Verifying Strategic Abilities in Multi-agent Systems with Private-Data Sharing

Extended Abstract

F. Belardinelli
Imperial College London and Université d'Evry
francesco.belardinelli@imperial.ac.uk

I. Boureanu
University of Surrey
i.boureanu@surrey.ac.uk

C. Dima
Université Paris-Est Créteil
dima@u-pec.fr

V. Malvone
Université d'Evry
vadim.malvone@univ-evry.fr

ACM Reference Format:
F. Belardinelli, I. Boureanu, C. Dima, and V. Malvone. 2019. Verifying Strategic Abilities in Multi-agent Systems with Private-Data Sharing. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 3 pages.

1 INTRODUCTION

Most formalisms for multi-agent systems (MAS) are not adept at explicitly expressing of data sharing between agents. Yet, disclosure and hiding of data amongst agents impacts on their strategic abilities, and so has a strong baring on non-classical logics that formally capture agents' coalitions, e.g., Alternating-time Temporal Logic (ATL) [1].

To this end, we devise *concurrent game structures with propositional control for atom-visibility (*vCGS*)*. In vCGS, agents a and b have an explicit endowment to see some of each others' variables, without other agents partaking in this. Second, we ascertain that the model checking problem for ATL with imperfect information and perfect recall on vCGSis undecidable. Third, we put forward a methodology to model check a formula φ in ATL^* on a vCGS M, by verifying a suitable translation of φ in a submodel of M.

2 AGENTS WITH VISIBILITY-CONTROL

We introduce a class of systems where each agent can change the truth value of atoms she controls, and can make them (in)visible to other agents. On the underlying concurrent game structure, we interpret ATL*. We consider finite sets Ag and AP of agents and atoms.

Definition 2.1 (Visibility Atom). Given atom $v \in AP$ and agent $a \in Ag$, $vis(v,a)$ denotes a *visibility atom* expressing intuitively that the value of v is visible to a. By VA, we denote the set of all *visibility atoms* $vis(v,a)$, for $v \in AP$ and $a \in Ag$. By $\mathsf{VA}_a = \{vis(v,a) \in \mathsf{VA} \mid v \in AP\}$, we denote the set of visibility atoms for agent a.

Definition 2.2 (Visibility-Controlling Agents: Syntax). Given set AP of atoms, an *agent* a is a tuple $a = \langle AP, V_a, GC_a \rangle$ such that

- $V_a \subseteq AP$ is the set of atoms controlled by agent a;
- GC_a is a finite set of *guarded commands*, which are of the form:

$$\gamma ::= \quad \varphi \rightsquigarrow \quad v_1 := \mathsf{tv}, \ldots, v_k := \mathsf{tv},$$
$$vis(v_{k+1}, a_1) := \mathsf{tv} \ldots, vis(v_{k+m}, a_m) := \mathsf{tv}$$

where each $v_i \in V_a$ is an atom controlled by a that occurs at most once, *guard* φ is a boolean formula over $AP \cup \mathsf{VA}_a$, all $a_1, \ldots, a_m$ are agents in Ag different from a, and tv is a truth value in $\{\mathtt{tt}, \mathtt{ff}\}$.

We denote with $\mathrm{g}(\gamma)$ and $asg(\gamma)$ the guard and assignment, respectively. Moreover, guarded commands can be of two disjoint types: `init`- or `update`-*type*. In the former, the guard is always equal to tt (and thus omitted) and the assignment contains $vis(v,a) ::= \mathtt{tt}$ for every atom $v \in V_a$ (i.e., an agent always has visibility of the atoms she controls).

Intuitively, Def. 2.2 says that (1) every agent a can change the value of the atoms in $V_a \subseteq AP$ through assignments $v := \mathsf{tv}$; (2) agent a can switch the visibility for some other agent a_i over some of a's atoms, by means of assignments $vis(v,a_i) := \mathsf{tv}$; (this is unlike [2, 18]); (3) since $a_1, \ldots, a_m$ are required to be different from a, agent a cannot remove visibility of her own atoms.

Hereafter, we assume that control is *exclusive*: for any two distinct agents a and b, $V_a \cap V_b = \emptyset$, i.e., the sets of controlled atoms are disjoint. Since control is exclusive, we often talk about the *owner* $own(v) \in Ag$ of an atom $v \in AP$.

