two worlds apart. That is, $(w_1, w_2) \notin \sim_G$ for any agent pattern G that contains the set $\{\{a\}, \{b\}\}$. Hence, $\mathcal{M}$ does not satisfy standard group knowledge.*

Lastly, notice, that K3 and K4 affect the formula $\mathsf{alive}(G)$ as well. For example, the following formula is κ-valid

$$\mathsf{alive}(G) \to \mathsf{alive}(\{B\}) \quad \text{if there exists } A \in G \text{ and } B \subseteq A.$$

Indeed, let $\mathcal{M} = (W, \sim, V)$ be a κ-model and assume $\mathcal{M}, w \Vdash \mathsf{alive}(G)$. By definition, there exists $v \in W$ with $w \sim_G v$ and $\mathcal{M}, v \Vdash \top$. By symmetry and transitivity of $\sim_G$ we obtain $w \sim_G w$. By Lemma 10 we have that $(w, w) \in \sim_{\{B\}}$ and thus $\mathcal{M}, w \Vdash \mathsf{alive}(\{B\})$.

Theorem 2 (Soundness). Syn^- *is sound with respect to κ-models.*

Proof. We only show (T), (NE), (Union), (Clo), (Mono), (Equiv) and ([G]-Nec).

1. (T): Consider a world $w \in W$ and assume that $\mathcal{M}, w \Vdash \mathsf{alive}(G)$ and $\mathcal{M}, w \Vdash [G]\phi$. By Lemma 9 we have $w \in \mathsf{Alive}(G)$, i.e., $w \sim_G w$. By the definition of truth we have $\mathcal{M}, w \Vdash \phi$.

2. (NE): Let $w \in W$ be arbitrary. By NE, there exists an agent pattern G such that $w \sim_G w$. By Lemma 9, $\mathcal{M}, w \Vdash \mathsf{alive}(G)$ and thus

$$\mathcal{M}, w \Vdash \bigvee_{G \subseteq \mathsf{Pow}(\mathsf{Ag}) \setminus \{\emptyset\}} \mathsf{alive}(G).$$

3. (Union): Assume $\mathcal{M}, w \Vdash \mathsf{alive}(G)$ and $\mathcal{M}, w \Vdash \mathsf{alive}(H)$. By Lemma 9, $w \in \mathsf{Alive}(G) \cap \mathsf{Alive}(H)$. By K1, $w \in \mathsf{Alive}(G \cup H)$ and by Lemma 9 $\mathcal{M}, w \Vdash \mathsf{alive}(G \cup H)$.

4. (Clo): Assume $\mathcal{M}, w \Vdash \mathsf{alive}(G)$ and let $A, B \in G$. By Lemma 9, we have $w \in \mathsf{Alive}(G)$. By K2, $w \in \mathsf{Alive}(\{A \cup B\})$ and by Lemma 9 we obtain $\mathcal{M}, w \Vdash \mathsf{alive}(\{A \cup B\})$.

5. (Mono): Assume that $G \subseteq H$ for arbitrary G, H. Let $w \in W$ be arbitrary such that $\mathcal{M}, w \Vdash [G]\phi$. By the definition of truth, $w \sim_G v$ implies $\mathcal{M}, v \Vdash \phi$ for all $v \in W$. By K3 we have that $\sim_H \subseteq \sim_G$, i.e., $w \sim_H v$ implies $w \sim_G v$. Thus, $\mathcal{M}, v \Vdash \phi$ whenever $w \sim_H v$. Therefore, it follows that $\mathcal{M}, w \Vdash [H]\phi$.

6. (Equiv): Assume that for an arbitrary G, there exists $A \in G$ with $B \subseteq A$. Further, let $w \in W$ be arbitrary such that $\mathcal{M}, w \Vdash [G \cup \{B\}]\phi$. Therefore, for all $v \in W$, $w \sim_{G \cup \{B\}} v$ implies $\mathcal{M}, v \Vdash \phi$ by assumption. By K4 we have that $\sim_G \subseteq \sim_{G \cup \{B\}}$ and thus $w \sim_G v$ implies $w \sim_{G \cup \{B\}} v$. Hence, $\mathcal{M}, v \Vdash \phi$ whenever $w \sim_G v$ and thus $\mathcal{M}, w \Vdash [G]\phi$.

Lastly, we show (G-Nec). Let $A \in \mathcal{L}$ and assume $\Vdash_\kappa A$. We need to show that $[G]A$ is κ-valid. Let $\mathcal{M} = (W, \sim, V)$ be an arbitrary κ-model. By assumption, $\mathcal{M}, w \Vdash A$ for all $w \in W$. Thus, for any $v \in W$ with $w \sim_G v$, it holds that $\mathcal{M}, v \Vdash A$. By the definition of truth, $\mathcal{M}, w \Vdash [G]A$, and since $w \in W$ was arbitrary, $\mathcal{M} \Vdash [G]A$. Moreover, due to $\mathcal{M}$ being arbitrary, $[G]A$ is κ-valid. $\square$

In what follows, we set up the machinery to prove completeness of $\mathcal{L}$ with respect to κ-models.

Definition 10. *Let G be an agent pattern and let $\Gamma \subseteq \mathcal{L}$, we define*

$$\Gamma \setminus [G] := \{\phi \mid [G]\phi \in \Gamma\}.$$

Definition 11 (Consistent set)**.** *A set $\Gamma \subseteq \mathcal{L}$ is consistent if and only if $\Gamma \nvdash \bot$. Γ is maximal consistent if none of its proper supersets is consistent.*

Definition 12 (Canonical Model)**.** *The canonical model $\mathcal{M}^c = (W^c, \sim^c, V^c)$ for $\mathcal{L}$ is defined as*

1. *$W^c := \{\Gamma \subseteq \mathcal{L} \mid \Gamma$ is a maximal consistent set$\}$ is the set of possible worlds;*

2. *$\sim^c$ is a function that assigns to each agent pattern G a relation*

$$\sim_G^c := \{(\Gamma, \Delta) \in W^c \times W^c \mid \Gamma \setminus [G] \subseteq \Delta\};$$

3. *$V^c : \mathsf{Prop} \to \mathsf{Pow}(W^c)$ is a function defined by*

$$V^c(p) := \{\Gamma \in W^c \mid p \in \Gamma\}.$$

Lemma 11. *Let $\mathcal{M}^c = (W^c, \sim^c, V^c)$ be the canonical model, G be an agent pattern, and $\Gamma \in W^c$, then*

$$\Gamma \in \mathsf{Alive}(G)_{\mathcal{M}^c} \quad \textit{iff} \quad \mathsf{alive}(G) \in \Gamma.$$

Proof. We first show that

$$\Gamma \in \mathsf{Alive}(G)_{\mathcal{M}^c} \quad \text{implies} \quad \mathsf{alive}(G) \in \Gamma.$$

Assume $\Gamma \in \mathsf{Alive}(G)_{\mathcal{M}^c}$. By Definition 8, $\Gamma \sim_G^c \Gamma$. Towards a contradiction, assume that $\mathsf{alive}(G) \notin \Gamma$, i.e., $[G]\bot \in \Gamma$ by the maximal consistency of Γ. Since $\Gamma \sim_G^c \Gamma$, i.e., $\Gamma \setminus [G] \subseteq \Gamma$, this yields $\bot \in \Gamma$, which contradicts the consistency of Γ. Thus $\mathsf{alive}(G) \in \Gamma$. We now show that

$$\mathsf{alive}(G) \in \Gamma \quad \text{implies} \quad \Gamma \in \mathsf{Alive}(G)_{\mathcal{M}^c}.$$

Assume $\mathsf{alive}(G) \in \Gamma$. We need to show $\Gamma \sim_G^c \Gamma$. Let $\phi \in \Gamma \setminus [G]$, i.e., $[G]\phi \in \Gamma$. Since
$$\mathsf{alive}(G) \to ([G]\phi \to \phi)$$

is an axiom of Syn^- and by the maximal consistency of Γ, it follows that $\phi \in \Gamma$, i.e., $\Gamma \sim_G^c \Gamma$. By Definition 8, we obtain $\Gamma \in \mathsf{Alive}(G)_{\mathcal{M}^c}$. $\square$

Lemma 12. $\mathcal{M}^c$ *is a κ-model.*

Proof. We show that $\mathcal{M}^c$ satisfies all properties of κ-models.

