Towards partial order reduction for model checking temporal epistemic logic

Alessio Lomuscio[1] and Wojciech Penczek[2] and Hongyang Qu[1]

[1] Department of Computing, Imperial College London, UK
{A.Lomuscio,Hongyang.Qu}@imperial.ac.uk
[2] Institute of Computer Science, PAS, and University of Podlasie, Poland
penczek@ipipan.waw.pl

Abstract. We introduce basic partial order reduction techniques in a temporal-epistemic setting. We analyse the semantics of interpreted systems with respect to the notions of trace-equivalence for the epistemic linear time logic LTLK$_{-X}$.

1 Introduction

In recent years there has been growing attention to the area of verification of multi-agent systems (MAS) by automatic model checking. Differently from standard reactive systems where plain temporal logics are often used, MAS are specified by using rich, intensional logics such as epistemic and deontic logics in combination with temporal logic. To accommodate for these needs several techniques for model checking have been suitably extended. For instance in [4, 20] OBDD-based techniques for temporal epistemic logic were introduced. Similar analysis were carried out previously for SAT-based approaches, including bounded and unbounded model checking [18, 10]. These approaches have now been implemented [4, 13, 1] and experimental results obtained in a variety of areas such as verification of security protocols, web-services, etc. Several extensions to other logics, including ATL, real-time, and others, have also been analysed.

It is surprising however that two mainstream techniques in symbolic verification, i.e., *predicate abstraction* and *partial order reduction* have not so far been applied to the verification of MAS logics. In this paper we begin the analysis of partial order reduction for temporal epistemic logic. Specifically, we look at the case of the linear temporal logic LTLK$_{-X}$ (i.e., the standard LTL [14] without the X next-time operator in which an epistemic modality is added [3]). The main contributions of this research note are the notions of weak and strong path equivalence defined on MAS semantics, the corresponding dependency relations, and a proof showing that these equivalences preserve the satisfaction of LTLK$_{-X}$ formulas.

The rest of the paper is organised as follows. In Section 2 we introduce syntax, semantics of our setting together with some basic notions. In Section 3 we present the definitions of path equivalence and dependency which are used in Theorem 1, the key result of the paper, showing that strongly equivalent paths preserve LTLK$_{-X}$ formulas. We exemplify the methodology in Section 4 while discussing an example, and present our conclusions in Section 5.

2 Preliminaries

We introduce here the basic technical background to the present paper. In particular we discuss the semantics of interpreted systems, properly augmented with suitable concepts for our needs, and the basic syntax we shall be using in the rest of the paper.

2.1 Interpreted Systems

The semantics of interpreted systems provides a setting to reason about MAS. Interpreted systems were originally developed independently by Parikh and Ramanujam [16], Halpern and Moses [8] and Rosenschein [21]. Their adoption as a semantics of choice for several MAS concept follows the publication of [3]. Although several valuable extensions have been proposed, in their basic settings interpreted systems offer a natural synchronous semantics for linear time and an external account of knowledge of the agents in the system. The following is a brief summary of the fundamental concepts needed for the rest of the paper; we refer to [3] for more details.

We begin by assuming a MAS to be composed of n agents $\mathcal{A} = \{1, \ldots, n\}$[3]. We associate a finite set of *possible local states* $L_i = \{l_i^1, l_i^2, \ldots, l_i^{nl_i}\}$ and *actions* $Act_i = \{a_i^1, a_i^2, \ldots, a_i^{na_i}\}$ to each agent $i \in \mathcal{A}$. In the interpreted systems model the actions of the agents are selected and performed synchronously according to each agent's *local protocol* $P_i : L_i \to 2^{Act_i}$; the local protocol effectively models the program the agent is executing. A *global state* $g = (l_1, \ldots, l_n)$ is a tuple of local states for all the agents in the MAS corresponding to an instantaneous snapshot of the system at a given time. Given a global state $g = (l_1, \ldots, l_n)$, we denote $g_i = l_i$ as the local component of agent $i \in \mathcal{A}$ in g. Global transitions are executed by means of joint actions on global states. In a nutshell, the *global evolution function* $t : G \times Act_1 \times \cdots \times Act_n \to G$ defines the target global state from a global state when a joint action $(a_1, \ldots, a_n) \in Act_1 \times \cdots \times Act_n$ is selected and performed by all agents in the system. More details can be found in [3].

In the following analysis we differ from the standard presentation by abstracting from the actual protocols and actions being performed and focus on the transitions only. For this reason we simply focus on the set of all possible *global transitions* $\mathcal{T} = \{(g, g') \mid \exists (a_1, \ldots, a_n) \in Act_1 \times \cdots \times Act_n \text{ such that } t(g, a_1, \ldots, a_n) = g'\}$. For simplicity we shall often use lower case letters $t_1, t_2, \ldots$ to denote elements of $\mathcal{T}$. Given the set $\mathcal{T}$ of global transitions we denote by $\mathcal{T}_i, i \in \mathcal{A}$, the set of all *local transitions* of the form $t_i = (l_i^k, l_i^{k+1})$ for an agent $i \in \mathcal{A}$. The set of all local transitions can be obtained by projecting $\mathcal{T}$ over the corresponding dimension for the agent in question; more formally $(l_i^k, l_i^{k+1}) \in \mathcal{T}_i$ if there exists a joint action $(a_1, \ldots, a_n)$ such that $t(g^k, a_1, \ldots, a_n) = g^{k+1}$, where the local component for agent i in g^k (respectively g^{k+1}) is l_i^k (respectively l_i^{k+1}). With slight abuse of notation for any global transition $t = (g, g') \in \mathcal{T}$ we write $t = (t_1, \ldots, t_n)$, where

[3] Note in the present study we do not consider the environment component. This may be added with no technical difficulty at the price of heavier notation.

each $t_i \in \mathcal{T}_i, i \in \mathcal{A}$ is such that $t_i(g_i, g'_i)$, and say that all $t_i, i = 1, \ldots, n$, are the local transitions in t.

