

Changing the Rules of the Game: Reasoning about Dynamic Phenomena in Multi-Agent Systems

AAAI Track

Rustam Galimullin
University of Bergen
Bergen, Norway
rustam.galimullin@uib.no

Munyuque Mittelmann
University of Naples Federico II
Naples, Italy
munyuque.mittelmann@unina.it

Maksim Gladyshev
Utrecht University
Utrecht, The Netherlands
m.gladyshev@uu.nl

Nima Motamed
Utrecht University
Utrecht, The Netherlands
n.motamed@uu.nl

ABSTRACT

The design and application of multi-agent systems (MAS) require reasoning about the effects of modifications on their underlying structure. In particular, such changes may impact the satisfaction of system specifications and the strategic abilities of their autonomous components. In this paper, we are concerned with the problem of verifying and synthesising modifications (or *updates*) of MAS. We propose an extension of the Alternating-Time Temporal Logic (ATL) that enables reasoning about the dynamics of model change, called the *Logic for ATL Model Building* (LAMB). We show how LAMB can express various intuitions and ideas about the dynamics of MAS, from normative updates to mechanism design. As the main technical result, we prove that, while being strictly more expressive than ATL, LAMB enjoys a P-complete model-checking procedure.

KEYWORDS

Strategy Logics; Model Change; Formal Verification

ACM Reference Format:

Rustam Galimullin, Maksim Gladyshev, Munyuque Mittelmann, and Nima Motamed. 2025. Changing the Rules of the Game: Reasoning about Dynamic Phenomena in Multi-Agent Systems: AAAI Track. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025)*, Detroit, Michigan, USA, May 19 – 23, 2025, IFAAMAS, 10 pages.

1 INTRODUCTION

Mechanism Design is a subfield of game theory concerned with the design of mathematical structures (i.e. *mechanisms*) describing the interaction of strategic agents that achieve desirable economic properties under the assumption of rational behavior [67]. Although it originated in economics, mechanism design provides an important foundation for the creation and analysis of multi-agent systems (MAS) [27, 70]. In numerous situations, the protocols and institutions describing interactions have already been designed and

implemented. When those do not comply with the designer’s objective (i.e. the economic properties) their complete redesign may not be feasible. For instance, a university with a selection procedure seen as *unfair*, would avoid implementing an entirely new procedure, but could be willing to adjust the existing one.

Although logic-based approaches have been widely used for the verification [25] and synthesis [28] of MAS, most research focuses on static or parametrised models and does not capture the dynamics of model change. One of the classic static approaches in the field is the Alternating-time Temporal Logic (ATL) [11]. ATL expresses conditions on the strategic abilities of agents interacting in a MAS, represented by a concurrent game model (CGM). In this paper, we are concerned with the problem of reasoning about the effects of modifications on CGMs. To tackle the problem, we extend ATL in two directions that have not been considered in the literature. First, we augment ATL with nominals and hybrid logic operators [13]. The resulting logic is called Hybrid ATL (HATL). HATL allows us to verify properties at states named by a given nominal, which cannot be captured by ATL. Second, we propose the *Logic for ATL Model Building* (LAMB), which enhances HATL with update operators that describe explicit modular modifications in the model.

We define three fundamental update operators in LAMB. First, we can *change the valuation* of some propositional variable in a particular state to the valuation of a given formula. Second, we can *switch the transition* from one state to another, which corresponds to modifying agents’ abilities in a given state. Finally, we can *add* a new state to the model and assign a fresh nominal to it. More complex operations, like adding a state, assigning it some propositional variable and adding incoming/outgoing transitions to it can be described in our language as sequences of primitive updates.

Our intuitions on model updates are guided by research in Dynamic Epistemic Logic (DEL) [77], where one can reason about the effects of epistemic events on agents’ knowledge. While epistemic and strategic reasoning are quite different domains, various DELs have been considered in the strategic setting (see, e.g., [6, 29, 62]). At the same time, research on updating models for strategic reasoning has been relatively sporadic and predominantly within the area of normative reasoning [9, 10].

We deem our contribution to be two-fold. On the *conceptual level*, we propose the exploration of general logic-based approaches to



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

dynamic phenomena in MAS, not confined to particular implementation areas. On the *technical level*, we propose a new formalism, LAMB, to reason about modifications on MAS that enables the verification of the strategic behavior of agents acting in a changing environment, and the synthesis of modifications on CGMs. We show how LAMB can express various intuitions and ideas about the dynamics of MAS from normative updates to mechanism design. We also formally study the logic (and some of its fragments) from the point of view of expressivity and model-checking complexity. Our results show that LAMB is strictly more expressive than HATL, which, in turn, is strictly more expressive than ATL. Finally, we present a P-complete algorithm for model checking LAMB.

2 REASONING ABOUT STRATEGIC ABILITIES IN CHANGING ENVIRONMENTS

2.1 Models

Let $Agt = \{1, \dots, n\}$ be a non-empty finite set of agents. We will call subsets $C \subseteq Agt$ *coalitions*, and complements \bar{C} of C *anti-coalitions*. Sometimes we also call Agt *the grand coalition*. Moreover, let $Prop = \{p, q, \dots\}$ and $Nom = \{\alpha, \beta, \dots\}$ be disjoint countably infinite sets of *atomic propositions* and *nominals* correspondingly. Finally, let $Act = \{a_1, \dots, a_m\}$ be a non-empty finite set of actions.

Definition 2.1 (Named CGM). Given a set of atomic propositions $Prop$, nominals Nom and agents Agt , a *named Concurrent Game Model* (nCGM) is a tuple $M = \langle S, \tau, L \rangle$, where:

- S is a non-empty finite set of states;
- $\tau : S \times Act^{Agt} \rightarrow S$ is a transition function that assigns the outcome state $s' = \tau(s, (a_1, \dots, a_n))$ to a state s and a tuple of actions (a_1, \dots, a_n) ;
- $L : Prop \cup Nom \rightarrow 2^S$ is a valuation function such that for all $\alpha \in Nom : |L(\alpha)| \leq 1$ and for all $s \in S$, there is some $\alpha \in Nom$ such that $L(\alpha) = \{s\}$.

We denote an nCGM M with a designated state s as M_s . Since all CGMs we are dealing with in this paper are named, we will abuse terminology and call nCGMs just CGMs.

Let $True(s) = \{p \in Prop \mid s \in L(p)\} \cup \{\alpha \in Nom \mid s \in L(\alpha)\}$ be the set of all propositional variables and nominals that are true in state s . We define the *size* of CGM M as $|M| = |Agt| + |Act| + |S| + |\tau| + \sum_{s \in S} |True(s)|$, where $|\tau| = |S| \cdot |Act|^{|Agt|}$. We call a CGM *finite*, if $|M|$ is finite. In this paper, we restrict our attention to finite models.

Our models differ from standard CGMs in two ways. First, similarly to hybrid logic [13], states of our models have names represented by nominals Nom . So, each nominal is assigned to at most one state, but each state may have multiple names. Observe that differently from hybrid logic, we allow nominals to have empty extensions, that is, to not be assigned to any state. In the next section we will use this property to ensure that once new states are introduced to a model, we always have names available for them. Finally, we also assume that our models are *properly named*, i.e. each state is assigned some nominal.

We also assume that all agents have the whole set Act of actions available to them. Such an assumption is relatively common in the strategy logics literature (see, e.g., [12, 20, 66]), and allows for a clearer presentation of our framework.

Definition 2.2 (Strategies). Given a coalition $C \subseteq Agt$, an *action profile for coalition C*, A_C , is an element of Act^C , and Act^{Agt} is the set of all *complete* action profiles, i.e. all tuples (a_1, \dots, a_n) for $Agt = \{1, \dots, n\}$.

Given an action profile A_C , we write $\tau(s, A_C)$ to denote a set $\{\tau(s, A) \mid A = A_C \cup A_{\bar{C}}, A_{\bar{C}} \in Act^{\bar{C}}\}$. Intuitively, $\tau(s, A_C)$ is the set of all states reachable by (complete) action profiles that extend a given action profile A_C ¹.

A (*memoryless*) *strategy profile* for C is a function $\sigma_C : S \times C \rightarrow Act$ with $\sigma_C(s, i)$ being an action agent i takes in s .

Given a CGM M , a *play* $\lambda = s_0 s_1 \dots$ is an infinite sequence of states in S such that for all $i \geq 0$, state s_{i+1} is a successor of s_i i.e., there exists an action profile $A \in Act^{Agt}$ s.t. $\tau(s_i, A) = s_{i+1}$. We will denote the i -th element of play λ by $\lambda[i]$. The set of all plays that can be realised by coalition C following strategy σ_C from some given state s , denoted by $\Lambda_{\sigma_C}^s$, is defined as

$$\{\lambda \mid \lambda[0] = s \text{ and } \forall i \in \mathbb{N} : \lambda[i+1] \in \tau(\lambda[i], \sigma_C(\lambda[i]))\}.$$

REMARK 1. Note that we consider only *memoryless (positional) strategies here*. One can also employ *memory-full (perfect recall) strategies*, where the choice of actions by a coalition depends not on the current state, but on the whole history of the system up to the given moment. We resort to the former for simplicity. Observe, however, that for ATL, the semantics based on these two types of strategies are equivalent [11, 50, 73]. We conjecture that this is also the case for LAMB, and leave it for future work.