1. Symmetry: Assume $\Gamma \setminus [G] \subseteq \Delta$. We need to show $\Delta \setminus [G] \subseteq \Gamma$. Let $\phi \in \Delta \setminus [G]$, i.e., $[G]\phi \in \Delta$, and assume towards a contradiction that $\phi \notin \Gamma$. Since Γ is a maximal consistent set and (B) is an axiom of Syn^-, we have $\neg\phi \in \Gamma$ as well as $[G]\neg[G]\neg\phi \in \Gamma$ by the maximal consistency of Γ. Therefore, $\neg[G]\neg\phi \in \Gamma \setminus [G]$ and thus $\neg[G]\neg\phi \in \Delta$. This is a contradiction because we assumed $[G]\phi \in \Delta$. Hence, we conclude $\phi \in \Gamma$ which shows that $\Delta \setminus [G] \subseteq \Gamma$.

2. Transitivity: Assume $\Gamma \sim_G^c \Delta$ and $\Delta \sim_G^c \Phi$. We need to show $\Gamma \sim_G^c \Phi$. By assumption, we have $\Gamma \setminus [G] \subseteq \Delta$ and $\Delta \setminus [G] \subseteq \Phi$. Let $\phi \in \Gamma \setminus [G]$, i.e., $[G]\phi \in \Gamma$. Since Γ is maximally consistent and because (4) is an axiom of Syn^-, we have $[G][G]\phi \in \Gamma$, and thus $[G]\phi \in \Delta$. Further, since $\Delta \sim_G^c \Phi$, we have $\phi \in \Phi$ and hence $\Gamma \setminus [G] \subseteq \Phi$.

3. K1: Assume $\Gamma \in \mathsf{Alive}(G)$ and $\Gamma \in \mathsf{Alive}(H)$. By Lemma 11 it follows that $\mathsf{alive}(G) \in \Gamma$ and $\mathsf{alive}(H) \in \Gamma$. Since Γ is a maximal consistent set and (Union) is an axiom of Syn^- it follows that $\mathsf{alive}(G \cup H) \in \Gamma$. By Lemma 11 we have $\Gamma \in \mathsf{Alive}(G \cup H)$.

4. K2: Assume $A, B \in G$ and let $\Gamma \in \mathsf{Alive}(G)$. By Lemma 11 it follows that $\mathsf{alive}(G) \in \Gamma$. Since Γ is a maximal consistent set and (Clo) is an axiom of Syn^- it follows that $\mathsf{alive}(\{A \cup B\}) \in \Gamma$ by the maximal consistency of Γ. By Lemma 11 we have $\Gamma \in \mathsf{Alive}(\{A \cup B\})$.

5. K3: Assume $(\Gamma, \Delta) \in \sim_H^c$, we need to show that $(\Gamma, \Delta) \in \sim_G^c$. Let $\phi \in \Gamma \setminus [G]$, i.e., $[G]\phi \in \Gamma$. Since Γ is maximally consistent and (Mono) is an axiom of Syn^- we have that $[H]\phi \in \Gamma$. Since we assumed $\Gamma \setminus [H] \subseteq \Delta$ we have $\phi \in \Delta$.

6. K4: Assume $(\Gamma, \Delta) \in \sim_G^c$, we need to show that $(\Gamma, \Delta) \in \sim_{G \cup \{B\}}^c$. Let $\phi \in \Gamma \setminus [G \cup \{B\}]$, i.e., $[G \cup \{B\}]\phi \in \Gamma$. Since Γ is maximally consistent and (Equiv) is an axiom of $\mathsf{Syn}^{\text{-}}$ we have that $[G]\phi \in \Gamma$. Since we assumed $\Gamma \setminus [G] \subseteq \Delta$ we have $\phi \in \Delta$.

7. NE: Let $\Gamma \in W^c$ be arbitrary. Since Γ is maximally consistent and (NE) is an axiom of $\mathsf{Syn}^{\text{-}}$, there exists an agent pattern G such that $\mathsf{alive}(G) \in \Gamma$. By Lemma 11, $\Gamma \sim_G^c \Gamma$. $\square$

The truth lemma is standard and completeness follows as usual.

Lemma 13 (Truth lemma). *Let $\mathcal{M}^c = (W^c, \sim^c, V^c)$ be the canonical model. For each world $\Gamma \in W^c$ and each formula $\phi \in \mathcal{L}$ we have*

$$\mathcal{M}^c, \Gamma \Vdash \phi \text{ iff } \phi \in \Gamma.$$

Theorem 3 (Completeness). $\mathsf{Syn}^{\text{-}}$ *is complete with respect to κ-models.*

5.2. δ-models

As demonstrated in Example 3, the previously introduced κ-models do not necessarily satisfy (D). In this section, we introduce δ-models which satisfy it. Furthermore, we show that $\mathsf{Syn}^{\text{-}}$ is sound and complete with respect to δ-models.

Definition 13 (δ-model). *A pre-model $\mathcal{M} = (W, \sim, V)$ is called a δ-model if and only if*

K2*: $\mathsf{Alive}(G)_\mathcal{M} \subseteq \mathsf{Alive}(\{A \cup B\})_\mathcal{M}$ for $A, B \in G$;*

K3*: $\sim_H \subseteq \sim_G$, if $G \subseteq H$;*

K4*: $\sim_G \subseteq \sim_{G \cup \{B\}}$ if there exists $A \in G$ with $B \subseteq A$;*

NE*: for all $w \in W$, there exists an agent pattern G such that $w \sim_G w$;*

D*: $\sim_G = \cap_{B \in G} \sim_{\{B\}}$.*

A formula ϕ is δ-valid, denoted by $\Vdash_\delta \phi$, if $\mathcal{M} \Vdash \phi$ for all δ-models $\mathcal{M}$. Moreover, observe that D implies K1 and thus, δ-models are κ-models. This implies that $\mathsf{Syn}^{\text{-}}$ is sound with respect to δ-models. For a pre-model $\mathcal{M} = (W, \sim, V)$, we call $F = (W, \sim)$ its frame and sometimes write $\mathcal{M} = (F, V)$ instead. Furthermore, we say that a frame is a κ-frame, or a δ-frame, in order to indicate which properties the frame satisfies. In what follows, we will show that $\mathsf{Syn}^{\text{-}}$ is complete with respect to δ-models.

Definition 14 (History). *Let $F = (W, \sim)$ be a frame. A history h is a non-empty and finite sequence of triples (w, G, v) where*

1. *$w \sim_G v$ and G is maximal under set inclusion. That means, there does not exist an agent pattern G' with $G \subsetneq G'$ and $w \sim_{G'} v$;*

2. *if (w', G', v') is the successor of (w, G, v), then $v = w'$.*

19

We write $\ell(h)$ to denote the last world of a history h. That is, if (w, G, v) is the last element of h, then $\ell(h) = v$. Furthermore, $h \parallel (\ell(h), G, v)$ denotes the extension of h with $(\ell(h), G, v)$. The set of all histories over a frame F is denoted by H_F.