With respect to the above we use the following notations. Given a local transition $t_i = (l_i, l'_i)$ we write $source(t_i) = l_i$ and $target(t_i) = l'_i$. Further, if $l_i = l'_i$, we denote t_i as ϵ. We use similar notation for global transitions too with obvious meaning in terms of source and target on global states. A sequence of global states $\rho = g^0 g^1 g^2 \ldots$ is called a path (or a run) if for every $g^k, g^{k+1} \in \rho$ ($k \geq 0$) we have that $(g^k, g^{k+1}) \in \mathcal{T}$. Given a path ρ we say $\rho|_i = g_i^0 g_i^1 g_i^2 \ldots$ is the local path for agent i in ρ. Given a path $\rho = g^0 g^1 g^2 \ldots$, $\rho(k) = g^k$, and $\rho\langle k \rangle = (g^k, g^{k+1}) = t^k$. Similarly, the k-th state and k-th transition in $\rho|_i$ are denoted as $\rho|_i(k)$ and $\rho|_i\langle k \rangle$ respectively. Let $\rho[0..k] = g^0 g^1 \ldots g^k$ (respectively $\rho|_i[0..k] = g_i^0 g_i^1 \ldots g_i^k$) be the prefix of ρ (respectively $\rho|_i$) and $\rho[k] = g^k g^{k+1} \ldots$ (respectively $\rho|_i[k] = g_i^k g_i^{k+1} \ldots$) the suffix. The set of paths originating from g is denoted as $\mathit{\Pi}(g)$.

We express synchronisation of transitions as follows. Local transitions are synchronised if they are always performed jointly by the system; this is formally expressed as follows.

Definition 1 (Synchronisation). *For any $i, j \in \mathcal{A}$ ($i \neq j$), a local transition t_i is said to be* semi-synchronised *to a local transition t_j if whenever t_i appears in a global transition $t = (t_1, \ldots, t_n)$ so does t_j. Two local transitions t_i, t_j are* synchronised *if t_i is semi-synchronised to t_j and t_j is semi-synchronised to t_i.*

We write $t_1 \rightarrow t_2$ to denote the fact that t_1 is semi-synchronised to t_2 and $t_1 \leftrightarrow t_2$ denote t_1 is synchronised to t_2. Figure 1 shows an interpreted system composed of three agents. The dotted lines represents synchronised transitions, i.e., the local transitions t_1^2 and t_2^2 are synchronised.

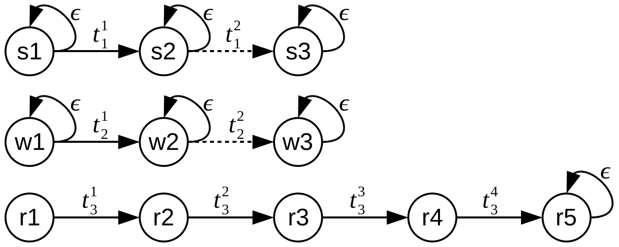

Fig. 1. A synchronous system.

Definition 2 (Interpreted Systems). *Given a set of atomic propositions P, an* interpreted system *(or simply a* model*) is a tuple $M = (G, G_0, \mathit{\Pi}, h)$, where G is a set of global states, $G_0 \subseteq G$ is a set of initial (global) states, $\mathit{\Pi} = \bigcup_{i \in G_0} \mathit{\Pi}(i)$ is the set of paths originating from all states in G_0, and $h : P \rightarrow 2^G$ is an interpretation for the atomic propositions. Particularly, we define a local atomic proposition p_i^j*

for each local state l_i^j of the agent $i \in \mathcal{A}$ such that $h(p_i^j) = \{g \mid g \in G$ and $g_i = l_i^j\}$. We assume G to be the set of states reachable from G_0 by any path in Π.

We can now define the syntax and interpretation of our language.

2.2 Syntax

Combinations of linear time and knowledge have long been used in the analysis of temporal epistemic properties of systems [3, 7]. In partial order reduction for LTL one typically excludes from the syntax the next time operator X as the preservation results [12] do not hold when X is present. Given this we consider LTLK$_{-X}$ in this paper.

Definition 3 (Syntax). *Let PV be set of atomic propositions to be interpreted over the global states of a system. The syntax of LTLK$_{-X}$ is defined by the following BNF grammar:*

$$\phi ::= true \mid false \mid p \mid \neg p \mid \phi \wedge \phi \mid \phi \vee \phi \mid \phi \mathcal{U} \phi \mid \phi \mathcal{R} \phi \mid K_i \phi \mid \overline{K}_i \phi,$$

where $p \in PV$.

The temporal operators $\mathcal{U}$ and $\mathcal{R}$ are named as usual *until* and *release* respectively. The formula $K_i\phi$ represents "agent i knows ϕ" and $\overline{K}_i\phi$ is the corresponding dual representing "agent i does not know whether or not ϕ holds". The epistemic modalities are defined by means of the following relations as standard.

Definition 4 (Epistemic relation). *For each agent $i \in \mathcal{A}$, $\sim_i \subseteq G \times G$ is an* epistemic indistinguishably *relation over global states defined by $g \sim_i g'$ if $g_i = g'_i$.*

Given a model $M = (G, G_0, \Pi, h)$, where $h(p)$ is the set of global states where p holds. Let $\overline{\Pi}$ denote the suffix-closure of Π, i.e., the set of all the paths in Π and their suffices. The formal semantics of an LTLK$_{-X}$ formula ϕ being satisfied by M and $\rho \in \overline{\Pi}$, denoted as $M, \rho \models \phi$, is recursively defined as follows.

Definition 5 (Satisfaction).

- $M, \rho \models true$ *for each* $\rho \in \overline{\Pi}$;
- $M, \rho \not\models false$ *for each* $\rho \in \overline{\Pi}$;
- $M, \rho \models p$ *iff* $\rho(0) \in h(p)$;
- $M, \rho \models \neg p$ *iff* $M, \rho \not\models p$;
- $M, \rho \models \phi_1 \wedge \phi_2$ *iff* $M, \rho \models \phi_1$ *and* $M, \rho \models \phi_2$;
- $M, \rho \models \phi_1 \vee \phi_2$ *iff* $M, \rho \models \phi_1$ *or* $M, \rho \models \phi_2$;
- $M, \rho \models \phi_1 \mathcal{U} \phi_2$ *iff* $M, \rho[k] \models \phi_2$ *for* some $k \geq 0$ *and* $M, \rho[j] \models \phi_1$ *for* all $0 \leq j < k$;
- $M, \rho \models \phi_1 \mathcal{R} \phi_2$ *iff either* $M, \rho[k] \models \phi_2$ *and* $M, \rho[k] \not\models \phi_1$ *for* all $k \geq 0$, *or* $M, \rho[k] \models \phi_1$ *for* some $k \geq 0$ *and* $M, \rho[j] \models \phi_2$ *for* all $0 \leq j \leq k$;
- $M, \rho \models K_i\phi$ *iff* all *paths* $\rho' \in \overline{\Pi}$ *we have that* $\rho'(0) \sim_i \rho(0)$ *implies* $M, \rho' \models \phi$.
- $M, \rho \models \overline{K}_i\phi$ *iff for* some *path* $\rho' \in \overline{\Pi}$ *we have that* $\rho'(0) \sim_i \rho(0)$ *and* $M, \rho' \models \phi$.