Example 2.3. Examples of CGMs are given in Figure 1, where an arrow labelled, for example, by ab denotes the action profile, where agent 1 takes action a and agent 2 takes action b . In a state s of model M , the two agents can make a transition to state t , if they synchronise on their actions, i.e. if they choose the same actions. In model N , on the other hand, a similar transition can be performed whenever they choose different actions.

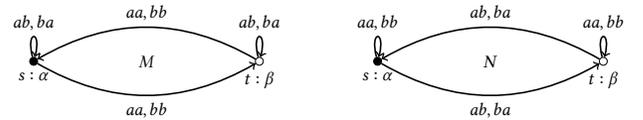


Figure 1: CGMs M and N for two agents and two actions. Propositional variable p is true in black states, and nominals α and β are true in their corresponding states.

2.2 Hybrid ATL

We start with an extension of ATL with hybrid logic features that allow us to express properties at states named by a given nominal.

Definition 2.4 (Syntax of HATL). The language of Hybrid ATL (HATL) logic is defined recursively as follows

$$\begin{aligned} \text{HATL} \ni \varphi ::= & p \mid \alpha \mid @_{\alpha} \varphi \mid \neg \varphi \mid (\varphi \wedge \varphi) \mid \langle\langle C \rangle\rangle X \varphi \mid \\ & \mid \langle\langle C \rangle\rangle \varphi \cup \varphi \mid \langle\langle C \rangle\rangle \varphi R \varphi \end{aligned}$$

¹Here, we here slightly abuse the notation and treat, whenever convenient, action profiles as sets rather than ordered tuples.

where $p \in Prop$, $\alpha \in Nom$ and $C \subseteq Agt$. The fragment of HATL without α and $@_\alpha\varphi$ corresponds to ATL².

Here, $\langle\langle C \rangle\rangle X\varphi$ means that a coalition C can ensure that φ holds in the neXt state. The operator $\langle\langle C \rangle\rangle\varphi U\psi$ (here U stands for *Until*) means that C has a strategy to enforce ψ while maintaining the truth of φ . $\langle\langle C \rangle\rangle\varphi R\psi$ means that C can maintain ψ until φ Releases the requirement for the truth of ψ . Derived operators $\langle\langle C \rangle\rangle F\varphi =_{def} \langle\langle C \rangle\rangle\top U\varphi$ and $\langle\langle C \rangle\rangle G\varphi =_{def} \langle\langle C \rangle\rangle\perp R\varphi$ mean “ C has a strategy to eventually make φ true” and “ C has a strategy to make sure that φ is always true” respectively. Hybrid logic operator $@_\alpha\varphi$ means “at state named α , φ is true”. This operator allows us to ‘switch’ our point of evaluation to the state labelled with nominal α in the syntax. Given a formula $\varphi \in$ HATL, the *size of φ* , denoted by $|\varphi|$, is the number of symbols in φ .

Definition 2.5 (Semantics of HATL). Let $M = \langle S, \tau, L \rangle$ be a CGM, $s \in S$, $p \in Prop$, $\alpha \in Nom$, and $\varphi, \psi \in$ HATL. The semantics of HATL is defined by induction as follows:

$$\begin{aligned} M_s \models p & \quad \text{iff } s \in L(p) \\ M_s \models \alpha & \quad \text{iff } s \in L(\alpha) \\ M_s \models @_\alpha\varphi & \quad \text{iff } \exists t \in S : \{t\} = L(\alpha) \text{ and } M_t \models \varphi \\ M_s \models \neg\varphi & \quad \text{iff } M_s \not\models \varphi \\ M_s \models \varphi \wedge \psi & \quad \text{iff } M_s \models \varphi \text{ and } M_s \models \psi \\ M_s \models \langle\langle C \rangle\rangle X\varphi & \quad \text{iff } \exists \sigma_C, \forall \lambda \in \Lambda_{\sigma_C}^s : M_{\lambda[1]} \models \varphi \\ M_s \models \langle\langle C \rangle\rangle\psi U\varphi & \quad \text{iff } \exists \sigma_C, \forall \lambda \in \Lambda_{\sigma_C}^s, \exists i \geq 0 : M_{\lambda[i]} \models \varphi \\ & \quad \text{and } M_{\lambda[j]} \models \psi \text{ for all } 0 \leq j < i \\ M_s \models \langle\langle C \rangle\rangle\psi R\varphi & \quad \text{iff } \exists \sigma_C, \forall \lambda \in \Lambda_{\sigma_C}^s, \forall i \geq 0 : M_{\lambda[i]} \models \varphi \\ & \quad \text{or } M_{\lambda[j]} \models \psi \text{ for some } 0 \leq j \leq i \end{aligned}$$

Observe that differently from standard hybrid logics [13], the truth condition for $@_\alpha\varphi$ also requires the state with name α to exist. This modification is necessary as we let some nominals to have empty denotations.

Example 2.6. Recall CGM M from Figure 1. It is easy to see that $M_s \models \alpha$ meaning that state s is named α ; $M_s \models \langle\langle \{1, 2\} \rangle\rangle X\neg p$, i.e. that from state s the grand coalition can force $\neg p$ to hold in the next state; and $M_s \models @_\beta \langle\langle \{1, 2\} \rangle\rangle G\beta$ meaning that in the state named β the grand coalition has a strategy to always remain in this state.

Expressivity of HATL. Even though the interplay between nominals and temporal modalities has been quite extensively studied (see, e.g., [18, 33, 42, 53]), to the best of our knowledge, the extension of ATL with nominals has never been considered. Our *Hybrid ATL* provides us with extra expressive power, compared to the standard ATL. We can, for instance, formulate safety and liveness properties in terms of *names*. For example, formula

$$\langle\langle C \rangle\rangle G \text{ safe} \wedge \langle\langle D \rangle\rangle \neg \text{crashed} \cup \alpha \wedge \langle\langle \emptyset \rangle\rangle F\beta$$

states that (1) coalition C can enforce that only safe states will be visited, (2) coalition D can avoid crashing until a state named α is visited, and (3) no matter what the agents Agt do, a state named β will eventually be visited. Both conjuncts (2) and (3) use nominals

²We define ATL over $\{\langle\langle C \rangle\rangle X\varphi, \langle\langle C \rangle\rangle\varphi U\varphi, \langle\langle C \rangle\rangle\varphi R\varphi\}$ as it is strictly more expressive than ATL defined over $\{\langle\langle C \rangle\rangle X\varphi, \langle\langle C \rangle\rangle\varphi U\varphi, \langle\langle C \rangle\rangle G\varphi\}$ [59].

to refer to state names in syntax, and are not expressible in ATL. Given that (as we will show later) the valuation of propositions in dynamic systems may change over time, as well as the transitions affecting reachability of some states, we believe that HATL allows for a more fine-grained way to capture safety and liveness.

Definition 2.7. Let L_1 and L_2 be two languages, and let $\varphi \in L_1$ and $\psi \in L_2$. We say that φ and ψ are *equivalent*, when for all CGMs M_s : $M_s \models \varphi$ if and only if $M_s \models \psi$.

If for every $\varphi \in L_1$ there is an equivalent $\psi \in L_2$, we write $L_1 \leq L_2$ and say that L_2 is *at least as expressive as L_1* . We write $L_1 < L_2$ iff $L_1 \leq L_2$ and $L_2 \not\leq L_1$, and we say that L_2 is *strictly more expressive than L_1* . Finally, if $L_1 \leq L_2$ and $L_2 \leq L_1$, we say that L_1 and L_2 are *equally expressive* and write $L_1 \approx L_2$.

Let us return to Figure 1, and assume that we have model N' , which is exactly like N with the only difference that $N'_s \models q$. Now, CGMs M and N' can be viewed as a disjoint union [17] $M \uplus N'$ (modulo renaming states in N' and assigning to them other nominals than α and β , like γ and δ correspondingly) of two isolated submodels M and N' . Note that no ATL formula φ that holds in M depends on the submodel N' as there are no transitions there. Hence, ATL cannot distinguish between M and $M \uplus N'$. At the same time, in contrast to state labels, we can use nominals in the syntax. So, formulas of HATL containing $@_\alpha$ operators can access states in the submodel N' , and hence can have different truth values in M and $M \uplus N'$. An example of such a formula would be $@_\gamma q$, which is false everywhere in M (since there is no state named γ), and true everywhere in $M \uplus N'$. This trivially implies that HATL is *strictly more expressive than ATL* ($ATL < HATL$).

2.3 Logic for ATL Model Building

Now we turn to the full logic with dynamic update operators.

Definition 2.8 (Syntax of LAMB). The language LAMB of *Logic for ATL Model Building* is defined as follows

$$\begin{aligned} \text{LAMB} \ni \varphi ::= & p \mid \alpha \mid @_\alpha\varphi \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\langle C \rangle\rangle X\varphi \mid \\ & \mid \langle\langle C \rangle\rangle\varphi U\varphi \mid \langle\langle C \rangle\rangle\varphi R\varphi \mid [\pi]\varphi \\ \pi ::= & (p\alpha := \psi) \mid \alpha \xrightarrow{A} \alpha \mid \bigcirc \end{aligned}$$

where $p \in Prop$, $\alpha \in Nom$, $C \subseteq Agt$, $A \in Act^{Agt}$. We will also write $[\pi_1; \dots; \pi_n]\varphi$ for $[\pi_1] \dots [\pi_n]\varphi$. Two important fragments of LAMB that we will also consider are the one without constructs $\alpha \xrightarrow{A} \alpha$ and \bigcirc , called *substitution LAMB (SLAMB)*; and the one without $p\alpha := \psi$ and \bigcirc , called *arrow LAMB (ALAMB)*.