Lemma 14 (Downwards closure). *Let $F = (W, \sim)$ be a κ-frame and consider a history $h \in H_F$ and let (w, U, v) be an element of h. If there exists $A \in U$ with $B \subseteq A$, then $B \in U$.*

Proof. Towards a contradiction, suppose that $B \notin U$. Consider the agent pattern $U' = U \cup \{B\}$. Clearly $U \subsetneq U'$ and $w \sim_{U'} v$ because of K4. This contradicts U being maximal under set inclusion and thus, $B \in U$. $\qquad\square$

Definition 15 ($\to_G$). *Let $F = (W, \sim)$ be a frame and consider $h, h' \in H_F$. For an agent pattern G, we define $\to_G \subseteq H_F \times H_F$ as follows*

$$h \to_G h' \quad iff \quad h' = h \parallel (\ell(h), U, \ell(h')) \ and \ G \subseteq U.$$

Lemma 15. *Let $F = (W, \sim)$ be a κ-frame and consider $h, h' \in H_F$ with $h \to_G h'$ for some agent pattern G. The following hold:*

1. *if $H \subseteq G$, then $h \to_H h'$;*

2. *if there exists $A \in G$ and $B \subseteq A$, then $h \to_{\{B\}} h'$;*

3. *if $h \to_H h'$, then $h \to_{G \cup H} h'$.*

Proof. By assumption, $h' = h \parallel (w, U, v)$ with $G \subseteq U$. For the first case, we have $H \subseteq G \subseteq U$ and thus $h \to_H h'$ by Definition 15. For the second case, we have $\{B\} \subseteq U$ by Lemma 14 and thus $h \to_{\{B\}} h'$ by Definition 15. Lastly, assume $h \to_G h'$ and $h \to_H h'$, i.e., $G \subseteq U$ and $H \subseteq U$. Hence $G \cup H \subseteq U$ by the properties of set union, and therefore $h \to_{G \cup H} h'$ by Definition 15. $\qquad\square$

Definition 16 ($U(F)$). *Let $F = (W, \sim)$ be a frame. We define the unravelled frame $U(F) = (H_F, \{\approx_G\}_{G \subseteq \mathsf{Pow}(\mathsf{Ag}) \setminus \{\emptyset\}})$ where $\approx_G$ is the transitive closure of the symmetric closure of $\to_G$, i.e., $\approx_G = (\to_G \cup \to_G^{-1})^*$.*

Definition 17 ($U(\mathcal{M})$). *Let $F = (W, \sim)$ be a frame and consider the pre-model $\mathcal{M} = (F, V)$. We call $U(\mathcal{M}) = (U(F), L)$ with*

$$h \in L(p) \quad iff \quad \ell(h) \in V(p)$$

the unravelled model of $\mathcal{M}$.

Definition 18. *(R-path) Let R be a relation on a set X. A R-path from x_1 to x_n is a sequence*

$$\tau = (x_1, x_2), (x_2, x_3), \cdots, (x_{n-2}, x_{n-1}), (x_{n-1}, x_n)$$

with $(x_i, x_{i+1}) \in R$ for $1 \leq i \leq n - 1$.

Definition 19. *Let R be a relation on a set X. We define*

$$R \circ R := \{(x, y) \in X \times X \mid \text{ There exists } z \in X \text{ with } (x, z) \in R \text{ and } (z, y) \in R\}.$$

We abbreviate $\overbrace{(\cdots (R \circ R) \circ R \cdots)}^{n-1 \ times} \circ R)$ to R^n.

Remark 5. *Let R be relation on a set X, then there is a R-path of length n from a to b if and only if $(a, b) \in R^n$.*

Corollary 4. *Let F be a κ-frame. The following two are equivalent*

1. *$(h, h') \in \approx_G$;*

2. *there exists a $(\to_G \cup \to_G^{-1})$-path τ from h to h'.*

   For brevity we refer to $(\to_G \cup \to_G^{-1})$-paths as $\to_G$-paths.

Remark 6. *If $F = (W, \sim)$ is a frame, then a $\to_G$-path from h to h' implies the existence of a $\sim_G$-path from $\ell(h)$ to $\ell(h')$. It follows by transitivity that $\ell(h) \sim_G \ell(h')$.*

Theorem 4. *Let $F = (W, \sim)$ be a κ-frame. Then $U(F)$ is a κ-frame.*

Proof. Observe that $\approx_G = (\to_G \cup \to_G^{-1})^*$ is transitive and symmetric, because the transitive closure of a symmetric relation is transitive and symmetric. Moreover, NE follows because histories are not empty.

- K1: Assume $(h, h) \in \approx_G$ and $(h, h) \in \approx_H$. By Remark 6, $\ell(h) \sim_G \ell(h)$ as well as $\ell(h) \sim_H \ell(h)$. Since F satisfies K1, it holds that $\ell(h) \sim_{G \cup H} \ell(h)$. Hence, $h^* = h \parallel (\ell(h), U, \ell(h))$ with $G \cup H \subseteq U$ is a valid history and $(h, h^*) \in \approx_{G \cup H}$. We obtain $(h, h) \in \approx_{G \cup H}$ by symmetry and transitivity.

- K2: Assume $(h, h) \in \approx_G$ and let $A, B \in G$. By assumption and Remark 6, $\ell(h) \sim_G \ell(h)$. Because F satisfies K2, we have $\ell(h) \sim_{\{A \cup B\}} \ell(h)$. Thus $h^* = h \parallel (\ell(h), U, \ell(h))$ with $\{A \cup B\} \subseteq U$ is a valid history which implies that $(h, h) \in \approx_{\{A \cup B\}}$ by symmetry and transitivity.

- K3: Assume $(h, h') \in \approx_H$ and $G \subseteq H$. By Corollary 4, there must exist a $\to_H$-path τ from h to h'. Let $(s, s') \in \tau$ be arbitrary. Since F satisfies K3, Lemma 15 implies that $s \approx_G s'$. Therefore, τ is a $\to_G$-path and $(h, h') \in \approx_G$ by Corollary 4.

- K4: Assume $(h, h') \in \approx_G$ and that there exists $A \in G$ with $B \subseteq A$. By Corollary 4 there exists a $\to_G$-path τ from h to h'. Let $(s, s') \in \tau$ be arbitrary. Since F satisfies K4, Lemma 15 implies that $s \approx_{G \cup \{B\}} s'$. Therefore, τ is a $\to_{G \cup \{B\}}$-path and $(h, h') \in \approx_{G \cup \{B\}}$ by Corollary 4. $\square$

We call a R-path from x_1 to x_n non-redundant if and only if $x_i \neq x_{i+2}$ for $1 \leq i < n - 1$. Further, we often write

$$\tau = x_1 R x_2 R x_3 \cdots x_{n-1} R x_n$$

instead of $\tau = (x_1, x_2), (x_2, x_3), \cdots, (x_{n-2}, x_{n-1}), (x_{n-1}, x_n)$. If $R = S \cup S^{-1}$, we will use S and S^{-1} in a path instead of R. We say that a $\to_G$-path has a *change of direction*, if $(x_i, x_{i+1}) \in \to_G$ (respectively $(x_i, x_{i+1}) \in \to_G^{-1}$) and $(x_{i+1}, x_{i+2}) \in \to_G^{-1}$ (respectively $(x_{i+1}, x_{i+2}) \in \to_G$). In what follows, we will write $h \leftarrow_G h'$ instead of $h \to_G^{-1} h'$ for better readability.

Lemma 16. *Let $F = (W, \sim)$ be a κ-frame and consider the unravelled frame $U(F)$. If $(h_1, h_n) \in \approx_G$, then h_1 and h_n have a common prefix.*

Proof. If one history is a prefix of the other, the claim follows trivially. Hence, we assume that neither of them is a prefix of the other. Corollary 4 ensures that there exists a $\to_G$-path $\tau = (h_1, h_2), \cdots, (h_{n-1}, h_n)$, which we can assume to be non-redundant. Moreover, due to our first assumption, τ must have at least one change of direction. In order to show that a common prefix of h_1 and h_n exists, we show that τ has exactly one change of direction, and that change is of the form $h_{i-1} \leftarrow h_i \to h_{i+1}$ with $i > 1$.

Let h_i be the history at which the first change of direction occurs. Notice, that this implies that $i > 1$. First, we observe that this change of direction cannot be of the form $h_{i-1} \to h_i \leftarrow h_{i+1}$, because this would imply $h_{i-1} = h_{i+1}$ and contradict τ being non-redundant. Therefore, the first change of direction must be of the form $h_{i-1} \leftarrow h_i \to h_{i+1}$. A subsequent change of direction would be of the form $h_{i+k-1} \to_G h_{i+k} \leftarrow_G h_{i+k+1}$ for $k \geq 1$ which contradicts that τ is non-redundant. Therefore, h_i is a common prefix. $\square$

Lemma 17. *Let $F = (W, \sim)$ be a κ-frame and consider the unravelled frame $U(F)$. For any two agent patterns G and H, the following holds*

$$(h, h') \in \approx_G \quad and \quad (h, h') \in \approx_H \quad implies \quad (h, h') \in \approx_{G \cup H}.$$

Proof. Assume $h \approx_G h'$ and $h \approx_H h'$. By Lemma 16 the following paths exist

- $h''_G \to_G \cdots \to_G h$,
- $h''_H \to_H \cdots \to_H h$, and
- $h''_G \to_G \cdots \to_G h'$,
- $h''_H \to_H \cdots \to_H h'$.