Given a global state g of M and an LTLK$_{-X}$ formula ϕ, we use the following notations:

- $M, g \models \phi$ iff $M, \rho \models \phi$ for all the paths $\rho \in \Pi(g)$.
- $M \models \phi$ iff $M, g \models \phi$ for all $g \in G_0$.
- $Props(\phi) \subseteq PV$ is the set of atomic propositions that appear in ϕ.

In order to define partial order reduction for LTLK$_{-X}$, we transform each formula $\neg p$ into a fresh atomic proposition q such that $h(q) = G \setminus h(p)$. Next, we present the main notions used for our reduction.

Definition 6 (Simple State Expression). *Let $I \subseteq \mathcal{A}$. A set $L_I \subseteq \bigcup_{i \in I} L_i$ is said to be* simple *if it contains exactly one element from each set L_i. Given a simple set L_I, a* simple state expression *$\mathcal{P}$ for an atomic proposition p is a Boolean formula of the form:*

$$\mathcal{P} = \bigwedge_{l_i^j \in L_I} p_i^j, \tag{1}$$

where p_i^j is the local atomic proposition corresponding to l_i^j and for all $g \in G$ and $i \in I$: $g_i \in L_I$ implies $g \in h(p)$.

In the above definition, each local atomic proposition in $\mathcal{P}$ denotes a local state which "forces" any global state in which it appears to satisfy p. Given any $I \subseteq \mathcal{A}$, let $[p]$ denote the set of all valid simple state expressions for p. Given an atomic proposition p, a set $I \subseteq \mathcal{A}$ and a simple state expression $\mathcal{P}$, we write $\overline{[\mathcal{P}]}$ for L_I and $\mathcal{A}|_{\mathcal{P}}$ for I.

Let $G|_{\mathcal{P}} \subseteq G$ be the set of global states in which $\mathcal{P}$ holds. Given two simple state expressions $\mathcal{P}_k, \mathcal{P}'_k \in [p]$, we write $\mathcal{P}_k \leq \mathcal{P}'_k$ iff $G|_{\mathcal{P}_k} \subseteq G|_{\mathcal{P}'_k}$ and $\mathcal{P}_k < \mathcal{P}'_k$ iff $\mathcal{P}_k \leq \mathcal{P}'_k$ and $\mathcal{P}_k \neq \mathcal{P}'_k$. Clearly, $([p], \leq)$ is a poset. Let $Max[p]$ be the set of the maximal elements in $[p]$. Note that the maximal elements intuitively correspond to the "smallest" simple state expressions.

Definition 7 (Full State Expression). *The* full state expression *E_p for an atomic proposition p is a Boolean formula of the form:*

$$E_p = \bigvee_{\mathcal{P} \in Max[p]} \mathcal{P}, \tag{2}$$

In other words, E_p encodes the set of global states where p holds, i.e., $h(p)$. In what follows we also use the following shortcuts: $\mathcal{A}|_p = \bigcup_{\mathcal{P} \in Max[p]} \mathcal{A}|_{\mathcal{P}}$ ($\mathcal{A}|_p$ denotes the set of agents appearing in the full state expression of p), and $\mathcal{A}|_\phi = \bigcup_{p \in Props(\phi)} \mathcal{A}|_p$.

3 Partial order reduction on interpreted systems

In the literature, partial order reduction has been studied intensively for asynchronous systems, e.g., [22, 6, 17, 9, 15, 5, 19, 11]. The technique permits the exploration of a portion of the state space when checking for satisfaction of a formula in a system. The basic idea consists in observing that two consecutive independent transitions in a path can sometimes be interchanged with no effect to the satisfaction of a formula. Because of this, the set of all the paths in a system can be partitioned into subsets, named *traces* [2]. In this section, we aim to define a dependency relation between transitions in order to be able to partition paths into traces. We begin with the notion of *stuttering* [12].

Definition 8. *The* stutter normal form *of a path ρ is a sequence $\#\rho$ such that each consecutive repetition of states in ρ is replaced by a single state. Two paths are said to be* equivalent up to stuttering *if they have the same stutter normal form.*

For example, two paths $g^1g^2g^2g^3g^3$ and $g^1g^2g^2g^2g^3$ are equivalent up to stuttering since their stutter normal form is $g^1g^2g^3$. The same definition applies to local paths $\rho|_i$.

Definition 9 (Weak equivalence). *Two paths ρ and ρ' are* weakly equivalent *iff $\rho|_i$ and $\rho|'_i$ are equivalent up to stuttering, for all agents $i \in \mathcal{A}$.*

Figure 2 and 3 display two weakly equivalent paths in the system of Figure 1 based on the above definition.

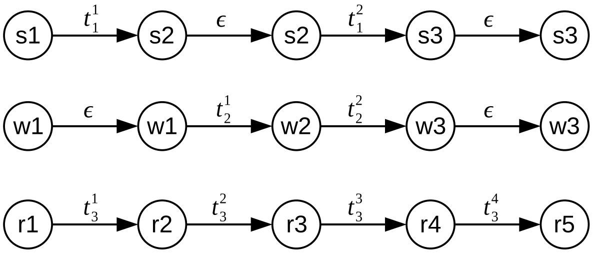

Fig. 2. A path ρ.

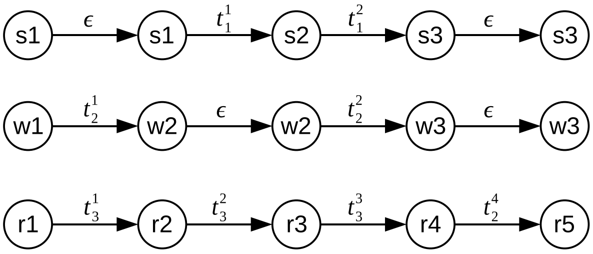

Fig. 3. A path weakly equivalent to ρ.

Observe that even if two paths are weakly equivalent, they may not satisfy the same LTLK$_{-X}$ formula. For example, consider the system in Figure 1 and two atomic propositions p and q such that p holds in all the global states containing $s1$ while q holds in all the global states containing $w2$. The formula

$$p\mathcal{U}q \tag{3}$$

holds in the path in Figure 3, but does not hold in the one in Figure 2.

Now we start to define dependency relations between transitions to strengthen weak equivalence in order to get strong equivalence preserving the LTLK$_{-X}$ formulae.