Definition 2.9 (Semantics of LAMB). Let $M = \langle S, \tau, L \rangle$ be a CGM, $s \in S$, $p \in Prop$, $\alpha \in Nom$, and $\varphi, \psi \in$ LAMB. The semantics of LAMB is defined as in Definition 2.5 with the following additional cases:

$$\begin{aligned} M_s \models [p\alpha := \psi]\varphi & \quad \text{iff } M_s^{p\alpha := \psi} \models \varphi \\ M_s \models [\bigcirc]\varphi & \quad \text{iff } M_s^{\bigcirc} \models \varphi \\ M_s \models [\alpha \xrightarrow{A} \beta]\varphi & \quad \text{iff } M_s^{\alpha \xrightarrow{A} \beta} \models \varphi \end{aligned}$$

where M^π with $\pi \in \{p\alpha := \psi, \alpha \xrightarrow{A} \beta, \bigcirc\}$ is called an *updated CGM*, and is defined as follows:

- $M_s^{p\alpha:=\psi} = \langle S, \tau, L^{p\alpha:=\psi} \rangle$, where, if $\exists t \in S$ such that $L(\alpha) = \{t\}$, then

$$L^{p\alpha:=\psi}(p) = \begin{cases} L(p) \cup \{t\} & \text{if } M_s \models \psi, \\ L(p) \setminus \{t\} & \text{if } M_s \not\models \psi, \end{cases}$$

and $M^{p\alpha:=\psi} = \langle S, \tau, L \rangle$ otherwise.

- $M^{\alpha \xrightarrow{A} \beta} = \langle S, \tau \xrightarrow{A} \beta, L \rangle$, where, if $\exists s, t \in S$ such that $L(\alpha) = s$ and $L(\beta) = t$, then

$$\tau \xrightarrow{A} \beta(s', A') = \begin{cases} t, & \text{if } s' = s \text{ and } A' = A \\ \tau(s', A') & \text{otherwise,} \end{cases}$$

and $M^{\alpha \xrightarrow{A} \beta} = \langle S, \tau, L \rangle$ otherwise.

- $M^{\odot} = \langle S^{\odot}, \tau^{\odot}, L^{\odot} \rangle$, where, if $L(\alpha) = \emptyset$, then $S^{\odot} = S \cup \{t\}$ with $t \notin S$, $\tau^{\odot} = \tau \cup \{(t, A, t) \mid \forall A \in Act^{Ag^t}\}$, and $L^{\odot}(\alpha) = L \cup \{(\alpha, \{t\})\}$. If $L(\alpha) \neq \emptyset$, then $M^{\odot} = \langle S, \tau, L \rangle$.

Intuitively, updating a given model with $p_\alpha := \psi$ assigns to propositional variable p in the state named α (if there is such a state) the truth value of ψ in the state of evaluation. Observe, however, that due to the presence of operators $@_\alpha$, we can also let the truth-value of p be dependent on the truth of ψ in any other state. Updating with $\alpha \xrightarrow{A} \beta$ results in redirecting the A -labelled arrow that starts in state named α from some state $\tau(L(\alpha), A)$ to state named β (if states named α and β exist). Finally, operator \odot adds a new state to the model and gives it name α . All propositional variables are false in the new state and all transitions are self-loops. The self-loops and the fact that propositional variables are false in the new state is a design choice that achieves two goals. First, adding new state results in a finite model (recall that to determine the size of a model we count all true propositions), and, second, we deem this to be the smallest meaningful change that is in line with the idea of modularity. All valuations and self-loops can be then further modified by the corresponding substitutions and arrow operators.

Example 2.10. Consider CGMs M and N from Figure 1 and an update $\alpha \xrightarrow{aa} \alpha$. Observe that such an update leaves N intact as there is already a self-loop labelled aa in the α -state. Model M , on the other hand, is changed and the resulting updated model $M^{\alpha \xrightarrow{aa} \alpha}$ is depicted in Figure 2.

Now it is easy to see, for example, that in the updated model the first agent can force the system to stay in the α -state by choosing action a , i.e. $M_s \models [\alpha \xrightarrow{aa} \alpha] \langle \{1\} \rangle X\alpha$, while it is not the case for $N^{\alpha \xrightarrow{aa} \alpha}$ (which is the same as N), i.e. $N_s \not\models [\alpha \xrightarrow{aa} \alpha] \langle \{1\} \rangle X\alpha$.

As a more complex update, consider $\pi_1 := \odot; p_\gamma := \top; fine_\gamma := \top$ and the resulting model M^{π_1} , where we have a new state named γ that satisfies propositions p and $fine$ (the intuition behind $fine$ is given in Section 3.1). Further, we can add redirection of some edges $\pi_2 := \pi_1; \alpha \xrightarrow{ab} \gamma; \alpha \xrightarrow{ba} \gamma; \gamma \xrightarrow{aa} \beta; \gamma \xrightarrow{bb} \beta$ to obtain model M^{π_2} . The final complex update of M with SN , also presented in Figure 2, is discussed in the next section in the context of normative updates.

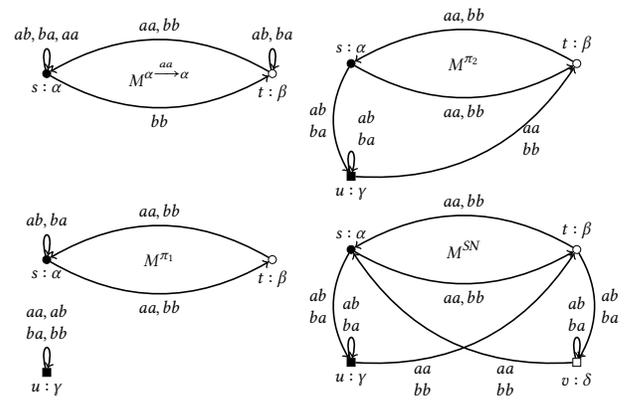


Figure 2: Updated CGMs $M^{\alpha \xrightarrow{aa} \alpha}$ (top left), M^{π_1} (bottom left), M^{π_2} (top right), and M^{SN} (bottom right). Proposition p is true in black states, and $fine$ is true in square states.

3 DYNAMIC MAS THROUGH THE LENS OF LAMB

In this section, we show how our dynamic approach to CGMs allows us to capture various ideas and intuitions from the MAS research.

3.1 Normative Updates on CGMs

Updates of CGMs play a crucial role in the area of normative MAS (see [9] for an overview). In the context of ATL, the general idea is to divide actions of agent's into those that are permitted to be executed in a current state, and into those that are not. We argue that LAMB, thanks to its dynamic features, captures various intuitions and technicalities of reasoning about norms in MAS.

As a use case of LAMB, we consider *norm-based updates* as presented in [22].³ The authors distinguish two types of norms for CGMs: regimenting norms and sanctioning norms. Regimenting norms prohibit certain transitions by looping them to the origin state. This corresponds to the intuition that a selected action profile has no noticeable effect and thus the system stays in the same state. Note that agents in this case always have available actions. This assumption is sometimes called *reasonableness* [4, 5, 7, 8, 40, 76].

While it is quite straightforward to model regimenting norms in LAMB using arrow updates, for the sake of an example we focus on more general and subtle sanctioning norms. Such norms put sanctions, or fines, on certain action profiles without explicitly prohibiting them. A bit more formally, a sanctioning norm is a triple $(\varphi, \mathcal{A}, \mathcal{S})$, where φ is a formula of a logic, \mathcal{A} is a set of action profiles, and \mathcal{S} is a set of sanctions, which is a subset of the set of propositional variables. Intuitively, sanctioning norm $(\varphi, \mathcal{A}, \mathcal{S})$ states that performing action profiles from \mathcal{A} from states satisfying φ imposes sanctions from \mathcal{S} .

The result of updating a CGM M with a norm $SN = (\varphi, \mathcal{A}, \mathcal{S})$ is a CGM M^{SN} such that for all \mathcal{A} -transitions from φ -states, we create copies of those states and make sanctioning atoms \mathcal{S} true in those copies. Non-sanctioned transitions from the copy-states have the same outcome as the original transitions in the initial model.

³We simplify the exposition for the sake of clarity.

Consider as an example the sanctioning norm $SN = (\top, \{ab, ba\}, \{fine\})$ and CGM M from Example 1. Recall that in M , agents can switch the current state only if they cooperate (i.e. choose the same actions). Hence, norm SN penalises non-cooperative behaviour with the sanction $fine$. We can directly model the implementation of SN on M in LAMB by translating SN into a complex update

$$SN = (\odot; p_Y := \top; fine_Y := \top; \alpha \xrightarrow{ab} \gamma; \alpha \xrightarrow{ba} \gamma; \gamma \xrightarrow{aa} \beta; \gamma \xrightarrow{bb} \beta; \odot; fine_\delta := \top; \beta \xrightarrow{ab} \delta; \beta \xrightarrow{ba} \delta; \delta \xrightarrow{aa} \alpha; \delta \xrightarrow{bb} \alpha).$$

The result of updating M with SN as well as some intermediate updates is presented in Figure 2. In M^{SN} , we create copies of states s and t (both satisfying \top), named u and v , which now also satisfy sanction $fine$. Then, undesirable action profiles $\{ab, ba\}$ originating in s and t lead to the sanctioned states u and v . At the same time, action profiles not from $\{ab, ba\}$ behave similarly to the initial model. It is easy to see that our LAMB update SN on M captures the effect of SN on M .