Observe that either $h''_G = h''_H$ or one of the histories is a proper prefix of the other. If $h''_G \neq h''_H$, let h'' be the longer history. If they are of the same length, fix either $h'' = h''_G$ or $h'' = h''_H$. We can write h and h' as

$$h = h'' \parallel (w_1, G_1, w_2) \parallel \cdots \parallel (w_{n-1}, G_n, w_n), \text{ and}$$

$$h' = h'' \parallel (w'_1, G'_1, w'_2) \parallel \cdots \parallel (w'_{m-1} G'_m, w'_m).$$

By Definition 16 we have $G \subseteq G_i, G'_j$ and $H \subseteq G_i, G'_j$ for all $i, j \geq 0$. Thus, by Lemma 15 and Corollary 4, we have $h \approx_{G \cup H} h'$. $\square$

Lemma 18. *Let $F = (W, \sim)$ be a κ-frame. $U(F)$ satisfies* D.

Proof. $\approx_G \subseteq \bigcap_{B \in G} \approx_{\{B\}}$ follows directly by K3. For the other direction, let $G = \{A_1, ..., A_n\}$ be an agent pattern and consider the sets

$$B_1 = \{A_1\} \text{ and } B_i = \{A_i\} \cup B_{i-1} \text{ for } 2 \le i \le n.$$

If $(u, v) \in \bigcap_{B \in G} \approx_{\{B\}}$, then $(u, v) \in \approx_{A_i}$ for $1 \le i \le n$. Applying Lemma 17 inductively yields $(u, v) \in \approx_{B_i}$, for $1 \le i \le n$. Since $B_n = G$, we obtain $(u, v) \in \approx_G$. $\qquad\square$

Corollary 5. *If F is a κ-frame, then $U(F)$ is a δ-frame.*

Definition 20 (Functional bisimulation). *Let $\mathcal{M} = (W_\mathcal{M}, \sim^\mathcal{M}, V_\mathcal{M})$, and $\mathcal{N} = (W_\mathcal{N}, \sim^\mathcal{N}, V_\mathcal{N})$ be two pre-models. $f : W_\mathcal{M} \to W_\mathcal{N}$ is a functional bisimulation if and only if*

1. **Atom**: *for all $w \in W_\mathcal{M}$, $V_\mathcal{M}(w) = V_\mathcal{N}(f(w))$.*

2. **Forth**: *for any agent pattern G, $w \sim_G^\mathcal{M} v$ implies $f(w) \sim_G^\mathcal{N} f(v)$.*

3. **Back**: *for any agent pattern G, $f(w) \sim_G^\mathcal{N} v'$ implies that there exists w' with $f(w') = v'$ such that $w \sim_G^\mathcal{M} w'$.*

Lemma 19. *If $f : W_\mathcal{M} \to W_\mathcal{N}$ is a functional bisimulation, then for any formula φ we have $\mathcal{M}, w \Vdash \varphi$ if and only if $\mathcal{N}, f(w) \Vdash \varphi$.*

Proof. Let f be a functional bisimulation. We show the claim by induction on the length of φ.

1. If $\varphi \in \mathsf{Prop}$, then the claim follows by **Atom**.

2. $\varphi = \neg\phi$ follows by the induction hypothesis.

3. $\varphi = \phi \wedge \psi$ follows by the induction hypothesis.

4. We will show equivalently that

$$\mathcal{M}, w \nVdash [G]\phi \quad \text{iff} \quad \mathcal{N}, f(w) \nVdash [G]\phi.$$

From left to right, assume $\mathcal{M}, w \nVdash [G]\phi$. Hence, there exists $v \in W_\mathcal{M}$ with $w \sim_G^\mathcal{M} v$ and $\mathcal{M}, v \nVdash \phi$. By **Forth** we have that $f(w) \sim_G^\mathcal{N} f(v)$ and by the induction hypothesis we have $\mathcal{N}, f(v) \nVdash \phi$. Therefore, it holds that $\mathcal{N}, f(w) \nVdash [G]\phi$. For the other direction, assume $\mathcal{N}, f(w) \nVdash [G]\phi$. Hence, there exists $v' \in W_\mathcal{N}$ with $f(w) \sim_G^\mathcal{N} v'$ and $\mathcal{N}, v' \nVdash \phi$. By **Back** there exists $w' \in W_\mathcal{M}$ such that $f(w') = v'$ and $w \sim_G^\mathcal{M} w'$. By the induction hypothesis we obtain $\mathcal{M}, w' \nVdash \phi$ and thus $\mathcal{M}, w \nVdash [G]\phi$. $\qquad\square$

Let $F = (W, \sim)$ be a κ-frame and consider a κ-model $\mathcal{M} = (W, \sim, V)$ as well as its unravelled model $U(\mathcal{M}) = (U(F), L)$. It is straightforward to show that the mapping $\mathsf{last} : H_F \to W$ that maps each $h \in H_F$ to $\ell(h) \in W$ is a functional bisimulation. We only show the claim for **Back**. Assume that $w \in H_F$ and $v' \in W$ such that $\ell(w) \sim_G v'$. It follows that, $w' = w \parallel (\ell(w), U, v')$ with $G \subseteq U$ is a valid history with $w' \approx_G w$ and $\ell(w') = v'$. Finally, the completeness of Syn^- with respect δ-models is an immediate consequence.

Theorem 5. Syn^- *is sound and complete with respect to δ-models.*

5.3. Completeness of Syn

In this section, we will show that Syn is sound and complete with respect to proper δ-models.

Definition 21. *Let* $\mathcal{M} = (W, \sim, V)$ *be a pre-model. For* $w \in W$*, we define*

$$\overline{w} := \{B \mid \exists G.B \in G \text{ and } w \in \mathsf{Alive}(G)_{\mathcal{M}}\}.$$

Definition 22 ($\equiv$)**.** *Let* $\mathcal{M} = (W, \sim, V)$ *be a pre-model. We define the relation* $\equiv$ *on* $W \times W$ *as*

$$w \equiv v \quad \text{iff} \quad \overline{w} = \overline{v} \text{ and } w \sim_{\overline{w}} v.$$

Definition 23 (Proper)**.** *Let* $\mathcal{M} = (W, \sim, V)$ *be a pre-model. We say that* $\mathcal{M}$ *is proper if and only if* $w \equiv v$ *implies* $w = v$.

If G is an agent pattern, we denote the set of maximal elements of G with $\max(G)$. Notice that $\max(\overline{w})$ always contains exactly one element.

Remark 7. *Let* $\mathcal{M} = (W, \sim, V)$ *be a* κ*-model. The following can be shown by using the properties of* κ*-models and the fact that* $\overline{w}$ *is maximal under subsets:*

1. *$\mathcal{M}, w \Vdash \mathsf{alive}(\overline{w})$ and $\mathcal{M}, w \Vdash \mathsf{dead}(\overline{w}^C)$ are always the case;*

2. *$\mathcal{M}, w \Vdash \mathsf{alive}(G) \wedge \mathsf{dead}(G^C)$ if and only if $\max(G) = \max(\overline{w})$;*

3. *if $\max(G) = \max(\overline{w})$, then $w \sim_G v$ and $\mathcal{M}, v \Vdash \mathsf{dead}(G^C)$ imply $\overline{w} = \overline{v}$.*

Definition 24. *Let* $\mathcal{M} = (W, \sim, V)$ *be a pre-model. We define the model* $\mathcal{M}^\rho = (W^\rho, \sim^\rho, V^\rho)$ *as*

1. *$W^\rho = W/\equiv$ is the set of equivalence classes of $\equiv$;*

2. *$[w] \sim^\rho_G [v]$ if and only if $w \sim_G v$;*

3. *for any $p \in$ Prop, $[w] \in V^\rho(p)$ if and only if $w \in V(p)$.*

$\mathcal{M}^\rho$ *is well-defined, if, for any two worlds* $w, v \in W$ *and* $p \in$ Prop,

$$w \equiv v \text{ implies } w \in V(p) \iff v \in V(p).$$

Lemma 20. *Let* $\mathcal{M} = (W, \sim, V)$ *be a pre-model such that* $\mathcal{M}^\rho$ *is well-defined. We find that:*

1. *$\mathcal{M}^\rho$ is proper;*

2. *$\mathcal{M}, w \Vdash \phi$ if and only if $\mathcal{M}^\rho, [w] \Vdash \phi$.*

Proof. In order to show that $\mathcal{M}^\rho$ is proper, observe that $[w] \equiv [v]$ implies $w \equiv v$. Indeed, since $[w] \sim^\rho_G [v]$ if and only if $w \sim_G v$, it holds that $\overline{[w]} = \overline{w}$ and $w \sim_{\overline{w}} v$. Hence, w and v belong to the same equivalence class, i.e., $[w] = [v]$.