Definition 2.3 (Visibility-Controlling Agents: Semantics). Given a set Ag of agents as in Def. 2.2, all defined on set AP of atoms, an *iCGS with propositional control for atom-visibility (*vCGS*)* is a tuple $\mathsf{M} = \langle Ag, AP, \{Act_a\}_{a \in Ag}, S, S_0, P, \tau, \{\sim_a\}_{a \in Ag}, \pi \rangle$ where:

- For every $a \in Ag$, $Act_a = GC_a$.
- $S = \{s \subseteq AP \cup \mathsf{VA} \mid \text{for every } a \in Ag, v \in V_a, vis(v,a) \in s\}$ is the *set of states.* For $s \in S$ and $a \in Ag$, $Vis(s,a) = \{v \in AP \mid vis(v,a) \in s\}$ is the set of atoms visible to a in state s. By def., $V_a \subseteq Vis(s,a)$ for every $s \in S$.
- $S_0 \subseteq S$ is the set of states $s_0 \in S$ such that for every $a \in Ag$, for some $\gamma \in \mathtt{init}(Act_a)$, for every $v \in V_a$, if $v := \mathtt{tt}$ occurs in $asg(\gamma)$, then $v \in s_0$; and if $v := \mathtt{ff}$ occurs in $asg(\gamma)$, then $v \notin s_0$. That is, atoms are initialised as either true or false only via an `init` command.
- For every state $s \in S$ and agent $a \in Ag$, the *protocol function* $P : S \times Ag \rightarrow 2^{\bigcup_{a \in Ag} Act_a}$ returns the set $P(s,a)$ of `update`-commands γ such that $AP(\mathrm{g}(\gamma)) \subseteq Vis(s,a)$ and $s \models \mathrm{g}(\gamma)$, where $AP(\phi)$ is the set of atoms occurring in formula ϕ. That is, all atoms appearing in the guard are visible to the agent and the guard is indeed true.
- The *transition function* $\tau : S \times ACT \rightarrow S$ is such that a transition $\tau(s, (\gamma_1, \ldots, \gamma_n)) = s'$ holds iff (1) for every $a \in Ag$, $\gamma_a \in P(s,a)$; (2) for every $v \in AP$ and $own(v) \in Ag$, $v \in s'$ iff either $asg(\gamma_{own(v)})$ contains an assignment $v := \mathtt{tt}$ or $v \in s$; whereas $v \notin s'$ iff either $asg(\gamma_{own(v)})$ contains an assignment $v := \mathtt{ff}$ or $v \notin s$; (3) for every $v \in AP$ and $own(v) \in Ag$, $vis(v,a) \in s'$ iff either $asg(\gamma_{own(v)})$ contains an assignment $vis(v,a) := \mathtt{tt}$ or $vis(v,a) \in s$; whereas $vis(v,a) \notin s'$ if either $asg(\gamma_{own(v)})$ contains an assignment $vis(v,a) := \mathtt{ff}$ or $vis(v,a) \notin s$.
- Let the set $R \subseteq S$ of *reachable states* be the transitive closure of S_0 under the transition function τ. The *indistinguishability relation*

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

is defined so that for every $s,s' \in R$, $s \sim_a s'$ iff $Vis(s,a) = Vis(s',a)$ and for every $v \in Vis(s,a) = Vis(s',a)$, $v \in s$ iff $v \in s'$; whereas for states in $S \backslash R$, each $\sim_a$ is the identity relation. We can easily check that if $s \sim_a s'$ then $P(s,a) = P(s',a)$.

- The labelling function $\pi : S \rightarrow 2^{AP}$ is the identity, i.e., each state is named with the atoms belonging to it.

ATL Syntax. State (φ) and path (ψ) formulas in ATL^* are defined as follows, where $q \in AP$ and $A \subseteq Ag$: $\varphi ::= q \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\langle A \rangle\rangle \psi$; $\psi ::= \varphi \mid \neg\psi \mid \psi \wedge \psi \mid X\psi \mid (\psi U \psi)$. Formulas in the ATL fragment of ATL^* are obtained by restricting path formulas ψ as follows, where φ is a state formula: $\psi ::= X\varphi \mid (\varphi U \varphi) \mid (\varphi R \varphi)$.

