

Practical Verification of Multi-Agent Systems against SLK Specifications

Petr Čermák^a, Alessio Lomuscio^a, Fabio Mogavero^b, Aniello Murano^c

^a*Imperial College London, UK*

^b*University of Oxford, UK*

^c*Università degli Studi di Napoli Federico II, Italy*

Abstract

We introduce Strategy Logic with Knowledge, a novel formalism to reason about knowledge and strategic ability in memoryless multi-agent systems with incomplete information. We exemplify its expressive power; we define the model checking problem for the logic and show that it is PSPACE-COMPLETE. We propose a labelling algorithm for solving the verification problem that we show is amenable to symbolic implementation. We introduce MCMAS_{SLK}, an extension of the open-source model checker MCMAS, implementing the proposed algorithm. We report the benchmarks obtained on a number of scenarios from the literature, including the dining cryptographers protocol.

Keywords: model checking, Strategy Logic, multi-agent systems, formal verification

1. Introduction

Multi-agent systems (MAS) are distributed systems whereby the components, or *agents*, display a high degree of autonomy and interact with their peers in a cooperative or adversarial way in order to maximise their private or common goals [1]. Over the years several logics have been put forward to reason about MAS, including epistemic logics [2], deontic logics [3] and formal languages accounting for the beliefs, desires and intentions of the agents [4]. These formalisms, once combined with temporal logic, are more expressive than logics used to reason about reactive systems, such as LTL and CTL [5]. This is because when reasoning about MAS, it is often not sufficient to establish whether a particular temporal statement is realised. Instead, there is an interest in establishing whether high-level properties of the agents hold in the system. These may involve the evolution of their beliefs, the intentions they want to bring about, what regulations they are subjected to and the interplay of all of these.

Knowledge and strategic ability are two particular aspects of agency that are of importance when reasoning about MAS. By means of epistemic specifications, we can, for example, reason about what the agents know about the world,

its evolution, their peers, their peers' knowledge, as well as epistemic group notions such as common knowledge [2]. By incorporating strategic abilities in the specifications, we can establish whether particular groups of agents have the ability to bring about certain temporal states of affairs.

There is a relatively long tradition in the development of verification technology, notably model checking, to verify MAS against temporal-epistemic specifications. This includes methods based on bounded [6, 7] verification, logical representation through Binary Decision Diagrams (*BDDs*, for short) [8, 9, 10], symmetry reduction [11] and abstraction [12, 13]. Verification techniques for validating systems against strategic abilities have also been put forward. For example, jMOCHA [14] is a model checker for the verification of systems against specifications in Alternating-time Temporal Logic (ATL). Proposals have also been made to devise methods supporting specifications that account for both the epistemic states of the agents as well as their strategic ability [15, 16, 17].

Also related to the present proposal is [18], where an extension of ATL for imperfect information games, namely CSLP, was introduced. CSLP can express sophisticated strategic and epistemic game properties of coalitions under uncertainty, including solution concepts. However, differently from the proposal made here, CSLP consider the agents' strategies implicitly, which are bound directly to the agents they refer to.

MCMAS [19, 20] is a BDD-based model checker supporting both ATL and epistemic specifications.

An important aspect in combining epistemic and strategic specifications is the information model the agents adhere to, namely whether they have either complete or incomplete information about the world and what form of memory they have. Epistemic analysis of MAS normally assumes that agents have private, incomplete information about the world. It is known that the model checking problem for ATL with incomplete information and perfect recall is undecidable [21]. Given this, the most widely adopted setting is incomplete information with memoryless local states. In turn, this implies that ATL modalities assume memoryless strategies. This raises further issues including whether local strategies should be uniform [22, 19, 17].

In this work we follow this tradition, but extend the strategic dimension of the analysis to a fragment of strategy logic [23]. A limitation of ATL is that while specifications relate to the strategies of the agents to achieve a certain state of affairs, the strategies themselves do not feature in the syntax explicitly; instead, they are treated implicitly through agent modalities that refer to coalitions. The logic SL was introduced to overcome this by introducing strategies as first-class citizens of the syntax and by allowing explicit quantification and binding over them.

In this paper we define a combination of SL with the standard epistemic modalities on a memoryless variant of interpreted systems, introduce an algorithm for model checking and present its implementation. A feature of the present work is the relatively low complexity of the model checking problem, which is shown to be PSPACE-COMPLETE. As discussed in more details below, this is achieved by limiting the scope of the epistemic operators in the language.

The rest of the paper is organised as follows. In Section 2 we introduce the specification language SLK, define its semantics on a variant of interpreted systems, illustrate its use, define the model checking problem and investigate its complexity. In Section 3 we present a labelling algorithm to solve the model checking problem for SLK specifications against interpreted systems. In Section 4 we present MCMAS_{SLK}, a symbolic model checker derived from MCMAS, implementing the labelling algorithm, and present experimental results. We conclude in Section 5 by discussing related work.

2. Strategy Logic with Knowledge

In this section, we introduce Strategy Logic with Knowledge [24] (SLK, for short) as an extension of the original *Strategy Logic* [25, 26] (SL, for short) introduced in [23]. Our aim is to define a formalism combining the ability of SL to express game-theoretic concepts with an epistemic framework for describing the agents’ knowledge in the context of incomplete information. In order to avoid the well-known undecidability result of the model-checking problem of multi-agent systems under incomplete information and perfect recall [27], we formalise the new logic by means of imperfect-recall semantics. This implies that agents have no memory of the past, including when planning their strategy to achieve a desired goal. SLK is defined *w.r.t.* interpreted systems [28], whereas SL has been introduced for concurrent game structures [27], since the former is normally used in the context of incomplete information.

In the rest of this section, we give formal definitions of SLK syntax and semantics, together with accessory concepts, such as strategy, profile and play. These provide us with a solid foundation for the development of the model-checking algorithm, which is theoretically presented in Section 3 and then developed as a software tool and benchmarked against several scalable scenarios in Section 4.

2.1. Syntax

SL syntactically extends liner-time temporal logic LTL [29] by introducing two *strategy quantifiers* $\langle\langle x \rangle\rangle$ and $\llbracket x \rrbracket$, and an *agent binding* (a, x) , where x is a variable and a an agent. Informally, these operators can be read as “*there exists a strategy x* ”, “*for all strategies x* ” and “*bind agent a to the strategy associated with x* ”, respectively. SLK further extends SL with *epistemic modalities* [28] representing *individual knowledge* K_a , *group knowledge* E_A , *distributed knowledge* D_A and *common knowledge* C_A , where a is an agent and A a set of agents. Therefore, this language can be seen as a unique blend of three well-established logic formalisms that allows us to reason about the temporal, strategic and epistemic aspects of a model in a unified way.

SLK formulas are defined as follows.

Definition 1 (Syntax). SLK formulas are built inductively from the sets of atomic propositions AP, variables Vr and agents Ag using the following context-free grammar, where $p \in \text{AP}$, $x \in \text{Vr}$, $a \in \text{Ag}$ and $A \subseteq \text{Ag}$:

$$\varphi ::= \perp \mid \top \mid p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid (\varphi \vee \varphi) \mid \mathbf{X}\varphi \mid (\varphi \mathbf{U}\varphi) \mid (\varphi \mathbf{R}\varphi) \mid \langle\langle x \rangle\rangle\varphi \mid \llbracket x \rrbracket\varphi \mid (a, x)\varphi \mid \mathbf{K}_a\varphi \mid \mathbf{E}_A\varphi \mid \mathbf{D}_A\varphi \mid \mathbf{C}_A\varphi$$

where the epistemic operators are applied to sentences only, i.e., $\text{free}(\varphi) = \emptyset$ in $\mathbf{K}_a\varphi$, $\mathbf{E}_A\varphi$, $\mathbf{D}_A\varphi$ and $\mathbf{C}_A\varphi$, where the function free is introduced in Definition 2. SLK denotes the set of formulas generated by the aforementioned rules.

As for SL , the *free agents and free variables* $\text{free}(\varphi)$ of an SLK formula φ are the subset of $\text{Ag} \cup \text{Vr}$ containing (i) all agents a for which there is no binding (a, x) before the occurrence of a temporal operator and (ii) all variables x for which a binding (a, x) is not in the scope of any quantification $\langle\langle x \rangle\rangle$ or $\llbracket x \rrbracket$.

Definition 2 (Free Agents and free Variables). *The set of free agents and free variables of an SLK formula is given by the function $\text{free}: \text{SLK} \rightarrow 2^{\text{Ag} \cup \text{Vr}}$ defined inductively as follows:*

1. $\text{free}(\perp), \text{free}(\top) \triangleq \emptyset$;
2. $\text{free}(p) \triangleq \emptyset$, where $p \in \text{AP}$;
3. $\text{free}(\neg\varphi) \triangleq \text{free}(\varphi)$;
4. $\text{free}(\varphi_1 \text{Op} \varphi_2) \triangleq \text{free}(\varphi_1) \cup \text{free}(\varphi_2)$, where $\text{Op} \in \{\wedge, \vee\}$;
5. $\text{free}(\mathbf{X}\varphi) \triangleq \text{Ag} \cup \text{free}(\varphi)$;
6. $\text{free}(\varphi_1 \text{Op} \varphi_2) \triangleq \text{Ag} \cup \text{free}(\varphi_1) \cup \text{free}(\varphi_2)$, where $\text{Op} \in \{\mathbf{U}, \mathbf{R}\}$;
7. $\text{free}(\mathbf{Qn}\varphi) \triangleq \text{free}(\varphi) \setminus \{x\}$, where $\mathbf{Qn} \in \{\langle\langle x \rangle\rangle, \llbracket x \rrbracket \mid x \in \text{Vr}\}$;
8. $\text{free}((a, x)\varphi) \triangleq \text{free}(\varphi)$, if $a \notin \text{free}(\varphi)$, where $a \in \text{Ag}$ and $x \in \text{Vr}$;
9. $\text{free}((a, x)\varphi) \triangleq (\text{free}(\varphi) \setminus \{a\}) \cup \{x\}$, if $a \in \text{free}(\varphi)$, where $a \in \text{Ag}$ and $x \in \text{Vr}$;
10. $\text{free}(\mathbf{K}_a\varphi), \text{free}(\mathbf{E}_A\varphi), \text{free}(\mathbf{D}_A\varphi), \text{free}(\mathbf{C}_A\varphi) \triangleq \emptyset$, where $a \in \text{Ag}$ and $A \subseteq \text{Ag}$.

A formula φ without free agents (resp. variables), i.e., with $\text{free}(\varphi) \cap \text{Ag} = \emptyset$ (resp. $\text{free}(\varphi) \cap \text{Vr} = \emptyset$), is called agent-closed (resp. variable-closed). If φ is both agent-closed and variable-closed, it is called a sentence.

Note that strategies can be *shared* among the agents. However, the interpreted systems on which we will later define the semantics allow agents to use different, possibly disjoint, sets of actions and thus also different sets of strategies. Therefore, to determine the set of strategies over which a variable can range, we need to know which agents are bound to it. We thus introduce the set of *sharing agents* that refers, intuitively, to the agents associated with the same given variable within a formula.

Definition 3 (Sharing Agents). *The set of sharing agents of an SLK formula w.r.t. a variable is given by the function $\text{shr}: \text{SLK} \times \text{Vr} \rightarrow 2^{\text{Ag}}$ defined inductively as follows:*

1. $\text{shr}(\perp, x), \text{shr}(\top, x) \triangleq \emptyset$;
2. $\text{shr}(p, x) \triangleq \emptyset$, where $p \in \text{AP}$;
3. $\text{shr}(\text{Op}\varphi, x) \triangleq \text{shr}(\varphi, x)$, where $\text{Op} \in \{\neg, \mathbf{X}\} \cup \{\langle\langle y \rangle\rangle, \llbracket y \rrbracket \mid y \in \text{Vr}\} \cup \{\mathbf{K}_a \mid a \in \text{Ag}\} \cup \{\mathbf{E}_A, \mathbf{D}_A, \mathbf{C}_A \mid A \subseteq \text{Ag}\}$;

4. $\text{shr}(\varphi_1 \text{ Op } \varphi_2, x) \triangleq \text{shr}(\varphi_1, x) \cup \text{shr}(\varphi_2, x)$, where $\text{Op} \in \{\wedge, \vee, \cup, \text{R}\}$;
5. $\text{shr}((a, x)\varphi, x) \triangleq \{a\} \cup \text{shr}(\varphi, x)$;
6. $\text{shr}((a, y)\varphi, x) \triangleq \text{shr}(\varphi, x)$, where $y \in \text{Vr} \setminus \{x\}$;

Note that, for conciseness, we assume throughout this work that every variable is quantified at most once in a given formula. This can be easily ensured by renaming variables which are not free in the formula (e.g., $(a_1, x)\langle\langle x \rangle\rangle(a_2, x)\text{X}p \equiv (a_1, x)\langle\langle y \rangle\rangle(a_2, y)\text{X}p \not\equiv (a_1, y)\langle\langle x \rangle\rangle(a_2, x)\text{X}p$).

2.2. Multi-Agent Model

As mentioned before, we provide a semantics for SLK based on interpreted systems [2]. Differently from standard interpreted systems, here we take an agent's local states to be composed of private, or *internal*, states and states resulting from portions of the environment that are visible to the agent in question. This is similar to variants of interpreted systems such as broadcasting systems [30].

Definition 4 (Interpreted Systems). *Let $\text{Ag} \triangleq \{1, \dots, n\} \cup \{\text{Env}\}$ be a set of agents, or players, where Env is a distinguished element representing the environment and $\Sigma \triangleq \{1, \dots, n\}$ is the set of proper agents, and AP a finite non-empty set of atomic propositions. An interpreted system is a tuple $\mathcal{I} \triangleq \langle (\text{St}_a, \text{Ac}_a, \text{P}_a, \text{tr}_a)_{a \in \text{Ag}}, \text{I}, \text{h} \rangle$ whose components are formally defined as follows:*

- St_a is a finite, non-empty set of local states of agent $a \in \text{Ag}$. For each proper agent $i \in \Sigma$, we assume that $\text{St}_i \triangleq \text{St}_i^{\text{P}} \times \text{St}_{\text{Env}}^{\text{vis}_i}$, where St_i^{P} is the set of internal states of agent i and $\text{St}_{\text{Env}}^{\text{vis}_i}$ is an image of the set of the environment states visible to agent i via the agent's visibility function $\text{vis}_i: \text{St}_{\text{Env}} \rightarrow \text{St}_{\text{Env}}^{\text{vis}_i}$. For conciseness, set $\text{St}_{\text{Env}}^{\text{P}} \triangleq \text{St}_{\text{Env}}$.

A tuple $s = (s_1^{\text{P}}, \dots, s_n^{\text{P}}, s_{\text{Env}}) \in \text{St} \triangleq \text{St}_1^{\text{P}} \times \dots \times \text{St}_n^{\text{P}} \times \text{St}_{\text{Env}}$ is called a global state. The symbols $\mathfrak{s}_i(s) \triangleq (s_i^{\text{P}}, \text{vis}_i(s_{\text{Env}}))$ and $\mathfrak{s}_i^{\text{P}}(s) \triangleq s_i^{\text{P}}$ represent the local and internal state of proper agent $i \in \Sigma$ in the global state s , respectively. Again, set $\mathfrak{s}_{\text{Env}}(s) \triangleq \mathfrak{s}_{\text{Env}}^{\text{P}}(s) \triangleq s_{\text{Env}}$ for conciseness.

- Ac_a is a finite non-empty set of actions that agent $a \in \text{Ag}$ can perform. By $\text{Dc} \triangleq \text{Ac}_1 \times \dots \times \text{Ac}_n \times \text{Ac}_{\text{Env}}$ we denote the set of decisions, or joint actions, of all agents. The symbol $c_a(\delta) \in \text{Ac}_a$ represents the action of agent $a \in \text{Ag}$ in the decision $\delta \in \text{Dc}$. Furthermore, $\text{Ac} \triangleq \bigcup_{a \in \text{Ag}} \text{Ac}_a$ and $\text{Ac}_A \triangleq \bigcap_{a \in A} \text{Ac}_a$ are the sets of total actions and shared actions of the agents in $A \subseteq \text{Ag}$.
- $\text{P}_a: \text{St}_a \rightarrow 2^{\text{Ac}_a} \setminus \{\emptyset\}$ is the protocol of agent $a \in \text{Ag}$, i.e., a function that associates each local state $s_a \in \text{St}_a$ with the non-empty set of actions $\text{P}_a(s_a)$ available to the agent a . The global protocol $\text{P}: \text{St} \rightarrow 2^{\text{Dc}}$ is defined as $\text{P}(s) \triangleq \{\delta \in \text{Dc} \mid \forall a \in \text{Ag}. c_a(\delta) \in \text{P}_a(\mathfrak{s}_a(s))\}$ for all global states $s \in \text{St}$.

- $\text{tr}_a : \text{St}_a \times \text{Dc} \rightarrow \text{St}_a^P$ is the transition function, or evolution function, which maps every local state $s_a \in \text{St}_a$ of an agent $a \in \text{Ag}$ and a decision $\delta \in \text{Dc}$ to a new internal state of the same agent. The resulting new local state of an agent is determined by both her internal state and the visible portion of the amended environment state. The evolution of the whole system is described by a global transition function $\text{tr} : \text{St} \times \text{Dc} \rightarrow \text{St}$ defined as follows: $\text{tr}(s, \delta) = s'$ iff, for all agents $a \in \text{Ag}$, it holds that $\text{tr}_a(\mathfrak{s}_a(s), \delta) = \mathfrak{s}_a^P(s')$.
- $I \subseteq \text{St}$ is a finite non-empty set of initial global states. W.l.o.g., we assume that St only contains the global states that are reachable from those in I by following the transitions prescribed by the global transition function tr .
- $\mathfrak{h} : \text{AP} \rightarrow 2^{\text{St}}$ is a valuation function that maps each atomic proposition $p \in \text{AP}$ to the set of global states $\mathfrak{h}(p)$ in which it is true.

Modelling agent knowledge in multi-agent systems is a critical aspect in numerous scenarios, since it allows to reason about the decision making process of individual agents, as well as the interactions within a group of agents [28]. The knowledge of (groups of) agents is represented using the four epistemic connectives K , E , D and C already considered for the syntax of SLK. Traditionally, the associated semantics is based on the concept of *epistemic accessibility* [31]. Intuitively, two states are *epistemically accessible* if they are indistinguishable by an agent. This is formalised by means of a binary epistemic accessibility relation on the states of the system. In the case of an interpreted system $\mathcal{I} = \langle (\text{St}_a, \text{Ac}_a, \text{P}_a, \text{tr}_a)_{a \in \text{Ag}}, I, \mathfrak{h} \rangle$, the individual epistemic accessibility relation $\sim_a \subseteq \text{St} \times \text{St}$ of an agent $a \in \text{Ag}$ is naturally induced by her local states in St_a : two global states $s_1, s_2 \in \text{St}$ are indistinguishable by agent a , in symbols $s_1 \sim_a s_2$, iff $\mathfrak{s}_a(s_1) = \mathfrak{s}_a(s_2)$, *i.e.*, iff her local states in s_1 and s_2 are the same. This relation provides the meaning for the modality K_a . The remaining modalities E_A , D_A and C_A base their semantics on suitable combinations of the knowledge that the agents in the set $A \subseteq \text{Ag}$ have:

- The *group* epistemic accessibility relation \sim_A^E is defined as $\sim_A^E \triangleq \bigcup_{a \in A} \sim_a$, *i.e.*, $s_1 \sim_A^E s_2$ iff the local states $\mathfrak{s}_a(s_1)$ and $\mathfrak{s}_a(s_2)$ of at least one agent $a \in A$ in the global states $s_1, s_2 \in \text{St}$ are the same.
- The *distributed* epistemic accessibility relation \sim_A^D is defined as $\sim_A^D \triangleq \bigcap_{a \in A} \sim_a$, *i.e.*, $s_1 \sim_A^D s_2$ iff the local states $\mathfrak{s}_a(s_1)$ and $\mathfrak{s}_a(s_2)$ of all agents $a \in A$ in the global states $s_1, s_2 \in \text{St}$ are the same.
- The *common* epistemic accessibility relation \sim_A^C is defined as $\sim_A^C \triangleq (\sim_A^E)^+$, *i.e.*, it is the transitive closure of the group accessibility relation.

As standard [28], we say that an agent *knows* a fact iff it is true at all the worlds the agent considers possible, *i.e.*, in all states that the agent cannot distinguish from the current one.