Definition 10 (Basic dependency relation). *For any agent $i \in \mathcal{A}$, the dependency relation D_i is the symmetric closure of the relation:*

$$\begin{aligned} d_i = \{(t_i, t'_i) \mid t_i, t'_i \in \mathcal{T}_i \text{ and (either } (t_i \neq \epsilon, t'_i \neq \epsilon) \text{ or} \\ (t_i \neq \epsilon \text{ and } \exists t_j \in \mathcal{T}_j, t_j \neq \epsilon, t'_i \to t_j \text{ or } t_j \to t'_i) \text{ or} \\ ((\exists t_j \in \mathcal{T}_j, t_j \neq \epsilon, t_i \to t_j \text{ or } t_j \to t_i) \text{ and } (\exists t_k \in \mathcal{T}_k, t_k \neq \epsilon, t'_i \to t_k \text{ or } t_k \to t'_i)))\}. \end{aligned}$$

The basic dependency relation relates two local transitions if either they cause an effective change of local states or they do not but they are (semi-)synchronised to other local transitions that do so.

Definition 11 (Dependency relation for synchronisation). *The dependency relation D_{syn} is the symmetric closure of the following relation:*

$$d_{syn} = \{(t_i, t_j) | t_i \in \mathcal{T}_i, t_j \in \mathcal{T}_j \text{ and } t_i \to t_j\}.$$

We now define the dependency relation for an LTLK$_{-X}$ formula. We begin with the dependency relation for an atomic proposition.

Definition 12 (Dependency relation for atomic propositions). *For an atomic proposition p with corresponding full state expression $E_p = \bigvee_{\mathcal{P} \in Max[p]} \mathcal{P}$, the dependency relation D_p for p is*

$$\begin{aligned} D_p = \{(t_i, t_j) | t_i \in \mathcal{T}_i, t_j \in \mathcal{T}_j, i \neq j, \mathcal{P} \in Max[p], \mathcal{P}' \in Max[p], \\ target(t_i) \in \overline{[\mathcal{P}]} \text{ and } t_i \neq \epsilon \text{ and } source(t_j) \in \overline{[\mathcal{P}']} \text{ and } t_j \neq \epsilon\}. \end{aligned}$$

D_p requires that each non-ϵ transition t_i entering a local state in $\overline{[\mathcal{P}]}$ is dependent on every non-ϵ transition t_j leaving a local state in any $\overline{[\mathcal{P}']}$. The reason for this is that p may become satisfied after t_i is executed and become unsatisfied after t_j is executed. For example, consider an atomic proposition p with full state expression $s2 \wedge r2$ (as shown in Figure 1). We have $D_p = \{(t_1^1, t_3^2), (t_3^1, t_1^2), (t_3^2, t_1^1), (t_1^2, t_3^1)\}$.

To define the dependency relation for an arbitrary LTLK$_{-X}$ formula ϕ, we need to preform some pre-processing on ϕ. Firstly, we need to make sure that each atomic proposition p occurs only once in ϕ. If there is more than one occurrence for p, we generate a fresh atomic proposition p' for each occurrence and define

$h(p') = h(p)$. It follows that $E_{p'} = E_p$. For example, we transform $\phi = K_i p \vee K_j p$ into $K_i p_1 \vee K_j p_2$ with $h(p_1) = h(p_2) = h(p)$. Secondly, we define the *epistemic nesting depth* $\{\psi\}_K$ for every sub-formula ψ of ϕ. The epistemic nesting of a sub-formula corresponds to the "epistemic depth" of a sub-formula in a formula. Intuitively, the "deeper" a sub-formula is in an epistemic formula the higher its nesting will be. To calculate the nesting we assign a level 0 of nesting to the whole formula and increase it by 1 every time we find an epistemic operator while exploring the parse tree of the formula. More formally, we proceed as follows.

Definition 13 (Epistemic nesting depth). *Given a formula ϕ, the epistemic nesting $\{\psi\}_K$ of a sub-formula ψ of ϕ is defined as follows.*

- *If $\psi = \phi$, then $\{\phi\}_K = \{\psi\}_K = 0$;*
- *If $\psi \in \{\psi_1 \wedge \psi_2, \psi_1 \vee \psi_2, \psi_1 \mathcal{U} \psi_2, \psi_1 \mathcal{R} \psi_2\}$, then $\{\psi_1\}_K = \{\psi_2\}_K = \{\psi\}_K$;*
- *If $\psi \in \{K_i \psi_1, \overline{K}_i \psi_1\}$, then $\{\psi_1\}_K = \{\psi\}_K + 1$;*
- *If $\psi = p$, then $\{p\}_K = \{\psi\}_K$.*

Let $|\phi|_K = \max\{\{p\}_K \mid p \in \mathit{Props}(\phi)\}$ be the maximum epistemic nesting depth of ϕ. Let AP_ϕ^m be the subset of $\mathit{Props}(\phi)$ such that for each $p \in AP_\phi^m$, $\{p\}_K = m$, and $AP_\phi = \bigcup_{0 \leq m \leq |\phi|_K} AP_\phi^m$. Assume $i_1, i_2, \ldots, i_m$ is the sequence of indexes for the epistemic modalities scoping p (e.g., for $\phi = K_1 q \wedge K_2(EF(K_1 p))$, the sequence of indexes for p is (2, 1)). Then we perform the following two steps on AP_ϕ:

1. For each $p \in AP_\phi^m$ for all $m > 0$, we generate the set of propositions

   $$\Sigma_p = \{p_{j_1, j_2, \ldots, j_m} \mid l_{i_1}^{j_1} \in L_{i_1}, \ldots, l_{i_m}^{j_m} \in L_{i_m} \text{and } E_{p_{j_1, j_2, \ldots, j_m}} = p_{i_1}^{j_1} \wedge \cdots \wedge p_{i_m}^{j_m} \wedge E_p\},$$

   where $p_{i_k}^{j_k}$ is the local atomic proposition for $l_{i_k}^{j_k}$. For example, consider $\phi = EF(K_2 p)$ with $E_p = s_2 \wedge r_2$ in the system of Figure 1. Since $\{p\}_K = 1$, we generate the propositions p_1, p_2, p_3 with $E_{p_1} = w_1 \wedge s_2 \wedge r_2$, $E_{p_2} = w_2 \wedge s_2 \wedge r_2$ and $E_{p_3} = w_3 \wedge s_2 \wedge r_2$. Let $AP'_\phi = \bigcup_{0 < m \leq |\phi|_K} (\bigcup_{p \in AP_\phi^m} \Sigma_p)$ be the set of the newly generated atomic propositions.
2. For each pair of atomic propositions p and q in $AP_\phi^0 \cup AP'_\phi$, we define a fresh atomic proposition r with $h(r) = h(p) \cup h(q)$. Let AP_ϕ^r be the set of atomic propositions generated in this step.

