

gc decides, and notifies oc , whether the conference will be run on site or online. In the former case, the epidemiologic risk is much higher, indicated by the atomic proposition $epid$. The AMAS and its model are shown in Figure 1.

3 REASONING ABOUT ABILITIES: ATL*

Let \mathcal{PV} be a set of propositions and $\mathbb{A}gt$ the set of all agents. The syntax of *alternating-time temporal logic* ATL* [4, 28] is defined as:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\langle A \rangle\rangle\gamma, \quad \gamma ::= \varphi \mid \neg\gamma \mid \gamma \wedge \gamma \mid X\gamma \mid \gamma U\gamma,$$

where $p \in \mathcal{PV}$, $A \subseteq \mathbb{A}gt$, X stands for “next”, and U for “strong until” ($\gamma_1 U \gamma_2$ denotes that γ_1 holds until γ_2 becomes true). The other operators, Boolean connectives, \top , and \perp are defined as usual.

A *positional imperfect information (ir) strategy* for i is a function $\sigma_i: L_i \rightarrow Evt_i$ such that $\sigma_i(l) \in R_i(l)$ for each $l \in L_i$ [28]. The set of such strategies is denoted by Σ_i^{ir} . Collective strategies Σ_A^{ir} for $A \subseteq \mathbb{A}gt$ are defined as usual. By $\sigma_A(g) = (\sigma_1(g), \dots, \sigma_m(g))$, we denote the tuple of selections of coalition $A = (1, \dots, m)$ at state g . An infinite sequence of global states and events $\pi = g_0\alpha_0g_1\alpha_1g_2\dots$ is called a *path* if $g_j \xrightarrow{\alpha_j} g_{j+1}$ for every $j \geq 0$. The set of all paths in model M starting at state g is denoted by $\Pi_M(g)$.

The *outcome* of strategy $\sigma_A \in \Sigma_A^{ir}$ in state $g \in St$ is the set $out_M(g, \sigma_A) \subseteq \Pi_M(g)$ such that $\pi = g_0\alpha_0g_1\alpha_1g_2\dots \in out_M(g, \sigma_A)$ iff $g_0 = g$, and $\forall j \geq 0 \alpha_j \in enabled(\pi[j], \sigma_A(\pi[j]))$.

A path π satisfies *concurrency-fairness (CF)* if there is no event α enabled in all states of π from $\pi[n]$ on and such that for every α_j actually executed in $\pi[j]$, $j = n, n+1, \dots$, we have $Agent(\alpha) \cap Agent(\alpha_j) = \emptyset$ [14]. Let $\Pi_M^{CF}(g)$ be the set of all such paths starting at g and $out_M^{CF}(g, \sigma_A) = out_M(g, \sigma_A) \cap \Pi_M^{CF}(g)$. The ir-semantics of ATL* [14] in asynchronous MAS is defined by the clause:

$$M, g \models_{ir} \langle\langle A \rangle\rangle\gamma \text{ iff there is a strategy } \sigma_A \in \Sigma_A^{ir} \text{ s.t. } out_M(g, \sigma_A) \neq \emptyset \text{ and } \forall \pi \in out_M(g, \sigma_A) \text{ we have } M, \pi \models_{ir} \gamma.$$

Moreover, the *concurrency-fair semantics* \models_{ir}^{CF} is obtained by replacing $out_M(g, \sigma_A)$ with $out_M^{CF}(g, \sigma_A)$ in the above clause.

Example 3.1. Clearly, formula $\langle\langle gc, oc \rangle\rangle G \neg epid$ holds in the conference model M_{conf} , in both \models_{ir} and \models_{ir}^{CF} semantics. To see that, fix $\sigma_{gc}(1) = online = \sigma_{oc}(0)$ in the collective strategy of $\{gc, oc\}$.

4 SEMANTIC PROBLEMS

We describe two kinds of problematic phenomena that follow from adding the concept of strategic ability to representations and models derived from concurrency theory, the way it was defined in [14].

4.1 Deadlock Strategies and Finite Paths

An automata network is typically required to produce no deadlock states. In case of AMAS, the situation is more delicate. Even if the AMAS as a whole produces no deadlocks, some strategies might, which makes the interpretation of strategic modalities cumbersome. We illustrate this on the following example.

Example 4.1. Consider the 3-agent AMAS and its model M_{conf} , which are depicted in Figure 1. Clearly, M_{conf} has no deadlock states. Let us now look at the collective strategies of coalition $\{gc, oc\}$, with agent sc serving as the opponent. It is easy to see that the coalition has no way to prevent the opening of the conference, i.e., it cannot

prevent the system from reaching state 101. However, the strategy depicted in Figure 1 produces only one *infinite* path, namely $(000 \text{ giveup } 002 \text{ giveup } \dots)$. Since the semantics in Section 3 disregards finite paths, we get $M_{conf}, 000 \models_{ir} \langle\langle gc, oc \rangle\rangle G \neg open$ and $M_{conf}, 000 \models_{ir}^{CF} \langle\langle gc, oc \rangle\rangle G \neg open$, which is counterintuitive.