We now recall the standard notions of *strategy*, *profile* and *play*. First observe that an interpreted system \mathcal{I} naturally induces a graph $\mathcal{G}(\mathcal{I}) \triangleq \langle \text{St}, Ed \rangle$, where the edge relation $Ed \triangleq \{(s, \text{tr}(s, \delta)) \in \text{St} \times \text{St} \mid \delta \in \text{Dc}\}$ is obtained by simply deleting all decisions on the transitions. A *path* $\pi \in \text{St}^\omega$ in \mathcal{I} is an infinite path in $\mathcal{G}(\mathcal{I})$. By π_i we will denote the i -th state along π . A *strategy* is a function $\sigma \in \text{Str} \triangleq \text{St} \rightarrow \text{Ac}$ prescribing which action has to be performed in a certain global state. We say that σ is *coherent w.r.t.* an agent $a \in \text{Ag}$ (a -coherent, for short) if, for each possible global state $s \in \text{St}$, the action that σ prescribes is available to a , *i.e.*, $\sigma(s) \in \text{P}_a(\mathfrak{s}_a(s))$. The notion of coherence can be lifted to a set of agents $A \subseteq \text{Ag}$ (A -coherent, for short) by simply requiring that σ is a -coherent, for every agent $a \in A$. By $\text{Str}_A \subseteq \text{Str}$ we denote the set of all A -coherent strategies. A strategy σ is also *uniform w.r.t.* the agent a (a -uniform, for short) if, for every pair of global states $s_1, s_2 \in \text{St}$ with $s_1 \sim_a s_2$, it holds that $\sigma(s_1) = \sigma(s_2)$. By $\text{UStr}_A \subseteq \text{Str}_A$ we denote the set of all a -coherent uniform strategies, for the agents $a \in A \subseteq \text{Ag}$. A *strategy profile* $\xi \in \text{Prf} \subseteq \text{Ag} \rightarrow \text{Str}$ specifies for each agent a coherent strategy. Given a profile $\xi \in \text{Prf}$ and an agent $a \in \text{Ag}$, the action $\xi(a)(s) \in \text{P}_a(\mathfrak{s}_a(s))$ determines which action a has chosen to perform in a global state $s \in \text{St}$. To identify, instead, the whole decision in s , we apply the flipping function $\hat{\xi} : \text{St} \rightarrow \text{Dc}$ defined as follows: $\hat{\xi}(s)(a) \triangleq \xi(a)(s)$. A path π is a *play w.r.t.* a profile ξ (ξ -*play*, for short) iff, for all $i \in \mathbb{N}$, there exists a decision $\delta \in \text{P}(\pi_i)$ such that $\delta = \hat{\xi}(\pi_i)$ and $\pi_{i+1} = \text{tr}(\pi_i, \delta)$, *i.e.*, π_{i+1} is the successor of π_i induced by the agent decision δ prescribed by the profile ξ in the same state. Note that, given a state s , the deterministic structure of the interpreted system ensures that there exists exactly one ξ -play π starting in s . Such a play is called (ξ, s) -*play* and is denoted by $\text{play}(\xi, s)$.

As an example, consider a generalised version of the cake-cutting problem [32], in which n agents share a cake of size d have to slice it in a fairly way. In particular, we would like to formalise a protocol allocating the slices fairly. The idea is a slight generalisation of the procedure proposed in [32] based on a *divide et impera* approach, where the agents take turns to slice the cake, while the environment responds by trying to ensure the cake is divided fairly. More precisely, we assume that, at each even round, the n agents concurrently choose the size of the piece they would each like to receive and, at each odd round, the environment makes the cuts and assigns each piece to a subset of the agents. The environment does this based on agent choices made in the previous turn. However, its action is not necessarily dictated by what agents want, since their willing may represent a non-fair partitioning of the cake. Due to the structure-by-rounds of our approach, the problem of cutting a cake of size d among n agents is divided into several simpler problems, in which pieces of size $d' < d$ have to be split among $n' < n$ agents. The resulting multi-player game terminates once each agent receives a different slice. For simplicity, we assume that we can only cut the cake into pieces of integer size. The formalisation of the corresponding interpreted system \mathcal{I} follows:

- The internal states of the environment have to maintain information about the pieces of cake already sliced, their assignment to the agents, as well

as the associated sizes. When it is the environment's turn to move, they also have to record the requests for the sizes of the pieces made by the agents in the previous turn. To distinguish between the turn of the agents and that of the environment, we make use of a flag \mathbf{A} that identifies those states in which the agents need to make a choice. Formally, we have

$$\text{St}_{\text{Env}} = \{((\mathbf{f}, \mathbf{g}), \alpha) \in \mathbf{S} \times \mathbf{A} \mid \alpha \neq \mathbf{A} \Rightarrow \forall a \in [1, n]. \alpha(a) \leq \mathbf{g}(\mathbf{f}(a))\}.$$

where $\mathbf{A} = \{\mathbf{A}\} \cup [1, n] \rightarrow [0, d]$, $\mathbf{S} = \{(\mathbf{f}, \mathbf{g}) \in \mathbf{Z} \mid \exists h \in [1, n]. \text{rng}(\mathbf{f}) = \text{dom}(\mathbf{g}) = [1, h] \wedge \sum_{i \in \text{dom}(\mathbf{g})} \mathbf{g}(i) \leq d\}$ and $\mathbf{Z} = ([1, n] \rightarrow [1, n]) \times ([1, n] \rightarrow [0, d])$. Intuitively, the function \mathbf{f} assigns to each proper agent $a \in [1, n]$ the $\mathbf{f}(a)$ -th piece of cake having size $\mathbf{g}(\mathbf{f}(a))$. Obviously, the sum of all sizes cannot be greater than d , but it can be strictly less than this value, since the environment can remove parts of the cake that cannot be split evenly among the agents. In addition, if it is not the agents' turn, *i.e.*, $\alpha \neq \mathbf{A}$, the function α returns, for each agent $a \in [1, n]$, her own desired size $\alpha(a)$ for the piece of cake to cut from the $\mathbf{f}(a)$ -th piece assigned to her. The actions of the environment are all possible splittings and assignments of the cake together with a nothing-to-do operation, *i.e.*, $\text{Ac}_{\text{Env}} = \mathbf{S} \cup \{\perp\}$. To complete the definition of the environment, we first need to define a partial order on the first component of a state, which allows to identify all possible continuations of a given state. Formally, for all $(\mathbf{f}, \mathbf{g}), (\mathbf{f}', \mathbf{g}') \in \mathbf{S}$, we write $(\mathbf{f}', \mathbf{g}') \preceq (\mathbf{f}, \mathbf{g})$ iff the following two conditions hold:

- If $\mathbf{f}'(a_1) = \mathbf{f}'(a_2)$, then $\mathbf{f}(a_1) = \mathbf{f}(a_2)$, *i.e.*, if agents $a_1, a_2 \in [1, n]$ share the same piece of cake in \mathbf{f}' then they do the same in \mathbf{f} .
- $\sum_{a \in \mathbf{f}^{-1}\{i\}} \mathbf{g}'(\mathbf{f}'(a)) \leq \mathbf{g}(i)$ for all $i \in \text{dom}(\mathbf{g})$, *i.e.*, the agents sharing the same piece i *w.r.t.* \mathbf{f} , have pieces *w.r.t.* \mathbf{f}' whose total size in \mathbf{g}' does not exceed that of i in \mathbf{g} .

Intuitively, $(\mathbf{f}', \mathbf{g}') \preceq (\mathbf{f}, \mathbf{g})$ if \mathbf{f}' is a refinement of \mathbf{f} , *i.e.*, \mathbf{f}' is a finer partition than \mathbf{f} , and \mathbf{g}' is coherent with the size of the pieces of the previous cut. At this point, we can formalise the protocol for Env as follows: $\text{P}_{\text{Env}}(((\mathbf{f}, \mathbf{g}), \mathbf{A})) = \{\perp\}$ and $\text{P}_{\text{Env}}(((\mathbf{f}, \mathbf{g}), \alpha)) = \{(\mathbf{f}', \mathbf{g}') \in \mathbf{S} \mid (\mathbf{f}', \mathbf{g}') \preceq (\mathbf{f}, \mathbf{g})\}$ if $\alpha \neq \mathbf{A}$. This means that, when it is the environment's turn, it can choose one of the possible finer continuations of a given state. Finally, for the transition function we have $\text{tr}(((\mathbf{f}, \mathbf{g}), \mathbf{A}), \delta) = ((\mathbf{f}, \mathbf{g}), \alpha)$, where $\alpha(a) = \mathbf{c}_a(\delta)$, for all $a \in [1, n]$, and $\text{tr}(((\mathbf{f}, \mathbf{g}), \alpha), \delta) = (\mathbf{c}_{\text{Env}}(\delta), \mathbf{A})$ when $\alpha \neq \mathbf{A}$. Informally, when it is the agents' turn, the environment just registers the agents' choices in its internal state. When it is its turn instead, it modifies its state by executing its own action $\mathbf{c}_{\text{Env}}(\delta)$.

- All information about the cake required by an agent $a \in [1, n]$ is already contained in the environment state and visible to to all agents, *i.e.*, $\text{vis}_a(s_{\text{Env}}) = s_{\text{Env}}$ for all environment states $s_{\text{Env}} \in \text{St}_{\text{Env}}$. Hence, the agent does not need any internal state, *i.e.*, $\text{St}_a^{\text{P}} = \{\mathbf{I}\}$. The actions are simply the dimensions of the possible pieces of cake together

with a nothing-to-do operation, *i.e.*, $Ac_a = [1, d] \cup \{\perp\}$. In the environment's turn, *i.e.*, when $\alpha \neq \mathbf{A}$, the only possible executable action is \perp , *i.e.*, $P_a((\mathbf{I}, ((\mathbf{f}, \mathbf{g}), \alpha))) = \{\perp\}$. Otherwise, we have $P_a((\mathbf{I}, ((\mathbf{f}, \mathbf{g}), \alpha))) = [0, \mathbf{g}(\mathbf{f}(a))]$, *i.e.*, a can choose a size from 0 to the current dimension of the $\mathbf{f}(a)$ -th piece. Because of the singleton internal state \mathbf{I} , we have $\text{tr}_a((\mathbf{I}, s_{\text{Env}}), \delta) = \mathbf{I}$ for all environment states $s_{\text{Env}} \in \text{St}_{\text{Env}}$ and decisions $\delta \in \text{Dc}$.

- The set of initial global states is the singleton $\mathbf{I} = \{(\mathbf{I}, \dots, \mathbf{I}, ((\mathbf{f}, \mathbf{g}), \mathbf{A}))\}$, where the entire cake is assigned to all agents, *i.e.*, $\mathbf{f}(a) = 1$ for all $a \in [1, n]$, and its size is $\mathbf{g}(1) = d$.
- The set of atomic propositions, which represent all possible sizes of a piece of cake an agent can receive, is $\text{AP} = [1, n] \times [1, d] \cup \{\#\}$. For example, $\langle a, i \rangle$ is true iff agent $a \in [1, n]$ receives a piece of cake of size $i \in [1, d]$, *i.e.*, $\mathbf{g}(\mathbf{f}(a)) = i$. In addition, the symbol $\#$ is used to label the terminal states in which each agent obtains a different piece. The labelling is defined as follows: $\mathbf{h}(\langle a, i \rangle) = \{(\mathbf{I}, \dots, \mathbf{I}, ((\mathbf{f}, \mathbf{g}), \alpha)) \in \text{St} \mid i = \mathbf{g}(\mathbf{f}(a))\}$ for all $\langle a, i \rangle \in \text{AP}$, $\alpha \in \mathbf{A}$ and $\mathbf{h}(\#) = \{(\mathbf{I}, \dots, \mathbf{I}, ((\mathbf{f}, \mathbf{g}), \alpha)) \in \text{St} \mid \forall a_1, a_2 \in [1, n]. a_1 \neq a_2 \Rightarrow \mathbf{f}(a_1) \neq \mathbf{f}(a_2)\}$ for all $\alpha \in \mathbf{A}$.

2.3. Semantics

We can now continue with the formalisation of SLK semantics. Similarly to first-order logic, the interpretation of a formula makes use of an assignment function which generally associates placeholders with some elements of the quantification domain. For SLK, as in SL, an *assignment* is a partial function $\chi \in \text{Asg} \triangleq (\text{Vr} \cup \text{Ag}) \rightarrow \text{Str}$ mapping variables and agents to strategies such that $\chi(a)$ is a -coherent for every agent $a \in \text{dom}(\chi) \cap \text{Ag}$. By Asg_{P} we denote the set of assignments defined over the placeholders in $\text{P} \subseteq \text{Ag} \cup \text{Vr}$. An assignment χ is *complete* iff it is defined on all agents, *i.e.*, $\text{Ag} \subseteq \text{dom}(\chi)$. In this case, it directly identifies the profile $\chi_{\upharpoonright \text{Ag}}$ given by the restriction of χ to Ag . In addition, $\chi[e \mapsto \sigma]$, with $e \in \text{Vr} \cup \text{Ag}$ and $\sigma \in \text{Str}$, is the assignment defined on $\text{dom}(\chi[e \mapsto \sigma]) = \text{dom}(\chi) \cup \{e\}$ which differs from χ only in the fact that e is associated with σ . Formally, $\chi[e \mapsto \sigma](e) \triangleq \sigma$ and $\chi[e \mapsto \sigma](e') \triangleq \chi(e')$ for all $e' \in \text{dom}(\chi) \setminus \{e\}$. Obviously, in case e is an agent, the previous operation is well-defined only if σ is e -coherent.

Given a complete assignment $\chi \in \text{Asg}_{\text{P}}$ over a set of placeholders $\text{P} \subseteq \text{Ag} \cup \text{Vr}$ and a state $s \in \text{St}$, we define the i -th translation of (χ, s) *w.r.t.* an index $i \in \mathbb{N}$ as the pair $(\chi, s)^i \triangleq (\chi, s') \in \text{Asg}_{\text{P}} \times \text{St}$, where $s' = \pi_i$ is the i -th state along the corresponding play $\pi = \text{play}(\chi_{\upharpoonright \text{Ag}}, s)$.

We can now formally introduce the semantics of SLK.

Definition 5 (Semantics). *Given an interpreted system \mathcal{I} , for all SLK formulas $\varphi \in \text{SLK}$, states $s \in \text{St}$ and assignments $\chi \in \text{Asg}$ such that $\text{free}(\varphi) \subseteq \text{dom}(\chi)$ and $\chi(x)$ is $\text{shr}(\varphi, x)$ -coherent for all $x \in \text{dom}(\chi) \cap \text{Vr}$, the satisfaction relation $\mathcal{I}, (\chi, s) \models \varphi$ is inductively defined as follows:*

1. $\mathcal{I}, (\chi, s) \models p$ if $s \in \mathbf{h}(p)$, with $p \in \text{AP}$.
2. Truth values and Boolean operators are interpreted as usual.
3. For all variables $x \in \text{Vr}$, we have that:
 - (a) $\mathcal{I}, (\chi, s) \models \langle\langle x \rangle\rangle\varphi$ if there exists a $\text{shr}(\varphi, x)$ -coherent uniform strategy $\sigma \in \text{UStr}_{\text{shr}(\varphi, x)}$ such that $\mathcal{I}, (\chi[x \mapsto \sigma], s) \models \varphi$;
 - (b) $\mathcal{I}, (\chi, s) \models \llbracket x \rrbracket\varphi$ if, for all $\text{shr}(\varphi, x)$ -coherent uniform strategies $\sigma \in \text{UStr}_{\text{shr}(\varphi, x)}$, it holds that $\mathcal{I}, (\chi[x \mapsto \sigma], s) \models \varphi$.
4. For an agent $a \in \text{Ag}$ and a variable $x \in \text{Vr}$, we have that: $\mathcal{I}, (\chi, s) \models (a, x)\varphi$ if $\mathcal{I}, (\chi[a \mapsto \chi(x)], s) \models \varphi$.
5. For an agent $a \in \text{Ag}$, we have that: $\mathcal{I}, (\chi, s) \models \mathbf{K}_a \varphi$ if, for all states $s' \in \text{St}$ with $s \sim_a s'$, it holds that $\mathcal{I}, (\emptyset, s') \models \varphi$.
6. For a set of agents $A \subseteq \text{Ag}$, we have that:
 - (a) $\mathcal{I}, (\chi, s) \models \mathbf{E}_A \varphi$ if, for all states $s' \in \text{St}$ with $s \sim_A^E s'$, it holds that $\mathcal{I}, (\emptyset, s') \models \varphi$;
 - (b) $\mathcal{I}, (\chi, s) \models \mathbf{D}_A \varphi$ if, for all states $s' \in \text{St}$ with $s \sim_A^D s'$, it holds that $\mathcal{I}, (\emptyset, s') \models \varphi$;
 - (c) $\mathcal{I}, (\chi, s) \models \mathbf{C}_A \varphi$ if, for all states $s' \in \text{St}$ with $s \sim_A^C s'$, it holds that $\mathcal{I}, (\emptyset, s') \models \varphi$.
7. Finally, if the assignment χ is also complete, we have that:
 - (a) $\mathcal{I}, (\chi, s) \models \mathbf{X} \varphi$ if $\mathcal{I}, (\chi, s)^1 \models \varphi$;
 - (b) $\mathcal{I}, (\chi, s) \models \varphi_1 \mathbf{U} \varphi_2$ if there is an index $i \in \mathbb{N}$ such that $\mathcal{I}, (\chi, s)^i \models \varphi_2$ and, for all indices $j \in \mathbb{N}$ with $j < i$, it holds that $\mathcal{I}, (\chi, s)^j \models \varphi_1$;
 - (c) $\mathcal{I}, (\chi, s) \models \varphi_1 \mathbf{R} \varphi_2$ if, for all indices $i \in \mathbb{N}$, it holds that $\mathcal{I}, (\chi, s)^i \models \varphi_2$ or there is an index $j \in \mathbb{N}$ with $j < i$ such that $\mathcal{I}, (\chi, s)^j \models \varphi_1$.

As the satisfaction of a sentence φ does not depend on the assignments, we omit them and write $\mathcal{I}, s \models \varphi$ for an arbitrary state $s \in \text{St}$, and $\mathcal{I} \models \varphi$ when $\mathcal{I}, s \models \varphi$ holds for all initial states $s \in \text{I}$ in \mathcal{I} .

SLK semantics differs from SL's one in the following two aspects, both of which are consequences of using incomplete information: (i) it is defined on uniform memoryless strategies; (ii) it supports epistemic operators. Note that, while epistemic modalities increase the expressive power of SLK *w.r.t.* SL, the imperfect recall feature limits the agents' behaviours.

As it is clear from the interpretation of the epistemic operators and the associated syntactic restriction on free placeholders (*i.e.*, all the epistemic operators can only predicate on sentences), we assume that agents are not aware of the current strategy assignment. In fact, in as far as the epistemic accessibility relation is concerned, they are not aware of the strategy chosen at a state. This assumption is implemented by the empty assignments in the inductive definitions of the epistemic cases.

Also observe that SLK strictly subsumes all logics in the ATL* hierarchy with incomplete information and imperfect recall imposed on all agents, including CTLK and ATLK. Moreover, it allows us to express properties that cannot be formalised in any of the previous logics, including the original SL. For example, given an interpreted system with agents $\text{Ag} = \{a, b\}$, the SLK sentence

$$\mathbf{E}_{\{a,b\}} \llbracket \epsilon \rrbracket \llbracket x \rrbracket (\text{Env}, e)(a, x)(b, x) [\mathbf{G} \mathbf{F} p \wedge \mathbf{G} \mathbf{F} \neg p]$$

expresses that a and b both know that if they use the same strategy, the atomic proposition p will be infinitely often true and infinitely often false.

Lastly, note that in the formalisation above all agents in the system are assumed to adhere to memoryless strategies. This is in contrast with memoryless ATL formalisms whereby agents the coalition in the formula are allowed to follow memoryfull strategies. We leave combinations of memoryless and memoryfull strategies for further work.

We now illustrate SLK's expressive power to express *Nash equilibria* by continuing our analysis of the cake cutting problem. Recall that we use atomic propositions $\langle a, i \rangle \in [1, n] \times [1, d]$ to indicate that agent a gets a piece of cake of size i . In addition, the symbol $\#$ is used to denote the terminal states in which a final division of the cake is reached. The existence of a protocol for the cake-cutting problem can be expressed by the following SLK specification

$$\varphi = \langle\langle x \rangle\rangle (\varphi_F \wedge \varphi_S)$$

where:

- $\varphi_F = \llbracket y_1 \rrbracket \cdots \llbracket y_n \rrbracket (\psi_{NE} \rightarrow \psi_E)$ ensures that the protocol x is *fair*, *i.e.*, all possible Nash equilibria (y_1, \dots, y_n) of the agents guarantee equality of splitting;
- $\varphi_S = \langle\langle y_1 \rangle\rangle \cdots \langle\langle y_n \rangle\rangle \psi_{NE}$ ensures that the protocol has a *solution*, *i.e.*, there is at least one Nash equilibrium;
- $\psi_{NE} = \flat \bigwedge_{a=1}^n \left(\bigwedge_{i=1}^d (\langle\langle z \rangle\rangle (a, z) p_a^i) \rightarrow \left(\bigvee_{j=i}^d p_a^j \right) \right)$ ensures that, if agent a has a strategy z that allows her to obtain a piece of cake of size i once the strategies of the other agents are fixed, she is already able to obtain a slice of size $j \geq i$ by means of her original strategy y_a ; in other words, (y_1, \dots, y_n) is a *Nash equilibrium*;
- $\psi_E = \flat \bigwedge_{a=1}^n p_a^{\lfloor d/n \rfloor}$ ensures that agent a is able to obtain a piece of size $\lfloor d/n \rfloor$;
- $\flat = (\text{Env}, x)(1, y_1) \cdots (n, y_n)$ and $p_a^i = F(\# \wedge \langle a, i \rangle)$ are auxiliary abbreviations.

Observe that we can perform the synthesis of cake-cutting protocols satisfying specific desired properties, by simply conjoining the corresponding SLK formalisation with the formulas φ_F and φ_S in φ .

Note, that, as above, when reasoning about Nash equilibria all agents in the system are assumed to follow memoryless strategies.

2.4. Model Checking

We now introduce the model-checking problem for SLK. In particular, we show that its imperfect-recall semantics allows us to obtain a decision procedure whose complexity is, in comparison to that of SL with complete information,

considerably simpler *w.r.t.* the length of specification, but notably harder *w.r.t.* the size of the system. Indeed, the perfect-recall semantics of SL [23] induces the related model checking problem to be polynomial in data complexity and non-elementary in formula complexity [26]. On the contrary, we prove that SLK model checking is PSPACE-COMPLETE *w.r.t.* both input dimensions.