ATL Semantics. Given a vCGS M, a *path* p is a sequence $s_1 s_2 \ldots$ of states such that for every $i > 1$ there exists a joint action $\vec{\alpha} \in ACT$ such that $\tau(s_i, \vec{\alpha}) = s_{i+1}$. A finite path $h \in S_0 \cdot S^*$ starting in an initial state is called a *history*. Hereafter, we extend the indistinguishability relation $\sim_a$ to histories in $S_0 \cdot S^*$ in a synchronous and pointwise manner, that is, $h \sim_a h'$ iff $|h| = |h'|$ and for every $i \leq |h|$, $h_i \sim_a h'_i$. A *uniform, memoryful strategy* for agent $a \in Ag$ is a function $f_a : S_0 \cdot S^* \rightarrow Act_a$ such that for all histories $h,h' \in S_0 \cdot S^*$, (i) $f_a(h) \in P(last(h),a)$; and (ii) if $h \sim_a h'$ then $f_a(h) = f_a(h')$. Given a *joint strategy* $F_A = \{f_a \mid a \in A\}$ for coalition $A \subseteq Ag$, and history $h \in S_0 \cdot S^*$, let $out(h, F_A)$ be the set of all infinite paths p starting from history h and compatible with F_A. More formally, we set $out(h,F_A) = \{p \mid p_{\leq |h|} = h$ and for all $i \geq |h|, p_{i+1} = \tau(p_i, \vec{\alpha})$, where for all $a \in A, \alpha_a = f_a(p_{\leq i})\}$.

The satisfaction relation $\models$ for a vCGS M, path p, index $i \in \mathbb{N}$, and ATL^* formula ϕ is defined as follows (clauses for Boolean operators and temporal operators are immediate and thus omitted): (1) $(M,p,i) \models q$ iff $q \in \pi(p_i)$; (2) $(M,p,i) \models \langle\langle A \rangle\rangle \psi$ iff for some joint strategy F_A, for all paths $p' \in out(p_{\leq i}, F_A)$, $(M,p') \models \psi$; (3) $(M,p) \models \varphi$ iff $(M,p,1) \models \varphi$.

A formula φ is *satisfied* by a vCGS M, or $M \models \varphi$, iff for all paths p starting in an initial state, $(M,p,1) \models \varphi$. Note that we adopt the *objective* interpretation of ATL^* [15], whereby strategy operator $\langle\langle A \rangle\rangle$ is evaluated against all paths $p \in out(h, F_A)$ starting from the present history h. The *model-checking problem* for vCGS asks for checking whether a given vCGS M satisfies a given ATL^* formula φ.

We state the main result of this section.

Theorem 2.4. *The model checking problem for ATL^* (resp. ATL) on* vCGS *is undecidable.*

Theorem 2.4 is proved by showing that model checking ATL on standard iCGS [15], which is known to be undecidable [10], is PTIME-reducible to the same problem on vCGS. A proof can be found in [3].

3 FORMULA-BASED MODEL REDUCTION

Now, we put forward a methodology to model check a formula φ in ATL^* on a vCGS M, by verifying a suitable translation of φ in a submodel of M. This reduction leads in general to a smaller state space and a less complex model checking instance. Under specific circumstances it might lead to decidable model checking.

Definition 3.1. Given a vCGS M and formula φ, we define by mutual recursion the sets $\Delta \subseteq AP$ of atoms and $\Gamma \subseteq Ag$ of agents:

$$\begin{aligned}
\Delta_0 &= AP(\varphi) \qquad \Gamma_0 = Own(\Delta_0) \\
\Delta_{n+1} &= \Delta_n \cup \{AP(g(\gamma)) \mid \text{some } v \in \Delta_n \text{ appears in } asg(\gamma)\} \cup \\
&\quad \{v \in AP \mid vis(v,b) \text{ appears in } \gamma \text{ for some } b \in \Gamma_n, \gamma \in Act\} \\
&\quad \cup \{v \in AP \mid vis(v,b) \in g(\gamma) \text{ for some } v' \in \Delta_n \text{ in } asg(\gamma)\} \\
\Gamma_{n+1} &= \Gamma_n \cup Own(\Delta_n)
\end{aligned}$$

where we recall that $AP(\phi) \subseteq AP$ is the set of atoms appearing in formula ϕ, and we take $Own(\Delta_i)$ to be $\{own(v) \mid v \in \Delta_i\} \subseteq Ag$. Then, let $\Delta = \bigcup_{n \in \mathbb{N}} \Delta_n \subseteq AP$ and $\Gamma = \bigcup_{n \in \mathbb{N}} \Gamma_n \subseteq Ag$.