For the second claim, we show the direction from right to left by induction on the length of ϕ. The other direction is symmetric. The base case follows because $\mathcal{M}^\rho$ is well defined. The only case left is $\phi = [G]\psi$. Assume $\mathcal{M}^\rho, [w] \Vdash [G]\psi$. We need to show $\mathcal{M}, w \Vdash [G]\psi$. Let $v \in W$ be such that $w \sim_G v$, i.e., $[w] \sim^\rho_G [v]$. By the definition of truth $\mathcal{M}^\rho, [v] \Vdash \psi$ and by the induction hypothesis $\mathcal{M}, v \Vdash \psi$, which concludes the proof. $\square$

Proving soundness and completeness of Syn with respect to proper κ-models simply requires us to show that proper κ-models satisfy (P) and that the canonical model is proper.

Lemma 21 (Soundness). Syn *is sound with respect to proper κ-models.*

Proof. We showed the cases for Syn⁻ in the proof of Theorem 2. Hence, we only need to show the case for (P). Let $\mathcal{M} = (W, \sim, V)$ be an arbitrary proper κ-model. Assume $\mathcal{M}, w \Vdash \mathsf{alive}(G) \wedge \mathsf{dead}(G^C) \wedge \phi$. By Remark 7, we find that $\max(G) = \max(\overline{w})$. This implies that for any $v \in W$ with $w \sim_G v$ such that $\mathcal{M}, v \Vdash \mathsf{dead}(G^C)$, we have $\overline{w} = \overline{v}$, i.e., $w \equiv v$ and by the properness of $\mathcal{M}$, it follows that $w = v$. Therefore $\mathcal{M}, v \Vdash \phi$. $\qquad\square$

Theorem 6. *The canonical model $\mathcal{M}^c$ for Syn is a proper κ-model.*

Proof. We showed the cases for Syn⁻ in the proof of Theorem 3. Hence, it suffices to show that $\mathcal{M}^c$ is proper. Let $\Gamma, \Delta \in W^c$ such that $\Gamma \equiv \Delta$, i.e., $\overline{\Gamma} = \overline{\Delta} = G$ and $\Gamma \sim_G^c \Delta$. We now show $\Gamma = \Delta$, i.e., for any $\phi \in \mathcal{L}$, $\phi \in \Gamma$ if and only if $\phi \in \Delta$. We show the direction from left to right. The other direction is symmetric. Assume $\phi \in \Gamma$. Since Γ is a maximal consistent set, it follows by Remark 7 and Lemma 11 that $\mathsf{alive}(G) \wedge \mathsf{dead}(G^C) \wedge \phi \in \Gamma$. Furthermore, by (P), we have $[G](\mathsf{dead}(G^C) \to \phi) \in \Gamma$ and thus $\mathsf{dead}(G^C) \to \phi \in \Delta$. By assumption, $\mathsf{dead}(G^C) \in \Delta$ and thus $\phi \in \Delta$, because Δ is a maximal consistent set. $\qquad\square$

Corollary 6. Syn *is sound and complete with respect to proper κ-models.*

Lemma 22 below shows that the construction of Definition 24 can be applied to the unravelled canonical model.

Lemma 22. *Let $\mathcal{M}^c$ be the canonical model for Syn, and let $U(\mathcal{M}^c) = (H, L)$ be the unravelled canonical model. For any $h, h' \in H$ with $h \equiv h'$ and $p \in$ Prop it holds that*
$$h \in L(p) \quad \textit{iff} \quad h' \in L(p).$$

Proof. Let $\mathcal{M}^c = (W^c, \sim^c, V^c)$ be the canonical model for Syn. Furthermore, consider two histories h, h' of $U(\mathcal{M}^c)$ with $h \equiv h'$, i.e., $\overline{h} = \overline{h'} = G$ and $h \approx_G h'$. Let $\Gamma = \ell(h)$ and $\Delta = \ell(h')$. By Remark 6, $\ell(h) \sim_G \ell(h')$, i.e., $\Gamma \sim_G^c \Delta$. We now show the direction from left to right. The other direction is symmetric. Let $p \in$ Prop with $h \in L(p)$, i.e., $\Gamma \in V^c(p)$. Since Γ is a maximal consistent set it follows by Remark 7 and Lemma 11 that $\mathsf{alive}(G) \wedge \mathsf{dead}(G^c) \wedge p \in \Gamma$. By (P), $[G](\mathsf{dead}(G^C) \to p) \in \Gamma$. Therefore, $\mathsf{dead}(G^C) \to p \in \Delta$. Lastly, by assumption and Remark 7, we have $\mathsf{dead}(G^C) \in \Delta$ and since Δ is a maximal consistent set, it follows that $p \in \Delta$, i.e., $\ell(h') \in V^c(p)$. $\qquad\square$

Corollary 7. *Let $\mathcal{M}^c$ be the canonical model for Syn. It holds that $U(\mathcal{M}^c)^\rho$ is proper and*
$$U(\mathcal{M}^c)^\rho, [h] \Vdash \phi \quad \textit{iff} \quad U(\mathcal{M}), h \Vdash \phi.$$

Corollary 8. Syn *is sound and complete with respect to proper δ-models.*

5.4. δ-translations

In this section, we show how any proper δ-model can be transformed to an equivalent simplicial model. Completeness of Syn with respect to simplicial models follows immediately.

Definition 25 (δ-translation). *Let $\mathcal{M} = (W, \sim, V)$ be a δ-model and consider a simplicial model $\mathcal{C} = (\mathbb{C}, L)$. $\mathcal{C}$ is δ-translation of $\mathcal{M}$ if and only if there exists a mapping $T : W \to \mathbb{C}$ such that for all $w, v \in W$*

1. *$G \subseteq (T(w) \cap T(v))^{\circ}$ iff $w \sim_G v$;*

2. *$T(w) \in L(p)$ iff $w \in V(p)$.*

Construction 1 δ-translation

1: **Input**
2:     A proper δ-frame $F = (W, \sim)$

3: **Initialisation**
4:     $w_i^* := \max\{A \subseteq \mathsf{Ag} \mid w_i \sim_{\{A\}} w_i\}$
5:     $S_i := \{(A, i) \mid A \subseteq w_i^*\}$ for $1 \leq i \leq n$

6: **Transformation**
7:     $i = 1$
8:     $j = 1$
9:     **while** w_i exists **do**
10:         $j \leftarrow i + 1$
11:         **while** w_j exists **do**
12:             **for each** $G \subseteq \mathsf{Pow}(\mathsf{Ag}) \setminus \{\emptyset\}$ **with** $w_i \sim_G w_j$ **do**
13:                 **for each** $B \in G$ **do**
14:                     $k \leftarrow \{l \mid w_l \sim_{\{B\}} w_i\}$
15:                     $S_j \leftarrow S_j \setminus \{(B, j)\}$
16:                     $S_j \leftarrow S_j \cup \{(B, k)\}$
17:             $j \leftarrow j + 1$
18:         $i \leftarrow i + 1$

19: **Output**
20:     $\mathbb{C} = \{S_i \mid w_i \in W\}$

Let $\mathcal{M} = (W, \sim, V)$ be a proper δ-model and $F = (W, \sim)$ be its δ-frame. We assume an arbitrary enumeration of worlds and write w_i for the i-th world. In what follows, we denote the simplicial image of w_i under a mapping T, i.e., $T(w_i)$, with S_i. The next lemmas show that there exists a simplicial model $\mathcal{C}$ that is a δ-translation of $\mathcal{M}$. Let $w_i^* \subseteq \mathsf{Ag}$ be the maximum set of all agents alive in w_i. Construction 1 on input F first initialises a simplex $S_i = \{(A, i) \mid A \subseteq w_i^*\}$ for each world w_i (lines 3 to 5). At this point, no two different simplices S_i and S_j are connected. Throughout the transformation phase (lines 6 to 16), Construction 1 glues related simplices together according to the relation $\sim$ of

F. It does so by iterating through all pairs (w_i, w_j) with $i < j$, which suffices by symmetry of $\sim_G$, and checking for all G, if $(w_i, w_j) \in \sim_G$. If so, for each $B \in G$, (B, j) is replaced with $(B, k) \in S_i$, where k is the smallest index such that $w_k \sim_{\{B\}} w_i$. After the replacement, the simplicies S_i and S_j are connected. Example 4 shows a simple execution.