Definition 6 (Model Checking). *Given an interpreted system \mathcal{I} , a state $s \in \text{St}$, an assignment χ and an SLK sentence $\varphi \in \text{SLK}$, the model-checking problem involves determining whether $\mathcal{I}, (\chi, s) \models \varphi$ holds.*

The original algorithmic procedure for SL follows an automata-theoretic approach [33], reducing the decision problem for the logic to the emptiness problem of a tree automaton, where the LTL properties are verified by means of an embedded word automaton obtained by a variation of the classic Vardi-Wolper construction [34]. The non-elementariness *w.r.t.* the length of the SL specification is due to the alternation of the memoryful strategy quantifiers that requires alternating projection operations on the automaton, each of which induces an exponential blow-up. This construction is, however, independent of the underlying model, a key feature that is crucial for the exhibited data complexity.

Unfortunately, due to the undecidability result proved for ATL with memoryful strategies under incomplete information [27], we cannot hope to use the same approach of standard SL when agents only have partial information about the global states. In fact, SLK avoids the undecidability of the model-checking problem by precisely exploiting its imperfect recall semantics. Effectively, we are removing agent memory and forcing them to choose their actions purely based on their current local states. In this way, the domain of strategies over which the quantifications have to range is necessarily finite, due to the finiteness of the underlying model. This fact allows us to apply a completely different approach, where one simply iterates over all possible memoryless strategies. Once the strategies are assigned to all corresponding variables in the SLK formula, we can then project them onto the interpreted structure obtaining a labelled graph, where the LTL property is verified recursively on the structure of the formula, by mimicking the semantics definition. It is quite immediate to observe that this procedure only requires a polynomial amount of space *w.r.t.* both the size of the model and the length of the SL specification. However, it is important to note that this approach is not independent of the model under verification. Even more, there is no way to avoid this unless PTIME = PSPACE. We indeed prove that the SLK model-checking problem is hard for the latter complexity class.

Since the model checking of ATL* with imperfect recall is known to be PSPACE-COMPLETE [35], one may be tempted to immediately conclude that the same holds for SLK. However, there is an important difference between ATL* and SLK semantics: the former assigns memoryless strategies only to the existentially quantified agents, whereas the latter assigns memoryless strategies to all agents. Informally, an ATL* expression $\langle\langle A \rangle\rangle\psi$ means that there exist memoryless strategies for agents in A such that, no matter how the other agents behave, the property ψ holds. The corresponding SLK formula $\{\langle\langle x_a \rangle\rangle(a, x_a)\}_{a \in A} \{\llbracket y_a \rrbracket\}$

$(a, y_a)\}_{a \in \text{Ag} \setminus A} \varphi$ has a slightly different meaning: there exist memoryless strategies for agents in A such that, for all memoryless strategies of the remaining agents, ψ holds. Intuitively, SLK restricts the universally quantified agents more than ATL^* does¹. Therefore, SLK cannot inherit the hardness result from ATL^* , which is, in its turn, inherited from that of LTL. Instead, we show a direct reduction from the satisfiability problem of quantified Boolean formulas (QBF, for short), which was proved to be complete for PSPACE in [36]. This reduction is loosely inspired by the one used to prove a similar result for Substructure Temporal Logic over finite models [37].

Before concluding this section with a proof of the model-checking result for SLK, we stress that although memoryless strategies are less powerful than the memoryful ones, they are more compact and, therefore, easier to handle. These features are desirable in several scenarios. Consider, for example, the situation in which we want to synthesise the correct behaviour of a simple hardware device with a fixed amount of memory. In this case, we can model each possible memory settings of the device as a local state and allow arbitrary transitions between them. Informally, we let the agent representing the device decide what she wants to remember. Thus, it is appropriate to treat this agent as being memoryless, since her real memory is already encoded in her local state space. We then model check the property describing the desired behaviour. If we succeed, we can extract from the procedure a memoryless strategy as a witness, which ensures that the fixed-size memory is sufficient for the agent to achieve the required goal.

Theorem 1 (Model Checking). *The model-checking problem for SLK is PSPACE-COMplete w.r.t. both the size of the model and the length of the specification.*

Proof. The proof is divided into two parts. In the first one we give the PSPACE decision procedure; in the second we provide the PSPACE hardness proof by means of a reduction from QBF satisfiability.

To verify that a given SLK formula φ is satisfied over an interpreted system \mathcal{I} at a state $s \in \text{St}$ under a memoryless assignments $\chi \in \text{Asg}$, we make use of the PSPACE computable recursive function $\text{Verify}_{\mathcal{I}} : \text{SL} \times \text{Asg} \times \text{St} \rightarrow \{\perp, \top\}$ defined as follows, for which it holds that $\mathcal{I}, (\chi, s) \models \varphi$ iff $\text{Verify}_{\mathcal{I}}(\varphi, \chi, s) = \top$:

- $\text{Verify}_{\mathcal{I}}(\varphi, \chi, s) \triangleq \varphi$, with $\varphi \in \{\perp, \top\}$.
- $\text{Verify}_{\mathcal{I}}(p, \chi, s) \triangleq \top$ iff $s \in h(p)$, with $p \in \text{AP}$.
- $\text{Verify}_{\mathcal{I}}(\neg\varphi, \chi, s) \triangleq \top$ iff $\text{Verify}_{\mathcal{I}}(\varphi, \chi, s) = \perp$.
- $\text{Verify}_{\mathcal{I}}(\varphi_1 \wedge \varphi_2, \chi, s) \triangleq \top$ iff $\text{Verify}_{\mathcal{I}}(\varphi_1, \chi, s) = \top$ and $\text{Verify}_{\mathcal{I}}(\varphi_2, \chi, s) = \top$.

¹Note that this subtle difference in the semantics has no effect on the relationship between SL and ATL^* with perfect recall since enforcing memoryful strategies on agents does not constrain their behaviour in any way.

- $\text{Verify}_{\mathcal{I}}(\varphi_1 \vee \varphi_2, \chi, s) \triangleq \top$ iff $\text{Verify}_{\mathcal{I}}(\varphi_1, \chi, s) = \top$ or $\text{Verify}_{\mathcal{I}}(\varphi_2, \chi, s) = \top$.
- $\text{Verify}_{\mathcal{I}}(\langle\langle x \rangle\rangle \varphi, \chi, s) \triangleq \top$ iff there exists a memoryless $\text{shr}(\varphi, x)$ -coherent uniform strategy $\sigma \in \text{UStr}_{\text{shr}(\varphi, x)}$ such that $\text{Verify}_{\mathcal{I}}(\varphi, \chi[x \mapsto \sigma], s) = \top$.
- $\text{Verify}_{\mathcal{I}}(\llbracket x \rrbracket \varphi, \chi, s) \triangleq \top$ iff for all memoryless $\text{shr}(\varphi, x)$ -coherent uniform strategies $\sigma \in \text{UStr}_{\text{shr}(\varphi, x)}$ it holds that $\text{Verify}_{\mathcal{I}}(\varphi, \chi[x \mapsto \sigma], s) = \top$.
- $\text{Verify}_{\mathcal{I}}((a, x)\varphi, \chi, s) \triangleq \top$ iff $\text{Verify}_{\mathcal{I}}(\varphi, \chi[a \mapsto \chi(x)], s) = \top$.
- $\text{Verify}_{\mathcal{I}}(\mathbf{K}_a \varphi, \chi, s) \triangleq \top$ iff $\text{Verify}_{\mathcal{I}}(\varphi, \emptyset, s') = \top$, for all states $s' \in \text{St}$ with $s \sim_a s'$.
- $\text{Verify}_{\mathcal{I}}(\mathbf{E}_A \varphi, \chi, s) \triangleq \top$ iff $\text{Verify}_{\mathcal{I}}(\varphi, \emptyset, s') = \top$, for all states $s' \in \text{St}$ with $s \sim_A^E s'$.
- $\text{Verify}_{\mathcal{I}}(\mathbf{D}_A \varphi, \chi, s) \triangleq \top$ iff $\text{Verify}_{\mathcal{I}}(\varphi, \emptyset, s') = \top$, for all states $s' \in \text{St}$ with $s \sim_A^D s'$.
- $\text{Verify}_{\mathcal{I}}(\mathbf{C}_A \varphi, \chi, s) \triangleq \top$ iff $\text{Verify}_{\mathcal{I}}(\varphi, \emptyset, s') = \top$, for all states $s' \in \text{St}$ with $s \sim_A^C s'$.
- $\text{Verify}_{\mathcal{I}}(\mathbf{X} \varphi, \chi, s) \triangleq \top$ iff $\text{Verify}_{\mathcal{I}}(\varphi, \chi, s') = \top$, where $s' = \text{tr}(s, \widehat{\chi \uparrow \text{Ag}}(s))$ is the successor of the state s following the decision $\widehat{\chi \uparrow \text{Ag}}(s)$ obtained by evaluating the flipping of the profile $\chi \uparrow \text{Ag}$ on s itself.
- To compute $\text{Verify}_{\mathcal{I}}(\varphi_1 \mathbf{U} \varphi_2, \chi, s)$, we use the auxiliary bounded-recursive function $\text{Verify}_{\mathcal{I}}^{\mathbf{U}}(\varphi_1 \mathbf{U} \varphi_2, \chi, s, C)$, where $C \subseteq \text{St}$ is the set of states already visited during the verification of the property $\varphi_1 \mathbf{U} \varphi_2$. We set $\text{Verify}_{\mathcal{I}}(\varphi_1 \mathbf{U} \varphi_2, \chi, s) \triangleq \text{Verify}_{\mathcal{I}}^{\mathbf{U}}(\varphi_1 \mathbf{U} \varphi_2, \chi, s, \emptyset)$ and define the latter as follows, by exploiting the classic one-step unfolding property of the until operator:
 - if $s \in C$, then $\text{Verify}_{\mathcal{I}}^{\mathbf{U}}(\varphi_1 \mathbf{U} \varphi_2, \chi, s, C) \triangleq \perp$;
 - if $s \notin C$ and $\text{Verify}_{\mathcal{I}}(\varphi_2, \chi, s) = \top$, then $\text{Verify}_{\mathcal{I}}^{\mathbf{U}}(\varphi_1 \mathbf{U} \varphi_2, \chi, s, C) \triangleq \top$;
 - if $s \notin C$ and $\text{Verify}_{\mathcal{I}}(\varphi_2, \chi, s) = \perp$, then $\text{Verify}_{\mathcal{I}}^{\mathbf{U}}(\varphi_1 \mathbf{U} \varphi_2, \chi, s, C) \triangleq \top$ iff $\text{Verify}_{\mathcal{I}}(\varphi_1, \chi, s) = \top$ and $\text{Verify}_{\mathcal{I}}^{\mathbf{U}}(\varphi_1 \mathbf{U} \varphi_2, \chi, s', C \cup \{s\}) = \top$, where $s' = \text{tr}(s, \widehat{\chi \uparrow \text{Ag}}(s))$ is the successor of the state s following the decision $\widehat{\chi \uparrow \text{Ag}}(s)$.
- Similarly to the previous case, to compute $\text{Verify}_{\mathcal{I}}(\varphi_1 \mathbf{R} \varphi_2, \chi, s)$, we use the auxiliary bounded-recursive function $\text{Verify}_{\mathcal{I}}^{\mathbf{R}}(\varphi_1 \mathbf{R} \varphi_2, \chi, s, C)$, where $C \subseteq \text{St}$ is the set of states already visited during the verification of the property $\varphi_1 \mathbf{R} \varphi_2$. We set $\text{Verify}_{\mathcal{I}}(\varphi_1 \mathbf{R} \varphi_2, \chi, s) \triangleq \text{Verify}_{\mathcal{I}}^{\mathbf{R}}(\varphi_1 \mathbf{R} \varphi_2, \chi, s, \emptyset)$ and define the latter as follows, by exploiting the classic one-step unfolding property of the release operator:

- if $s \in C$, then $\text{Verify}_{\mathcal{I}}^R(\varphi_1 R \varphi_2, \chi, s, C) \triangleq \top$;
- if $s \notin C$ and $\text{Verify}_{\mathcal{I}}(\varphi_2, \chi, s) = \perp$, then $\text{Verify}_{\mathcal{I}}^R(\varphi_1 R \varphi_2, \chi, s, C) \triangleq \perp$;
- if $s \notin C$ and $\text{Verify}_{\mathcal{I}}(\varphi_2, \chi, s) = \top$, then $\text{Verify}_{\mathcal{I}}^R(\varphi_1 R \varphi_2, \chi, s, C) \triangleq \perp$ iff $\text{Verify}_{\mathcal{I}}(\varphi_1, \chi, s) = \perp$ and $\text{Verify}_{\mathcal{I}}^R(\varphi_1 R \varphi_2, \chi, s', \text{CU}\{s\}) = \perp$, where $s' = \text{tr}(s, \widehat{\chi}_{\uparrow \text{Ag}}(s))$ is the successor of the state s following the decision $\widehat{\chi}_{\uparrow \text{Ag}}(s)$.

The correctness of the construction can be shown by induction on the syntactic structure of the SLK formulas, by observing that a formula $\varphi_1 U \varphi_2$ is not satisfied on \mathcal{I} starting at s under χ , when the same state is found twice on the path induced by χ without proving that φ_2 holds somewhere on it before φ_1 is falsified. A similar reasoning can be used to establish the case for $\varphi_1 R \varphi_2$ as well. As far as the complexity is concerned, a single $\text{Verify}_{\mathcal{I}}(\varphi, \chi, s)$ function call (excluding recursion) requires $|\varphi| + |\text{St}| \cdot \lceil \log_2 |\text{Ac}| \rceil \cdot |\text{free}(\varphi)| + \lceil \log_2 |\text{St}| \rceil$ space to maintain its parameters. Similar reasoning can be done for $\text{Verify}_{\mathcal{I}}^U(\varphi_1 U \varphi_2, \chi, s, C)$ and $\text{Verify}_{\mathcal{I}}^R(\varphi_1 R \varphi_2, \chi, s, C)$ function calls, which require $|\varphi| + |\text{St}| \cdot \lceil \log_2 |\text{Ac}| \rceil \cdot |\text{free}(\varphi)| + \lceil \log_2 |\text{St}| \rceil + |\text{St}|$ space. Furthermore, in each recursive call (of any of the three checking functions), either φ is decomposed or C is augmented. Hence, the recursion depth cannot be more than $O(|\varphi| \cdot |\text{St}|)$ at any point in time. It follows that the model checking procedure requires only $O(|\varphi|^2 \cdot |\text{St}|^2 \cdot \lceil \log_2 |\text{Ac}| \rceil)$ space.

We now turn to the PSPACE lower bound, which is proved by means of a direct reduction from QBF satisfiability. Let $\varphi = \wp \psi$ be a QBF formula over the Boolean variables in $X = \{x_1, \dots, x_k\}$, where \wp is the quantification prefix of the form $\text{Qn}_1 x_1 \cdots \text{Qn}_k x_k$, with $\text{Qn}_a \in \{\exists, \forall\}$ for $a \in [1, k]$, and the matrix $\psi = \bigwedge_{i=1}^m D_i$ is a Boolean formula in conjunctive normal form over the variables in X , where each clause D_i can be seen as a subset of the literals in $\text{AP} = \{x, \bar{x} \mid x \in X\}$. We transform φ into a suitable SLK formula $\tilde{\varphi}$ over the set of atomic propositions AP , where, for each $x \in X$, we need two atomic propositions, one for the positive literal x and one for the negative literal $\neg x$. In the formula $\tilde{\varphi}$, the k propositional quantifications are replaced by corresponding strategy quantifications for k different proper agents. These quantifications are applied to a suitable encoding $\tilde{\psi}$ of the matrix ψ , where the verification of the conjunction of clauses is assigned to a distinguished agent, the $k+1$ -th, while that of each clause to the agent embodying the environment. Formally, we have $\tilde{\varphi} = \text{Qn}'_1 x_1 \cdots \text{Qn}'_k x_k \llbracket x_{k+1} \rrbracket \langle\langle x_{\text{Env}} \rangle\rangle (1, x_1) \cdots (k, x_k) (k+1, x_{k+1}) (\text{Env}, x_{\text{Env}}) \tilde{\psi}$, where the LTL formula $\tilde{\psi}$ is defined later in the proof and $\text{Qn}'_a x_a = \langle\langle x_a \rangle\rangle$, if $\text{Qn}_a = \exists$, and $\text{Qn}'_a x_a = \llbracket x_a \rrbracket$ otherwise. The formula $\tilde{\varphi}$ is then checked against the interpreted system \mathcal{I} described in the following paragraph, which ensures that φ is satisfiable iff $\mathcal{I} \models \tilde{\varphi}$. The correctness of this construction can be shown by induction on the number k of quantifications occurring in φ .

The idea behind the construction of \mathcal{I} is to have a global state for each clause D_i , each variable in X and each of its two possible Boolean valuations. Every proper agent a from 1 to k determines the valuation of her associated

variable x_a . Agent $k+1$ determines which clause D_i we have to focus on for the verification task. Finally, the environment agent Env tries to satisfy the chosen clause, by selecting the variable x_a with the appropriate valuation. In Figure 1, we exemplify the reduction for the QBF formula $\varphi = \forall p \exists q \exists r . (p \vee q) \wedge (\neg q \vee r) \wedge (\neg r \vee \neg p)$, by showing the internal graph structure of the environment. The formal description of \mathcal{I} follows:

- The system consists of the $k+1$ proper agents in $\Sigma = \{1, \dots, k+1\}$ together with the environment Env.
- The set of internal states of the environment is $\text{St}_{\text{Env}} = \{\mathbf{I}\} \cup \{D_1, \dots, D_m\} \cup \{x_1, \dots, x_k\}$, where \mathbf{I} represents the initial state. Since the task assigned to Env is to verify a given clause D_i by selecting one of its literals, we have $\text{Ac}_{\text{Env}} = \{\perp\} \cup \{1, \dots, k\}$ as set of actions, where \perp represents the idle action. Accordingly, Env employs the following protocol: $\text{P}_{\text{Env}}(s_{\text{Env}}) = \{a \in \{1, \dots, k\} \mid \{x_a, \bar{x}_a\} \cap D_i \neq \emptyset\}$, if $s_{\text{Env}} = D_i$ for some $i \in \{1, \dots, m\}$, and $\text{P}_{\text{Env}}(s_{\text{Env}}) = \{\perp\}$ otherwise. In this way, we are sure that Env can only select one of the indices a of the variables occurring in the chosen clause D_i . Finally, the transition function is defined as follows: $\text{tr}_{\text{Env}}(\mathbf{I}, \delta) = D_{c_{k+1}(\delta)}$, $\text{tr}_{\text{Env}}(D_i, \delta) = x_{c_{\text{Env}}(\delta)}$ and $\text{tr}_{\text{Env}}(x_j, \delta) = x_j$, for $j \in [1, k]$. Intuitively, when Env is at the initial state, it transits to the clause of index $c_{k+1}(\delta)$ chosen by the agent $k+1$. If the state is a clause, the environment can decide to move to the variable of index $c_{\text{Env}}(\delta)$. Once a state representing a variable is reached, it remains there forever.
- Agent $k+1$ only needs to decide which clause will be verified. Therefore, she has only one internal state, $\text{St}_{k+1}^{\text{P}} = \{\mathbf{I}\}$, sees the whole environment state, $\text{vis}_{k+1}(s_{\text{Env}}) = s_{\text{Env}}$, her transition function is trivial, $\text{tr}_{k+1}(s_{k+1}, \delta) = \mathbf{I}$, her actions are $\text{Ac}_{k+1} = \{\perp\} \cup \{1, \dots, m\}$ and the protocol is set as follows: $\text{P}_{k+1}(s_{k+1}) = \{1, \dots, m\}$, if $s_{k+1} = (\mathbf{I}, \mathbf{I})$, and $\text{P}_{k+1}(s_{k+1}) = \{\perp\}$ otherwise.
- Each agent $a \in \{1, \dots, k\}$ has a set of internal states equal to $\text{St}_a^{\text{P}} = \{\mathbf{I}, 0, 1\}$, where \mathbf{I} represents the initial state, while 0 and 1 are the Boolean valuations to be assigned to the variable x_a . The agent sees the whole environment state, *i.e.*, $\text{vis}_a(s_{\text{Env}}) = s_{\text{Env}}$ for all $s_{\text{Env}} \in \text{St}_{\text{Env}}$. The set of actions is $\text{Ac}_a = \{\perp, 0, 1\}$, where the idle action \perp is used by a when she does not need to set the value of the variable. The protocol and the transition function are defined as follows: $\text{P}_a(s_a) = \{0, 1\}$ and $\text{tr}_a(s_a, \delta) = c_a(\delta)$, if $s_a = (\mathbf{I}, x_a)$, and $\text{P}_a(s_a) = \{\perp\}$ and $\text{tr}_a(s_a, \delta) = s_a$, if $s_a \neq (\mathbf{I}, x_a)$. Intuitively, agent a can choose the value for her own variable x_a only when she is in her initial state and the environment is in its internal state associated with the variable x_a .
- There is only one global initial state represented by the $(k+2)$ -tuple whose components are all equal to \mathbf{I} , $\mathbf{I} = \{(\mathbf{I}, \dots, \mathbf{I})\}$.