Things can get even trickier. For the ir-semantics, it may happen that the outcomes of some (or even all) strategies of a coalition are empty, which leads to situations where the intuitive meaning of a strategic formula differs significantly from its formal interpretation.

Example 4.2. Let us add the transition $0 \xrightarrow{proceed} 0$ in agent oc , and remove the transitions labeled with *giveup* in agent sc . The resulting model M'_{conf} has no deadlock states, yet all the joint strategies of $\{gc, oc\}$ produce only finite runs. Since finite paths are not included in the outcome sets, and the semantics in Section 3 rules out strategies with empty outcomes, we get that $\neg \langle\langle gc, oc \rangle\rangle F \top$, which seems definitely wrong.

Notice that removing the non-emptiness requirement from the semantic clause in Section 3 does not help. In that case, any joint strategy of $\{gc, oc\}$ could be used to demonstrate that $\langle\langle gc, oc \rangle\rangle G \perp$.

4.2 Strategies in Asymmetric Interaction

In this section, we point out that AMAS is too restricted to model the strategic aspects of asymmetric synchronization (e.g., a sender sending a message to a receiver) in a natural way.

Example 4.3. Consider the global state 101 of the conference model M_{conf} , i.e., the state where it has just been decided to proceed with the conference. In that state, we have $\langle\langle gc \rangle\rangle G \neg epid$, meaning that the GC chair can make sure that the epidemic risk is always low. This is achieved by gc selecting *online* at its local state 1. Then, the next transition can be obtained only if the oc module synchronizes with gc on event *online*. On the other hand, we also have that $\langle\langle oc \rangle\rangle F epid$ holds at $M_{conf}, 101$, which is obtained by the OC’s strategy selecting *onsite* at 0. That is rather odd, in particular it violates the standard postulate of superadditivity [24].

The problem arises because the repertoire functions in AMAS are based on the assumption that the agent can choose any single event in $R_i(l_i)$. This does not allow for a natural specification of the situation when the transition is determined by another agent.

5 CONCLUSIONS

In this paper, we reconsider the asynchronous semantics of strategic ability for multi-agent systems, proposed recently in [14]. We show that adding strategic reasoning on top of the modeling machinery, inherited from distributed systems, leads to counterintuitive interpretation of formulas. We identify two main sources of problems. First, the execution semantics does not handle reasoning about deadlock-inducing strategies well. Secondly, the class of representations lacks constructions to resolve the tension between the asymmetry imposed by strategic operators on the one hand, and the asymmetry of interaction, e.g. between communicating parties.

ACKNOWLEDGMENTS

We acknowledge the support of NCBR Poland and FNR Luxembourg, under the PolLux/FNR-CORE project STV (POLLUX-VII/1/2019).