Remark 8. *Construction 1 could formally be given by a corecursion. However, we think that the presentation as an algorithm makes it easier to understand.*

We claim that the model $\mathcal{C} = (\mathbb{C}, L)$, where $\mathbb{C}$ is the complex returned by Construction 1 (line 20), and L is a valuation such that $S_i \in L(p)$ if and only if $w_i \in V(p)$, is a δ-translation of $\mathcal{M}$.

Example 4. *Let $\mathsf{Ag} = \{a, b\}$ and consider the proper δ-model $\mathcal{M} = (W, \sim, V)$ depicted in Figure 4. Construction 1 first initialises*

$$S_1 = \left\{ \begin{matrix} ab1 \\ a1, b1 \end{matrix} \right\} , \; S_2 = \left\{ \begin{matrix} ab2 \\ a2, b2 \end{matrix} \right\}, \; and \; S_3 = \left\{ \begin{matrix} ab3 \\ a3, b3 \end{matrix} \right\}.$$

During the transformation phase, when $i = 1$ and $j = 2$, Construction 1 replaces $(a, 2)$ and $(b, 2)$ with $(a, 1)$ and $(b, 1)$ due to $w_1 \sim_{\{a\}} w_2$ and $w_1 \sim_{\{b\}} w_2$. Moreover, it replaces $(b, 3) \in S_3$ with $(b, 1)$. Observe that if Line 14 was missing, i.e., $k = i$, then, for $i = 2$ and $j = 3$, the Construction would add $(b, 2)$ to S_3 as well, which would make it an ill-formed simplex. The resulting complex is

$$\mathbb{C} = \left\{ \left\{ \begin{matrix} ab1 \\ a1, b1 \end{matrix} \right\}, \left\{ \begin{matrix} ab2 \\ a1, b1 \end{matrix} \right\}, \left\{ \begin{matrix} ab3 \\ a3, b1 \end{matrix} \right\} \right\}.$$

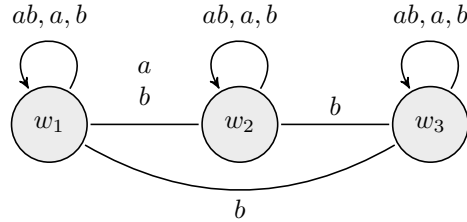Figure 4: The δ-model for Example 4. Only arrows for maximal agent patterns are shown.

The next Lemma ensures that we can safely assume the existence of an unique and non-empty S_i for each w_i throughout our proofs.

Lemma 23 (Uniqueness)**.** *Let $F = (W, \sim)$ be a proper δ-frame. After the initialisation of Construction 1 on input F and after each execution of Line 16, the following holds for all $w_i, w_j \in W$:*

1. *$S_i \neq \emptyset$;*

2. *$S_i = S_j$ if and only if $i = j$.*

Proof. $S_i \neq \emptyset$ follows immediately because δ-frames satisfy NE and thus, each S_i is initialised to some non-empty set. Furthermore, since each element that is removed gets replaced, we conclude that $S_i \neq \emptyset$ for all $w_i \in W$ after Line 16. For the second claim, the direction from right to left follows immediately. For the other direction, observe that if $S_i = S_j$, then $\overline{w_i} = \overline{w_j}$ and $w_i \sim_{\overline{w_i}} w_j$, i.e., $w_i \equiv w_j$. By the properness of F, it follows that $w_i = w_j$, i.e., $i = j$. $\square$

Since the smaller index gets precedence, some elements of a simplex may never be replaced. For example, the simplex S_1 remains unchanged throughout Construction 1. Lemma 24 shows that those simplices are correct.

Lemma 24. *Let $F = (W, \sim)$ be a proper δ-frame and consider Construction 1 on input F. After the initialisation (lines 3 to 5), for all $w_j \in W$, it holds that $(A, j) \in S_j$ iff $w_j \sim_{\{A\}} w_j$.*

Proof. From left to right, we have $(A, j) \in S_j$ if and only if $A \subseteq w_j^*$. Since $w_j \sim_{\{w_j^*\}} w_j$ by definition, we obtain $w_j \sim_{\{A\}} w_j$ by Lemma 10. Regarding the other direction, assume $w_j \sim_{\{A\}} w_j$. By the definition of w_j^* it follows that $A \subseteq w_j^*$ which implies that $(A, j) \in S_j$ after the initialisation phase. $\square$

Lemma 25. *Let $F = (W, \sim)$ be a proper δ-frame and consider $i, j, k \in \mathbb{N}$ such that $i \neq j, i \neq k$, and $j \neq k$. There does not exist $B \subseteq \mathsf{Ag}$ such that, while running Construction 1 on input F, (B, i) is replaced with (B, j) and (B, j) is replaced with the pair (B, k).*

Proof. Towards a contradiction, assume that there exists such a $B \subseteq \mathsf{Ag}$, i.e., for some S_i, (B, i) is replaced with (B, j) and (B, j) is replaced with (B, k). By assumption, $i < j < k$ because Construction 1 replaces (B, l) with (B, m) only if $l > m$. Since lines 15 and 16 are executed at least twice, there exist $G, G' \subseteq \mathsf{Pow}(\mathsf{Ag}) \setminus \{\emptyset\}$ such that $B \in G$ and $B \in G'$ with $w_i \sim_G w_j$ and $w_j \sim_{G'} w_k$. Since $\{B\} \subseteq G$ and $\{B\} \subseteq G'$, we have by K3 that $w_i \sim_{\{B\}} w_j$ and $w_j \sim_{\{B\}} w_k$. Further, $w_i \sim_{\{B\}} w_k$ follows by transitivity of $\sim_{\{B\}}$. But this means, that $i = k$ because (B, k) was replaced by (B, i) prior which contradicts that $i < k$. $\square$

Lemma 26. *Let $F = (W, \sim)$ be a proper δ-frame. The output $\mathbb{C}$ of Construction 1 on input F satisfies the following two properties*

T*1: Let $S \in \mathbb{C}$. If $(A, j) \in S$ and $(A, k) \in S$, then $j = k$.*

T*2: Let $w_i, w_j \in W$ and $G \subseteq \mathsf{Pow}(\mathsf{Ag}) \setminus \{\emptyset\}$, then*

$$\forall B \in G. \exists k \in \mathbb{N}. (B, k) \in S_i \wedge (B, k) \in S_j \quad \text{iff} \quad w_i \sim_G w_j.$$

Proof. T1 follows by construction. Regarding T2, we start by showing the direction from left to right. Assume

$$\forall B \in G. \exists k \in \mathbb{N}. (B, k) \in S_i \wedge (B, k) \in S_j.$$

Without loss of generality, we fix $i \leq j$. We have that $(B, k) \in S_i \cap S_j$ only if there exists G' with $w_i \sim_{G'} w_j$ and $B \in G'$. By K3 we get $w_i \sim_{\{B\}} w_j$. Since B is arbitrary, D implies that $w_i \sim_G w_j$.