- Finally, the valuation function for the atomic propositions is set as follows:

$$\begin{aligned}
h(x_a) &= \{s \in \text{St} \mid [\exists i \in \{1, \dots, m\} \cdot s_{\text{Env}}(s) = D_i \wedge x_a \in D_i] \vee \\
&\quad [s_{\text{Env}}(s) = x_a \wedge s_a^p(s) = 1]\} \quad \text{for all } x_a \in \text{AP} \cap X \\
h(\bar{x}_a) &= \{s \in \text{St} \mid [\exists i \in \{1, \dots, m\} \cdot s_{\text{Env}}(s) = D_i \wedge \bar{x}_a \in D_i] \vee \\
&\quad [s_{\text{Env}}(s) = x_a \wedge s_a^p(s) = 0]\} \quad \text{for all } \bar{x}_a \in \text{AP} \setminus X
\end{aligned}$$

Intuitively, we label with x_a (*resp.* \bar{x}_a) all global states corresponding to clauses containing this proposition. We also label those states corresponding to the valuation of the variable set to 1 (*resp.* 0).




Figure 1: Structure underlying the environment.

Now note that the number of reachable global states in this interpreted system is $1+m+3k$. Indeed, at the initial global state (I, \dots, I, I) , the only agent that can execute a non-idle action is agent $(k+1)$, choosing one of the m possible clauses D_i . In the resulting global state (I, \dots, I, D_i) , the environment is the only agent capable of changing its internal state, thereby reaching one of the k variables, say x_j . The global state resulting from this action is (I, \dots, I, x_j) . In this situation, only agent j can change her own state, by choosing one of the two possible actions $\alpha \in \{0, 1\}$, which represents the valuation for the variable x_j . Consequently, the system reaches the global state $(I, \dots, \alpha, \dots, I, x_j)$, where α is the j -th component of the tuple. The system will then remain in this state forever. Observe that, since only one agent at a time can perform a non-idle action, the number of possible transitions in the system, *i.e.*, the number of decisions declared in the global protocol, is linear in the size of the input QBF formula as well, precisely $m + \sum_{i=1}^m |D_i| + 2k$.

To conclude the reduction, we have to define the LTL formula $\tilde{\psi}$. The idea here is to ensure that, when the environment chooses a given variable to satisfy

a clause D_i selected by agent $k + 1$, the polarity of the variable in the clause and that of its valuation agree. In other words, for each proposition x_a , we have to ensure that either the proposition itself, or its dual \bar{x}_a occurs along the play identified by the strategies. Therefore, to ensure the above, we take $\tilde{\psi} = \bigwedge_{a=1}^k \neg((F x_a) \wedge (F \bar{x}_a))$. \square

3. Labelling Algorithm for Model Checking

In this section we describe our novel labelling algorithm for SLK which solves the model-checking problem described in Definition 6.

3.1. Algorithm

The model checking algorithm $\text{Check}_{\mathcal{I}}$ for SLK, which calculates the set of global states in which a given formula is true, is an extension of the existing ones for temporal logics. It differs from those in two ways:

1. It has an extra input parameter which represents the *binding* of agents to variables. When model checking a formula, we start with an empty binding and augment it whenever an agent binding operator $(a, x)\varphi$ is encountered.
2. Unlike the original algorithm, which merely returns the set of states in which a given formula holds, the present algorithm returns a set of *extended states*. Intuitively, an extended state is a pair of (i) a global state and (ii) a variable assignment (mapping variables to strategies) subject to which the formula holds in that state.

We will now define the aforementioned concepts of a binding and variable assignment more formally. For simplicity, we will fix Vr to always be the set of variables quantified in the SLK formula we are considering (*e.g.*, if the formula to be checked is $\varphi = \langle\langle x \rangle\rangle[y](a, x)(b, y)\mathbf{X}p$, then we set $\text{Vr} = \{x, y\}$). This will allow us to define variable assignments as total functions, which will make the theory and proofs much simpler.

Definition 7 (Bindings). *Let \mathcal{I} be an interpreted system. Then a binding is a partial function $b: \text{Ag} \rightarrow \text{Vr}$ which maps agents in its domain to variables. $\text{Bnd} \triangleq \text{Ag} \rightarrow \text{Vr}$ denotes the set of all bindings.*

Definition 8 (Variable Assignments). *Let \mathcal{I} be an interpreted system. Then a variable assignment is a function $v: \text{Vr} \rightarrow \text{UStr}$ which maps variables in its domain to uniform shared memoryless strategies. $\text{VAsg} \triangleq \text{Vr} \rightarrow \text{UStr}$ denotes the set of all variable assignments.*

Note that variable assignments are also assignments, *i.e.*, $\text{VAsg} \subseteq \text{Asg}$. As an example of an extended state, consider the tuple $(s_1, \{(x, \sigma_x), (y, \sigma_y)\})$, representing the global state s_1 in which the agents bound to variables x and y act according to the strategies σ_x and σ_y , respectively. We formalise this below.

Definition 9 (Extended States). *Let \mathcal{I} be an interpreted system, $s \in \text{St}$ a global state and $v \in \text{VAsg}$ a variable assignment. Then an extended state is a pair $\langle s, v \rangle \in \text{St} \times \text{VAsg}$. $\text{Ext} \triangleq \text{St} \times \text{VAsg}$ denotes the set of all extended states.*

Intuitively, an extended state $\langle s, v \rangle \in \text{Ext}$ guarantees a formula $\varphi \in \text{SLK}$ iff all assignments which agree with v make the formula true in the state s .

Definition 10 (Guarantees). *Let \mathcal{I} be an interpreted system, $\langle s, v \rangle \in \text{Ext}$ an extended state, $b \in \text{Bnd}$ a binding and $\varphi \in \text{SLK}$ an SLK formula with $\text{free}(\varphi) \cap \text{Ag} \subseteq \text{dom}(b)$. We say that $\langle s, v \rangle$ guarantees φ in \mathcal{I} under b iff for the assignment $\chi_v = v \cup \{(a, v(b(a))) \mid a \in \text{dom}(b)\} \in \text{Asg}$, we have $\mathcal{I}, (\chi_v, s) \models_{\text{SLK}} \varphi$.*

Now that we have covered the basic structures, we can define the concepts of negation and predecessors of extended states. These will be necessary for the model checking algorithm presented at the end of this subsection. Let us start with negation. Assume that we have calculated the set of extended states E which guarantee the formula φ under a binding b . We want to find the set E' of extended states which guarantee the formula $\neg\varphi$. Intuitively, E' should contain all extended states that somehow disagree with E . E' is calculated as follows:

$$E' = \text{Ext} \setminus E$$

We will now prove that our claim is correct.

Lemma 1. *Let \mathcal{I} be an interpreted system, $b \in \text{Bnd}$ a binding and $\varphi \in \text{SLK}$ an SLK formula with $\text{free}(\varphi) \cap \text{Ag} \subseteq \text{dom}(b)$. Let $E \subseteq \text{Ext}$ be the set of all extended states which guarantee φ in \mathcal{I} under b . Then $\text{Ext} \setminus E$ is the set of all extended states which guarantee $\neg\varphi$ in \mathcal{I} under b .*

Proof. Let $E' \subseteq \text{Ext}$ be the set of all extended states which guarantee $\neg\varphi$ in \mathcal{I} under b . We show that $E' = \text{Ext} \setminus E$:

- \Rightarrow : Take an arbitrary extended state $\langle s, v \rangle \in E'$. Since $\langle s, v \rangle$ guarantees $\neg\varphi$, we have $\mathcal{I}, (\chi_v, s) \models_{\text{SLK}} \neg\varphi$. By SLK semantics (Definition 5), we have $\mathcal{I}, (\chi_v, s) \not\models_{\text{SLK}} \varphi$. Thus, $\langle s, v \rangle \notin E$, so we have $\langle s, v \rangle \in \text{Ext} \setminus E$.
- \Leftarrow : Take an arbitrary extended state $\langle s, v \rangle \in \text{Ext} \setminus E$. Since $\langle s, v \rangle$ does not guarantee φ (otherwise, we would have $\langle s, v \rangle \in E$), we have $\mathcal{I}, (\chi_v, s) \not\models_{\text{SLK}} \varphi$. By SLK semantics (Definition 5), $\mathcal{I}, (\chi_v, s) \models_{\text{SLK}} \neg\varphi$. Thus, $\langle s, v \rangle \in \text{Ext}$ guarantees $\neg\varphi$, so we have $\langle s, v \rangle \in E'$. \square

Having covered negation, it remains to define how to calculate the set of previous extended states, *i.e.*, given a set of extended states $E \subseteq \text{Ext}$ which guarantee φ under a binding b , determine the set of extended states $E' \subseteq \text{Ext}$ which guarantee $X\varphi$ under the same binding.

We first define the transition relation on global states implied by a binding and a variable assignment. Intuitively, there is a transition from state $s_1 \in \text{St}$ to a state $s_2 \in \text{St}$ (given the binding and variable assignment) if there exists a decision between them such that each agent acts according to the strategy assigned to the variable she is bound to.

Definition 11 (Implied Transition Relation). *Let \mathcal{I} be an interpreted system, $s_1, s_2 \in \text{St}$ two global states, $b \in \text{Bnd}$ a binding and $v \in \text{VAsg}$ a variable assignment. Then the transition relation $\rightarrow_v^b \subseteq \text{St} \times \text{St}$ implied by b and v is defined by $s_1 \rightarrow_v^b s_2$, iff $\text{dom}(b) = \text{Ag}$ and there exists a decision $\delta \in \text{Dc}$ such that $\text{tr}(s_1, \delta) = s_2$ and, for all agents $a \in \text{Ag}$, it holds that $c_a(\delta) = v(b(a))(s_1)$.*

We are now ready to explain how the set of previous extended states is calculated. Intuitively, given an extended state $\langle s, v \rangle \in \text{Ext}$, the previous extended states are pairs $\langle s', v \rangle$ where s is the successor of s' when all agents act according to their strategies in v .

Definition 12 (Previous Extended States). *Let \mathcal{I} be an interpreted system, $E \subseteq \text{Ext}$ a set of extended states and $b \in \text{Bnd}$ a binding such that $\text{dom}(b) = \text{Ag}$. Then the function $\text{pre}: 2^{\text{Ext}} \times \text{Bnd} \rightarrow 2^{\text{Ext}}$ returning the set of previous extended states is defined as $\text{pre}(E, b) \triangleq \{ \langle s, v \rangle \in \text{Ext} \mid \exists s' \in \text{St}. \langle s', v \rangle \in E \wedge s \rightarrow_v^b s' \}$.*

Again, we will show that the function pre is correct.

Lemma 2. *Let \mathcal{I} be an interpreted system, $b \in \text{Bnd}$ a binding with $\text{dom}(b) = \text{Ag}$ and $\varphi \in \text{SLK}$ an SLK formula. Let $E \subseteq \text{Ext}$ be the set of all extended states which guarantee φ in \mathcal{I} under b . Then $\text{pre}(E) \subseteq \text{Ext}$ is the set of all extended states which guarantee $X\varphi$ in \mathcal{I} under b .*

Proof. Let $E' \subseteq \text{Ext}$ be the set of all extended states which guarantee $X\varphi$ in \mathcal{I} under b . We show that $\text{pre}(E) = E'$:

\Rightarrow : Take an arbitrary extended state $\langle s, v \rangle \in \text{pre}(E, b)$. By construction, there is a global state $s' \in \text{St}$ such that $\langle s', v \rangle \in E$ and $s \rightarrow_v^b s'$. Since $\langle s', v \rangle \in E$ guarantees φ , we have $\mathcal{I}, (\chi_v, s') \models_{\text{SLK}} \varphi$. The implied transition relation implies that there exists a decision $\delta \in \text{Dc}$ such that $s' = \text{tr}(s, \delta)$ where $c_a(\delta) = v(b(a))(s)$ for all agents $a \in \text{Ag}$. This can be also rewritten as $s' = \text{tr}(s, \langle \chi_v(a)(s) : a \in \text{Ag} \rangle)$ because $\chi_v(a) = v(b(a))$ for all agents $a \in \text{dom}(b) = \text{Ag}$ (see Definition 10).

We want to show that $\langle s, v \rangle$ guarantees $X\varphi$, i.e., $\mathcal{I}, (\chi_v, s) \models_{\text{SLK}} X\varphi$. This is the case iff $\mathcal{I}, (\chi_v, \pi_1) \models_{\text{SLK}} \varphi$ where $\pi = \text{play}(\chi_v, s)$ (see Definition 5). From the definition of a play (see Subsection 2.1), we obtain $\pi_1 = \text{tr}(s, \langle \chi_v(a)(s) : a \in \text{Ag} \rangle)$, so $\pi_1 = s'$. Since we have already shown that $\mathcal{I}, (\chi_v, s') \models_{\text{SLK}} \varphi$, we have $\mathcal{I}, (\chi_v, s) \models_{\text{SLK}} X\varphi$, so $\langle s, v \rangle \in E'$ as required.

\Leftarrow : Take an arbitrary extended state $\langle s, v \rangle \in E'$. Thus, we have $\mathcal{I}, (\chi_v, s) \models_{\text{SLK}} X\varphi$. By Definition 5, this means that $\mathcal{I}, (\chi_v, \pi_1) \models_{\text{SLK}} \varphi$ where $\pi = \text{play}(\chi_v, s)$. From the definition of a play (see Subsection 2.1), we obtain $\pi_1 = \text{tr}(s, \langle \chi_v(a)(s) : a \in \text{Ag} \rangle)$. This can be equivalently written as $\pi_1 = \text{tr}(s, \delta)$ for some $\delta \in \text{Dc}$ where $c_a(\delta) = \chi_v(a)(s) = v(b(a))(s)$ for all agents $a \in \text{Ag}$.

As $\text{dom}(b) = \text{Ag}$ by assumption, we have $s \rightarrow_v^b \pi_1$. Moreover, since $\mathcal{I}, (\chi_v, \pi_1) \models_{\text{SLK}} \varphi$, we have $\langle \pi_1, v \rangle \in E$ (because it guarantees φ). Therefore, we have $\langle s, v \rangle \in \text{pre}(E, b)$ as required. \square

Finally, we have all the ingredients to define the model checking algorithm $\text{Check}_{\mathcal{I}}(\cdot, \cdot)$ for SLK.

Definition 13 (SLK Model Checking Algorithm). *Let \mathcal{I} be an interpreted system, $\varphi \in \text{SLK}$ an SLK formula and $b \in \text{Bnd}$ a binding, such that $\text{free}(\varphi) \cap \text{Ag} \subseteq \text{dom}(b)$. Then the model checking function $\text{Check}_{\mathcal{I}} : \text{SLK} \times \text{Bnd} \rightarrow 2^{\text{Ext}}$ is inductively defined as follows:*

1. $\text{Check}_{\mathcal{I}}(\top, b) \triangleq \text{Ext}$.
2. $\text{Check}_{\mathcal{I}}(p, b) \triangleq \{\langle s, v \rangle \mid s \in \mathbf{h}(p)\}$, with $p \in \text{AP}$.
3. For all formulas $\varphi, \varphi_1, \varphi_2 \in \text{SLK}$, it is defined as:
 - (a) $\text{Check}_{\mathcal{I}}(\neg\varphi, b) \triangleq \text{Ext} \setminus \text{Check}_{\mathcal{I}}(\varphi, b)$;
 - (b) $\text{Check}_{\mathcal{I}}(\varphi_1 \wedge \varphi_2, b) \triangleq \text{Check}_{\mathcal{I}}(\varphi_1, b) \cap \text{Check}_{\mathcal{I}}(\varphi_2, b)$;
 - (c) $\text{Check}_{\mathcal{I}}(\varphi_1 \vee \varphi_2, b) \triangleq \text{Check}_{\mathcal{I}}(\varphi_1, b) \cup \text{Check}_{\mathcal{I}}(\varphi_2, b)$.
4. For an agent $a \in \text{Ag}$, a variable $x \in \text{Vr}$ and a formula $\varphi \in \text{SLK}$, $\text{Check}_{\mathcal{I}}((a, x)\varphi, b) \triangleq \text{Check}_{\mathcal{I}}(\varphi, b[a \mapsto x])$.
5. For a variable $x \in \text{Vr}$ and an SLK formula $\varphi \in \text{SLK}$, it is defined as:
 - (a) $\text{Check}_{\mathcal{I}}(\langle\langle x \rangle\rangle\varphi, b) \triangleq \{\langle s, v \rangle \in \text{Ext} \mid \exists f \in \text{UStr}_{\text{shr}(\varphi, x)}. \langle s, v[x \mapsto f] \rangle \in \text{Check}_{\mathcal{I}}(\varphi, b)\}$;
 - (b) $\text{Check}_{\mathcal{I}}(\llbracket x \rrbracket\varphi, b) \triangleq \{\langle s, v \rangle \in \text{Ext} \mid \forall f \in \text{UStr}_{\text{shr}(\varphi, x)}. \langle s, v[x \mapsto f] \rangle \in \text{Check}_{\mathcal{I}}(\varphi, b)\}$.
6. For all formulas $\varphi, \varphi_1, \varphi_2 \in \text{SLK}$, it is defined as:
 - (a) $\text{Check}_{\mathcal{I}}(\text{X}\varphi, b) \triangleq \text{pre}(\text{Check}_{\mathcal{I}}(\varphi, b), b)$;
 - (b) $\text{Check}_{\mathcal{I}}(\text{F}\varphi, b) \triangleq \text{Check}_{\mathcal{I}}(\top \text{U}\varphi, b)$;
 - (c) $\text{Check}_{\mathcal{I}}(\text{G}\varphi, b) \triangleq \text{Check}_{\mathcal{I}}(\neg\text{F}\neg\varphi, b)$;
 - (d) $\text{Check}_{\mathcal{I}}(\varphi_1 \text{U}\varphi_2, b) \triangleq \text{lfp}_{\text{X}}[\text{Check}_{\mathcal{I}}(\varphi_2, b) \cup (\text{Check}_{\mathcal{I}}(\varphi_1, b) \cap \text{pre}(\text{X}, b))]$ where $\text{lfp}_{\text{X}}f$ is the least fixed point of function f ;
 - (e) $\text{Check}_{\mathcal{I}}(\varphi_1 \text{R}\varphi_2, b) \triangleq \text{Check}_{\mathcal{I}}(\neg((\neg\varphi_1) \text{U}(\neg\varphi_2)), b)$.
7. For an agent $a \in \text{Ag}$, a set of agents $A \subseteq \text{Ag}$ and a formula $\varphi \in \text{SLK}$, it is defined as:
 - (a) $\text{Check}_{\mathcal{I}}(\text{K}_a\varphi, b) \triangleq \text{acc}(\varphi, \sim_a)$;
 - (b) $\text{Check}_{\mathcal{I}}(\text{E}_A\varphi, b) \triangleq \text{acc}(\varphi, \sim_A^{\text{E}})$;
 - (c) $\text{Check}_{\mathcal{I}}(\text{D}_A\varphi, b) \triangleq \text{acc}(\varphi, \sim_A^{\text{D}})$;
 - (d) $\text{Check}_{\mathcal{I}}(\text{C}_A\varphi, b) \triangleq \text{acc}(\varphi, \sim_A^{\text{C}})$;
where $\text{acc}(\varphi, R) \triangleq \{\langle s, v \rangle \in \text{Ext} \mid \forall \langle s', v' \rangle \in \text{Check}_{\mathcal{I}}(\neg\varphi, \emptyset). \neg R(s', s)\}$.

The correctness of the algorithm is asserted in the following theorem.

Theorem 2. *Let \mathcal{I} be an interpreted system and $\varphi \in \text{SLK}$ an SLK sentence. Then the set of all states at which φ holds is such that:*

$$\{s \in \text{St} \mid \mathcal{I}, s \models_{\text{SLK}} \varphi\} = \{s \in \text{St} \mid \exists v \in \text{VAsg}. \langle s, v \rangle \in \text{Check}_{\mathcal{I}}(\varphi, \emptyset)\}.$$

Proof (Sketch). We prove by induction that for an arbitrary interpreted system \mathcal{I} , SLK formula $\varphi \in \text{SLK}$ and binding $b \in \text{Bnd}$ such that $\text{free}(\varphi) \subseteq \text{dom}(b)$, $\text{Check}_{\mathcal{I}}(\varphi, b)$ is the set of all extended states that guarantee φ in \mathcal{I} under b .

Lemma 1 and Lemma 2 report the proofs for negation and the temporal operators, respectively.

Since the topmost formula φ is a sentence, it will either hold, or not hold in each state (regardless of the variable assignment). Therefore, we existentially quantify over variable assignments. \square

Observe that all cases of the SLK model checking algorithm are well-defined because the sets UStr, Vr, VAsg, Ag, Dc, Bnd, St and, consequently, Ext are finite. An efficient symbolic implementation of the algorithm using BDDs is presented in Subsection 3.3.