3.2 Synthesis

Assume that you have a model M_s of a MAS that does not satisfy some safety requirement φ . One way to go about it would be to create a completely new model from scratch, or to try to manually fix the existing one. However, in case of large models and complex safety requirements, both options may not be feasible. This is exactly where the problem of *synthesis* (or *modification synthesis*) comes in. The full language of LAMB is expressive enough to capture the (bounded) synthesis problem from a given specification and starting model. In such a way, the modifications of the model are presented as updates operators of LAMB.

Definition 3.1. For a given CGM M_s , formula φ , and natural number n , the *bounded (modification) synthesis existence problem* decides whether there is an update $\pi_\varphi := [\pi_1, \dots, \pi_k]$, with the size at most n , such that $M_s \models [\pi_\varphi]\varphi$.

The bounded synthesis existence problem is important, as it tackles the synthesis of *compact* modifications. Indeed, oftentimes designers of (a model of) MAS are interested in a modification that achieves the goal without significantly altering the initial model. Think of an example of a software update, where we would like to extend the functionality of the software without rewriting most of its code. Moreover, the bounded synthesis problem is also important in cases where changes in the given model are costly, and we want to avoid wasting resources as much as possible.

In Section 5, we show that the complexity of the model checking problem for LAMB is P-complete. With this in mind, it is easy to see that the complexity of the bounded synthesis existence problem is NP-complete.

PROPOSITION 3.2. *The bounded synthesis problem for LAMB is NP-complete.*

PROOF. To see that the problem is in NP, let M_s be a CGM, φ be a formula, and n be a natural number. We guess a π_φ of size at most n . It follows that $[\pi_\varphi]\varphi$ is of polynomial size w.r.t $|\varphi| + n$. Then we can model check $M_s \models [\pi_\varphi]\varphi$ in polynomial time (Theorem 5.1).

For the NP-hardness, we employ the reduction from 3-SAT problem. Let $\varphi := \bigwedge_{1 \leq i \leq k} (\psi_{i,1} \vee \psi_{i,2} \vee \psi_{i,3})$, where ψ 's are literals, be an

instance of 3-SAT, and let $P^\varphi = \{p^1, \dots, p^m\}$ be the set of propositions appearing in φ . We construct model M^φ over one agent consisting of a single state s with the name α , and such that $V^\varphi(p^i) = \emptyset$ for all $p^i \in P^\varphi$. All transitions for the agent are self-loops. Now it is easy to see that φ is satisfiable iff there is a π_φ consisting only of substitutions $p_\alpha^i := \top$, and thus of the size linear in $|P^\varphi|$, such that $M_s^\varphi \models [\pi_\varphi]\varphi$. In other words, if there is an assignment that makes φ true, we can explicitly simulate it in M_s^φ using constructs $p_\alpha^i := \top$ for variables $p^i \in P^\varphi$ that should be assigned 'true'. \square

The fact that LAMB captures the modification synthesis is significant, because a *constructive* solution to the problem would produce a step-by-step recipe, or an instruction, of how to modify a given model to make it satisfy some desirable property φ . In this section we studied the computational complexity of checking *whether* a required modification of certain size exists. In future work, we will focus on constructive solutions to the problem, i.e. providing algorithms that automatically construct the required LAMB update.

3.3 Mechanisms for Social Choice

The key advantage of synthesising mechanisms from logical specifications is that, as a declarative approach, the designer is not required to construct a complete solution for the problem of interest; instead, she can describe the desired mechanism in terms of its rules and desirable properties. Some recent approaches use model checking and satisfiability procedures for Strategy Logic [63, 64]. We, however, focus on the dynamic LAMB approach.

Let us assume a mechanism encoded as a CGM, and a finite set of alternatives Alt . Such a mechanism may represent, for instance, a single-winner election or a resource allocation protocol. In such CGM, we let the atomic proposition $pref(i,j)_a$ denote that agent a prefers the alternative i to j (e.g., she prefers that i is elected over j). We also let $dislike_a(i)$ indicate that i is *disliked* by agent a , and $chosen_i$ denote that the alternative i was chosen.

We now illustrate how to capture two classic concepts from game theory: individual rationality and Pareto optimality. Individual rationality expresses the idea that each agent can ensure nonnegative utility [67]. In our setting, this can be seen as avoiding disliked candidates and expressed with the following formula

$$\bigwedge_{a \in Agt, i \in Alt} \langle\langle \{a\} \rangle\rangle G \neg (chosen_i \wedge dislike_a(i))$$

which states that each agent has a strategy to enforce that none of the disliked alternatives are ever chosen.

A mechanism is Pareto optimal if any change of outcome that is beneficial to one agent is detrimental to at least one of the other agents. The formula

$$pareto_i := \bigwedge_{a \in Agt, j \in Alt \setminus \{i\}} ((pref_a(j,i) \rightarrow \bigvee_{b \in Agt} pref_b(i,j)))$$

expresses that an alternative i is Pareto optimal whenever for any other alternative j , if j is preferred to an agent a , then there is an agent b that prefers i over j .

We can then express that there is a strategy for the coalition C to ensure that any chosen alternative is Pareto optimal, with the formula $\varphi_{po} := \langle\langle C \rangle\rangle G \bigwedge_{i \in Alt} chosen_i \rightarrow pareto_i$. If we let $C = \emptyset$,

this formula requires choices to be Pareto optimal for any possible behavior of the agents.

Once desirable mechanism properties are expressed as ATL formulas, one can use LAMB to verify whether CGM M_s has such property or would have it in case a sequence of modifications $\pi_1; \dots; \pi_n$ was performed, i.e. whether $M_s \models [\pi_1; \dots; \pi_n] \varphi_{po}$ for Pareto optimality. Further, employing the ideas of synthesis, one could automatically obtain the required modifications.

4 EXPRESSIVITY

In this section, we compare the expressive power of LAMB and its fragments. In particular, we show that, interestingly enough, substitutions on their own do not add expressive power compared to the base HATL (Theorem 4.1). The ability to move arrows, on the other hand, leads to an increase in expressivity (Theorem 4.2).

First, we show that SLAMB and HATL are equally expressive. This is quite intriguing since it means that adding substitutions to the base logic HATL does not allow us to express anything that we could not express in HATL. We prove this result by providing a truth-preserving translation from formulas of SLAMB into formulas of HATL. While expressivity results that use translation schemas (usually based on *reduction axioms*) are quite ubiquitous in DEL (see [77] for an overview, and [54, 55] for substitution specific translations), to the best of our knowledge, this is the first time reduction has been discussed in the context of dynamics for logics of strategic abilities.

THEOREM 4.1. HATL \approx SLAMB

PROOF. Consider the reduction axioms below.

- A1 $[p_\alpha := \psi]p \leftrightarrow (\alpha \rightarrow \psi) \wedge (\neg\alpha \rightarrow p)$
- A2 $[p_\alpha := \psi]q \leftrightarrow q$ for $p \neq q$
- A3 $[p_\alpha := \psi]\beta \leftrightarrow \beta$
- A4 $[p_\alpha := \psi]@_\beta\varphi \leftrightarrow @_\beta[p_\alpha := \psi]\varphi$
- A5 $[p_\alpha := \psi]\neg\varphi \leftrightarrow \neg[p_\alpha := \psi]\varphi$
- A6 $[p_\alpha := \psi](\varphi \wedge \chi) \leftrightarrow [p_\alpha := \psi]\varphi \wedge [p_\alpha := \psi]\chi$
- A7 $[p_\alpha := \psi]\langle\langle C \rangle\rangle X\varphi \leftrightarrow \langle\langle C \rangle\rangle X[p_\alpha := \psi]\varphi$
- A8 $[p_\alpha := \psi]\langle\langle C \rangle\rangle\varphi U\chi \leftrightarrow \langle\langle C \rangle\rangle[p_\alpha := \psi]\varphi U[p_\alpha := \psi]\chi$
- A9 $[p_\alpha := \psi]\langle\langle C \rangle\rangle\varphi R\chi \leftrightarrow \langle\langle C \rangle\rangle[p_\alpha := \psi]\varphi R[p_\alpha := \psi]\chi$

Soundness of the axioms can be shown by application of the definition of semantics (Definition 2.9). As an example, consider axiom A1. Let us assume that for some CGM M_s it holds that $M_s \models [p_\alpha := \psi]p$. According to the definition of $M_s^{p_\alpha := \psi}$ (Definition 2.9), we consider two cases. First, if $M_s \models \alpha$, then the truth-value of p is equivalent to whether $M_s \models \psi$. Hence, the first conjunct of the right-hand part of A1: $M_s \models \alpha \rightarrow \psi$. Second, if $M_s \not\models \alpha$, then the truth-value of p remains intact, i.e. $M_s \models \neg\alpha \rightarrow p$. Combining two parts together, we have that $M_s \models [p_\alpha := \psi]p$ if and only if $M_s \models (\alpha \rightarrow \psi) \wedge (\neg\alpha \rightarrow p)$.