For the other direction, let $w_i \sim_G w_j$ and let $B \in G$. By construction, $(B, j) \in S_j$ is replaced by (B, k) with $k \leq i$. Let k be the smallest index such that $w_k \sim_G w_i$, i.e., if $k < i$, then $(B, i) \in S_i$ was replaced with (B, k) before. Lemma 24 ensures that $(B, k) \in S_k$. Hence, if $k < i$, $(B, j) \in S_j$ and $(B, i) \in S_i$ are both replaced by (B, k) in the same iteration. If $k = i$, only (B, j) is replaced by (B, i), and $(B, i) \in S_i$ at the end of the Construction due to the symmetry of $\sim_G$ and k being minimal. By Lemma 25 there cannot be any further replacements and thus $(B, k) \in S_i$ and $(B, k) \in S_j$ at the end of Construction 1. $\qquad\square$

A consequence of T2 is that

$$S_i \sim_G S_j \text{ if and only if } w_i \sim_G w_j. \tag{12}$$

Lemma 27. *Let $F = (W, \sim)$ be a proper δ-frame. Let $\mathbb{C}$ be the complex returned by Construction 1 on input F. For each $w_i \in W$, $S_i \in \mathbb{C}$ is a well formed simplex.*

Proof. We show that S_i satisfies S1, S2, and S3.

1. S_i satisfies S1. Let $(A, j) \in S_i$ and $(B, k) \in S_i$ be two distinct maximal elements. First, we note that $A \neq B$ due to T1. By T2, $w_i \in \mathsf{Alive}(\{A\})$ and $w_i \in \mathsf{Alive}(\{B\})$. By K3, $w_i \in \mathsf{Alive}(\{A\} \cup \{B\})$. Further, by K2, $w_i \in \mathsf{Alive}(\{A \cup B\})$. Hence, $(A \cup B, l) \in S_i$ for some l. But since $A \subsetneq A \cup B$, this contradicts that (A, i) is the maximal element of S_i.

2. S_i satisfies S2. Let $(B, j) \in S_i$ and $C \subseteq B$, we need to show that there exists a unique k such that $(C, k) \in S_i$. By Lemma 10 we have $\sim_{\{B\}} \subseteq \sim_{\{C\}}$. Since we assume that $(B, j) \in S_i$, we get $w_i \sim_{\{B\}} w_i$ by T2. Due to $\sim_{\{B\}} \subseteq \sim_{\{C\}}$, we have that $w_i \sim_{\{C\}} w_i$. Thus, by T2, there exists $k \in \mathbb{N}$ such that $(C, k) \in S_i$. Condition T1 ensures that k is unique.

3. S_i satisfies S3. Let $\max(S_i) = (B, k)$ and suppose that $(A, j) \in S_i$ for some $A \nsubseteq B$. We have $w_i \sim_{\{A\}} w_i$ and $w_i \sim_{\{B\}} w_i$ by T2. By K3, it follows that $w_i \in \mathsf{Alive}(\{A\} \cup \{B\})$ and by K2 we have that $w_i \in \mathsf{Alive}(\{A \cup B\})$. Thus $(A \cup B, l) \in S_i$ for some l which contradicts the maximality of (B, k). $\qquad\square$

Lemma 28. *Let $F = (W, \sim)$ be a proper δ-frame and consider Construction 1 on input F. The output set $\mathbb{C} = \{S_i \mid w_i \in W\}$ is a complex.*

Proof. In order to show that $\mathbb{C}$ is a complex, we need to show that it satisfies Condition C. Consider the simplices $S_m, S_n \in \mathbb{C}$ and assume that there exists $A \subseteq \mathsf{Ag}$ and i with $(A, i) \in S_m$ and $(A, i) \in S_n$. We need to show that for all $B \subseteq A$ and all $j \in \mathbb{N}$ it holds that

$$(B, j) \in S_m \quad \text{iff} \quad (B, j) \in S_n.$$

Since S_m and S_n are arbitrary, it is enough to show only one direction. Assume that $(B, j) \in S_m$. By T2, we have $w_m \sim_{\{A\}} w_n$ and by Lemma 10, it holds that $w_m \sim_{\{B\}} w_n$. Since $(B, j) \in S_m$, T2 implies $w_m \sim_{\{B\}} w_j$. Further, by symmetry and transitivity we obtain $w_j \sim_{\{B\}} w_n$. Hence, it holds that $j \leq n$. Therefore, (B, j) replaced (B, n) in S_n. By Lemma 25, no more replacements of that pair can happen during the execution of Construction 1 and we conclude $(B, j) \in S_n$. $\square$

Theorem 7. *For each proper δ-model, there exists a δ-translation.*

Proof. Let $\mathcal{M} = (W, \sim, V)$ be a proper δ-model. Further, consider $\mathcal{C} = (\mathbb{C}, L)$, where $\mathbb{C}$ is the output of Construction 1 on input $(W, \sim)$ and $L : \mathbb{C} \to \mathsf{Prop}$ is a valuation such that

$$S_i \in L(p) \quad \text{iff} \quad w_i \in V(p). \tag{13}$$

By Lemma 28, $\mathcal{C}$ is a simplicial model. It follows from (12) and (13) that $\mathcal{C}$ is a δ-translation of $\mathcal{M}$. $\square$

Theorem 8. *Let $\mathcal{M} = (W, \sim, V)$ be a proper δ-model and let $\mathcal{C} = (\mathbb{C}, L)$ be a δ-translation of $\mathcal{M}$. It holds that*

$$\mathcal{M}, w_i \Vdash \phi \text{ if and only if } \mathcal{C}, S_i \Vdash_\sigma \phi.$$

Proof. By induction on the length of formulas.

1. Let $\phi \equiv \mathsf{alive}(G)$ for some $G \subseteq \mathsf{Pow}(\mathsf{Ag}) \setminus \{\emptyset\}$. We have $\mathcal{M}, w_i \Vdash \mathsf{alive}(G)$ iff $w_i \sim_G w_i$ iff $S_i \sim_G S_i$ iff $\mathcal{C}, S_i \Vdash_\sigma \mathsf{alive}(G)$.

2. Let $\phi \in \mathsf{Prop}$. We have $\mathcal{M}, w_i \Vdash \phi$ iff $w_i \in V(p)$ iff $S_i \in L(p)$ iff $\mathcal{C}, S_i \Vdash_\sigma \phi$ (by the definition of L).

3. Let $\phi \equiv \neg \psi$. Follows by the induction hypothesis.

4. Let $\phi \equiv \psi \wedge \varphi$. Follows by the induction hypothesis.

5. Let $\phi \equiv [G]\psi$. We equivalently show

$$\mathcal{M}, w_i \nVdash [G]\psi \text{ iff } \mathcal{C}, S_i \nVdash_\sigma [G]\psi.$$

It holds that $\mathcal{M}, w_i \nVdash [G]\psi$ iff there exists $w_j \in W$ with $w_i \sim_G w_j$ and $\mathcal{M}, w_j \nVdash \psi$ iff $G \subseteq (S_i \cap S_j)^\circ$ (12) and $\mathcal{C}, S_j \nVdash_\sigma \psi$ (by hypothesis) iff $\mathcal{C}, S_i \nVdash_\sigma [G]\psi$ by definition. $\square$

Hence, if $\nvdash \varphi$, then there exits a proper δ-model $\mathcal{M}$ (Corollary 8) such that $\mathcal{M} \nVdash \varphi$. By Theorem 8, there must exist a simplicial model $\mathcal{C}$ such that $\mathcal{C} \nVdash_\sigma \varphi$. Thus, Syn is sound and complete with respect to simplicial models.

Corollary 9. $\Vdash_\sigma \varphi$ *if and only if* $\vdash \varphi$.

6. Conclusion

This work presented the normal modal logic Syn, a logic in which we can reason about what a group of agents knows beyond their pooled individual knowledge. The new modality $[G]$ allows us to express the knowledge of a group while also accounting for relations between its members. Hence, the knowledge of two seemingly equal groups may differ due to different relationships among their agents. Moreover, we provided simplicial semantics for Syn based on semi-simplicial sets and demonstrated its use by analysing consensus objects and the dining cryptographers problem. To round off our work on synergistic knowledge, we showed that Syn is sound and complete with respect to our simplicial models.