3.2. Strategy Synthesis

One of the features of SLK is that it may encode *non-behavioural strategies* [38], where an agent's action in a particular scenario may depend on actions in counterfactual scenarios. Consequently, SLK strategies are difficult to synthesise. To see why this is the case, consider the SLK sentence $\varphi = \llbracket x \rrbracket \llbracket y \rrbracket \langle\langle z \rangle\rangle \psi$. Assume that φ holds at a particular state $s \in \text{St}$ in an interpreted system \mathcal{I} , i.e., $\mathcal{I}, s \models_{\text{SLK}} \varphi$. We would now like to synthesise a uniform shared strategy $f_z : \text{St} \rightarrow \text{Ac}_{\text{shr}(\varphi, z)}$ for the variable z depending on the uniform shared strategies $f_x : \text{St} \rightarrow \text{Ac}_{\text{shr}(\varphi, x)}$ and $f_y : \text{St} \rightarrow \text{Ac}_{\text{shr}(\varphi, y)}$ for the variables x and y , respectively. If SLK strategies were behavioural, there *would* exist a mapping m_1 (or an *elementary dependence map* as in [38]) from the next actions of f_x and f_y in s to the next action of f_z in s :

$$m_1 : \text{St} \rightarrow \underbrace{(\text{Ac}_{\text{shr}(\varphi, x)} \times \text{Ac}_{\text{shr}(\varphi, y)} \rightarrow \text{Ac}_{\text{shr}(\varphi, z)})}_{(s, f_x(s), f_y(s)) \mapsto f_z(s)}$$

There are at most $|\text{Ac}_{\text{shr}(\varphi, x)}| \times |\text{Ac}_{\text{shr}(\varphi, y)}|$ possible inputs to m_1 to determine $f_z(s)$ as it depends only on $f_x(s)$ and $f_y(s)$. Unfortunately, such a mapping does not exist in general because SLK strategies are non-behavioural. Instead, a more general mapping m_2 (also referred to as *dependence map* [38]) from strategies f_x and f_y to the strategy f_z must be considered:

$$m_2 : \underbrace{(\text{St} \rightarrow \text{Ac}_{\text{shr}(\varphi, x)}) \times (\text{St} \rightarrow \text{Ac}_{\text{shr}(\varphi, y)})}_{(f_x, f_y, s) \mapsto f_z(s)} \rightarrow (\text{St} \rightarrow \text{Ac}_{\text{shr}(\varphi, z)})$$

Informally, determining $f_z(s)$ requires the same amount of information as constructing the whole strategy f_z . In order to synthesise the action $f_z(s)$ or the strategy f_z , we possibly need to know the complete strategies f_x and f_y . More importantly, the maximum number of *entries* in the mappings m_1 and m_2 are:

$$\begin{aligned} m_1 &: |\text{St}| \times |\text{Ac}_{\text{shr}(\varphi, x)}| \times |\text{Ac}_{\text{shr}(\varphi, y)}| \\ m_2 &: |\text{St}| \times |\text{Ac}_{\text{shr}(\varphi, x)}|^{|\text{St}|} \times |\text{Ac}_{\text{shr}(\varphi, y)}|^{|\text{St}|} \end{aligned}$$

Assume that the interpreted system \mathcal{I} is relatively small; e.g., assume $|\text{St}| = 10$ global states and $|\text{Ac}_{\text{shr}(\varphi, x)}| = |\text{Ac}_{\text{shr}(\varphi, y)}| = |\text{Ac}_{\text{shr}(\varphi, z)}| = 10$ actions. While

the mapping m_1 for *behavioural* strategies would require at most 1000 entries of the form $(s, f_x(s), f_y(s)) \mapsto f_z(s)$, the mapping m_2 for *non-behavioural* strategies might have up to 10^{21} entries of the form $(f_x, f_y, s) \mapsto f_z(s)$. Furthermore, if we encode each output of f_z using only $\lceil \log_2 |\text{Ac}_{\text{shr}(\varphi, z)}| \rceil = 4$ bits and store the whole mapping in a large array, m_1 will use at most 500 *bytes* while m_2 might need up to 434 *exabytes*.

While SLK strategy synthesis is infeasible in general, it can be performed efficiently on certain types of formulas. We will now explain the concepts of *witness* and *counterexample strategies* and describe how these can be synthesised using the model checking algorithm discussed in Subsection 3.1. Let \mathcal{I} be an interpreted system, $s \in \text{St}$ a global state and $\varphi_w = \langle\langle x \rangle\rangle \psi_w$ and $\varphi_c = \llbracket y \rrbracket \psi_c$ two SLK sentences. Furthermore, assume that φ_w holds at s while φ_c does not, *i.e.*, $\mathcal{I}, (\emptyset, s) \models_{\text{SLK}} \varphi_w$ and $\mathcal{I}, (\emptyset, s) \not\models_{\text{SLK}} \varphi_c$. By SLK semantics (see Definition 5), there is a memoryless uniform shared strategy f_w for x which makes ψ_w true at s . Conversely, there must be a memoryless shared strategy f_c for y which makes ψ_c false at s . f_w and f_c are referred to as a *witness* and a *counterexample strategy*, respectively. Intuitively, the strategy f_w is a “witness” to φ_w being true at s while f_c is a “counterexample” for φ_c at s . A slightly more general form of the two concepts is provided in the following definition.

Definition 14 (Witness and Counterexample Strategies). *Let \mathcal{I} be an interpreted system, $s \in \text{St}$ a global state and $\varphi_w = \langle\langle x_0 \rangle\rangle \dots \langle\langle x_{m-1} \rangle\rangle \psi_w$ and $\varphi_c = \llbracket y_0 \rrbracket \dots \llbracket y_{n-1} \rrbracket \psi_c$ two SLK sentences. Then:*

- *Memoryless uniform shared strategies $f_{w0}, \dots, f_{w(m-1)}$ are witness strategies for φ_w at s iff (i) $f_{wi} \in \text{UStr}_{\text{shr}(\psi_w, x_i)}$ for all $0 \leq i < m$ and (ii) $\mathcal{I}, (\chi_w, s) \models_{\text{SLK}} \psi_w$ where $\chi_w = \{(x_i, f_{wi}) \mid 0 \leq i < m\}$.*
- *Memoryless uniform shared strategies $f_{c0}, \dots, f_{c(n-1)}$ are counterexample strategies for φ_c at s iff (i) $f_{ci} \in \text{UStr}_{\text{shr}(\psi_c, y_i)}$ for all $0 \leq i < n$ and (ii) $\mathcal{I}, (\chi_c, s) \not\models_{\text{SLK}} \psi_c$ where $\chi_c = \{(y_i, f_{ci}) \mid 0 \leq i < n\}$.*

The two concepts are *duals* of each other in the sense that if f is a witness strategy for $\langle\langle x \rangle\rangle \psi$ at s , then it is a counterexample strategy for $\llbracket x \rrbracket \neg \psi$ at s (and vice versa).

Lemma 3. *Let \mathcal{I} be an interpreted system, $s \in \text{St}$ a global state, ψ an agent-closed SLK formula with $\text{free}(\psi) = \{x_0, \dots, x_{n-1}\}$ and f_0, \dots, f_{n-1} memoryless uniform shared strategies such that $f_i \in \text{UStr}_{\text{shr}(x_i, \psi)}$ for $0 \leq i < n$. Then f_0, \dots, f_{n-1} are witness strategies of $\varphi_w = \langle\langle x_0 \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle \psi$ at s iff f_0, \dots, f_{n-1} are counterexample strategies of $\varphi_c = \llbracket x_0 \rrbracket \dots \llbracket x_{n-1} \rrbracket \neg \psi$ at s .*

Proof. We prove both directions of the equivalence separately:

\Rightarrow : Assume that f_0, \dots, f_{n-1} are witness strategies of φ_w at s . By Definition 14, we have $f_i \in \text{UStr}_{\text{shr}(\psi, x_i)}$ for all $0 \leq i < n$ and $\mathcal{I}, (\chi, s) \models \psi$ where $\chi = \{(x_i, f_i) \mid 0 \leq i < n\}$. By Definition 3, we have $\text{shr}(\neg \varphi, x) = \text{shr}(\varphi, x)$ for all SLK formulas $\varphi \in \text{SLK}$ and variables $x \in \text{Vr}$ so $f_i \in \text{UStr}_{\text{shr}(\neg \psi, x_i)}$

for all $0 \leq i < n$. Since $\mathcal{I}, (\chi, s) \models_{\text{SLK}} \psi$, $\mathcal{I}, (\chi, s) \not\models_{\text{SLK}} \neg\psi$ holds by SLK semantics (Definition 5). Hence, by Definition 14, f_0, \dots, f_{n-1} are counterexample strategies for φ_c at s .

\Leftarrow : Assume that f_0, \dots, f_{n-1} are counterexample strategies of φ_c at s . By Definition 14, we have $f_i \in \text{UStr}_{\text{shr}(\neg\psi, x_i)}$ for all $0 \leq i < n$ and $\mathcal{I}, (\chi, s) \not\models_{\text{SLK}} \neg\psi$ where $\chi = \{(x_i, f_i) \mid 0 \leq i < n\}$. Again, we have $f_i \in \text{UStr}_{\text{shr}(\psi, x_i)}$ for all $0 \leq i < n$ by Definition 3. Since $\mathcal{I}, (\chi, s) \not\models_{\text{SLK}} \neg\psi$, $\mathcal{I}, (\chi, s) \models_{\text{SLK}} \psi$ holds by SLK semantics (Definition 5). Hence, by Definition 14, f_0, \dots, f_{n-1} are witness strategies of φ_w at s . \square

This duality will allow us to focus on witness strategies and their synthesis only. Before describing how witness strategies can be retrieved, we need to show that they always exist when an SLK formula holds.

Lemma 4. *Let \mathcal{I} be an interpreted system, ψ an agent-closed SLK formula such that $\text{free}(\psi) = \{x_0, \dots, x_{n-1}\}$ and $\varphi = \langle\langle x_0 \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle \psi$ an SLK sentence. Then the following holds: $\mathcal{I}, (\emptyset, s) \models_{\text{SLK}} \varphi$ iff there exist witness strategies f_0, \dots, f_{n-1} for φ at s .*

Proof. We prove both directions of the equivalence separately:

\Rightarrow : Assume that $\mathcal{I}, (\emptyset, s) \models_{\text{SLK}} \varphi$. By SLK semantics (Definition 5), there exist strategies f_0, \dots, f_{n-1} such that $f_i \in \text{UStr}_{\text{shr}(\langle\langle x_{i+1} \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle \psi, x_i)}$ for all $0 \leq i < n$ and $\mathcal{I}, (\chi, s) \models_{\text{SLK}} \psi$ where $\chi = \{(x_i, f_i) \mid 0 \leq i < n\}$. By Definition 3, we have $\text{shr}(\langle\langle x_{i+1} \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle \psi, x_i) = \text{shr}(\psi, x_i)$ for all $0 \leq i < n$. Hence f_0, \dots, f_{n-1} satisfy both conditions for being witness strategies of φ at s (see Definition 14).

\Leftarrow : Assume that there exist witness strategies f_0, \dots, f_{n-1} for φ at s . By Definition 14, we have $f_i \in \text{UStr}_{\text{shr}(\psi, x_i)}$ for all $0 \leq i < n$ and $\mathcal{I}, (\chi, s) \models_{\text{SLK}} \psi$ where $\chi = \{(x_i, f_i) \mid 0 \leq i < n\}$. By Definition 3, we have $\text{shr}(\psi, x_i) = \text{shr}(\langle\langle x_{i+1} \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle \psi, x_i)$. Hence, for all $0 \leq i < n$, we get $f_i \in \text{UStr}_{\text{shr}(\langle\langle x_{i+1} \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle \psi, x_i)}$. Therefore, $\mathcal{I}, (\emptyset, s) \models_{\text{SLK}} \varphi$ by SLK semantics (Definition 5). \square

It remains to explain how witness strategies can be synthesised using the SLK model checking algorithm we introduced in Subsection 3.1. Consider an SLK sentence $\varphi = \langle\langle x_0 \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle \psi$ that holds at a global state $s \in \text{St}$ in some interpreted system \mathcal{I} . Since $\mathcal{I}, (\emptyset, s) \models_{\text{SLK}} \varphi$, by Lemma 4, there must be some witness strategies f_0, \dots, f_{n-1} for φ at s , which we want to synthesise. The corresponding assignment $\chi \in \text{Asg}$ on the variables x_0, \dots, x_{n-1} satisfies $\mathcal{I}, (\chi, s) \models_{\text{SLK}} \psi$. Let $\text{E} = \text{Check}_{\mathcal{I}}(\psi, \emptyset)$ be the set of extended states which guarantee ψ in \mathcal{I} under the empty binding. E contains all possible extended states $\langle s', v' \rangle$ such that $\mathcal{I}, (s', v') \models_{\text{SLK}} \psi$. Hence, it must be the case that $\langle s, v \rangle \in \text{E}$ where v is some variable assignment which extends χ .

Therefore, in order to synthesise witness strategies for φ at s , it suffices to pick an arbitrary extended state $\langle s, v \rangle \in \text{Check}_{\mathcal{I}}(\psi, \emptyset)$. The witness strategies for φ at s are then $v(x_0), \dots, v(x_{n-1})$.

Lemma 5. *Let \mathcal{I} be an interpreted system, ψ an agent-closed SLK formula such that $\text{free}(\psi) = \{x_0, \dots, x_{n-1}\}$ and $\varphi = \langle\langle x_0 \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle \psi$ an SLK sentence. Then the following properties hold:*

1. *For all variable assignments $\langle s, v \rangle \in \text{Check}_{\mathcal{I}}(\psi, \emptyset)$, $v(x_0), \dots, v(x_{n-1})$ are witness strategies for φ at s .*
2. *If there exist witness strategies f_0, \dots, f_{n-1} for φ at s , then there exists a variable assignment $v \in \text{VAsg}$ such that $v(x_i) = f_i$ for $0 \leq i < n$ and $\langle s, v \rangle \in \text{Check}_{\mathcal{I}}(\psi, \emptyset)$.*

Proof. We prove both properties separately:

1. Take an arbitrary extended state $\langle s, v \rangle \in \text{Check}_{\mathcal{I}}(\psi, \emptyset)$. Observe that the SLK model checking algorithm (see item 5 in Definition 13) ensures uniformity of v wrt x_0, \dots, x_{n-1} , *i.e.*, $v(x_i) \in \text{UStr}_{\text{shr}(\psi, x_i)}$ for $0 \leq i < n$. As in the proof of Theorem 2, $\langle s, v \rangle$ guarantees ψ in \mathcal{I} under \emptyset . By Definition 10, $\mathcal{I}, (v, s) \models_{\text{SLK}} \psi$. By SLK semantics (Definition 5), we have $\mathcal{I}, (\chi, s) \models_{\text{SLK}} \psi$ where $\chi = \{(x_i, v(x_i)) \mid 0 \leq i < n\}$ since $\text{free}(\psi) \subseteq \{x_0, \dots, x_{n-1}\} \subseteq \text{dom}(v)$. Hence, $v(x_0), \dots, v(x_{n-1})$ satisfy both conditions for being witness strategies for φ at s (see Definition 14).
2. Assume that f_0, \dots, f_{n-1} are witness strategies for φ at s . By Definition 14, $f_i \in \text{UStr}_{\text{shr}(\psi, x_i)}$ for all $0 \leq i < n$ and $\mathcal{I}, (\chi, s) \models_{\text{SLK}} \psi$ where $\chi = \{(x_i, f_i) \mid 0 \leq i < n\}$. Since $\text{dom}(\chi) = \text{free}(\psi)$, by SLK semantics (Definition 5), we have $\mathcal{I}, (v, s) \models_{\text{SLK}} \psi$ for all $v \in \text{VAsg}$ such that $\chi \subseteq v$. There must exist at least one such variable assignment v because $\text{dom}(\chi) \subseteq \text{Vr}$ (simply set $v(x) = \chi(x)$ for $x \in \text{dom}(\chi)$ and assign arbitrary strategies to variables $y \in \text{Vr} \setminus \text{dom}(\chi)$). By Definition 10, $\langle s, v \rangle \in \text{Ext}$ guarantees ψ in \mathcal{I} under \emptyset . Therefore, $\langle s, v \rangle \in \text{Check}_{\mathcal{I}}(\psi, \emptyset)$ (see proof of Theorem 2). \square

Informally, the first property in Lemma 5 expresses *soundness* of the approach, *i.e.*, that it will return only witness strategies for φ at s . Conversely, the second property asserts *completeness* of the approach, *i.e.*, that it will return witness strategies for φ at s , if they exist.

3.3. Symbolic Implementation

In this subsection we discuss how the algorithm presented in Subsection 3.1 can be implemented symbolically using binary decision diagrams. BDDs are an efficient representation for Boolean formulas and are used by many existing model checkers including MCMAS for model checking temporal logic formulas [39]. We here present a modification of the relevant algorithms for SLK.

We start by representing the parameters of the interpreted system by means of Boolean formulas [10]. Given an interpreted system:

$$\mathcal{I} = \left\langle (\text{St}_a, \text{Ac}_a, \text{Pa}, \text{tr}_a)_{a \in \text{Ag}}, \text{I}, \mathbf{h} \right\rangle$$

we can represent global states and decisions as follows [39, 10]:

- For every agent $a \in \text{Ag}$, we can encode her set of internal states St_a^{P} with $\text{nv}(a) = \lceil \log_2 |\text{St}_a| \rceil$ Boolean variables. Thus, a *global state* $s = (s_1^{\text{P}}, \dots, s_n^{\text{P}}, s_{\text{Env}}) \in \text{St}$ can be encoded as a conjunction $s[\bar{v}]$ of the variables in a Boolean vector $\bar{v} = \bar{v}_1^{\text{P}} \dots \bar{v}_n^{\text{P}} \bar{v}_{\text{Env}} = (v_0, \dots, v_{N-1})$, where $N = \sum_{a \in \text{Ag}} \text{nv}(a)$.
- For every agent $a \in \text{Ag}$, we can encode her set of actions Ac_a with $\text{na}(a) = \lceil \log_2 |\text{Ac}_a| \rceil$ Boolean variables. Thus, a *decision* $\delta = (c_1, \dots, c_n, c_{\text{Env}}) \in \text{Dc}$ can be encoded as a conjunction $\delta[\bar{w}]$ of the variables in a Boolean vector $\bar{w} = \bar{w}_1 \dots \bar{w}_n \bar{w}_{\text{Env}} = (w_0, \dots, w_{M-1})$, where $M = \sum_{a \in \text{Ag}} \text{na}(a)$.

3.3.1. Running Example

Consider a variant of the classical Nim game [40], where two players take turns to remove one or two objects from a single heap which initially contains four objects. The player who removes the last object wins the game.

We model the scenario as an interpreted system \mathcal{I}_{Nim} with agents $\text{Ag} = \{\text{A}, \text{B}, \text{Env}\}$, where the proper agents $\Sigma = \{\text{A}, \text{B}\}$ correspond to the two players:

- We fix the *environment states* $\text{St}_{\text{Env}} = \{s_{rq}^{\text{Env}} \mid 0 \leq r \leq 4 \wedge q \in \Sigma\}$, where r is the number of objects remaining on the heap and q is the active player, whose turn it is to remove objects. Both players have a single internal state, $\text{St}_{\text{A}}^{\text{P}} = \text{St}_{\text{B}}^{\text{P}} = \{\text{I}\}$, and see the whole environment state, *i.e.*, $\text{vis}_p(s_{\text{Env}}) = s_{\text{Env}}$ for all $p \in \Sigma$, $s_{\text{Env}} \in \text{St}_{\text{Env}}$. Hence:

$$\begin{aligned} \text{St}_p &= \{s_{rq}^p = (\text{I}, s_{rq}^{\text{Env}}) \mid s_{rq}^{\text{Env}} \in \text{St}_{\text{Env}}\} \quad \text{for } p \in \Sigma \\ \text{St} &= \{s_{rq} = (\text{I}, \text{I}, s_{rq}^{\text{Env}}) \mid s_{rq}^{\text{Env}} \in \text{St}_{\text{Env}}\} \end{aligned}$$

- A player can perform three possible *actions*: do nothing, remove one object, or remove two objects, which is formally represented as $\text{Ac}_{\text{A}} = \text{Ac}_{\text{B}} = \{\perp, 1, 2\}$. The environment always does nothing, *i.e.*, $\text{Ac}_{\text{Env}} = \{\perp\}$.
- The *protocols* of the agents require that the active player removes one or two objects:

$$\begin{aligned} \text{P}_p(s_{rq}^p) &= \begin{cases} \{1, 2\} & \text{if } p = q \wedge r > 1 \\ \{1\} & \text{if } p = q \wedge r > 0 \\ \{\perp\} & \text{otherwise} \end{cases} \quad \text{for } p \in \Sigma, s_{rq}^p \in \text{St}_p \\ \text{P}_{\text{Env}}(s_{rq}^{\text{Env}}) &= \{\perp\} \quad \text{for } s_{rq}^{\text{Env}} \in \text{St}_{\text{Env}} \end{aligned}$$

- The *evolution functions* simply reflect the active player's action for $\delta \in \text{Dc}$:

$$\begin{aligned} \text{tr}_p(s_{rq}^p, \delta) &= \text{I} \quad \text{for } p \in \Sigma, s_{rq}^p \in \text{St} \\ \text{tr}_{\text{Env}}(s_{rq}^{\text{Env}}, \delta) &= \begin{cases} s_{(r-c_{\text{A}}(\delta))\text{B}}^{\text{Env}} & \text{if } q = \text{A} \wedge r > 0 \\ s_{(r-c_{\text{B}}(\delta))\text{A}}^{\text{Env}} & \text{if } q = \text{B} \wedge r > 0 \\ s_{rq}^{\text{Env}} & \text{otherwise} \end{cases} \quad \text{for } s_{rq}^{\text{Env}} \in \text{St}_{\text{Env}} \end{aligned}$$

- The set of *initial states* is $I = \{s_{4A}\} = \{(I, I, s_{4A}^{\text{Env}})\}$, *i.e.*, the heap initially contains 4 objects and player A plays first.
- Finally, we use two *atomic propositions* $\text{AP} = \{\text{win}_A, \text{win}_B\}$ and reflect the winning condition in the *valuation function*: $\text{h}(\text{win}_A) = \{s_{0B}\}$ and $\text{h}(\text{win}_B) = \{s_{0A}\}$.