As another example, consider A7. Assume that for an arbitrary CGM M_s we have that $M_s \models [p_\alpha := \psi]\langle\langle C \rangle\rangle X\varphi$. By the definition of semantics, the latter is equivalent to $M_s^{p_\alpha := \psi} \models \langle\langle C \rangle\rangle X\varphi$, which in turn is equivalent to the fact that $\exists \sigma_C, \forall \lambda \in \Lambda_{\sigma_C}^s : M_{\lambda[1]}^{p_\alpha := \psi} \models \varphi$. Now, by the semantics of substitutions, we have that $\exists \sigma_C, \forall \lambda \in \Lambda_{\sigma_C}^s : M_{\lambda[1]} \models [p_\alpha := \psi]\varphi$, which in turn is equivalent to $M_s \models \langle\langle C \rangle\rangle X[p_\alpha := \psi]\varphi$. Other cases can be shown in a similarly.

Now, having a set of sound reduction axioms, it is straightforward to define a truth-preserving translation $tr : \text{SLAMB} \rightarrow \text{HATL}$. Such a translation would work from inside-out, taking the innermost occurrence of a substitution and pushing it closer to propositions and nominals, eventually translating them away completely. In particular, axioms A2 and A3 state that changing the valuation of p does not affect other propositional variables and nominals. Axiom A1 intuitively means that if the current state is named α , then the valuation of p is the same as the valuation ψ in this state. And if the current state is not α , then p in the *current* state retains its valuation. See [54, 55, 77] for translations in the context of DEL. \square

An observant reader may notice, that the presented translation of a SLAMB formula may result in an exponentially larger formula of HATL (due to repetitions in A6, A8, and A9). Or, equivalently, that the initial formula of SLAMB is exponentially *more succinct* than its translation. Such a blow-up is quite natural for reduction-based translations in DEL [35, 61]. In the next section we will show that despite this, model checking SLAMB is P-complete.

Now we turn to ALAMB and show that the ability to move arrows grants us additional expressivity.

THEOREM 4.2. HATL $<$ ALAMB and SLAMB $<$ ALAMB.

PROOF. The fact that HATL \leq ALAMB follows trivially as HATL is a fragment of ALAMB. To see that ALAMB $\not\leq$ HATL recall Example 2.10. It is easy to see that models M_s and N_s cannot be distinguished by any HATL formula⁴. Indeed, in both models states s and t agree on their corresponding nominals and propositional variables, and, moreover in all states none of the agents can force a transition on their own. The fact that such a transition requires different action profiles by $\{1, 2\}$ in different models cannot be captured by formulas of HATL as they do not have the access to particular actions, rather just to abilities of the agent.

At the same time, as shown in Example 2.10, the ALAMB formula $[\alpha \xrightarrow{aa} \alpha] \langle\langle \{1\} \rangle\rangle X\alpha$ holds in M_s and is false in N_s , and, therefore, HATL $<$ ALAMB. The fact that SLAMB $<$ ALAMB follows by transitivity of the expressivity relation from HATL \approx SLAMB. \square

From Theorem 4.2 it follows that HATL $<$ LAMB and SLAMB $<$ LAMB. For future work, we leave open the question of whether LAMB is strictly more expressive than ALAMB, and conjecture that it is indeed the case. Figure 3 summarises the expressivity results.

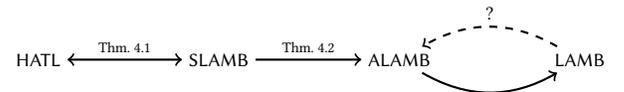


Figure 3: Overview of the expressivity results. An arrow from L_1 to L_2 means $L_1 \leq L_2$. If there is no symmetric arrow, then $L_1 < L_2$. This relation is transitive, and we omit transitive arrows in the figure. The dashed arrow with the question mark denotes the open question.

⁴In fact, there is an alternating bisimulation [3] between the models. The discussion of an appropriate notion of bisimulation for LAMB is outside of the scope of this paper and is left for future work.

5 MODEL CHECKING

The model checking problem for LAMB consists in determining, for a CGM M_s and a formula φ , whether $M_s \models \varphi$. We show that despite the increase in expressivity, the complexity of the model checking problem for the full language of LAMB is still P-complete.

The complexity of model checking ATL is known to be P-complete [11], and it is easy to see that it remains the same also for HATL. The algorithm for ATL uses function $Pre(M, C, Q)$ that computes for a given CGM M , coalition $C \subseteq Agt$ and a set $Q \subseteq S$, the set of states, from which coalition C can force the outcome to be in one of the Q states. Function Pre can be computed in polynomial time. We can use exactly the same algorithm for computing Pre as for the standard ATL, however in our case we will compute Pre for not only the original model M , but its updated versions as well.

Algorithm 1 An algorithm for model checking LAMB

```

1: procedure MC( $M, s, \varphi$ )
2:   case  $\varphi = \alpha$ 
3:     return  $s \in L(\alpha)$ 
4:   case  $\varphi = @_\alpha\psi$ 
5:     if  $L(\alpha) \neq \emptyset$  then
6:       return MC( $M, L(\alpha), \psi$ )
7:     else
8:       return false
9:   case  $\varphi = [\pi]\psi$  with  $\pi \in \{p_\alpha := \psi, \alpha \xrightarrow{A} \beta, \textcircled{\alpha}\}$ 
10:    return MC(UPDATE( $M, s, \pi$ ),  $s, \psi$ )
11: end procedure

```

The model checking algorithm for LAMB (Algorithm 1) is similar to the one for ATL when it comes to temporal modalities and Boolean cases, and thus we omit them for brevity (see the full algorithm in [39]). Apart from them, we have hybrid cases $\varphi = \alpha$ and $\varphi = @_\alpha\psi$, and the dynamic case $\varphi = [\pi]\psi$ with $\pi \in \{p_\alpha := \psi, \alpha \xrightarrow{A} \beta, \textcircled{\alpha}\}$. Regarding the $\varphi = @_\alpha\psi$, we evaluate ψ at state named α , if the state with such a name exists. If the denotation of name α is empty, then $@_\alpha\psi$ is false.

The dynamic case $\varphi = [\pi]\psi$ is a bit more involved as the algorithm evaluates ψ in a new updated model M^π . Procedure UPDATE for constructing updated models is captured by Algorithm 2.

Algorithm 2 An algorithm for computing updated models

```

1: procedure UPDATE( $M, s, \pi$ )
2:   case  $\pi = p_\alpha := \psi$ 
3:     if  $L(\alpha) \neq \emptyset$  then
4:       if MC( $M, s, \psi$ ) then
5:          $L^\pi(p) = L(p) \cup L(\alpha)$ 
6:       else
7:          $L^\pi(p) = L(p) \setminus L(\alpha)$ 
8:       return  $M^\pi = \langle S, \tau, L^\pi \rangle$ 
9:     else
10:      return  $M$ 
11:   case  $\pi = \alpha \xrightarrow{A} \beta$ 
12:     if  $L(\alpha) \neq \emptyset$  and  $L(\beta) \neq \emptyset$  then
13:        $\tau^\pi = \tau \setminus \{(L(\alpha), A, \tau(L(\alpha), A))\} \cup \{(L(\alpha), A, L(\beta))\}$ 
14:       return  $M^\pi = \langle S, \tau^\pi, L \rangle$ 
15:     else
16:       return  $M$ 
17:   case  $\pi = \textcircled{\alpha}$ 
18:     if  $L(\alpha) = \emptyset$  then

```

```

19:        $S^\pi = S \cup \{t\}$ , where  $t$  is fresh
20:        $\tau^\pi = \tau \cup \{(t, A, t) \mid A \in Act^{Agt}\}$ 
21:        $L^\pi = L \cup \{(\alpha, \{t\})\}$ 
22:       return  $M^\pi = \langle S^\pi, \tau^\pi, L^\pi \rangle$ 
23:     else
24:       return  $M$ 
25: end procedure

```

In the procedure, an updated model is constructed according to Definition 2.9. For the case of substitutions, we first check whether state named α exists, and if it does, we update the valuation function L based on whether $M_s \models \psi$. For arrows $\alpha \xrightarrow{A} \beta$, if both states named α and β exist, we substitute in τ transition $(L(\alpha), A, \tau(L(\alpha), A))$ (i.e. transition from state named α via A to whatever state is assigned according to $\tau(L(\alpha), A)$) by the required transition $(L(\alpha), A, L(\beta))$ (line 13). Finally, to add a new state with name α , we first check whether the name is not used, and then extend S , τ , and L of the original model accordingly.

Model checking for LAMB is done then recursively by a combination of procedures MC, which decides whether a given formula is true in a given model, and UPDATE, which computes required updated models. MC calls UPDATE when it needs to perform an update (line 10, Alg. 1), and UPDATE calls MC when it needs to compute the valuation of ψ for case $\pi = p_\alpha := \psi$ (line 4, Alg. 2).

We run MC(N, t, ψ) for at most $|\varphi|$ formulas ψ and at most $|\varphi|$ models N . Each run, similarly to the algorithm for ATL, is done in polynomial time with respect to $|M|$. Hence, procedure MC is used by the model checking algorithm for a polynomial amount of time.

At the same time, we run UPDATE(N, t, π) for at most $|\varphi|$ models N and at most $|\varphi|$ formulas ψ (the substitution case). The sizes of updated models are bounded by $|\varphi| \cdot |M|$ (the case of adding a new state). Thus, each run of UPDATE takes polynomial time, and hence we spend a polynomial amount of time in the procedure while performing model checking.