Definition 14 (Dependency relation for an LTLK$_{-X}$ formula ϕ). *The dependency relation D_ϕ for ϕ is defined as follows:*

$$D_\phi = \bigcup_{p \in AP_\phi^0 \cup AP'_\phi \cup AP_\phi^r} D_p.$$

Consider the example $\phi = EF(K_2 p)$ with $E_p = s_2 \wedge r_2$ again. D_ϕ is the symmetric closure of the following set: $\{(t_2^1, t_1^1), (t_2^1, t_3^1), (t_1^1, t_3^2), (t_3^1, t_1^2), (t_2^2, t_1^1), (t_2^2, t_3^1), (t_2^1, t_1^2), (t_2^1, t_3^2), (t_2^2, t_1^2), (t_2^2, t_3^2)\}$. The above dependency relation is used to avoid inconsistencies among weakly equivalent paths where a formula holds in one path but does not hold in the other. For example, the paths in Figure 2 and Figure 3 can be distinguished now with respect to Formula (3). Since $D_{p\mathcal{U}q} = \{(t_1^1, t_2^1), (t_2^1, t_1^1)\}$, t_1^1 and t_2^1 are not interchangeable and the execution order between them has an impact on the satisfaction of the formula.

Definition 15 (Extended Formula). *For any LTLK$_{-X}$ formula ϕ, an* extended formula *ϕ' for ϕ is defined by replacing each subformula $\psi = K_i\varphi$ with*

$$\psi' = K_i((p_i^1 \wedge \varphi) \vee \ldots \vee (p_i^{nl_i} \wedge \varphi)),$$

where p_i^j is the local atomic proposition corresponding to l_i^j $(1 \leq j \leq nl_i)$. The substitution is carried out bottom-up in the parse tree.

Note that obviously $D_\phi = D_{\phi'}$. So in what follows we assume to be dealing with extended formulae only.

Given an LTLK$_{-X}$ formula ϕ, let

$$D = (\bigcup_{i \in \mathcal{A}} D_i) \cup D_{syn} \cup D_\phi. \tag{4}$$

For a path ρ containing two specific occurrences t_i and t_j $(i, j \in \mathcal{A})$ of local transitions, we write $t_i <_\rho t_j$ if t_i happens earlier than t_j in ρ. We write $t_i =_\rho t_j$ if they are executed together in a global transition. We use $t_i \leq_\rho t_j$ to denote either $t_i <_\rho t_j$ or $t_i =_\rho t_j$.

Now we are ready to present the main result of this note. To this aim we first define strong equivalence, and then show that it preserves the LTLK$_{-X}$ formulae.

Definition 16 (Strong equivalence). *Two paths ρ and ρ' are* strongly equivalent *with respect to an LTLK$_{-X}$ formula ϕ iff the following two conditions hold:*

(1) ρ and ρ' are weakly equivalent,
(2) for any two occurrences t and t' of local transitions in ρ and $(t, t') \in D$, $t <_\rho t'$ implies $t <_{\rho'} t'$, and $t =_\rho t'$ implies $t =_{\rho'} t'$.

Given the above equivalence, we formulate two auxiliary lemmas.

Lemma 1. *The following two conditions hold:*

A) For a path ρ and an LTLK$_{-X}$ *formula ϕ, if $M, \rho \models \phi$ and $M, \rho[1] \not\models \phi$, then there exists $p \in Props(\phi)$ such that $M, \rho \models p$ and $M, \rho[1] \not\models p$,*
B) if $M, \rho \not\models \phi$ and $M, \rho[1] \models \phi$, we can find an atomic proposition $p \in Props(\phi)$ such that $M, \rho \not\models p$ and $M, \rho[1] \models p$.

Proof. We prove A) by induction on the structure of ϕ. The condition B) can be shown similarly.

1. $\phi = p$. This case is obvious.
2. $\phi = \psi_1 \wedge \psi_2$. We have $M, \rho \models \psi_1 \wedge \psi_2$ and $M, \rho[1] \not\models \psi_1 \wedge \psi_2$. If $M, \rho[1] \not\models \psi_1$, given that $M, \rho \models \psi_1$, it follows that there exists an atomic proposition p in ψ_1 such that $M, \rho \models p$ and $M, \rho[1] \not\models p$.

3. $\phi = \psi_1 \vee \psi_2$. This case is similar to the previous one.
4. $\phi = \psi_1 \mathcal{U} \psi_2$. We have $M, \rho \models \psi_1 \mathcal{U} \psi_2$ and $M, \rho[1] \not\models \psi_1 \mathcal{U} \psi_2$. So $M, \rho \models \psi_2$ and $M, \rho[1] \not\models \psi_2$. Therefore, by induction the case holds.
5. $\phi = \psi_1 \mathcal{R} \psi_2$. We have $M, \rho \models \psi_1 \mathcal{R} \psi_2$ and $M, \rho[1] \not\models \psi_1 \mathcal{R} \psi_2$. If ψ_2 holds in all states in ρ and ψ_1 does not holds in any states, then $M, \rho[1] \models \phi$. Thus there exists k such that ψ_1 holds in $\rho(k)$ and ψ_2 holds in $\rho(j)$ for all $0 \leq j \leq k$. Similarly to the $\mathcal{U}$ case, $k = 0$, and ψ_1 or ψ_2 does not hold in $\rho(1)$. Then there exists p in ψ_1 or ψ_2 satisfying the lemma.
6. $\phi = K_i \psi$. We have $M, \rho \models K_i \psi$ and $M, \rho[1] \not\models K_i \psi$. So $\rho|_i(0) \neq \rho|_i(1)$. Since ϕ is an extended formula, we know that $M, \rho \models K_i((p_i^1 \wedge \psi) \vee ... \vee (p_i^{nl_i} \wedge \psi))$, and there exists a $1 \leq j \leq nl_1$ such that p_i^j is the local atomic proposition corresponding to $\rho|_i(0)$. We have $M, \rho \models p_i^j$ and $M, \rho[1] \not\models p_i^j$.
7. $\phi = \overline{K}_i \psi$. This case is similar to the one above. □

Lemma 2. *Let ϕ be an* LTLK$_{-X}$ *formula and paths $\rho, \rho' \in \Pi$ be strongly equivalent. Then there exist $k, k' \geq 0$ such that the following two conditions hold:*

A) If $M, \rho[k] \models \phi$, then $M, \rho'[k'] \models \phi$;
B) There exists an $i \in \mathcal{A}|_\phi$ such that the paths $\rho|_i[0..k]$ and $\rho'|_i[0..k']$ are equivalent up to stuttering, and if $M, \rho[k-1] \not\models \phi$ and $M, \rho[k] \models \phi$, then $\rho|_i\langle k \rangle \neq \epsilon$.