To encode the global states St , the symbolic implementation uses a Boolean vector $\bar{v} = (v_0, v_1, v_2, v_3)$ with $\text{nv}(A) + \text{nv}(B) + \text{nv}(\text{Env}) = 0 + 0 + 4 = 4$ variables. Variables v_0, v_1, v_2 are used to encode the number of remaining objects on the heap as a binary number and variable v_3 is true iff B is the active player. For example, $s_{1A}[\bar{v}] = \neg v_0 \wedge \neg v_1 \wedge v_2 \wedge \neg v_3$.

Similarly, the decisions Dc are represented using a Boolean vector $\bar{w} = (w_0, w_1, w_2, w_3)$ with $\text{na}(A) + \text{na}(B) + \text{na}(\text{Env}) = 2 + 2 + 0 = 4$ variables. Variables w_0, w_1 and w_2, w_3 encode the actions of players A and B, respectively, using binary encoding (treating \perp as zero). For example, $(2, \perp, \perp)[\bar{w}] = w_0 \wedge \neg w_1 \wedge \neg w_2 \wedge \neg w_3$.

3.3.2. Encoding of Extended States

In order to implement the SLK model checking algorithm, we need to represent sets of *extended states* (see Definition 9), which consist of a global state and a variable assignment. We thus represent the variable assignment explicitly using Boolean variables as well.

Since the number of strategy variables (and hence the domain of a variable assignment) depends on the SLK formula we are checking, the total number of Boolean variables also depends on the formula. Let $\varphi \in \text{SLK}$ be an SLK sentence and $\text{vars}(\varphi) \subseteq \text{Vr}$ the set of variables quantified in φ . For each variable $x \in \text{vars}(\varphi)$, we represent the strategy associated with x using a number of Boolean variables.

Let us now consider an arbitrary variable $x \in \text{vars}(\varphi)$. In the SLK model checking algorithm (Definition 13), the variable quantifies over strategies $f \in \text{UStr}_{\text{shr}(\varphi, x)}$. The domain and the range of a strategy $f \in \text{UStr}_A$ are St and Ac_A , respectively (see Subsection 2.3). Hence, we could represent it using $|\text{St}| \times \lceil \log_2 |\text{Ac}_A| \rceil$ Boolean variables by encoding the action associated with each global state. By considering the protocols² of all agents in A, we can reduce the number of Boolean variables to $\sum_{s \in \text{St}} \lceil \log_2 |\bigcap_{a \in A} \text{P}_a(\mathbf{s}_a(s))| \rceil$.

Furthermore, we can exploit the fact that the strategies are *uniform*. Intuitively, if we have a strategy $f \in \text{UStr}_{\{a\}}$ for an agent $a \in \text{Ag}$, then for all states $s_1, s_2 \in \text{St}$, $s_1 \sim_a s_2$ implies $f(s_1) = f(s_2)$. Hence, there is no point in storing both actions $f(s_1)$ and $f(s_2)$. More generally, the epistemic accessibility relations \sim_a with $a \in A$ induce regions of the global state space, to which the uniform strategies $f \in \text{UStr}_A$ must assign the same action. It turns out that these regions are equivalence classes with respect to the common epistemic accessibility relation \sim_A^{C} .

² $|\bigcap_{a \in A} \text{P}_a(\mathbf{s}_a(s))| \leq |\text{Ac}_A|$ for each global state $s \in \text{St}$.

Lemma 6. *Let \mathcal{I} be an interpreted system and $A \subseteq \text{Ag}$ a set of agents. Then a memoryless strategy $f : \text{St} \rightarrow \text{Ac}_A$ is uniform iff for each set of global states in the quotient set $S \in \text{St}/\sim_A^C$, we have $f(s_1) = f(s_2)$ for all $s_1, s_2 \in S$.*

Proof. We will prove both directions of the equivalence separately:

\Rightarrow : Assume that f is uniform and take an arbitrary set $S \in \text{St}/\sim_A^C$. Furthermore, take arbitrary global states $s_1, s_2 \in S$. By the definition of quotient set, we have $s_1 \sim_A^C s_2$. There are two cases:

- $s_1 = s_2$. Trivially, $f(s_1) = f(s_2)$.
- By the definition of common epistemic accessibility relation (see Subsection 2.3), there is a chain of global states $s'_1, s'_2, \dots, s'_n \in \text{St}$ and agents $i_1, i_2, \dots, i_{n+1} \in A$ with $n \geq 0$ such that $s_1 \sim_{i_1} s'_1 \sim_{i_2} s'_2 \dots s'_n \sim_{i_{n+1}} s_2$. By uniformity of f , we get $f(s_1) = f(s'_1) = \dots = f(s'_n) = f(s_2)$.

\Leftarrow : Assume that for each $S \in \text{St}/\sim_A^C$, we have $f(s_1) = f(s_2)$ for all $s_1, s_2 \in S$. Take an arbitrary agent $a \in A$ and global states $s_1, s_2 \in \text{St}$ such that $s_1 \sim_a s_2$. To prove uniformity of f , we need to show that $f(s_1) = f(s_2)$. Since $s_1 \sim_a s_2$, we also have $s_1 \sim_A^C s_2$. Hence, s_2 belongs to the equivalence class $[s_1]_{\sim_A^C}$. By definition of an equivalence class, it must be the case that $s_1 \in [s_1]_{\sim_A^C}$ and $[s_1]_{\sim_A^C} \in \text{St}/\sim_A^C$. Since $s_1, s_2 \in [s_1]_{\sim_A^C}$ and $[s_1]_{\sim_A^C} \in \text{St}/\sim_A^C$, we get $f(s_1) = f(s_2)$ by the initial assumption as required. \square

This allows us to present an even more compact representation of a strategy $f \in \text{UStr}_A$ with $A \subseteq \text{Ag}$ in an interpreted system \mathcal{I} . We only need to store one action for each *shared local state* $S \in \text{St}/\sim_A^C$. Thus, we can represent f using $\sum_{S \in \text{St}/\sim_A^C} \lceil \log_2 \left| \bigcap_{s \in S} \bigcap_{a \in A} P_a(s_a(s)) \right| \rceil$ Boolean variables. Finally, a variable assignment $v \in \text{VAsg}$ for a formula $\varphi \in \text{SLK}$ can be represented using a Boolean vector $\bar{u} = (u_0, \dots, u_{K-1})$ such that³:

$$K = \sum_{x \in \text{vars}(\varphi)} \sum_{S \in \text{St}/\sim_{\text{shr}(\varphi, x)}^C} \left\lceil \log_2 \left| \bigcap_{s \in S} \bigcap_{a \in \text{shr}(\varphi, x)} P_a(s_a(s)) \right| \right\rceil$$

as follows:

$$v[\bar{u}] := \bigwedge_{x \in \text{vars}(\varphi)} \bigwedge_{S \in \text{St}/\sim_{\text{shr}(\varphi, x)}^C} v(x)(S)[\bar{u}_{S, x}]$$

³Assuming that each strategy and action is stored separately in \bar{u} , the provided representation is optimal. If we relaxed this assumption, the number of Boolean variables could be reduced to $\left\lceil \log_2 \left| \prod_{x \in \text{vars}(\varphi)} \prod_{S \in \text{St}/\sim_{\text{shr}(\varphi, x)}^C} \bigcap_{s \in S} \bigcap_{a \in \text{shr}(\varphi, x)} P_a(s_a(s)) \right| \right\rceil$. However, we do not pursue this approach any further because it would complicate the symbolic implementation of the labelling algorithm.

where $v(x)(S)[\bar{u}_{S,x}]$ is the Boolean formula representing that the single action assigned by the uniform strategy mapped to variable x in all states $s \in S$ is $v(x)(s)$ for any $s \in S$.

Note that despite both optimisations, the worst case still remains $K = |\text{vars}(\varphi)| \times |\text{St}| \times \lceil \log_2(\max_{a \in \text{Ag}} |\text{Ac}_a|) \rceil$, *i.e.*, we need polynomially many Boolean variables with respect to both the size of the model $|\mathcal{I}|$ and the number of strategy variables in the formula $|\text{vars}(\varphi)|$.

An extended state $\langle s, v \rangle \in \text{Ext}$ can be represented by the conjunction $\langle s, v \rangle[\bar{v}, \bar{u}] := s[\bar{v}] \wedge v[\bar{u}]$ over the vectors \bar{v} and \bar{u} . A set of extended states $E \subseteq \text{Ext}$ can in turn be expressed as the disjunction $E[\bar{v}, \bar{u}] := \bigvee_{\langle s, v \rangle \in E} s[\bar{v}] \wedge v[\bar{u}]$.

Consider the SLK sentence $\varphi_{\text{win}_A} = \langle\langle x \rangle\rangle [\![y]\!] [\![e]\!] (A, x)(B, y)(\text{Env}, e) F \text{win}_A$, which specifies that there is a strategy for player A to win regardless of the strategies followed by the other player and the environment. Our Boolean encoding represents variable assignments for φ_{win_A} in \mathcal{I}_{Nim} using a Boolean vector \bar{u} with $K = 6$ variables:

$$\begin{aligned}
K &= \left\lceil \log_2 \left| \mathbf{P}_{\text{Env}}(s_{0A}^{\text{Env}}) \right| \right\rceil + \dots + \left\lceil \log_2 \left| \mathbf{P}_{\text{Env}}(s_{4A}^{\text{Env}}) \right| \right\rceil + \\
&\quad \left\lceil \log_2 \left| \mathbf{P}_{\text{Env}}(s_{0B}^{\text{Env}}) \right| \right\rceil + \dots + \left\lceil \log_2 \left| \mathbf{P}_{\text{Env}}(s_{4B}^{\text{Env}}) \right| \right\rceil + \\
&\quad \left\lceil \log_2 \left| \mathbf{P}_A(s_{0A}^A) \right| \right\rceil + \dots + \left\lceil \log_2 \left| \mathbf{P}_A(s_{4A}^A) \right| \right\rceil + \\
&\quad \left\lceil \log_2 \left| \mathbf{P}_A(s_{1B}^A) \right| \right\rceil + \dots + \left\lceil \log_2 \left| \mathbf{P}_A(s_{4B}^A) \right| \right\rceil + \\
&\quad \left\lceil \log_2 \left| \mathbf{P}_B(s_{0A}^B) \right| \right\rceil + \dots + \left\lceil \log_2 \left| \mathbf{P}_B(s_{4A}^B) \right| \right\rceil + \\
&\quad \left\lceil \log_2 \left| \mathbf{P}_B(s_{0B}^B) \right| \right\rceil + \dots + \left\lceil \log_2 \left| \mathbf{P}_B(s_{4B}^B) \right| \right\rceil \\
&= \left\lceil \log_2 |\{\perp\}| \right\rceil + \dots + \left\lceil \log_2 |\{\perp\}| \right\rceil + \\
&\quad \left\lceil \log_2 |\{\perp\}| \right\rceil + \dots + \left\lceil \log_2 |\{\perp\}| \right\rceil + \\
&\quad \left\lceil \log_2 |\{\perp\}| \right\rceil + \left\lceil \log_2 |\{1\}| \right\rceil + \left\lceil \log_2 |\{1, 2\}| \right\rceil + \left\lceil \log_2 |\{1, 2\}| \right\rceil + \left\lceil \log_2 |\{1, 2\}| \right\rceil + \\
&\quad \left\lceil \log_2 |\{\perp\}| \right\rceil + \dots + \left\lceil \log_2 |\{\perp\}| \right\rceil + \\
&\quad \left\lceil \log_2 |\{\perp\}| \right\rceil + \dots + \left\lceil \log_2 |\{\perp\}| \right\rceil + \\
&\quad \left\lceil \log_2 |\{\perp\}| \right\rceil + \left\lceil \log_2 |\{1\}| \right\rceil + \left\lceil \log_2 |\{1, 2\}| \right\rceil + \left\lceil \log_2 |\{1, 2\}| \right\rceil + \left\lceil \log_2 |\{1, 2\}| \right\rceil \\
&= 6
\end{aligned}$$

Intuitively, this is because there are only 6 states in which some strategy has two actions to choose from: strategy $f_x = v(x)$ in states s_{2A} , s_{3A} , s_{4A} and strategy $f_y = v(y)$ in states s_{2B} , s_{3B} , s_{4B} . Each Boolean variable corresponds to one of these situations ($u_0 \rightsquigarrow (x, s_{2A})$, $u_1 \rightsquigarrow (x, s_{3A})$, $u_2 \rightsquigarrow (x, s_{4A})$, $u_3 \rightsquigarrow (y, s_{2B})$, $u_4 \rightsquigarrow (y, s_{3B})$, $u_5 \rightsquigarrow (y, s_{4B})$). If a variable is true, then the relevant strategy prescribes removing two items from the heap in the relevant state. For example, if u_0 is false, then $f_x(s_{2A}) = 1$. Conversely, if u_4 is true, then $f_y(s_{3B}) = 2$. Furthermore, a sample extended state $e = (s_{0A}, \{x \mapsto f_x, y \mapsto f_y, e \mapsto f_e\})$ with strategies:

$$\begin{array}{lll}
f_x(s_{0*}) = \perp & f_y(s_{0*}) = \perp & f_e(s_{**}) = \perp \\
f_x(s_{1A}) = 1 & f_y(s_{1B}) = 1 & \\
f_x(s_{2A}) = 2 \quad [u_0] & f_y(s_{2B}) = 1 \quad [\neg u_3] & \\
f_x(s_{3A}) = 1 \quad [\neg u_1] & f_y(s_{3B}) = 2 \quad [u_4] & \\
f_x(s_{4A}) = 1 \quad [\neg u_2] & f_y(s_{4B}) = 2 \quad [u_6] & \\
f_x(s_{*B}) = \perp & f_y(s_{*A}) = \perp &
\end{array}$$

is symbolically encoded as $e[\bar{v}, \bar{u}] = \neg v_0 \wedge \neg v_1 \wedge \neg v_2 \wedge \neg v_3 \wedge u_0 \wedge \neg u_1 \wedge \neg u_2 \wedge \neg u_3 \wedge u_4 \wedge u_6$.

3.3.3. Encoding of the Algorithm

Given a binding $b \in \text{Bnd}$ such that $\text{dom}(b) = \text{Ag}$, we define a formula $S^b(e, a)$, where $e \in \text{Ext}$ and $\delta \in \text{Dc}$, representing the *strategy restrictions* of the implied transition relation (see Definition 11). The formula asserts that all agents act according to their strategies:

$$S^b(\langle s, v \rangle, \delta) \triangleq \forall a \in \text{Ag}. v(b(a))(s) = c_a(\delta)$$

Let \bar{v} , \bar{w} and \bar{u} be the Boolean vectors for representing current global states, decisions and variable assignments, respectively. We can encode the strategy restrictions as a Boolean formula $S^b[\bar{v}, \bar{w}, \bar{u}]$:

$$S^b[\bar{v}, \bar{w}, \bar{u}] := \bigwedge_{a \in \text{Ag}} \bigvee_{s_a \in \text{St}_a} \left(s_a[\bar{v}_a] \wedge \bigvee_{c \in \text{P}_a(s_a)} c[\bar{w}_a] \wedge c[\bar{u}_{s_a, b(a)}] \right)$$

where:

- $s_a[\bar{v}_a]$ is the Boolean formula representing that the local state⁴ of agent $a \in \text{Ag}$ is $s_a \in \text{St}_a$;
- $c[\bar{w}_a]$ is the Boolean formula representing that the action of agent $a \in \text{Ag}$ is $c \in \text{Ac}_a$;
- $c[\bar{u}_{s_a, b(a)}]$ is the Boolean formula representing that the action in local state⁵ s_a assigned by the strategy mapped to variable $b(a)$ is $c \in \text{Ac}_a$.

Given the binding $b = (A \mapsto x, B \mapsto y, \text{Env} \mapsto e)$ in φ_{win_A} , the strategy re-

⁴ Note that for a proper agent $a \in \Sigma$, the local state is a combination of a private state and an image of the environment state, *i.e.*, $s_a = (s_a^p, \text{vis}_a(s_{\text{Env}}))$ for some $s_a^p \in \text{St}_a^p$, $s_{\text{Env}} \in \text{St}_{\text{Env}}$. This is reflected in the encoding: $s_a[\bar{v}_a] = (s_a^p, \text{vis}_a(s_{\text{Env}}))[\bar{v}_a] := s_a^p[\bar{v}_a^p] \wedge \text{vis}_a(s_{\text{Env}})[\bar{v}_{\text{Env}}]$.

Typically, vis_a is encoded as a filter over \bar{v}_{Env} , *e.g.*, **Obsvars** and **Lobsvars** sections of the ISPL file format used by MCMAS [20]. Intuitively, only some Boolean variables encoding the environment state are visible to agent a .

⁵Strictly speaking, the strategy maps *global* states to actions. However, as explained earlier, there exists a set $S \subseteq \text{St} / \sim_{\text{shr}(\varphi, b(a))}^c$ such that, for all global states $s \in \text{St}$, if $\mathbf{s}_a(s) = s_a$ then $s \in S$. Due to uniformity, the strategy mapped to $b(a)$ must assign the same action to all global states in S . Thus, we can also interpret the strategy as a mapping from *local* states to actions.

restrictions for our running example are encoded as follows:

$$\begin{aligned}
S^b[\bar{v}, \bar{w}, \bar{u}] = & (v0000 \wedge w00-- \wedge u----- \vee v0001 \wedge w00-- \wedge u----- \vee \\
& v0010 \wedge w01-- \wedge u----- \vee v0011 \wedge w00-- \wedge u----- \vee \\
& v0100 \wedge (w01-- \wedge u0----- \vee w10-- \wedge u1-----) \vee v0101 \wedge w00-- \wedge u----- \vee \\
& v0110 \wedge (w01-- \wedge u0----- \vee w10-- \wedge u1-----) \vee v0111 \wedge w00-- \wedge u----- \vee \\
& v1000 \wedge (w01-- \wedge u--0--- \vee w10-- \wedge u--1---) \vee v1001 \wedge w00-- \wedge u-----) \wedge \\
& (v0000 \wedge w--00 \wedge u----- \vee v0001 \wedge w--00 \wedge u----- \vee \\
& v0010 \wedge w--00 \wedge u----- \vee v0011 \wedge w--01 \wedge u----- \vee \\
& v0100 \wedge w--00 \wedge u----- \vee v0101 \wedge (w--01 \wedge u--0-- \vee w--10 \wedge u--1--) \vee \\
& v0110 \wedge w--00 \wedge u----- \vee v0111 \wedge (w--01 \wedge u--0- \vee w--10 \wedge u--1-) \vee \\
& v1000 \wedge w--00 \wedge u----- \vee v1001 \wedge (w--01 \wedge u-----0 \vee w--10 \wedge u-----1)) \wedge \\
& (v0000 \wedge w---- \wedge u----- \vee v0001 \wedge w---- \wedge u----- \vee \\
& v0010 \wedge w---- \wedge u----- \vee v0011 \wedge w---- \wedge u----- \vee \\
& v0100 \wedge w---- \wedge u----- \vee v0101 \wedge w---- \wedge u----- \vee \\
& v0110 \wedge w---- \wedge u----- \vee v0111 \wedge w---- \wedge u----- \vee \\
& v1000 \wedge w---- \wedge u----- \vee v1001 \wedge w---- \wedge u-----)
\end{aligned}$$

Above we replaced conjunctions of variables within a Boolean vector with a sequential representation of their sign or absence. For example $w--01$ is a shorthand for $\neg w_2 \wedge w_3$.

Furthermore, let \bar{v}' be the Boolean vector representing successor (global) states. We can then represent the global protocol and the evolution function as Boolean formulas $P[\bar{v}, \bar{w}]$ and $\text{tr}[\bar{v}, \bar{w}, \bar{v}']$ by taking the conjunctions of Boolean formulas representing the individual agents' protocols P_a and the evolution functions tr_a for $a \in \text{Ag}$:

$$\begin{aligned}
P[\bar{v}, \bar{w}] & := \bigwedge_{a \in \text{Ag}} \bigvee_{s_a \in \text{St}_a} \left(s_a[\bar{v}_a] \wedge \bigvee_{c \in P_a(s_a)} c[\bar{w}_a] \right) \\
\text{tr}[\bar{v}, \bar{w}, \bar{v}'] & := \bigwedge_{a \in \text{Ag}} \bigvee_{s_a \in \text{St}_a} \left(s_a[\bar{v}_a] \wedge \bigvee_{\delta \in \text{Dc}} \delta[\bar{w}] \wedge \text{tr}_a(s_a, \delta)[\bar{v}'^P_a] \right)
\end{aligned}$$

The evolution function for our running example is encoded as follows:

$$\begin{aligned}
\text{tr}[\bar{v}, \bar{w}, \bar{v}'] = & (v0000 \wedge w0000 \wedge v'0000 \vee v0001 \wedge w0000 \wedge v'0001 \vee \\
& v0010 \wedge w0100 \wedge v'0001 \vee v0011 \wedge w0100 \wedge v'0000 \vee \\
& v0100 \wedge (w0100 \wedge v'0011 \vee w1000 \wedge v'0001) \vee \\
& v0101 \wedge (w0001 \wedge v'0010 \vee w0010 \wedge v'0000) \vee \\
& v0110 \wedge (w0100 \wedge v'0101 \vee w1000 \wedge v'0011) \vee \\
& v0111 \wedge (w0001 \wedge v'0100 \vee w0010 \wedge v'0010) \vee \\
& v1000 \wedge (w0100 \wedge v'0111 \vee w1000 \wedge v'0101) \vee \\
& v1001 \wedge (w0001 \wedge v'0110 \vee w0010 \wedge v'0100))
\end{aligned}$$

The Boolean formula $R_t^b[\bar{v}, \bar{w}, \bar{v}']$ for the *implied transition relation* $\rightarrow_v^b \subseteq \text{St} \times \text{St}$ (see Definition 11) is then constructed from the conjunction of the Boolean formulas representing the global protocol, the global evolution function

and the strategy restrictions⁶:

$$R_t^b[\bar{v}, \bar{v}', \bar{u}] := \exists \bar{w}. \mathbf{P}[\bar{v}, \bar{w}] \wedge \text{tr}[\bar{v}, \bar{w}, \bar{v}'] \wedge S^b[\bar{v}, \bar{w}, \bar{u}]$$

Note that we quantify over actions, encoded as \bar{w} , but we keep the variable assignment in the extra parameter \bar{u} .⁷ Quantification over the variable assignment is performed when a strategy quantifier ($\langle\langle x \rangle\rangle$, $\llbracket x \rrbracket$) is encountered. Also note that the strategy restrictions S^b depend on the binding b and thus have to be recomputed when the binding is updated, *i.e.*, when the agent binding operator (a, x) is encountered. This is not the case for the other Boolean formulas ($\mathbf{P}[\bar{v}, \bar{w}]$ and $\text{tr}[\bar{v}, \bar{w}, \bar{v}']$), which are constant for a given interpreted system.