Both procedures, MC and UPDATE, take polynomial time to run, and, therefore, model checking for LAMB can be done in polynomial time. The lower bound follows straightforwardly from P-completeness of ATL model checking.

THEOREM 5.1. *The model checking problem for LAMB is P-complete.*

REMARK 2. *To make the language of LAMB more succinct, we can extend it with constructs $[\pi \cup \rho]\varphi$, meaning ‘whichever update we implement, π or ρ , φ will be true (in both cases)’. The model checking of the resulting logic is PSPACE-complete. Details about the extension and proofs can be found in section A Note On Succinctness in [39].*

6 RELATED WORK

Strategic Reasoning. From the perspective of strategic reasoning, our work is related to the research on rational verification and synthesis. The first is the problem of checking whether a temporal goal is satisfied in some (or all) game-theoretic equilibria of a CGM [1, 45]. Rational synthesis consists in the automated construction of such a model [26, 32]. In this direction, [49] investigated the problem of finding *incentives* by manipulating the weights of atomic propositions to persuade agents to act towards a temporal goal.

Recent work has also investigated the use of formal methods to verify and synthesize mechanisms for social choice using model checking and satisfiability procedures for variants of Strategy Logic

[63, 64]. While being able to analyse MAS with respect to complex solution concepts, all these works face high complexity issues. In particular, key decision problems for rational verification with temporal specifications are known to be 2EXPTIME-complete [45] and model checking Strategy Logic is NONELEMENTARY for memoryfull agents [65]. Compared to these approaches, LAMB offers relatively high expressivity while maintaining the P-completeness of its model checking problem.

The recently introduced *obstruction ATL* [23, 24] (OATL) allows reasoning about agents’ strategic abilities while being hindered by an external force, called the *Demon*. Being inspired by sabotage modal logic [74], in this logic the Demon is able to disable some transitions and thus impact the strategic abilities of the agents in a system. This is somewhat related to normative updates that we covered in Section 3.1 and the *module checking problem* [51, 56], where agents interact with a non-deterministic environment that may inhibit access to certain paths of the computation tree.

Notice that LAMB is significantly more general than the presented approaches as it allows to not only restrict transitions or access, but also change it in a more nuanced way by redirecting arrows (and, e.g., granting access to a state). Moreover, LAMB also allows adding *new* states, as well as changing the valuations of propositions. Moreover, updates in LAMB are explicitly present in the syntax that enables explicit synthesis of model modifications.

Nominals. Nominals are an integral part of *hybrid logic* [13] and is a common tool whenever one needs to refer to particular states on the syntax level. For example, nominals and other hybrid modalities are ubiquitous in the research on *logics for social networks* (see [69, Chapter 3] for a comprehensive overview).

In the setting of DEL, tools and methods of hybrid logic have been used, for example, to relax the assumption of common knowledge of agents’ names [79], to study the interplay between public announcements and distributed knowledge [46], and to tackle the information and intentions dynamics in interactive scenarios [72]. Moreover, nominals were used to provide an axiomatisation of a hybrid variant of *sabotage modal logic* [75], which extends the standard language of modal logic with constructs $\blacklozenge\varphi$ meaning ‘after removing some edge in the model, φ holds’ [14, 74].

Nominals were also used in *linear- and branching-time temporal logics* to refer to particular points in computation (see, e.g., [18, 19, 33, 34, 42, 52, 53, 58]). In the framework of strategic reasoning, [48] used some ideas from hybrid logic, but neither $@_\alpha$ nor nominals themselves. Hence, in terms of novelty, to the best of our knowledge, the *Hybrid ATL* (HATL) proposed in this paper is the first attempt to combine nominals with the ATL-style strategic reasoning.

The Interplay Between DEL and Strategic Reasoning. As we mentioned in the introduction, albeit DEL and various strategic logics being very different formalism, some avenues of DEL research has incorporated ideas from logics for strategic reasoning. Examples include the exploration of *concurrent DEL games* [62], *alternating-time temporal DEL* [29], *coalitions announcements* [6, 36] and other forms of *strategic multi-agent communication* (see, e.g., [2, 41]).

To the best of our knowledge, DEL updates for CGMs, up until now, were considered only in [37, 38], where the authors capture granting and revoking actions of singular agents as well as updates

based on *action models* [15]. Both works are limited to the neXt-time fragment of ATL (so-called *coalition logic* [68]). Moreover, they do not support such expressive features of LAMB as adding *new* states and changing the valuation of propositional variables. Additionally, our arrow-redirecting operators allow for greater flexibility while dealing with agents’ strategies.

7 DISCUSSION & CONCLUSION

We proposed LAMB, a logic for updating CGMs that combines ideas from both the strategy logics tradition (ATL in our case) and the DEL tradition. We have argued that LAMB can be useful for reasoning about a variety of dynamic phenomena in MAS thanks to the modular nature of its primitive update operators. Finally, we have explored the expressivity hierarchy of LAMB and its fragments, and demonstrated that the model checking problem for LAMB is P-complete.

As we have just scratched the surface of dynamic updates for CGMs, there is a plethora of open questions. One of the immediate ones is to explore the satisfiability problem for LAMB. Another one is to assign costs to different types of model changes. In such a way, we will be able to generalise the bounded modification synthesis to the scenarios, where some changes are more costly to implement and thus are less optimal. Moreover, having costs associated with model changes will allow for a direct comparison of (this generalised version of) LAMB and Obstruction ATL (see Section 6). We conjecture that in such a scenario, LAMB will subsume OATL.

Apart from that, in Section 3.2, we mentioned the exploration of constructive solutions to the synthesis problem, both bounded and unbounded versions, as a promising area of further research. One way to go about it is, perhaps, taking intuitions from the constructive approaches to the ATL satisfiability [43, 44, 78]. Those solutions can be then embedded into some of the existing model checkers, like MCMAS [60] and STV [57], so that if a model checker returns FALSE for a given model and property φ , the tool automatically constructs an update that can fix the model so that it satisfies φ .

In a more general setting of CGM updates, one can also consider modifications that cannot be captured by LAMB. For example, we can explore the effects of granting or revoking actions to/from certain agents, changing the number of agents, or any combinations thereof with the LAMB updates. As mentioned in Section 6, some preliminary work on changing the actions available to agents has been done for the neXt-time fragment of ATL [37, 38].

Finally, as LAMB is based on ATL, we find it particularly interesting to consider more expressive base languages, like, ATL^* or variations of strategy logic SL [65]. Of particular interest is the *simple goal fragment of SL* [16], which is strictly more expressive than ATL and yet allows for P-time model checking. We would also like to consider the ideas of model updates in the setting of STIT logics [21, 47]. Additionally, we believe that ideas from Separation Logics [30, 31, 71], which were proposed to verify programs with mutable data structures, could also provide insights on how to reason about separation and composition-based modifications of MAS.

ACKNOWLEDGMENTS

This project has received funding from the European Union’s H2020 Marie Skłodowska-Curie project with grant agreement No 101105549.