Intuitively, Δ is the set of all atoms that are relevant to determine the atoms in formula φ, the visibility of agents in Γ, or of action guards that influence the truth of atoms in Δ (including $AP(\varphi)$); whereas Γ is the set of owners of the atoms in Δ.

Definition 3.2. Given Δ and Γ as in Def.3.1 and an agent $a \in \Gamma$, we define a new agent $a' = (\Delta, V'_a, GC'_a)$ such that
(1) $V'_a = V_a \cap \Delta$;
(2) $GC'_a = \{\gamma' ::= g(\gamma) \rightsquigarrow asg(\gamma)_{|\Delta,\Gamma} \mid \gamma \in GC_a, AP(g(\gamma))$
$\cup \{v \in AP \mid vis(v,a) \text{ appears in } g(\gamma)\} \subseteq \Delta\}$;
where $asg(\gamma)_{|\Delta,\Gamma}$ is the restriction of $asg(\gamma)$ to Δ and Γ.

We now state the main result of this section. First of all, given a model M, we write $M_{|\Delta,\Gamma}$ for the model generated by using agents as in Def.3.2 and restricted over Δ. Further, given a formula φ, we write $\varphi_{|\Gamma}$ for the formula generated by substituting every sub-formula $\langle\langle A \rangle\rangle \psi$ of φ with $\langle\langle A \cap \Gamma \rangle\rangle \psi$. Also, given a path p, $p_{|\Delta,\Gamma} = (p_1)_{|\Delta,\Gamma}, (p_2)_{|\Delta,\Gamma}, \ldots$ is the component-wise restriction of p to Δ and Γ.

Theorem 3.3. *Given a* vCGS M *and a formula φ, we have that*

$$(M,p,i) \models \varphi \quad \textit{iff} \quad (M_{|\Delta,\Gamma}, p_{|\Delta,\Gamma}, i) \models \varphi_{|\Gamma}$$

By Theorem 3.3 we obtain the following corollary.

Corollary 3.4. *Given a* vCGS M *with reduction* $M_{|\Delta,\Gamma}$, *and a formula* $\langle\langle A \rangle\rangle \varphi$, *if* $\Gamma \subseteq A$ *then it is decidable whether* $M \models \varphi$.

4 RELATED WORK AND CONCLUSIONS

Coalition logic based on propositional control [20] has been extended with transfer of control in [4, 11, 12, 19]. Yet, these mainly deal with *coalition logic*, which is the "next" $\langle\langle A \rangle\rangle X$ fragment of ATL, while assuming perfect information. Meanwhile, we analyse the case of imperfect information and for the whole of ATL^*. This is also unlike verification of just game-theoretic equilibria under imperfect information [13, 14], in reactive modules with guarded commands [2]. Similarly, other works miss the strategic-ability edge that we have, yet they focus on more expressivity, e.g., at the epistemic level. This is the case of [16] where (propositional) visibility is also analysed but employing modal operators for visibility. Only limited type of strategic reasoning, stemming from no local "implementation" of actions, is also offered by another semantics similar to ours, i.e., dynamic epistemic logic and epistemic planning [21]. Finally, various restrictions on iCGS to retain decidability of model checking under imperfect information and perfect recall has been recently explored in [5–9, 17].

We put forward a formalism for the explicit expression of private-data sharing in multi-agent systems. On these "MAS with 1-to-1 private-channels", we ascertain that the model checking problem for Alternating-time Temporal Logic under imperfect information and perfect recall is, as expected, undecidable. Yet, we put forward a methodology to model check a formula φ in ATL^* on a vCGS M, by verifying a suitable translation of φ in a submodel of M. As future work, we aim to find general classes of vCGS for which the model checking of ATL becomes decidable, and show-case vCGS in modelling ICT problems.

Acknowledgements. F. Belardinelli acknowledges the support of ANR JCJC Project SVeDaS (ANR-16-CE40-0021).

REFERENCES

[1] R. Alur, T.A. Henzinger, and O. Kupferman. 2002. Alternating-time temporal logic. *J. ACM* 49, 5 (2002), 672–713.