The implied transition relation for our running example is encoded as follows:

$$\begin{aligned} R_t^b[\bar{v}, \bar{v}', \bar{u}] = & (v0000 \wedge v'0000 \vee v0001 \wedge v'0001 \vee \\ & v0010 \wedge v'0001 \vee v0011 \wedge v'0000 \vee \\ & v0100 \wedge (v'0011 \wedge u0----- \vee v'0001 \wedge u1-----) \vee \\ & v0101 \wedge (v'0010 \wedge u---0-- \vee v'0000 \wedge u---1--) \vee \\ & v0110 \wedge (v'0101 \wedge u-0---- \vee v'0011 \wedge u-1----) \vee \\ & v0111 \wedge (v'0100 \wedge u----0- \vee v'0010 \wedge u----1-) \vee \\ & v1000 \wedge (v'0111 \wedge u--0--- \vee v'0101 \wedge u--1---) \vee \\ & v1001 \wedge (v'0110 \wedge u-----0 \vee v'0100 \wedge u-----1)) \end{aligned}$$

The *epistemic accessibility relations* R_a^K , R_A^E , R_A^D , R_A^C for an agent $a \in \text{Ag}$ and a set of agents $A \subseteq \text{Ag}$ are also constant for a given interpreted system and are encoded in a similar fashion [39]:

$$\begin{aligned} R_a^K[\bar{v}, \bar{v}'] & := \bigvee_{(s, s') \in \sim_a} s[\bar{v}] \wedge s'[\bar{v}'] \\ R_A^E[\bar{v}, \bar{v}'] & := \bigvee_{a \in A} R_a^K[\bar{v}, \bar{v}'] \\ R_A^D[\bar{v}, \bar{v}'] & := \bigwedge_{a \in A} R_a^K[\bar{v}, \bar{v}'] \\ R_A^C[\bar{v}, \bar{v}'] & := \text{lf}_{\text{p}_\rho} (R_A^E[\bar{v}, \bar{v}'] \vee (\exists v^\circ. \rho[\bar{v}, v^\circ] \wedge \rho[v^\circ, \bar{v}'])) \end{aligned}$$

Finally, the algorithm $\text{Check}_{\mathcal{I}}: \text{SLK} \times \text{Bnd} \rightarrow 2^{\text{Ext}}$ can be translated into operations on BDDs representing sets of extended states.

To check the SLK sentence φ_{win_A} in \mathcal{I}_{Nim} , we perform the following steps (see Definition 13):

1. We calculate the set of extended states which guarantee win_A :

$$\text{Check}_{\mathcal{I}_{\text{Nim}}}(\text{win}_A, b)[\bar{v}, \bar{w}] = \mathbf{h}(\text{win}_A)[v] = v0001$$

⁶Observe that, for all bindings b and Boolean vectors $\bar{u}, \bar{v}, \bar{w}$, if $S^b[\bar{v}, \bar{w}, \bar{u}]$ is true, then so is $\mathbf{P}[\bar{v}, \bar{w}]$. In other words, the protocol is already included in the strategy restrictions. Hence, the encoding of $R_t^b[\bar{v}, \bar{v}', \bar{u}]$ could be simplified to $\exists \bar{w}. \text{tr}[\bar{v}, \bar{w}, \bar{v}'] \wedge S^b[\bar{v}, \bar{w}, \bar{u}]$.

⁷Observe that we do not need to allocate BDD variables \bar{u}' for any successor states as these are not affected by a temporal transition (see Definition 12).

2. We evaluate the least fixed point of $F[\bar{v}, \bar{u}] = \mathbf{Check}_{\mathcal{I}_{\text{Nim}}}(\text{win}_A, b)[\bar{v}, \bar{u}] \vee \exists \bar{v}'. R_t^b[\bar{v}, \bar{u}, \bar{v}'] \wedge F[\bar{v}', \bar{u}]$ to find the set of extended states from which win_A can be eventually guaranteed:

$$\begin{aligned}
F^0[\bar{v}, \bar{u}] &= v0001 \\
F^1[\bar{v}, \bar{u}] &= F^0[\bar{v}, \bar{u}] \vee v0010 \vee v0100 \wedge u1----- \\
F^2[\bar{v}, \bar{u}] &= F^1[\bar{v}, \bar{u}] \vee v0101 \wedge u---0-- \vee v0111 \wedge u----1- \vee \\
&\quad v0111 \wedge u1---0- \vee v1001 \wedge u1-----1 \\
F^3[\bar{v}, \bar{u}] &= F^2[\bar{v}, \bar{u}] \vee v0110 \wedge u-0-0-- \vee v1000 \wedge u--10-- \vee \\
&\quad v1000 \wedge u--0-1- \vee v1000 \wedge u1-0-0- \\
F^4[\bar{v}, \bar{u}] &= F^3[\bar{v}, \bar{u}] \vee v1001 \wedge u-0-0-0 = F^5[\bar{v}, \bar{u}] \\
\mathbf{Check}_{\mathcal{I}_{\text{Nim}}}(\text{F win}_A, b)[\bar{v}, \bar{u}] &= v0001 \vee v0010 \vee v0100 \wedge u1----- \vee v0101 \wedge u---0-- \vee \\
&\quad v0111 \wedge (u----1- \vee u1---0-) \vee v0110 \wedge u-0-0-- \vee \\
&\quad v1000 \wedge (u--10-- \vee u--0-1- \vee u1-0-0-) \vee \\
&\quad v1001 \wedge (u1-----1 \vee u-0-0-0)
\end{aligned}$$

3. We universally quantify over the Boolean variables associated with strategy variables y and e to find the set of extended states which guarantee $\psi_{\text{win}_A} = \llbracket y \rrbracket \llbracket e \rrbracket (A, x)(B, y)(\text{Env}, e)\text{F win}_A$:

$$\begin{aligned}
\mathbf{Check}_{\mathcal{I}_{\text{Nim}}}(\psi_{\text{win}_A}, \emptyset)[\bar{v}, \bar{u}] &= \forall u_3 u_4 u_5. \mathbf{Check}_{\mathcal{I}_{\text{Nim}}}(\text{F win}_A, b)[\bar{v}, \bar{u}] \\
&= v0001 \vee v0010 \vee v0100 \wedge u1----- \vee \\
&\quad v0111 \wedge u1----- \vee v1000 \wedge u1-0---
\end{aligned}$$

4. We existentially quantify over the Boolean variables associated with the strategy variable x to find the set of extended states which guarantee $\varphi_{\text{win}_A} = \langle\langle x \rangle\rangle \psi_{\text{win}_A}$:

$$\begin{aligned}
\mathbf{Check}_{\mathcal{I}_{\text{Nim}}}(\varphi_{\text{win}_A}, \emptyset)[\bar{v}, \bar{u}] &= \exists u_0 u_1 u_2. \mathbf{Check}_{\mathcal{I}_{\text{Nim}}}(\psi_{\text{win}_A}, \emptyset)[\bar{v}, \bar{u}] \\
&= v0001 \vee v0010 \vee v0100 \vee v0111 \vee v1000
\end{aligned}$$

5. Finally, to determine whether φ_{win_A} holds in all initial states of \mathcal{I}_{Nim} , we check if I is a subset of $\mathbf{Check}_{\mathcal{I}_{\text{Nim}}}(\varphi_{\text{win}_A}, \emptyset)$ symbolically:

$$\begin{aligned}
(I \cap \mathbf{Check}_{\mathcal{I}_{\text{Nim}}}(\varphi_{\text{win}_A}, \emptyset))[\bar{v}, \bar{u}] &= I[\bar{v}] \wedge \mathbf{Check}_{\mathcal{I}_{\text{Nim}}}(\varphi_{\text{win}_A}, \emptyset)[\bar{v}, \bar{u}] \\
&= v0001 \wedge (v0001 \vee v0010 \vee v0100 \vee v0111 \vee v1000) \\
&= v0001 \\
&= I[\bar{v}]
\end{aligned}$$

We conclude that φ_{win_A} holds in \mathcal{I}_{Nim} , *i.e.*, $\mathcal{I}_{\text{Nim}} \models_{\text{SLK}} \varphi_{\text{win}_A}$.

3.3.4. Strategy Synthesis

Synthesising *witness* and *counterexample strategies* (see Definition 14) using the symbolic implementation is straightforward because our representation of variable assignments ensures that all strategies are uniform (see Subsection 2.3).

To synthesise the witness strategies for an SLK sentence $\varphi = \langle\langle x_0 \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle \psi$ at a global state $s \in \text{St}$, we proceed as follows. Since $\mathcal{I}, (\emptyset, s) \models_{\text{SLK}} \varphi$, there exist witness strategies for φ at s (see Lemma 4). We calculate $E = \mathbf{Check}_{\mathcal{I}}(\psi, \emptyset)$ and pick an arbitrary extended state $\langle s, v \rangle \in E$. Note that all strategies in v are uniform due to our representation. By Lemma 5, there exists at least one

such extended state $\langle s, v \rangle$ (because there exist witness strategies for φ at s) and for all such extended states, $v(x_0), \dots, v(x_{n-1})$ are witness strategies for φ at s . These steps can be implemented symbolically as follows:

1. Calculate $E = \text{Check}_{\mathcal{I}}(\psi, \emptyset)$. This is performed using the symbolic implementation presented earlier in this subsection.
2. Remove from the set E all the extended states with a different underlying global state $E' = \{(s', v') \in E \mid s' = s\}$. The symbolic representation of this operation is $E'[\bar{v}, \bar{u}] := E[\bar{v}, \bar{u}] \wedge s[\bar{v}]$.
3. Pick an arbitrary extended state $\langle s, v \rangle \in E'$. This is equivalent to selecting one conjunct (also referred to as *minterm*) $C[\bar{v}, \bar{u}]$ from $E'[\bar{v}, \bar{u}]$. BDD packages usually provide a built-in function for this operation⁸.
4. The conjunct $C[\bar{v}, \bar{u}]$ encodes the extended state $\langle s, v \rangle$, where v contains the witness strategies. To get the next action $f_i(s')$ of a strategy $f_i = v(x_i)$ for $0 \leq i < n$ at a global state $s' \in \text{St}$, we find a possible action $c \in \bigcap_{a \in \text{shr}(\psi, x_i)} P_a(s_a(s'))$ such that $C[\bar{v}, \bar{u}] \wedge c[\bar{u}_{s', x_i}]$ is *not* equivalent to false, where $c[\bar{u}_{s', x_i}]$ is the Boolean formula representing the fact that the next action of the strategy mapped to variable x_i at the global state s' is c .

To finish our running example, we show how to synthesise a witness strategy $f_x: \text{St} \rightarrow \text{Ac}_A$ for x at the initial state s_{4A} in $\varphi_{\text{win}_A} = \langle\langle x \rangle\rangle \psi_{\text{win}_A}$:

1. $E[\bar{v}, \bar{u}] = \text{Check}_{\mathcal{I}_{\text{Nim}}}(\psi_{\text{win}_A}, \emptyset)[\bar{v}, \bar{u}] = v0001 \vee v0010 \vee v0100 \wedge u1\text{-----} \vee v0111 \wedge u1\text{-----} \vee v1000 \wedge u1\text{-0---}$ (we have already done this).
2. $E'[\bar{v}, \bar{u}] = E[\bar{v}, \bar{u}] \wedge s_{4A}[\bar{v}] = v1000 \wedge u1\text{-0---}$.
3. $C[\bar{v}, \bar{u}] = v1000 \wedge u1\text{-0---}$ (there is only one conjunct in $E'[\bar{v}, \bar{u}]$).
4. The variable assignment in $C[\bar{v}, \bar{u}] = v1000 \wedge u0 \wedge \neg u_2$ encodes the witness strategy f_x :

$$\begin{array}{ll} f_x(s_{*B}) = \perp & f_x(s_{2A}) = 2 \quad [u_0] \\ f_x(s_{0*}) = \perp & f_x(s_{3A}) = 1 \text{ or } 2 \quad [u_1 \vee \neg u_1] \\ f_x(s_{1A}) = 1 & f_x(s_{4A}) = 1 \quad [\neg u_2] \end{array}$$

3.3.5. Complexity

As we have just shown, the symbolic encoding of the labelling algorithm has polynomial size and requires polynomially many Boolean variables. Therefore, the algorithm runs in *exponential time* with respect to both the size of the model $|\mathcal{I}|$ and the number of quantified variables $|\text{vars}(\varphi)|$.

Theorem 3. *Let \mathcal{I} be an arbitrary interpreted and $\varphi \in \text{SLK}$ be an SLK sentence. The worst case time complexity of the symbolic implementation of the model checking algorithm $\text{Check}_{\mathcal{I}}(\varphi, \emptyset)$ is:*

$$O(|\varphi| \times |\text{Ag}|) \times 2^{O(|\text{St}| \times |\text{vars}(\varphi)| \times \log_2 |\text{Ac}|)}$$

⁸If no such function is available, we can skip step 3 ($C[\bar{v}, \bar{u}] := E'[\bar{v}, \bar{u}]$) and refine the conjunct upon each lookup in step 4 ($C[\bar{v}, \bar{u}] := C[\bar{v}, \bar{u}] \wedge c[\bar{u}_{s', x_i}]$).

where Ag , St and Ac are the sets of agents, reachable global states and total actions of \mathcal{I} , respectively.

Proof (Sketch). The symbolic implementation uses:

- $O(\log_2 |\text{St}|)$ Boolean variables to represent the *current global state*;
- $O(\log_2 |\text{St}|)$ Boolean variables to represent the *next global state*;
- $O(\log_2 |\text{Ac}|)$ Boolean variables to represent the *decisions*;
- $O(|\text{St}| \times |\text{vars}(\varphi)| \times \log_2 |\text{Ac}|)$ Boolean variables to represent the *variable assignment*.

The total number of Boolean variables is thus $O(|\text{St}| \times |\text{vars}(\varphi)| \times \log_2 |\text{Ac}|)$. Consequently, any BDDs built on top of these variables will have at most $2^{O(|\text{St}| \times |\text{vars}(\varphi)| \times \log_2 |\text{Ac}|)}$ nodes. Since individual BDD operations take polynomial time with respect to the size of the relevant BDDs [41], the worst case time complexity of each BDD operation is $2^{O(|\text{St}| \times |\text{vars}(\varphi)| \times \log_2 |\text{Ac}|)}$. As the recursive algorithm presented in Definition 13 will perform at most

$$\underbrace{O(|\varphi| \times |\text{St}|)}_{\text{check}_{\mathcal{I}} \text{ calls}} \times \underbrace{O(|\text{St}| \times |\text{Ag}| \times |\text{Ac}|)}_{\text{size of } R_t^b \text{ and } R_a^K \text{ defs}} \times \underbrace{O(|\text{St}| \times |\text{vars}(\varphi)| \times \log_2 |\text{Ac}|)}_{\text{Boolean variables}}$$

BDD operations, our claim follows. \square

The proof shows that the complexity of the algorithm is dominated by the encoding of the variable assignments.

Observe that the symbolic implementation of the algorithm $\text{check}_{\mathcal{I}}$ can use more than a polynomial amount of space because it performs operations on sets of extended states. While the procedure described in the proof of Theorem 1 uses only a polynomial amount of space (as it operates on individual states and assignments), we believe it is unlikely that such an explicit approach would outperform the symbolic algorithm in practice [42].

4. Implementation and Experimental Results

The model checker $\text{MCMAS}_{\text{SLK}}$ [43], first introduced in [24], implements the state labelling algorithm presented in Section 3. The tool is based on the existing open-source model checker MCMAS [20], which supports the verification of CTL and ATL formulas with fairness constraints and epistemic and deontic modalities. Both checkers are written in C++ and based on the CUDD BDD package [44].

$\text{MCMAS}_{\text{SLK}}$ takes as input an ISPL file [45] containing a description of an interpreted system (see Definition 4) and a list of SLK specifications. For each specification, the tool calculates the associated set of reachable extended states, encoded as BDDs, and verifies the specification using the symbolic implementation of the state labelling algorithm presented in Section 3. If requested, the

checker also generates witnesses, counterexamples and strategies for each specification (see Definition 14). The latter is an enhancement of MCMAS as it provides the means to automatically synthesise agents' behaviour satisfying an SLK specification.

The usage of the SLK extension is the same as that of the original tool. Given an ISPL file `system.ispl`, the following command checks the SLK specifications against the interpreted system:

```
$ ./mcmas system.ispl
```

A witness/counterexample execution with strategies can be requested (where possible):

```
$ ./mcmas -c 1 system.ispl
```

We present here the experimental results obtained using $\text{MCMAS}_{\text{SLK}}$ on several scalable real-life scenarios. The experiments were run on an Intel Core i7-3770 CPU 3.40GHz machine with 16GB RAM running Linux kernel version 3.8.0-35-generic. We measure the amount of time and memory used by the tool and compare it with the performance of the original tool, MCMAS, on CTL and ATL (where possible). The examples in this section also demonstrate the expressiveness of SLK.

4.1. Dining Cryptographers

The dining cryptographers protocol [46, 45] is a commonly studied anonymity protocol. It is suited as a case study as anonymity can naturally be expressed as lack of knowledge [47]. We model the protocol with $n \geq 3$ cryptographers as an interpreted system (see Definition 4) with agents $\text{Ag} = \{\text{Env}, c_1, \dots, c_n\}$. Table 1 reports the results obtained when verifying the dining cryptographers protocol against the following CTLK, ATLK and SLK specifications:

$$\varphi_{\text{CTLK}} = \text{AG } \psi \tag{1}$$

$$\varphi_{\text{ATLK}} = \langle\langle \emptyset \rangle\rangle \mathbf{G} \psi \tag{2}$$

$$\varphi_{\text{SLK}} = \llbracket x_e \rrbracket \llbracket x_1 \rrbracket \cdots \llbracket x_n \rrbracket (\text{Env}, x_e)(c_1, x_1) \cdots (c_n, x_n) \mathbf{G} \psi$$

where:

$$\psi = (\text{odd} \wedge \neg \text{paid}_1) \rightarrow \underbrace{\left[\mathbf{K}_{c_1} \bigvee_{i=2}^n \text{paid}_i \right]}_{\text{cryptographer } c_1 \text{ knows that another cryptographer paid}} \wedge \underbrace{\left[\bigwedge_{i=2}^n \neg \mathbf{K}_{c_1} \text{paid}_i \right]}_{\text{cryptographer } c_1 \text{ does not know which cryptographer paid}}$$

φ_{CTLK} and φ_{ATLK} are the usual epistemic specifications for the protocol [48, 45] and φ_{SLK} is its natural extension where strategies are quantified. The results indicate that the checker can verify reasonably large state spaces. The performance depends on the number of Boolean variables required to represent the

extended states. In the case of SLK specifications, the number of Boolean variables is proportional to the number of states. The last six columns of Table 1 show that the tool’s performance drops considerably faster when verifying SLK formulas compared to CTLK and ATLK ones. This is because neither CTLK nor ATLK verification requires assignments; hence extended states collapse to plain states. In contrast, both CTLK and ATLK performance is dominated by the reachable state space computation. We found that splitting the extended state space generation across multiple threads has almost negligible impact on performance; this is due to the fact that the CUDD BDD package used by MCMAS does not support concurrency.

4.2. Cake Cutting

The cake-cutting problem, a well-known mathematical puzzle in which n agents take turns to slice a cake of size d and the environment responds by trying to ensure the cake is divided fairly, was already described in Section 2 together with the SLK specification $\varphi = \langle\langle x \rangle\rangle (\varphi_F \wedge \varphi_S)$ asserting the existence of a solution. This example demonstrates not only the ability of SLK to express *Nash equilibria*, but also the fact that the checker can synthesise protocols which achieve them.

We were able to verify the formula φ defined above on a system with $n = 2$ agents and a cake of size $d = 2$. Moreover, the checker automatically synthesised a witness strategy f_x for the environment as well as the Nash equilibrium (f_{y_1}, f_{y_2}) for the agents. For more details, see [49, Subsection 6.4.2].

We were unable to verify larger examples due to out of memory errors. For example, with $n = 2$ agents and $d = 3$ slices, there are 29 reachable states and the encoding requires 105 Boolean variables, 54 of which represent the strategy f_x we wish to synthesise. The intermediate BDDs represent 8.09×10^{20} possible *extended* states using in the order of 10^9 nodes. A single intermediate BDD thus requires approximately 32GB of memory and the whole algorithm would need to use over 10^{11} bytes to complete the calculation. This is to be expected given the theoretical difficulty of the problem.