REFERENCES

- [1] Alessandro Abate, Julian Gutierrez, Lewis Hammond, Paul Harrenstein, Marta Kwiatkowska, Muhammad Najib, Giuseppe Perelli, Thomas Steeples, and Michael J. Wooldridge. 2021. Rational verification: game-theoretic verification of multi-agent systems. *Applied Intelligence* 51, 9 (2021), 6569–6584. <https://doi.org/10.1007/S10489-021-02658-Y>
- [2] Thomas Ágotnes, Philippe Balbiani, Hans van Ditmarsch, and Pablo Seban. 2010. Group announcement logic. *Journal of Applied Logic* 8, 1 (2010), 62–81. <https://doi.org/10.1016/j.jal.2008.12.002>
- [3] Thomas Ágotnes, Valentin Goranko, and Wojciech Jamroga. 2007. Alternating-time temporal logics with irrevocable strategies. In *Proceedings of the 11th TARK*, Dov Samet (Ed.), 15–24. <https://doi.org/10.1145/1324249.1324256>
- [4] Thomas Ágotnes, Wiebe van der Hoek, Juan A. Rodríguez-Aguilar, Carles Sierra, and Michael Wooldridge. 2007. On the Logic of Normative Systems. In *Proceedings of the 20th IJCAI*, Manuela M. Veloso (Ed.), 1175–1180.
- [5] Thomas Ágotnes, Wiebe van der Hoek, and Michael J. Wooldridge. 2010. Robust normative systems and a logic of norm compliance. *Logic Journal of the IGPL* 18, 1 (2010), 4–30. <https://doi.org/10.1093/jigpal/jzp070>
- [6] Thomas Ágotnes and Hans van Ditmarsch. 2008. Coalitions and announcements. In *Proceedings of the 7th AAMAS*, Lin Padgham, David C. Parkes, Jörg P. Müller, and Simon Parsons (Eds.), IFAAMAS, 673–680.
- [7] Thomas Ágotnes and Michael J. Wooldridge. 2010. Optimal social laws. In *Proceedings of the 9th AAMAS*, Wiebe van der Hoek, Gal A. Kaminka, Yves Lespérance, Michael Luck, and Sandip Sen (Eds.), IFAAMAS, 667–674.
- [8] Natasha Alechina, Mehdi Dastani, and Brian Logan. 2013. Reasoning about Normative Update. In *Proceedings of the 23rd IJCAI*, Francesca Rossi (Ed.), 20–26.
- [9] Natasha Alechina, Mehdi Dastani, and Brian Logan. 2018. Norm Specification and Verification in Multi-agent Systems. *Journal of Applied Logics* 5, 2 (2018), 457–491.
- [10] Natasha Alechina, Giuseppe De Giacomo, Brian Logan, and Giuseppe Perelli. 2022. Automatic Synthesis of Dynamic Norms for Multi-Agent Systems. In *Proceedings of the 19th KR*, Gabriele Kern-Isberner, Gerhard Lakemeyer, and Thomas Meyer (Eds.),
- [11] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. 2002. Alternating-time temporal logic. *J. ACM* 49 (2002), 672–713. <https://doi.org/10.1145/585265.585270>
- [12] Benjamin Aminof, Marta Kwiatkowska, Bastien Maubert, Aniello Murano, and Sasha Rubin. 2019. Probabilistic Strategy Logic. In *Proceedings of the 28th IJCAI*, Sarit Kraus (Ed.), 32–38. <https://doi.org/10.24963/IJCAI.2019/5>
- [13] Carlos Areces and Balder ten Cate. 2007. Hybrid logics. In *Handbook of Modal Logic*, Patrick Blackburn, Johan Van Benthem, and Frank Wolter (Eds.), Studies in Logic and Practical Reasoning, Vol. 3. Elsevier, 821–868.
- [14] Guillaume Aucher, Johan van Benthem, and Davide Grossi. 2018. Modal logics of sabotage revisited. *Journal of Logic and Computation* 28, 2 (2018), 269–303. <https://doi.org/10.1093/LOGCOM/EXX034>
- [15] Alexandru Baltag, Lawrence S. Moss, and Slawomir Solecki. 2022. Logics for Epistemic Actions: Completeness, Decidability, Expressivity. *CoRR* abs/2203.06744 (2022). <https://doi.org/10.48550/ARXIV.2203.06744>
- [16] Francesco Belardinelli, Wojciech Jamroga, Damian Kurpiewski, Vadim Malvone, and Aniello Murano. 2019. Strategy Logic with Simple Goals: Tractable Reasoning about Strategies. In *Proceedings of the 28th IJCAI*, Sarit Kraus (Ed.), 88–94. <https://doi.org/10.24963/IJCAI.2019/13>
- [17] Patrick Blackburn, Maarten de Rijke, and Yde Venema. 2001. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science, Vol. 53. CUP.
- [18] Patrick Blackburn and Miroslava Tzakova. 1999. Hybrid Languages and Temporal Logic. *Logic Journal of the IGPL* 7, 1 (1999), 27–54. <https://doi.org/10.1093/JIGPAL/7.1.27>
- [19] Laura Bozzelli and Ruggero Lanotte. 2010. Complexity and succinctness issues for linear-time hybrid logics. *Theoretical Computer Science* 411, 2 (2010), 454–469. <https://doi.org/10.1016/j.tcs.2009.08.009>
- [20] Laura Bozzelli, Aniello Murano, and Loredana Sorrentino. 2020. Alternating-time temporal logics with linear past. *Theoretical Computer Science* 813 (2020), 199–217.
- [21] Jan M. Broersen and Andreas Herzog. 2015. Using STIT Theory to Talk About Strategies. In *Models of Strategic Reasoning - Logics, Games, and Communities*, Johan van Benthem, Sujata Ghosh, and Rineke Verbrugge (Eds.), LNCS, Vol. 8972. Springer, 137–173. https://doi.org/10.1007/978-3-662-48540-8_5
- [22] Nils Bulling and Mehdi Dastani. 2016. Norm-based mechanism design. *Artificial Intelligence* 239 (2016), 97–142.
- [23] Davide Catta, Jean Leneutre, and Vadim Malvone. 2023. Obstruction Logic: A Strategic Temporal Logic to Reason About Dynamic Game Models. In *Proceedings of the 26th ECAI (Frontiers in Artificial Intelligence and Applications, Vol. 372)*, Kobi Gal, Ann Nowé, Grzegorz J. Nalepa, Roy Fairstein, and Roxana Radulescu (Eds.), IOS Press, 365–372. <https://doi.org/10.3233/FAIA230292>
- [24] Davide Catta, Jean Leneutre, Vadim Malvone, and Aniello Murano. 2024. Obstruction Alternating-time Temporal Logic: A Strategic Logic to Reason about Dynamic Models. In *Proceedings of the 23rd AAMAS*, Mehdi Dastani, Jaime Simão Sichman, Natasha Alechina, and Virginia Dignum (Eds.), IFAAMAS, 271–280. <https://doi.org/10.5555/3635637.3662875>
- [25] Edmund M. Clarke, Orna Grumberg, Daniel Kroening, Doron A. Peled, and Helmut Veith. 2018. *Model checking, 2nd Edition*. MIT Press.
- [26] Rodica Condurache, Emmanuel Filiot, Raffaella Gentilini, and Jean-François Raskin. 2016. The Complexity of Rational Synthesis. In *Proceedings of the 43rd ICALP (LIPIcs, Vol. 55)*, Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi (Eds.), Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 121:1–121:15. <https://doi.org/10.4230/LIPICS.ICALP.2016.121>
- [27] Rajdeep K. Dash, Nicholas R. Jennings, and David C. Parkes. 2003. Computational-Mechanism Design: A Call to Arms. *IEEE Intelligent Systems* 18, 6 (2003), 40–47. <https://doi.org/10.1109/MIS.2003.1249168>
- [28] Cristina David and Daniel Kroening. 2017. Program synthesis: challenges and opportunities. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 375, 2104 (2017), 20150403.
- [29] Tiago de Lima. 2014. Alternating-time temporal dynamic epistemic logic. *Journal of Logic and Computation* 24, 6 (2014), 1145–1178. <https://doi.org/10.1093/LOGCOM/EXS061>
- [30] Stéphane Demri and Morgan Deters. 2015. Separation logics and modalities: a survey. *J. Appl. Non Class. Logics* 25, 1 (2015), 50–99.
- [31] Stéphane Demri and Raul Fervari. 2019. The power of modal separation logics. *J. Log. Comput.* 29, 8 (2019), 1139–1184.
- [32] Dana Fisman, Orna Kupferman, and Yoav Lustig. 2010. Rational Synthesis. In *Proceedings of the 16th TACAS (LNCS, Vol. 6015)*, Javier Esparza and Rupak Majumdar (Eds.), Springer, 190–204. https://doi.org/10.1007/978-3-642-12002-2_16
- [33] Massimo Franceschet and Maarten de Rijke. 2006. Model checking hybrid logics (with an application to semistructured data). *Journal of Applied Logic* 4, 3 (2006), 279–304. <https://doi.org/10.1016/J.JAL.2005.06.010>
- [34] M. Franceschet, B. Schlingloff, and M. Rijke. 2003. Hybrid Logics on Linear Structures: Expressivity and Complexity. In *Proceedings of the 10th TIME*, 166–173. <https://doi.org/10.1109/TIME.2003.1214893>
- [35] Tim French, Wiebe van der Hoek, Petar Iliev, and Barteld P. Kooi. 2013. On the succinctness of some modal logics. *Artificial Intelligence* 197 (2013), 56–85. <https://doi.org/10.1016/J.ARTINT.2013.02.003>
- [36] Rustam Galimullin. 2021. Coalition and relativised group announcement logic. *Journal of Logic, Language and Information* 30, 3 (2021), 451–489. <https://doi.org/10.1007/s10849-020-09327-2>
- [37] Rustam Galimullin and Thomas Ágotnes. 2021. Dynamic Coalition Logic: Granting and Revoking Dictatorial Powers. In *Proceedings of the 8th LORI (LNCS, Vol. 13039)*, Sujata Ghosh and Thomas Icard (Eds.), Springer, 88–101. https://doi.org/10.1007/978-3-030-88708-7_7
- [38] Rustam Galimullin and Thomas Ágotnes. 2022. Action Models for Coalition Logic. In *Proceedings of the 4th DaLi (LNCS, Vol. 13780)*, Carlos Areces and Diana Costa (Eds.), Springer, 73–89. https://doi.org/10.1007/978-3-031-26622-5_5
- [39] Rustam Galimullin, Maksim Gladyshev, Munyque Mittelman, and Nima Motamed. 2025. Changing the Rules of the Game: Reasoning about Dynamic Phenomena in Multi-Agent Systems. *CoRR* abs/2502.11785 (2025). <https://doi.org/10.48550/arXiv.2502.11785>
- [40] Rustam Galimullin and Louwe B. Kuijter. 2024. Synthesizing Social Laws with ATL Conditions. In *Proceedings of the 23rd AAMAS*, Mehdi Dastani, Jaime Simão Sichman, Natasha Alechina, and Virginia Dignum (Eds.), IFAAMAS, 2270–2272. <https://doi.org/10.5555/3635637.3663130>
- [41] Rustam Galimullin and Fernando R. Velázquez-Quesada. 2024. Topic-Based Communication Between Agents. *Studia Logica* (2024). <https://doi.org/10.1007/s11225-024-10119-z>
- [42] Valentin Goranko. 2000. Temporal Logics with Reference Pointers and Computation Tree Logics. *Journal of Applied Non-Classical Logics* 10, 3-4 (2000), 221–242. <https://doi.org/10.1080/11663081.2000.10510998>
- [43] Valentin Goranko and Dmitry Shkatov. 2009. Tableau-based decision procedures for logics of strategic ability in multiagent systems. *ACM Transactions on Computational Logic* 11, 1 (2009), 3:1–3:51. <https://doi.org/10.1145/1614431.1614434>
- [44] Valentin Goranko and Govert van Drimmlen. 2006. Complete axiomatization and decidability of Alternating-time temporal logic. *Theoretical Computer Science* 353, 1-3 (2006), 93–117. <https://doi.org/10.1016/J.TCS.2005.07.043>
- [45] Julian Gutierrez, Muhammad Najib, Giuseppe Perelli, and Michael J. Wooldridge. 2023. On the complexity of rational verification. *Annals of Mathematics and Artificial Intelligence* 91, 4 (2023), 409–430.
- [46] Jens Ulrik Hansen. 2011. A Hybrid Public Announcement Logic with Distributed Knowledge. In *Int. Workshop on Hybrid Logic and Applications, HyLo 2010*.
- [47] John F. Horty. 2001. *Agency and deontic logic*. OUP.
- [48] Xiaowei Huang and Ron Van Der Meyden. 2018. An Epistemic Strategy Logic. *ACM Trans. Comput. Logic* 19, 4, Article 26 (nov 2018), 45 pages. <https://doi.org/10.1145/3233769>
- [49] David Hyland, Munyque Mittelman, Aniello Murano, Giuseppe Perelli, and Michael Wooldridge. 2024. Incentive Design for Rational Agents. In *Proceedings of the 21st KR*, Pierre Marquis, Magdalena Ortiz, and Maurice Pagnucco (Eds.), 464–474. <https://doi.org/10.24963/kr.2024/44>