REFERENCES

- [1] R. Alur, L. de Alfaro, R. Grossu, T.A. Henzinger, M. Kang, C.M. Kirsch, R. Majumdar, F.Y.C. Mang, and B.-Y. Wang. 2001. jMocha: A Model-Checking Tool that Exploits Design Structure. In *Proceedings of ICSE*. IEEE Computer Society Press, 835–836.
- [2] R. Alur, T. Henzinger, F. Mang, S. Qadeer, S. Rajamani, and S. Tasiran. 1998. MOCHA: Modularity in Model Checking. In *Proceedings of CAV (Lecture Notes in Computer Science, Vol. 1427)*. Springer, 521–525.
- [3] R. Alur, T. A. Henzinger, and O. Kupferman. 1997. Alternating-Time Temporal Logic. In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE Computer Society Press, 100–109.
- [4] R. Alur, T. A. Henzinger, and O. Kupferman. 2002. Alternating-Time Temporal Logic. *J. ACM* 49 (2002), 672–713.
- [5] F. Belardinelli, A. Lomuscio, A. Murano, and S. Rubin. 2017. Verification of Broadcasting Multi-Agent Systems against an Epistemic Strategy Logic. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*. 91–97.
- [6] F. Belardinelli, A. Lomuscio, A. Murano, and S. Rubin. 2017. Verification of Multi-agent Systems with Imperfect Information and Public Actions. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2017, São Paulo, Brazil, May 8-12, 2017*. 1268–1276.
- [7] S. Busard, C. Pecheur, H. Qu, and F. Raimondi. 2014. Improving the Model Checking of Strategies under Partial Observability and Fairness Constraints. In *Formal Methods and Software Engineering*. Lecture Notes in Computer Science, Vol. 8829. Springer, 27–42.
- [8] P. Cermák, A. Lomuscio, F. Mogavero, and A. Murano. 2014. MCMAS-SLK: A Model Checker for the Verification of Strategy Logic Specifications. In *Proc. of CAV'14 (Lecture Notes in Computer Science, Vol. 8559)*. Springer, 525–532.
- [9] P. Cermák, A. Lomuscio, and A. Murano. 2015. Verifying and Synthesising Multi-Agent Systems against One-Goal Strategy Logic Specifications. In *Proceedings of the Twenty-Ninth AAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA*. 2038–2044.
- [10] T. Chen, V. Forejt, M. Kwiatkowska, D. Parker, and A. Simaitis. 2013. PRISM-games: A Model Checker for Stochastic Multi-Player Games. In *Proceedings of TACAS (Lecture Notes in Computer Science, Vol. 7795)*. Springer, 185–191.
- [11] R. Fagin. J. Y. Halpern, Y. Moses, and M. Y. Vardi. 1995. *Reasoning about Knowledge*. MIT Press.
- [12] X. Huang and R. van der Meyden. 2014. Symbolic Model Checking Epistemic Strategy Logic. In *Proceedings of AAI*. 1426–1432.
- [13] W. Jamroga, M. Knapik, D. Kurpiewski, and L. Mikulski. 2019. Approximate Verification of Strategic Abilities under Imperfect Information. *Artificial Intelligence* 277 (2019).
- [14] W. Jamroga, W. Penczek, P. Dembiński, and A. Mazurkiewicz. 2018. Towards Partial Order Reductions for Strategic Ability. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. IFAAMAS, 156–165.
- [15] W. Jamroga, W. Penczek, and T. Sidoruk. 2020. Strategic Abilities of Asynchronous Agents: Semantic Paradoxes and How to Tame Them. *CoRR* abs/2003.03867 (2020). arXiv:2003.03867 <https://arxiv.org/abs/2003.03867>
- [16] W. Jamroga, W. Penczek, T. Sidoruk, P. Dembiński, and A. W. Mazurkiewicz. 2020. Towards Partial Order Reductions for Strategic Ability. *J. Artif. Intell. Res.* 68 (2020), 817–850. <https://doi.org/10.1613/jair.1.11936>
- [17] M. Kacprzak, A. Niewiadomski, and W. Penczek. 2020. SAT-Based ATL Satisfiability Checking. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning, KR 2020, Rhodes, Greece, September 12-18, 2020*, Diego Calvanese, Esra Erdem, and Michael Thielscher (Eds.). 539–549. <https://doi.org/10.24963/kr.2020/54>
- [18] M. Kacprzak and W. Penczek. 2004. Unbounded Model Checking for Alternating-Time Temporal Logic. In *Proceedings of AAMAS-04*. IEEE Computer Society, 646–653.
- [19] D. Kurpiewski, W. Jamroga, and M. Knapik. 2019. STV: Model Checking for Strategies under Imperfect Information. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems AAMAS 2019*. IFAAMAS, 2372–2374.
- [20] A. Lomuscio, W. Penczek, and H. Qu. 2010. Partial Order Reductions for Model Checking Temporal-epistemic Logics over Interleaved Multi-agent Systems. *Fundam. Inform.* 101, 1-2 (2010), 71–90.
- [21] A. Lomuscio, H. Qu, and F. 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.
- [22] A. Lomuscio and F. Raimondi. 2006. MCMAS : A Model Checker for Multi-Agent Systems. In *Proceedings of TACAS (Lecture Notes in Computer Science, Vol. 4314)*. Springer, 450–454.
- [23] A. Niewiadomski, M. Kacprzak, D. Kurpiewski, M. Knapik, W. Penczek, and W. Jamroga. 2020. MsATL: A Tool for SAT-Based ATL Satisfiability Checking. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '20, Auckland, New Zealand, May 9-13, 2020*, Amal El Fallah Seghrouchni, Gita Sukthankar, Bo An, and Neil Yorke-Smith (Eds.). International Foundation for Autonomous Agents and Multiagent Systems, 2111–2113.
- [24] M. Pauly. 2002. A Modal Logic for Coalitional Power in Games. *Journal of Logic and Computation* 12, 1 (2002), 149–166.
- [25] W. Penczek. 2018. Improving Efficiency of Model Checking for Variants of Alternating-time Temporal Logic. In *Proceedings of the 27th International Workshop on Concurrency, Specification and Programming, Berlin, Germany, September 24-26, 2018 (CEUR Workshop Proceedings, Vol. 2240)*, Bernd-Holger Schlingloff and Samira Akili (Eds.). CEUR-WS.org.
- [26] J. Pilecki, M.A. Bednarczyk, and W. Jamroga. 2014. Synthesis and Verification of Uniform Strategies for Multi-Agent Systems. In *Proceedings of CLIMA XV (Lecture Notes in Computer Science, Vol. 8624)*. Springer, 166–182.
- [27] L. Priese. 1983. Automata and concurrency. *Theoretical Computer Science* 25, 3 (1983), 221 – 265.
- [28] P. Y. Schobbens. 2004. Alternating-Time Logic with Imperfect Recall. *Electronic Notes in Theoretical Computer Science* 85, 2 (2004), 82–93.