Proof. A) By induction on the structure of ϕ.
The base case: $\phi = p$.
Assume $M, \rho[k] \models p$ for some $k \geq 0$. Given that $\rho(k) \in h(p)$, we have that there exists a simple state expression $\mathcal{P} \in Max[p]$ for some simple set L_I, $I \subseteq A$ and $\rho(k) \in G|_{\mathcal{P}}$. For any $i \in I$, consider the shortest and longest prefixes of the projections of ρ' onto i that are equivalent to $\rho|_i[0..k]$ up to stuttering. Call $\rho'|_i[0..j_i]$ the shortest and $\rho'|_i[0..\overline{j}_i]$ the longest. Given ρ and ρ' are strongly equivalent, they are weakly equivalent and therefore, we have $\rho'|_i(j_i) = \rho'|_i(\overline{j}_i) = \rho|_i(k)$. Consider the following two cases, which may arise.

1. $\bigcap\limits_{i \in I} [j_i, \overline{j}_i] \neq \emptyset$. Then, there is a $k' \geq 0$ such that $k' \in \bigcap\limits_{i \in I} [j_i, \overline{j}_i]$. Given that $\rho'|_i(k') = \rho|_i(k)$ for all $i \in I$, we have that $M, \rho'[k'] \models p$.
2. $\bigcap\limits_{i \in I} [j_i, \overline{j}_i] = \emptyset$. Then, there must exist $x, y \in I$ such that $j_x > \overline{j}_y$. This implies that the transitions $t_x^{j_x - 1}$ and $t_y^{\overline{j}}$ are dependent. However, by the inductive hypothesis ρ, ρ' are strongly equivalent and therefore we have $t_y^{\overline{j}} \leq_{\rho'} t_x^{j_x - 1}$. This is a contradiction. So, we have $\bigcap\limits_{i \in I} [j_i, \overline{j}_i] \neq \emptyset$.

The induction steps.

1. $\phi = \psi_1 \wedge \psi_2$. Assume $M, \rho[k] \models \psi_1 \wedge \psi_2$, therefore $M, \rho[k] \models \psi_1$ and $M, \rho[k] \models \psi_2$. By the inductive assumption there exist $k', k'' \geq 0$ such that $M, \rho'[k'] \models \psi_1$ and $M, \rho'[k''] \models \psi_2$. If $k' = k''$, then $M, \rho'[k'] \models \psi_1 \wedge \psi_2$. So, we are done. Without loss of generality, assume now that $k' < k''$. Let $\bar{k}' \geq k'$ be the biggest natural number such that $M, \rho'[j] \models \psi_1$ for $k' \leq j \leq \bar{k}'$ and $M, \rho'[\bar{k}' + 1] \not\models \psi_1$.

Similarly let $\bar{k}''$ be the smallest natural number such that $M, \rho'[j] \models \psi_2$ for $\bar{k}'' \leq j \leq k''$ and $M, \rho'[\bar{k}''-1] \not\models \psi_2$. If $\bar{k}'' \leq \bar{k}'$ then there exists a k''' such that $M, \rho'[k'''] \models \psi_1 \wedge \psi_2$.
Otherwise, we have $\bar{k}' < \bar{k}''$. By Lemma 1, there exists an atomic proposition p in ψ_1 such that $M, \rho'[\bar{k}'] \models p$ and $M, \rho'[\bar{k}'+1] \not\models p$, and an atomic proposition q in ψ_2 such that $M, \rho'[\bar{k}''-1] \not\models q$ and $M, \rho'[\bar{k}''] \models q$. Assume p is satisfied by the simple state expression $\mathcal{P}_1 \in Max[p]$ for some $I \subseteq \mathcal{A}$ and q by $\mathcal{P}_2 \in Max[q]$ for some $I' \subseteq \mathcal{A}$. Therefore, there exist an agent $i \in I$ such that $t_i^{\bar{k}'} = \rho'|_i\langle \bar{k}' \rangle$ ($t_i^{\bar{k}'} \neq \epsilon$) leaves the local state $\rho'|_i(\bar{k}') \in \overline{[\mathcal{P}_1]}$, and an agent $j \in I'$ such that $t_j^{\bar{k}''-1} = \rho'|_j\langle \bar{k}''-1 \rangle$ ($t_j^{\bar{k}''-1} \neq \epsilon$) enters the local state $\rho'|_j(\bar{k}'') \in \overline{[\mathcal{P}_2]}$ (note that $i \neq j$, otherwise we would have $k' = k''$.). According to the construction of D_ϕ, $t_i^{\bar{k}'}$ and $t_j^{\bar{k}''-1}$ are dependent. So we have $t_i^{\bar{k}'} \leq_{\rho'} t_j^{\bar{k}''-1}$ and $t_j^{\bar{k}''-1} <_\rho t_i^{\bar{k}'}$. But ρ and ρ' are strongly equivalent by the inductive hypothesis, so we get a contradiction.

2. $\phi = \psi_1 \vee \psi_2$. This case is immediate.
3. $\phi = \psi_1 \mathcal{U} \psi_2$. Assume $M, \rho[k] \models \psi_1 \mathcal{U} \psi_2$. By definition we have that there exists a $k' \geq k$ such that $M, \rho[k'] \models \psi_2$ and $M, \rho[j] \models \psi_1$ for $k \leq j \leq k'$. Then by induction, we have that there exists a k'' such that $M, \rho'[k''] \models \psi_2$. So we have $M, \rho'[k''] \models \psi_1 \mathcal{U} \psi_2$.
4. $\phi = \psi_1 \mathcal{R} \psi_2$. According to the semantics of $\mathcal{R}$, we know that $M, \rho[k] \models \psi_2$ and thus there exists k' such that $M, \rho'[k'] \models \psi_2$. If for all $j > k$, $M, \rho[j] \not\models \psi_1$, then for all $j' > k'$, $M, \rho'[j'] \not\models \psi_1$ (otherwise, there exists $\bar{j} > k$ such that $M, \rho[\bar{j}] \models \psi_1$). If there exists j ($j \geq k$), $M, \rho[j] \models \psi_1$, then $\psi_1 R \psi_2 = \psi_2 U (\psi_1 \wedge \psi_2)$ and the case may be shown similarly to the above.
5. $\phi = K_i \psi$. Assume $M, \rho[k] \models K_i \psi$. Since ρ, ρ' are strongly equivalent, $\rho|_i[0..k]$ and $\rho'|_i[0..k']$ are equivalent up to stuttering for some k'. So $\rho|_i(k) = \rho'|_i(k')$. Therefore $M, \rho'[k'] \models \phi$.
6. $\phi = \overline{K}_i \psi$. It is the same as the K_i case.

B) A proof of this condition follows from the above proof. □

Strong equivalence for an LTLK$_{-X}$ formula ϕ naturally partitions Π into traces of strongly equivalent paths. We have the following theorem.