4.3. Scheduler

Lastly, we consider a *preemptive scheduler system* [49, Subsection 6.4.3] composed of $n \in \mathbb{N}$ processes and an arbiter which ensures mutual exclusion of a shared resource while preventing starvation, to compare the performance of CTL, ATL and SLK model checking on safety and fairness properties.

The desired properties of the system can be expressed using the following SLK specifications:

1. **Mutual exclusion.** The following specification asserts that at most one process owns the resource at any given point in time:

$$\varphi_{\text{ME}} = \llbracket x \rrbracket \llbracket y_1 \rrbracket \cdots \llbracket y_n \rrbracket (\text{Env}, x)(1, y_1) \cdots (n, y_n) \text{G} \neg \bigvee_{i=1}^n \bigvee_{j=i+1}^n \langle \text{rs}, i \rangle \wedge \langle \text{rs}, j \rangle$$

crypts n	possible states	reachable states	reachability time (s)	CTLK		ATLK		SLK					
								1 thread		2 threads		4 threads	
				time (s)	mem (MB)	time (s)	mem (MB)	time (s)	mem (MB)	time (s)	mem (MB)	time (s)	mem (MB)
3	82944	128	0.00	0.00	2.69	0.00	2.69	0.00	2.79	0.01	4.91	0.02	9.23
4	1.99×10^6	320	0.02	0.00	2.71	0.00	2.70	0.00	2.91	0.02	4.97	0.04	9.24
5	4.78×10^7	768	0.04	0.00	2.75	0.00	2.75	0.01	3.07	0.03	5.10	0.05	9.67
6	1.15×10^9	1792	0.08	0.00	2.79	0.00	2.77	0.02	3.41	0.04	5.37	0.07	9.94
7	2.75×10^{10}	4096	0.21	0.00	3.03	0.00	2.97	0.05	4.47	0.08	5.88	0.12	10.71
8	6.60×10^{11}	9216	0.40	0.00	3.19	0.00	3.27	0.18	6.37	0.21	7.42	0.24	12.56
9	1.58×10^{13}	20480	0.43	0.00	3.38	0.00	3.42	0.38	9.59	0.38	9.82	0.46	15.96
10	3.80×10^{14}	45056	4.19	0.28	13.02	0.16	12.75	2.29	18.62	2.23	22.33	2.43	28.21
11	9.13×10^{15}	98304	1.69	0.04	5.49	0.01	4.75	5.01	23.55	4.37	27.51	4.77	39.59
12	2.19×10^{17}	212992	2.32	0.01	4.40	0.01	4.20	11.61	39.25	9.61	40.36	10.32	60.82
13	5.26×10^{18}	458752	2.05	0.10	6.50	0.03	5.62	32.63	81.95	26.30	79.75	28.39	118.27
14	1.26×10^{20}	983040	1.96	0.08	6.87	0.03	5.00	89.46	169.95	67.99	169.68	73.57	243.80
15	3.03×10^{21}	2.10×10^6	21.35	0.35	12.73	0.08	9.29	163.77	363.66	139.42	354.18	156.57	500.96
16	7.27×10^{22}	4.46×10^6	6.66	0.09	6.65	0.04	4.83	422.03	758.43	345.83	736.02	377.50	977.79
17	1.74×10^{24}	9.44×10^6	9.10	0.13	6.67	0.08	6.04	734.44	1360.59	643.03	1270.93	705.44	1359.73
18	4.19×10^{25}	1.99×10^7	65.94	0.50	12.66	0.14	12.67	2654.46	3550.66	2129.54	3680.89	<i>process killed</i>	

Table 1: Experimental model checking results for the dining cryptographers protocol.

This property can be equivalently expressed in CTL and ATL using the AG and $\langle\langle\emptyset\rangle\rangle\mathbf{G}$ operators, respectively (see Equations 1 and 2). The results in Table 2 show that the equivalent CTL and ATL formulas are checked almost instantaneously. Similarly to the dining cryptographers scenario (see Subsection 4.1), the SLK model checking performance drops quickly for $n \geq 8$ processes.

2. **Absence of starvation.** The following specification asserts the existence of an arbiter strategy which ensures that every request is eventually satisfied:

$$\varphi_{\text{AS}} = \langle\langle x \rangle\rangle [[y_1]] \dots [[y_n]] (\text{Env}, x)(1, y_1) \dots (n, y_n) \mathbf{G} \bigwedge_{i=1}^n (\langle \text{wt}, i \rangle \rightarrow \mathbf{F} \neg \langle \text{wt}, i \rangle)$$

This property cannot be expressed in CTL because it refers to the capability of an agent to enforce a property. Moreover, it cannot be expressed in ATL either, because it contains nested temporal operators (\mathbf{G} and \mathbf{F}) which depend on the same strategies. The model checking results for the formulas with $2 \leq n \leq 3$ processes are shown in Table 2. In this case MCMAS_{SLK} ran out of physical memory for $n > 3$ processes.

4.4. Analysis

In general, MCMAS_{SLK} has significantly lower performance than the original tool on CTL and ATL. We expected this outcome given the high model checking complexity of SLK. However, the checker can verify specifications which cannot be checked by the original tool. We have discussed two examples of such properties: (i) Nash equilibria (see Subsection 4.2) and (ii) fairness conditions (see Subsection 4.3). Moreover, the checker supports witness strategy synthesis, which can be used to automatically generate agents' behaviour. In other words, the aim of MCMAS_{SLK} is not to replace MCMAS, but to add support for completely new features. Therefore, the results presented in this section should be regarded as an illustration of the new functionality.

The experimental results presented in this section confirm that the main performance bottleneck of MCMAS_{SLK} is the BDD encoding of the extended states, which allocates separate BDD variables for each shared local state of a strategy (see Subsection 3.3). Nevertheless, the checker can handle reasonably large state spaces for certain specifications as we have seen in the dining cryptographers scenario (see Subsection 4.1). We were able to slightly improve the tool's performance by using multiple threads (see Table 1) despite the fact that the CUDD BDD package used by MCMAS has no built-in support for concurrency. This suggests that further speedup could be achieved by using a different BDD package with a higher degree of parallelism.

5. Conclusions and Related Work

As argued in the introduction, considerable progress has been made over the past 15 years in the development of methodologies for the verification of

procs n	possible states	reach. states	reach. time (s)	Mutual exclusion φ_{ME}						Absence of starvation φ_{AS}	
				CTL		ATL		SLK		SLK	
				time (s)	mem (MB)	time (s)	mem (MB)	time (s)	mem (MB)	time (s)	mem (MB)
2	72	9	0.00	0.00	2.64	0.00	2.64	0.00	2.67	0.00	2.67
3	432	21	0.00	0.00	2.72	0.00	2.72	0.00	2.81	116.78	124.90
4	2592	49	0.01	0.00	2.81	0.00	2.81	0.00	3.11	<i>out of memory</i>	
5	15552	113	0.02	0.00	3.02	0.00	3.03	0.02	4.20	<i>out of memory</i>	
6	93312	257	0.02	0.00	3.40	0.00	3.40	0.11	8.98	<i>out of memory</i>	
7	559872	577	0.08	0.00	4.47	0.00	4.48	0.86	23.18	<i>out of memory</i>	
8	3.36×10^6	1281	0.17	0.00	6.73	0.00	6.72	4.88	91.26	<i>out of memory</i>	
9	2.02×10^7	2817	0.20	0.00	11.62	0.00	11.62	36.71	561.69	<i>out of memory</i>	
10	1.21×10^8	6145	0.10	0.00	22.23	0.00	22.23	222.36	2769.64	<i>out of memory</i>	

Table 2: Experimental model checking results for the preemptive scheduler system.

MAS against agent-based specifications, with particular emphasis on epistemic specifications. Statements capturing the strategic interplay between agents in a system have also been supported, but they have traditionally been limited to ATL specifications.

It has been argued that the intrinsic limitations of ATL with respect to direct naming and binding of strategies severely restrict the expressiveness of the specifications [50, 23, 51]. These limitations are particularly severe when reasoning about game theoretic concepts. In response to this criticism, *Strategy logic* (SL), among others, was put forward [23]. Sophisticated concepts such as Nash equilibria, which cannot be expressed by ATL, can naturally be encoded in SL.

Given this, the question that arises is whether *automatic and efficient* verification methodologies for MAS against SL specifications can be devised. To answer this question, we have introduced and studied SLK, a variant of epistemic logic incorporating some SL concepts defined on a memoryless variant of interpreted systems. We have shown that SLK admits an efficient labelling algorithm solving the associated model checking problem and introduced MCMAS_{SLK}, an extension of the model checker MCMAS [20], for the verification of MAS against SLK specifications.

A notable feature of MCMAS_{SLK} is that it allows for automatic verification of game concepts such as various forms of equilibria, including Nash equilibria and subgame perfect Nash equilibria. Since SLK also supports a limited form of epistemic modalities, the proposed approach further enables us to express specifications concerning individual and group knowledge of cooperation properties. Furthermore, since SLK subsumes ATLK*, MCMAS_{SLK} is also the first tool enabling the verification of multi-agent systems against specifications given in ATLK*.

Related Work. The present work builds on [19, 20] where a symbolic approach for the verification of a variant of ATL was put forward and implemented in MCMAS. While the present work reuses parts of the labelling algorithm and its implementation, it supports SL modalities which were not covered in [20].

A number of investigations have explored the complexity of the model checking problem for combinations of plain ATL and epistemic operators under a number of assumptions [15, 16, 17]. These identify lower and upper bounds for the verification problem; but no implementation is reported.

In [52, 53] a combination of epistemic logic and a logic supporting strategic concepts was introduced together with a BDD-based toolkit. Similarly to the present approach, [52] assumes incomplete information and memoryless strategies. This work, however, shares ATL’s implicit notion of strategies. In contrast, here we follow SL closely and adopt an explicit notion of strategies which can be bound to and shared by different agents.

In [54] a verification methodology for a subset of SL, namely its one-goal fragment, was introduced. However the underlying assumptions are quite different as perfect recall and complete information is assumed. Furthermore, the logic does not support epistemic operators.

Future Work. The experiments that we reported show that the perfor-

mance of $\text{MCMAS}_{\text{SLK}}$ on SLK is comparable to that of MCMAS on equivalent ATL and CTLK formulas. This is because we adopted an approach in which the colouring with strategies is specification-dependent and is only performed after the set of reachable states is computed. We found that the main impediment to better performance of the tool is the size of the BDDs required to encode sets of extended states. Future efforts will be devoted to mitigate this problem as well as to support logics stronger than SLK, including relaxing the assumption on epistemic sentences.

Acknowledgements. This research was partly funded by the EPSRC (grant EP/I00520X), the Regione Campania (Embedded System Cup project B25B09090 100007), the EU (FP7 project 600958-SHERPA), and the MIUR (ORCHESTRA project).

Bibliography

- [1] M. Wooldridge, An introduction to MultiAgent systems, second edition Edition, Wiley, 2009.
- [2] R. Fagin, J. Y. Halpern, Y. Moses, M. Y. Vardi, Reasoning about Knowledge, MIT Press, Cambridge, 1995.
- [3] A. Lomuscio, M. Sergot, Deontic interpreted systems, *Studia Logica* 75 (1) (2003) 63–92.
- [4] A. S. Rao, Decision procedures for propositional linear-time Belief-Desire-Intention logics, in: M. Wooldridge, J. P. Müller, M. Tambe (Eds.), *Intelligent Agents II (LNAI 1037)*, Springer-Verlag: Heidelberg, Germany, 1996, pp. 33–48.
- [5] A. Pnueli, Applications of temporal logic to the specification and verification of reactive systems: A survey of current trends., in: J. W. de Bakker, W. P. de Roever, G. Rozenberg (Eds.), *Current Trends in Concurrency, Overviews and Tutorials, LNCS 224*, Springer, 1986, pp. 510–584.
- [6] W. Penczek, A. Lomuscio, Verifying epistemic properties of multi-agent systems via bounded model checking, *Fundamenta Informaticae* 55 (2) (2003) 167–185.
- [7] M. Kacprzak, W. Nabialek, A. Niewiadomski, W. Penczek, A. Pólrola, M. Szreter, B. Wozna, A. Zbrzezny, Verics 2007 - a model checker for knowledge and real-time, *Fundamenta Informaticae* 85 (1-4) (2008) 313–328.
- [8] R. Bryant, Graph-Based Algorithms for Boolean Function Manipulation., *Transactions on Computers* 35 (8) (1986) 677–691.
- [9] P. Gammie, R. van der Meyden, MCK: Model checking the logic of knowledge, in: *Proceedings of 16th International Conference on Computer Aided Verification (CAV'04)*, LNCS 3114, Springer, 2004, pp. 479–483.

- [10] F. Raimondi, A. Lomuscio, Automatic verification of multi-agent systems by model checking via OBDDs, *Journal of Applied Logic* 5 (2) (2005) 235–251.
- [11] M. Cohen, M. Dam, A. Lomuscio, H. Qu, A symmetry reduction technique for model checking temporal-epistemic logic, in: *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI'09)*, Pasadena, USA, 2009, pp. 721–726.
- [12] M. Cohen, M. Dam, A. Lomuscio, F. Russo, Abstraction in model checking multi-agent systems, in: *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'09)*, IFAAMAS Press, Budapest, Hungary, 2009, pp. 945–952.
- [13] A. Lomuscio, J. Michaliszyn, Verification of multi-agent systems via predicate abstraction against ATLK specifications, in: *Proc. of the 15th Int. Conference on Autonomous Agents and Multiagent Systems (AAMAS'16)*, 2016, pp. 662–670.
- [14] R. Alur, L. de Alfaro, R. Grosu, T. Henzinger, A. Thomas, M. Kang, C. Kirsch, R. Majumdar, F. Mang, B.-Y. Wang, jMocha: A model checking tool that exploits design structure, in: *Proceedings of the 23rd International Conference on Software Engineering (ICSE'01)*, IEEE, 2001, pp. 835–836.
- [15] W. Hoek, M. Wooldridge, Cooperation, knowledge, and time: Alternating-time temporal epistemic logic and its applications, *Studia Logica* 75 (1) (2003) 125–157.
- [16] W. Jamroga, Some remarks on alternating temporal epistemic logic, in: B. Dunin-Kęplicz, R. Verbrugge (Eds.), *Proceedings of the International Workshop on Formal Approaches to Multi-Agent Systems (FAMAS'03)*, 2004, pp. 133–140.
- [17] T. Ågotnes, V. Goranko, W. Jamroga, M. Wooldridge, Knowledge and ability, in: *Handbook of Logics for Knowledge and Belief*, College Publications, 2015.
- [18] N. Bulling, W. Jamroga, Rational play and rational beliefs under uncertainty, in: *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'16)*, IFAAMAS, 2009, pp. 257–264.
- [19] A. Lomuscio, F. Raimondi, Model checking knowledge, strategies, and games in multi-agent systems, in: *Proceedings of the 5th International Joint Conference on Autonomous agents and Multi-Agent Systems (AAMAS'06)*, ACM Press, 2006, pp. 161–168.
- [20] A. Lomuscio, H. Qu, F. Raimondi, MCMAS: A model checker for the verification of multi-agent systems, *Software Tools for Technology Transfer*.

- [21] R. Alur, T. A. Henzinger, O. Kupferman, Alternating-time temporal logic, *Journal of the ACM* 49 (5) (2002) 672–713.
- [22] G. Jonker, Feasible strategies in alternating-time temporal epistemic logic, Master’s thesis, University of Utrecht, The Netherlands (2003).
- [23] F. Mogavero, A. Murano, M. Vardi, Reasoning About Strategies., in: *Foundations of Software Technology and Theoretical Computer Science’10, LIPIcs 8*, Leibniz-Zentrum fuer Informatik, 2010, pp. 133–144.
- [24] P. Čermák, A. Lomuscio, F. Mogavero, A. Murano, MCMAS-SLK: A model checker for the verification of strategy logic specifications, in: *Proceedings of the 26th International Conference on Computer Aided Verification (CAV’14)*, LNCS 8559, Springer, 2014, pp. 525–532.
- [25] K. Chatterjee, T. Henzinger, N. Piterman, Strategy Logic., *Information and Computation* 208 (6) (2010) 677–693.
- [26] F. Mogavero, A. Murano, G. Perelli, M. Vardi, Reasoning About Strategies: On the Model-Checking Problem., *Transactions On Computational Logic* 15 (4) (2014) 34:1–42.
- [27] R. Alur, T. Henzinger, O. Kupferman, Alternating-Time Temporal Logic., *Journal of the ACM* 49 (5) (2002) 672–713.
- [28] R. Fagin, J. Halpern, Y. Moses, M. Vardi, Reasoning about Knowledge., MIT Press, 1995.
- [29] A. Pnueli, The Temporal Logic of Programs., in: *Foundation of Computer Science’77*, IEEE Computer Society, 1977, pp. 46–57.
- [30] A. Lomuscio, R. Meyden, M. Ryan, Knowledge in multi-agent systems: Initial configurations and broadcast, *ACM Transactions of Computational Logic* 1 (2) (2000) 246–282.
- [31] A. Lomuscio, Knowledge sharing among ideal agents, Ph.D. thesis, School of Computer Science, University of Birmingham, Birmingham, UK (Jun. 1999).
- [32] S. Even, A. Paz, A Note on Cake Cutting., *Discrete Applied Mathematics* 7 (1984) 285–296.
- [33] O. Kupferman, M. Vardi, P. Wolper, An Automata Theoretic Approach to Branching-Time Model Checking., *Journal of the ACM* 47 (2) (2000) 312–360.
- [34] M. Vardi, P. Wolper, An Automata-Theoretic Approach to Automatic Program Verification., in: *Logic in Computer Science’86*, IEEE Computer Society, 1986, pp. 332–344.

- [35] N. Bulling, J. Dix, W. Jamroga, Model Checking Logics of Strategic Ability: Complexity., in: Specification and Verification of Multi-Agent Systems'10, Springer, 2010, pp. 125–159.
- [36] L. Stockmeyer, A. Meyer, Word Problems Requiring Exponential Time (Preliminary Report)., in: Symposium on Theory of Computing'73, Association for Computing Machinery, 1973, pp. 1–9.
- [37] M. Benerecetti, F. Mogavero, A. Murano, Reasoning About Substructures and Games., Transactions On Computational Logic 16 (3) (2015) 25:1–46.
- [38] F. Mogavero, A. Murano, L. Sauro, On the boundary of behavioral strategies, in: Proceedings of the 28th Annual IEEE/ACM Symposium on Logic in Computer Science (LICS'13), IEEE, 2013, pp. 263–272.
- [39] A. Lomuscio, F. Raimondi, Model checking knowledge, strategies, and games in multi-agent systems, in: Proceedings of the 5th international joint conference on Autonomous agents and multiagent systems (AAMAS'06), ACM Press, 2006, pp. 161–168.
- [40] A. Morgenstern, Symbolic controller synthesis for LTL specifications, Ph.D. thesis (2010).
- [41] M. Huth, M. Ryan, Logic in Computer Science: Modelling and Reasoning about Systems, Cambridge University Press, New York, NY, USA, 2004.
- [42] J. R. Burch, E. M. Clarke, K. L. McMillan, D. L. Dill, L. J. Hwang, Symbolic model checking: 10^{20} states and beyond, Information and Computation 98 (2) (1992) 142–170.
- [43] MCMAS-SLK - A Model Checker for the Verification of Strategy Logic Specifications., <http://vas.doc.ic.ac.uk/software/tools/>.
- [44] F. Somenzi, CUDD: CU decision diagram package, <http://vlsi.colorado.edu/~fabio/CUDD/>.
- [45] A. Lomuscio, H. Qu, F. Raimondi, MCMAS: A model checker for the verification of multi-agent systems, in: Proceedings of the 21th International Conference on Computer Aided Verification (CAV'09), LNCS 5643, Springer, 2009, pp. 682–688.
- [46] D. Chaum, The dining cryptographers problem: Unconditional sender and recipient untraceability, Journal of Cryptology 1 (1) (1988) 65–75.
- [47] W. van der Hoek, A. Lomuscio, A logic for ignorance, Electronic Notes in Theoretical Computer Science 85 (2) (2004) 117–133.
- [48] R. van der Meyden, K. Su, Symbolic model checking the knowledge of the dining cryptographers, in: Proceedings of the 17th IEEE Computer Security Foundations Workshop (CSFW'04), IEEE Computer Society, Washington, DC, USA, 2004, pp. 280–291.

- [49] P. Čermák, A model checker for strategy logic, Master's thesis, Department of Computing, Imperial College London, UK (2014).
- [50] K. Chatterjee, T. A. Henzinger, N. Piterman, Strategy logic, *Inf. Comput.* 208 (6) (2010) 677–693.
- [51] A. D. C. Lopes, F. Laroussinie, N. Markey, ATL with strategy contexts: Expressiveness and model checking, in: *Proceedings of the 30th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'10)*, LIPIcs 8, 2010, pp. 120–132.
- [52] X. Huang, R. v. Meyden, Symbolic model checking epistemic strategy logic, in: *Proceedings of the 28th Conference on Artificial Intelligence (AAAI'14)*, AAAI, 2014, pp. 1426–1432.
- [53] X. Huang, R. v. Meyden, A temporal logic of strategic knowledge, in: *Proceedings of the 14th International Conference on Principles of Knowledge Representation and Reasoning (KR'14)*, AAAI, 2014, pp. 418–427.
- [54] P. Čermák, A. Lomuscio, A. Murano, Verifying and synthesising multi-agent systems against one-goal strategy logic specifications, in: *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI'15)*, AAAI Press, 2015, pp. 2038–2044.