Semi-simplicial sets were initially studied in order to represent improper Kripke models and to explore new types of group knowledge (cf. van Ditmarsch et al. [5]). While synergistic knowledge definitely is a new form of distributed knowledge, we still must require that Syn contains the axiom (P) because for any simplex S, S° is a valid agent pattern. Thus, Syn⁻ can describe improper Kripke models. For example, the δ-model shown in Figure 5 is a valid model for Syn⁻, but not for Syn. Applying Construction 1 would result in a complex with only one world, which means that this complex cannot be a δ-translation of the model shown in Figure 5. However, if we restrict agent patterns to sets, we might be able to represent improper Kripke models. For example, when treating agent patterns as sets, the complex depicted in Figure 1 corresponds to an improper Kripke model.
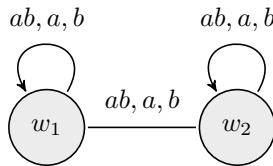

Figure 5: An improper δ-model. Only the maximal agent patterns are shown.

7. Future Work

In this work, we presented a candidate for an indistinguishability relation. However, there may be other reasonable options. For example, another possible notion of indistinguishability is one that takes the connectivity of an agent pattern into account. It is an intuitive option for settings where adding an unrelated agent to a pattern must not result in new knowledge. For example, adding an agent without any communication links to a network does not strengthen the network's reasoning power. In other words, an agent pattern is only as strong as its weakest connected component. Unfortunately, the concept of indistinguishability so far cannot capture this. To see why this is the case, consider the model depicted in Figure 6, where the pattern $G = \{a, b\}$ cannot distinguish X and Y.

But if we add the completely unrelated agent c to G, then, the updated pattern can suddenly distinguish X and Y. Hence, we would like a definition in which $\{a, b, c\}$ cannot distinguish X from Y. Let G be the pattern $\{a, bc\}$. We could 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. Hence, $\{a, b, c\}$ cannot tell the difference between X and Y because a and b cannot do so individually. Moreover, c does not have the means to tell a or b why it can distinguish the two worlds. A further investigation of this kind of indistinguishability seems to be promising for network topologies. We think that it is certainly fruitful to have communication pattern logic (cf. Castañeda et al. [3]) in mind when studying this notion of indistinguishability. A first analysis of this topic can be found in the preliminary version of this work [2].
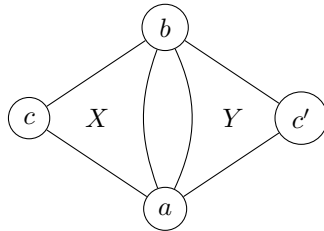

Figure 6: A model in which two agents a and b can together distinguish between the worlds X and Y. However, they cannot do so individually.

Additionally, we would like to explore the notion of synergy in distributed computing. Especially for tasks which require agents to commit to certain roles. For example, a smart contract that assigns buyers to sellers can be seen as a synergistic primitive for two parties. Indeed, without a buyer, the seller cannot sell and without the seller, the buyer cannot buy. Hence, this task needs two roles: a seller and a buyer. The model depicted in Figure 7 on the left shows how simplicial models can capture this. The buyer B offers an amount o and the seller S sells for a price p. Individually, they do not know if a trade can happen. That is, they do not know whether $p \leq o$ (i.e., deal) or $p > o$ (i.e., no deal). Only after querying the previously mentioned smart contract, they know if an exchange can take place. Lastly, it is also reasonable to analyse this from a dynamic standpoint in which querying the smart contract eliminates multiple edges. Figure 7 shows an update for the model on the left, in the case that $p \leq o$. It would be interesting to formulate updates that eventually transform models with parallel faces to the usual simplicial complexes. However, it is currently not clear how such update rules could transfer to distributed computing.
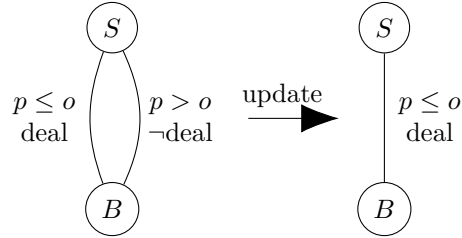
32

Figure 7: Without accessing a buyer-seller smart contract, the buyer B and the seller S do not know whether a trade can happen. We can represent the model after accessing the smart contract by simply removing the parallel edge, i.e., global state, that is not the case.

References

[1] Cachin, C., Lehnherr, D., Studer, T., 2023a. Modal and justification logics for multi-agent systems (invited talk), in: Herzig, A., Luo, J., Pardo, P. (Eds.), Logic and Argumentation - 5th International Conference, CLAR 2023, Hangzhou, China, September 10-12, 2023, Proceedings, Springer. pp. 3–8. URL: https://doi.org/10.1007/978-3-031-40875-5_1, doi:10.1007/978-3-031-40875-5_1.

[2] Cachin, C., Lehnherr, D., Studer, T., 2023b. Synergistic knowledge, in: Dolev, S., Schieber, B. (Eds.), Stabilization, Safety, and Security of Distributed Systems - 25th International Symposium, SSS 2023, Jersey City, NJ, USA, October 2-4, 2023, Proceedings, Springer. pp. 552–567. URL: https://doi.org/10.1007/978-3-031-44274-2_41, doi:10.1007/978-3-031-44274-2_41.

[3] Castañeda, A., van Ditmarsch, H., Rosenblueth, D.A., Velázquez, D.A., 2023. Communication pattern logic: Epistemic and topological views. Journal of Philosophical Logic 52, 1445–1473. doi:10.1007/s10992-023-09713-8.

[4] Chaum, D., 1988. The dining cryptographers problem: Unconditional sender and recipient untraceability. J. Cryptol. 1, 65–75. URL: https://doi.org/10.1007/BF00206326, doi:10.1007/BF00206326.

[5] van Ditmarsch, H., Goubault, É., Ledent, J., Rajsbaum, S., 2022. Knowledge and simplicial complexes, in: Lundgren, B., Nuñez Hernández, N.A. (Eds.), Philosophy of Computing, Springer International Publishing, Cham. pp. 1–50.

[6] van Ditmarsch, H., Kuznets, R., Randrianomentsoa, R., 2023a. On two- and three-valued semantics for impure simplicial complexes, in: GandALF, pp. 50–66. doi:10.4204/EPTCS.390.4.

[7] van Ditmarsch, H., Kuznets, R., Randrianomentsoa, R.F., 2023b. Impure simplicial complexes: Complete axiomatization. Log. Methods Comput. Sci. 19. URL: `https://doi.org/10.46298/lmcs-19(4:3)2023`, doi:`10.46298/LMCS-19(4:3)2023`.

[8] Fagin, R., Halpern, J.Y., Moses, Y., Vardi, M.Y., 1995. Reasoning About Knowledge. MIT Press. URL: `https://doi.org/10.7551/mitpress/5803.001.0001`, doi:`10.7551/mitpress/5803.001.0001`.

[9] Goubault, É., Kniazev, R., Ledent, J., Rajsbaum, S., 2023a. Semi-simplicial set models for distributed knowledge, in: LICS, pp. 1–13. URL: `https://doi.org/10.1109/LICS56636.2023.10175737`, doi:`10.1109/LICS56636.2023.10175737`.

[10] Goubault, É., Kniazev, R., Ledent, J., Rajsbaum, S., 2023b. Simplicial models for the epistemic logic of faulty agents. `arXiv:2311.01351`.

[11] Goubault, É., Ledent, J., Rajsbaum, S., 2021. A simplicial complex model for dynamic epistemic logic to study distributed task computability. Inf. Comput. 278, 104597. URL: `https://doi.org/10.1016/j.ic.2020.104597`, doi:`10.1016/j.ic.2020.104597`.

[12] Goubault, É., Ledent, J., Rajsbaum, S., 2022. 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), Schloss Dagstuhl - Leibniz-Zentrum für Informatik. pp. 33:1–33:20. URL: `https://doi.org/10.4230/LIPIcs.STACS.2022.33`, doi:`10.4230/LIPIcs.STACS.2022.33`.

[13] Halpern, J.Y., Moses, Y., 1984. 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, ACM. pp. 50–61. URL: `https://doi.org/10.1145/800222.806735`, doi:`10.1145/800222.806735`.

[14] Herlihy, M., 1991. Wait-free synchronization. ACM Trans. Program. Lang. Syst. 13, 124–149. URL: `https://doi.org/10.1145/114005.102808`, doi:`10.1145/114005.102808`.

[15] Herlihy, M., Kozlov, D.N., Rajsbaum, S., 2013. Distributed Computing Through Combinatorial Topology. Morgan Kaufmann. URL: `https://store.elsevier.com/product.jsp?isbn=9780124045781`.