Theorem 1. *For any LTLK$_{-X}$ ϕ and any two strongly equivalent paths $\rho, \rho' \in \Pi$, we have $M, \rho \models \phi$ iff $M, \rho' \models \phi$.*

Proof. By induction on the structure of ϕ.
The base case $\phi = p$ is obvious given $\rho(0) = \rho'(0)$.
The induction steps $\phi = \psi_1 \wedge \psi_2$, $\phi = \psi_1 \vee \psi_2$, $\phi = K_i \psi$ and $\phi = \overline{K}_i \psi$ can be obtained similarly. In the following, we prove the case $\phi = \psi_1 \mathcal{U} \psi_2$. A similar proof can be obtained for $\phi = \psi_1 \mathcal{R} \psi_2$.
$\phi = \psi_1 \mathcal{U} \psi_2$. Assume $M, \rho \models \psi_1 \mathcal{U} \psi_2$. If $M, \rho \models \psi_2$, then $M, \rho' \models \psi_2$ and therefore $M, \rho' \models \phi$. Assume there exists a $k \geq 0$ such that $M, \rho[k] \models \psi_2$ and $M, \rho[j] \models \psi_1$ for $0 \leq j < k$. By Lemma 2, there exists a smallest $k' > 0$ such that $M, \rho'[k'] \models \psi_2$; we need to show that $M, \rho'[j] \models \psi_1$ for all $0 \leq j < k'$. Assume that $M, \rho'[j] \not\models \psi_1$ for the smallest $0 \leq j < k'$. Note that $M, \rho'[0] \models \psi_1$; so this implies that $M, \rho'[j-$

$1] \models \psi_1$. So there must exist a set of agents $I \subseteq \mathcal{A}$ such that $\rho'|_i(j-1) \neq \rho'|_i(j)$ for all $i \in I$. Similarly observe there exists a set of agents $I' \subseteq \mathcal{A}$ such that $\rho'|_i(k') \neq \rho'|_i(k'-1)$ for all $i \in I'$. So by observing there are atomic propositions changing values from $\rho'(k'-1)$ to $\rho'(k')$ and from $\rho'(j-1)$ to $\rho'(j)$, and reasoning similarly to the case of conjunction in the proof of Lemma 2, we can reach a contradiction with hypothesis of ρ, ρ' being strongly equivalent. □

Theorem 1 implies that partial order reduction based on the relation of strong equivalence preserves $LTLK_{-X}$ properties.

4 Example

We exemplify the technique above on the system of three agents $\mathcal{A} = \{1, 2, 3\}$ of Figure 1 with respect to the formula

$$\phi = \Diamond K_3\ p.$$

We assume p is an atomic proposition that holds in the global state $(s2, w2, r5)$, i.e., its full state expression is

$$E_p = s2 \wedge w2 \wedge r5.$$

Before we start to explore the state space, we need to generate the dependency relation according to the formula 4.

- The basic dependency relation is defined as follows.

  $$\begin{aligned} D_1 &= \{(t_1^1, t_1^1), (t_1^2, t_1^2), (t_1^1, t_1^2), (t_1^2, t_1^1)\} \\ D_2 &= \{(t_2^1, t_2^1), (t_2^2, t_2^2), (t_2^1, t_2^2), (t_2^2, t_2^1)\} \\ D_3 &= \{(t_3^1, t_3^1), (t_3^2, t_3^2), (t_3^3, t_3^3), (t_3^4, t_3^4), (t_3^1, t_3^2), (t_3^1, t_3^3), (t_3^1, t_3^4), (t_3^2, t_3^3), (t_3^2, t_3^4), \\ &\qquad (t_3^3, t_3^4), (t_3^2, t_3^1), (t_3^3, t_3^1), (t_3^4, t_3^1), (t_3^3, t_3^2), (t_3^4, t_3^2), (t_3^4, t_3^3)\} \end{aligned}$$

- The dependency relation for synchronisation is as follows.

  $$D_{syn} = \{(t_1^2, t_2^2), (t_2^2, t_1^2)\}$$

- The dependency relation for atomic propositions is as follows.

  $$D_p = \{(t_1^2, t_3^4), (t_2^2, t_3^4), (t_3^4, t_1^2), (t_3^4, t_2^2)\}$$

- The dependency relation for the formula is defined as follows. For $K_3\ p$, we construct a new atomic proposition p' such that

  $$\begin{aligned} E_{p'} = &\bigvee_{l_3 \in L_3} (l_3 \wedge s2 \wedge w2 \wedge r5) \\ &= (r1 \wedge s2 \wedge w2 \wedge r5) \vee (r2 \wedge s2 \wedge w2 \wedge r5) \vee (r3 \wedge s2 \wedge w2 \wedge r5) \vee \\ &\qquad (r4 \wedge s2 \wedge w2 \wedge r5) \vee (r5 \wedge s2 \wedge w2 \wedge r5) \vee \\ &= s2 \wedge w2 \wedge r5 \\ &= E_p. \end{aligned}$$

  Therefore, we have $D_\phi = D_p$.

By means of the technique discussed, to check the validity of the formula above we do not need to explore the full state space shown in Figure 4. Since p does not hold in the state $(s1, w2, r5)$ (nor in $(s1, w1, r5)$, $(s2, w1, r5)$, $(s3, w3, r5)$) and $(s1, w2, r5) \sim_3 (s2, w2, r5)$, $K_3\, p$ does not hold in the model.

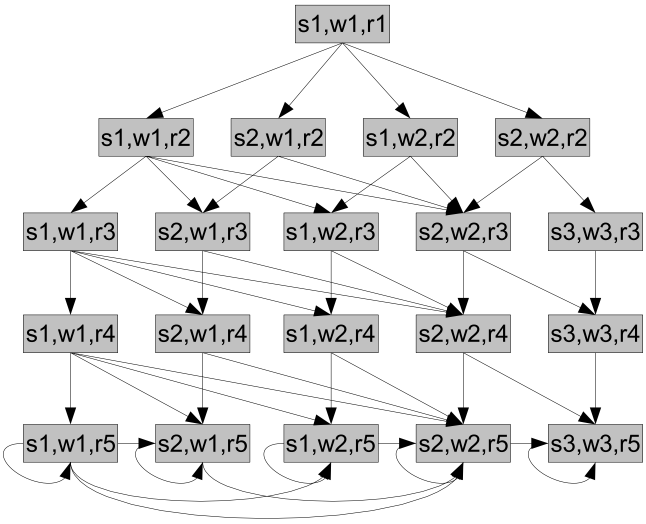

Fig. 4. The full state space.

After applying partial order reduction, we are able to check that $K_3\, p$ does not hold. Figure 5 illustrates the reduced state space, clearly showing the potential of this technique.