- [50] Wojciech Jamroga. 2015. *Logical Methods for Specification and Verification of Multi-Agent Systems*. Institute of Computer Science, Polish Academy of Sciences.
- [51] Wojciech Jamroga and Aniello Murano. 2015. Module Checking of Strategic Ability. In *Proceedings of the 14th AAMAS*, Gerhard Weiss, Pinar Yolum, Rafael H. Bordini, and Edith Elkind (Eds.). ACM, 227–235.
- [52] Ahmet Kara, Volker Weber, Martin Lange, and Thomas Schwentick. 2009. *On the Hybrid Extension of CTL and CTL**. LNCS, Vol. 5734. Springer, 427–438. https://doi.org/10.1007/978-3-642-03816-7_37
- [53] Daniel Kernberger and Martin Lange. 2020. On the expressive power of hybrid branching-time logics. *Theoretical Computer Science* 813 (2020), 362–374. <https://doi.org/10.1016/j.tcs.2020.01.014>
- [54] Barteld Kooi. 2007. Expressivity and completeness for public update logics via reduction axioms. *Journal of Applied Non-Classical Logics* 17, 2 (2007), 231–253. <https://doi.org/10.3166/JANCL.17.231-253>
- [55] Louwe B. Kuijter. 2014. The Expressivity of Factual Change in Dynamic Epistemic Logic. *Review of Symbolic Logic* 7, 2 (2014), 208–221. <https://doi.org/10.1017/S1755020313000324>
- [56] Orna Kupferman, Moshe Y Vardi, and Pierre Wolper. 2001. Module checking. *Information and Computation* 164, 2 (2001), 322–344.
- [57] Damian Kurpiewski, Wojciech Jamroga, and Michal Knapik. 2019. STV: Model Checking for Strategies under Imperfect Information. In *Proceedings of the 18th AAMAS*, Edith Elkind, Manuela Veloso, Noa Agmon, and Matthew E. Taylor (Eds.). IFAAMAS, 2372–2374.
- [58] Martin Lange. 2009. Model Checking for Hybrid Logic. *Journal of Logic, Language and Information* 18, 4 (2009), 465–491. <https://doi.org/10.1007/s10849-009-9088-7>
- [59] Francois Laroussinie, Nicolas Markey, and Ghassan Oreiby. 2008. On the Expressiveness and Complexity of ATL. *Logical Methods in Computer Science* Volume 4, Issue 2 (May 2008).
- [60] Alessio Lomuscio, Hongyang Qu, and Franco Raimondi. 2017. MCMAS: an open-source model checker for the verification of multi-agent systems. *International Journal on Software Tools for Technology Transfer* 19, 1 (2017), 9–30. <https://doi.org/10.1007/s10009-015-0378-x>
- [61] Carsten Lutz. 2006. Complexity and succinctness of public announcement logic. In *Proceedings of the 5th AAMAS*, Hideyuki Nakashima, Michael P. Wellman, Gerhard Weiss, and Peter Stone (Eds.). ACM, 137–143. <https://doi.org/10.1145/1160633.1160657>
- [62] Bastien Maubert, Sophie Pinchinat, François Schwarzentruber, and Silvia Stranieri. 2020. Concurrent Games in Dynamic Epistemic Logic. In *Proceedings of the 29th IJCAI*, Christian Bessière (Ed.). 1877–1883. <https://doi.org/10.24963/IJCAI.2020/260>
- [63] Munyque Mittelmann, Bastien Maubert, Aniello Murano, and Laurent Perrussel. 2023. Formal Verification of Bayesian Mechanisms. In *Proceedings of the 37th AAAI*, Brian Williams, Yiling Chen, and Jennifer Neville (Eds.). AAAI Press, 11621–11629. <https://doi.org/10.1609/AAAI.V37I10.26373>
- [64] Munyque Mittelmann, Bastien Maubert, Aniello Murano, and Laurent Perrussel. 2025. Formal verification and synthesis of mechanisms for social choice. *Artif. Intell.* 339 (2025), 104272.
- [65] Fabio Mogavero, Aniello Murano, Giuseppe Perelli, and Moshe Y. Vardi. 2014. Reasoning About Strategies: On the Model-Checking Problem. *ACM Trans. Comput. Log.* 15, 4 (2014), 34:1–34:47.
- [66] Fabio Mogavero, Aniello Murano, and Moshe Y. Vardi. 2010. Reasoning About Strategies. In *Proceedings of the 30th FSTTCS (LIPIcs, Vol. 8)*, Kamal Lodaya and Meena Mahajan (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 133–144. <https://doi.org/10.4230/LIPICS.FSTTCS.2010.133>
- [67] Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani (Eds.). 2007. *Algorithmic Game Theory*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511800481>
- [68] Marc Pauly. 2002. A modal logic for coalitional power in games. *Journal of Logic and Computation* 12, 1 (2002), 149–166. <https://doi.org/10.1093/logcom/12.1.149>
- [69] Mina Young Pedersen. 2024. *Malicious Agents and the Power of the Few : On the Logic of Abnormality in Social Networks*. Ph.D. Dissertation. University of Bergen, Norway. <https://bora.uib.no/bora-xmlui/handle/11250/3151733>
- [70] Steve Phelps, Peter McBurney, and Simon Parsons. 2010. Evolutionary mechanism design: a review. *Autonomous agents and multi-agent systems* 21, 2 (2010), 237–264.
- [71] John C. Reynolds. 2002. Separation Logic: A Logic for Shared Mutable Data Structures. In *LICS*. IEEE Computer Society, 55–74.
- [72] Olivier Roy. 2009. A dynamic-epistemic hybrid logic for intentions and information changes in strategic games. *Synthese* 171, 2 (2009), 291–320. <https://doi.org/10.1007/S11229-009-9644-3>
- [73] Pierre-Yves Schobbens. 2003. Alternating-time logic with imperfect recall. In *Proceedings of the 1st LCMAS (ENTCS, Vol. 85)*, Wiebe van der Hoek, Alessio Lomuscio, Erik P. de Vink, and Michael J. Wooldridge (Eds.). Elsevier, 82–93. [https://doi.org/10.1016/S1571-0661\(05\)82604-0](https://doi.org/10.1016/S1571-0661(05)82604-0)
- [74] Johan van Benthem. 2005. An Essay on Sabotage and Obstruction. In *Mechanizing Mathematical Reasoning (LNCS, Vol. 2605)*, Dieter Hutter and Werner Stephan (Eds.). Springer, 268–276. https://doi.org/10.1007/978-3-540-32254-2_16
- [75] Johan van Benthem, Lei Li, Chenwei Shi, and Haoxuan Yin. 2023. Hybrid sabotage modal logic. *Journal of Logic and Computation* 33, 6 (2023), 1216–1242. <https://doi.org/10.1093/LOGCOM/EXAC006>
- [76] Wiebe van der Hoek, Mark Roberts, and Michael Wooldridge. 2007. Social laws in alternating time: effectiveness, feasibility and synthesis. *Synthese* 156 (2007), 1–19.
- [77] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. 2008. *Dynamic Epistemic Logic*. Synthese Library, Vol. 337. Springer.
- [78] Dirk Walther, Carsten Lutz, Frank Wolter, and Michael J. Wooldridge. 2006. ATL Satisfiability is Indeed EXPTIME-complete. *Journal of Logic and Computation* 16, 6 (2006), 765–787. <https://doi.org/10.1093/LOGCOM/EXL009>
- [79] Yanjing Wang and Jeremy Seligman. 2018. When Names Are Not Commonly Known: Epistemic Logic with Assignments. In *Proceedings of the 12th AiML*, Guram Bezhaniashvili, Giovanna D’Agostino, George Metcalfe, and Thomas Studer (Eds.). College Publications, 611–628.