[2] R. Alur and T. A. Henzinger. 1999. Reactive Modules. *Formal Methods in System Design* 15, 1 (1999), 7–48. https://doi.org/10.1023/A:1008739929481

[3] F. Belardinelli, I. Boureanu, C. Dima, and V. Malvone. 2019. *Model Checking ATL* on vCGS.* https://arxiv.org/abs/1903.04350

[4] F. Belardinelli, U. Grandi, A. Herzig, D. Longin, E. Lorini, A. Novaro, and L. Perrussel. 2017. Relaxing Exclusive Control in Boolean Games. In *TARK17*. 43–56.

[5] F. Belardinelli, A. Lomuscio, A. Murano, and S. Rubin. 2017. Verification of Broadcasting Multi-Agent Systems against an Epistemic Strategy Logic. In *IJCAI 2017*, ijcai.org, 91–97. https://doi.org/10.24963/ijcai.2017/14

[6] F. Belardinelli, A. Lomuscio, A. Murano, and S. Rubin. 2017. Verification of Multi-agent Systems with Imperfect Information and Public Actions. In *AAMAS 17*. 1268–1276.

[7] F. Belardinelli, A. Lomuscio, A. Murano, and S. Rubin. 2018. Decidable Verification of Multi-agent Systems with Bounded Private Actions. In *AAMAS 2018*. 1865–1867.

[8] R. Berthon, B. Maubert, and A. Murano. 2017. Decidability Results for ATL* with Imperfect Information and Perfect Recall. In *AAMAS 2017*, ACM, 1250–1258. http://dl.acm.org/citation.cfm?id=3091299

[9] R. Berthon, B. Maubert, A. Murano, S. Rubin, and M. Y. Vardi. 2017. Strategy logic with imperfect information. In *LICS 2017*. IEEE Computer Society, 1–12. https://doi.org/10.1109/LICS.2017.8005136

[10] C. Dima and F.L. Tiplea. 2011. Model-checking ATL under Imperfect Information and Perfect Recall Semantics is Undecidable. *CoRR* abs/1102.4225 (2011). http://arxiv.org/abs/1102.4225

[11] J. Gerbrandy. 2006. Logics of propositional control. In *AAMAS 2006*. 193–200. https://doi.org/10.1145/1160633.1160664

[12] D. Grossi, E. Lorini, and F. Schwarzentruber. 2015. The Ceteris Paribus Structure of Logics of Game Forms. *JAIR* 53 (2015), 91–126. https://doi.org/10.1613/jair.4666

[13] J. Gutierrez, P. Harrenstein, and M. Wooldridge. 2017. From model checking to equilibrium checking: Reactive modules for rational verification. *Artif. Intell.* 248 (2017), 123–157.

[14] J. Gutierrez, G. Perelli, and M. Wooldridge. 2016. Imperfect Information in Reactive Modules Games. In *KR'15*. 390–400. http://www.aaai.org/ocs/index.php/KR/KR16/paper/view/12848

[15] W. Jamroga and W. van der Hoek. 2004. Agents that Know How to Play. *Fund. Inf.* 62 (2004), 1–35.

[16] F. Maffre. 2016. *Ignorance is bliss : observability-based dynamic epistemic logics and their applications*. Ph.D. Dissertation. Université Paul Sabatier - Toulouse III.

[17] V. Malvone, A. Murano, and L. Sorrentino. 2017. Hiding Actions in Multi-Player Games. In *AAMAS 2017*. ACM Press, 1205–1213.

[18] W. van der Hoek, A. Lomuscio, and M. Wooldridge. 2006. On the complexity of practical ATL model checking knowledge, strategies, and games in multi-agent systems. In *AAMAS 2006*. ACM Press, 201–208.

[19] W. van der Hoek, D. Walther, and M. Wooldridge. 2010. Reasoning About the Transfer of Control. *JAIR* 37 (2010), 437–477. https://doi.org/10.1613/jair.2901

[20] W. van der Hoek and M. Wooldridge. 2005. On the logic of cooperation and propositional control. *Artificial Intelligence* 164, 1-2 (2005), 81–119. https://doi.org/10.1016/j.artint.2005.01.003

[21] H. van Ditmarsch, W. van der Hoek, and B. Kooi. 2007. *Dynamic Epistemic Logic*. Synthese Library, Vol. 337. Springer.