It is easy to see that any path in Figure 4 has a strongly equivalent path in Figure 5. For example, the path

$$(s1, w1, r1)(s1, w1, r2)(s2, w2, r3)(s2, w2, r4)(s2, w2, r5)(s3, w3, r5)$$

is equivalent to

$$(s1, w1, r1)(s1, w1, r2)(s1, w1, r3)(s1, w1, r4)(s1, w1, r5)(s2, w2, r5)(s3, w3, r5).$$

We can use similar considerations to check any LTLK$_{-X}$ formulae effectively.

5 Conclusions

In this research note we have extended a partial order reduction technique to a basic logic for knowledge and linear time. Our main result concerns the preservation of satisfaction of LTLK$_{-X}$ formulae on equivalent paths on synchronous interpreted systems semantics.

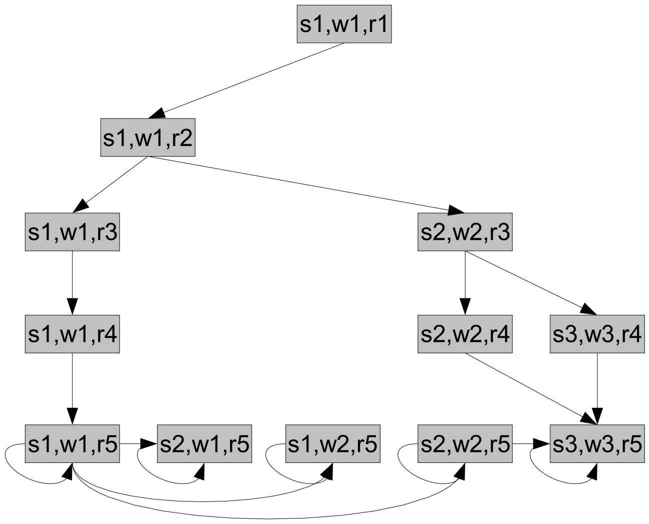

Fig. 5. The reduced state space.

The dependency relation we defined is quite general, as we do not impose any restrictions on the underlying models. While this makes it easier to design an algorithm ans test its effectiveness, we believe we can further enhance its effectiveness by exploring particular properties in the temporal epistemic logic.

We are currently investigating the feasibility of an algorithm to verify satisfiability on reduced traces and plan to test its implementation against known results for temporal epistemic specification available in the multi-agent systems literature.

Acknowledgements.

The research described in this paper is partly supported by the European Commission Framework 6 funded project CONTRACT (IST Project Number 034418) and by the Polish Ministry of Science and Higher Education under grant 3T11C01128.

References

1. P. Dembiński, A. Janowska, P. Janowski, W. Penczek, A. Pólrola, M. Szreter, B. Woźna, and A. Zbrzezny. VerICS: A tool for verifying Timed Automata and Estelle specifications. In *Proc. of the 9th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'03)*, volume 2619 of *LNCS*, pages 278–283. Springer-Verlag, 2003.
2. V. Diekert and G. Rozemberg, editors. *The Book of Traces.* World Scientific Publishing Co. Pte. Ltd., 1995.
3. R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge.* MIT Press, Cambridge, 1995.

4. P. Gammie and R. van der Meyden. MCK: Model checking the logic of knowledge. In *Proceedings of 16th International Conference on Computer Aided Verification (CAV'04)*, volume 3114 of *LNCS*, pages 479–483. Springer-Verlag, 2004.
5. R. Gerth, R. Kuiper, D. Peled, and W. Penczek. A partial order approach to branching time logic model checking. *Information and Computation*, 150:132–152, 1999.
6. P. Godefroid. Using partial orders to improve automatic verification methods. In E. M. Clarke and R. P. Kurshan, editors, *Proceedings of the 2nd International Conference on Computer Aided Verification (CAV'90)*, volume 3 of *ACM/AMS DIMACS Series*, pages 321–340, 1991.
7. J. Halpern, R. van der Meyden, and M. Y. Vardi. Complete axiomatisations for reasoning about knowledge and time. *SIAM Journal on Computing*, 33(3):674–703, 2003.
8. J. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587, 1990. A preliminary version appeared in *Proc. 3rd ACM Symposium on Principles of Distributed Computing*, 1984.
9. G. Holzmann and D. Peled. Partial order reduction of the state space. In *First SPIN Workshop*, Montréal, Quebec, 1995.
10. M. Kacprzak, A. Lomuscio, and W. Penczek. From bounded to unbounded model checking for temporal epistemic logic. *Fundamenta Informaticae*, 63(2,3):221–240, 2004.
11. M. E. Kurbán, P. Niebert, H. Qu, and W. Vogler. Stronger reduction criteria for local first search. In *ICTAC*, LNCS 4281, pages 108–122. Springer, 2006.
12. L. Lamport. What good is temporal logic? In *IFIP Congress*, pages 657–668, 1983.
13. A. Lomuscio and F. Raimondi. MCMAS: A model checker for multi-agent systems. In *Proceedings of TACAS 2006*, volume 3920, pages 450–454. Springer Verlag, 2006.
14. Z. Manna and A. Pnueli. *The temporal logic of reactive and concurrent systems*, volume 1. Springer-Verlag, Berlin/New York, 1992.
15. K. L. McMillan. A technique of a state space search based on unfolding. *Formal Methods in System Design*, 6(1):45–65, 1995.
16. R. Parikh and R. Ramanujam. Distributed processes and the logic of knowledge. In *Logic of Programs*, volume 193 of *Lecture Notes in Computer Science*, pages 256–268. Springer, 1985.
17. D. Peled. All from one, one for all: On model checking using representatives. In *Proceedings of the 5th International Conference on Computer Aided Verification (CAV'93)*, volume 697 of *LNCS*, pages 409–423. Springer-Verlag, 1993.
18. W. Penczek and A. Lomuscio. Verifying epistemic properties of multi-agent systems via bounded model checking. *Fundamenta Informaticae*, 55(2):167–185, 2003.
19. W. Penczek, M. Szreter, R. Gerth, and R. Kuiper. Improving partial order reductions for universal branching time properties. *Fundamenta Informaticae*, 43:245–267, 2000.
20. F. Raimondi and A. Lomuscio. Automatic verification of multi-agent systems by model checking via OBDDs. *Journal of Applied Logic*, 2007. To appear in Special issue on Logic-based agent verification.
21. S.J. Rosenschein. Formal theories of ai in knowledge and robotics. *New Generation Computing*, 3:345–357, 1985.
22. A. Valmari. A stubborn attack on state explosion. In *Proceedings of the 2nd International Conference on Computer Aided Verification (CAV'90)*, volume 531 of *LNCS*, pages 156–165. Springer-Verlag, 1990.