Uniform Strategies under Bounded Resources

Natasha Alechina
University of Nottingham
nza@cs.nott.ac.uk

Mehdi Dastani
Utrecht University
m.m.dastani@uu.nl

Brian Logan
University of Nottingham
bsl@cs.nott.ac.uk

Abstract

We consider a logic of strategic ability that extends Alternating Time Temporal Logic (ATL) with syntactic epistemic modalities and explicit bounds on the cost of strategies. We show that in a special kind of models, which have an explicit communication step before every action selection, the model checking problem for our logic is decidable under perfect recall uniform strategies. This result does not contradict known results on the undecidability of ATL with perfect recall uniform strategies, since the class of models is restricted to allow communication between agents in a coalition.

1 Introduction

This paper addresses the question of verifying the existence of uniform strategies in multi-agent systems where the agents act under resource constraints and imperfect information. *Uniform strategies* are strategies where agents select the same actions in all states where they have the same information available to them. Uniform strategies are important because real agents can only select actions based on the information they have. The model-checking problem for Alternating-Time Temporal Logic (ATL) under imperfect information and with uniform perfect recall strategies is known to be undecidable [3], even without taking into account resource constraints. In [1], we introduced a logic called Resource-Bounded Alternating Time Syntactic Epistemic Logic (RB±ATSEL). RB±ATSEL has a decidable model-checking problem for so-called *coalition-uniform* strategies, and an algorithm for verifying existence of coalition-uniform (rather than uniform) strategies for RB±ATSEL is given in [1]. A strategy is coalition-uniform if agents in a coalition select the same joint action in all states where the knowledge of the *coalition* is the same. This decidability result holds for any notion of coalition knowledge, where coalition-indistinguishability between states is decidable. For example, this result applies for uniformity with respect to the distributed knowledge of a coalition.

However, the notion that a coalition can select actions based on, say, its distributed knowledge, presupposes free and unbounded communication between the agents in the coalition before every action selection. This rather goes against the grain of the resource-bounded setting in [1]. In this paper, we investigate the question of when and how we can convert a coalition-uniform strategy in the setting of [1] to a uniform strategy by adding explicit (and explicitly costed) communication before each action selection. At first sight, this problem seems trivial: just add a communication step, which basically converts the distributed knowledge of the coalition into individual knowledge of each agent, and add corresponding costs to the computation of the strategy. On closer inspection, however, there are several complications. For example, inserting an extra communication step before every 'real' action requires changing the semantics of the temporal next state operator. Even with such a change, and ignoring the resource bounds, the set of formulas true under the requirement that strategies are coalition uniform changes when we require uniform strategies and introduce an explicit communication step. Adding extra knowledge to agents' states to ensure uniform strategies makes some negative epistemic formulas false. This in turn affects formulas ascribing the agents the ability to achieve some goal while *not* knowing whether p is true (e.g., a coalition of agents performing a task to acquire a surprise present for one of the agents in the coalition, without that agent knowing about the existence of the present). In what follows, we present precise changes that need to be made to the semantics of RB±ATSEL to incorporate explicit communication before every action selection.

2 Background

In this section, we recap definitions from [1] where Resource Bounded Alternating Time Temporal Logic with Syntactic Knowledge (RB±ATSEL) was introduced.

The language of RB±ATSEL is built from the following components:

- $Agt = \{a_1, \ldots, a_n\}$ a set of n agents,
- $Res = \{res_1, \ldots, res_r\}$ a set of r resources,
- Π a set of propositions.

The set of possible resource bounds or resource allocations is $B = Agt \times Res \to \mathbb{N}_\infty$, where $\mathbb{N}_\infty = \mathbb{N} \cup \{\infty\}$.

Formulas of the language $\mathcal{L}$ of RB±ATSEL are defined by the following syntax

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \psi \mid \langle\!\langle A^b \rangle\!\rangle \bigcirc \varphi \mid \langle\!\langle A^b \rangle\!\rangle \varphi \,\mathcal{U}\, \psi \mid \langle\!\langle A^b \rangle\!\rangle \Box \varphi \mid K_a \varphi$$

where $p \in \Pi$ is a proposition, $A \subseteq Agt$, $b \in B$ is a resource bound and $a \in Agt$.

The meaning of RB±ATSEL formulas is as follows:

- $\langle\!\langle A^b \rangle\!\rangle \bigcirc \psi$ means that a coalition A has a strategy executable within resource bound b to ensure that the next state satisfies ψ.
- $\langle\!\langle A^b \rangle\!\rangle \psi_1 \,\mathcal{U}\, \psi_2$ means that A has a strategy executable within resource bound b to ensure ψ_2 while maintaining the truth of ψ_1.
- $\langle\!\langle A^b \rangle\!\rangle \Box \psi$ means that A has a strategy executable within resource bound b to ensure that ψ is always true.
- $K_a \phi$ means that formula ϕ is in agent a's knowledge base. Note that this is a syntactic knolwedge definition.

Definition 1 *A model of RB±ATSEL is a structure $M = (\Phi, Agt, Res, S, \Pi, Act, d, c, \delta)$ where:*

- *Φ is a finite set of formulas of $\mathcal{L}$ (possible contents of the local states of the agents).*
- *Agt is a non-empty set of n agents.*
- *Res is a non-empty set of r resources.*
- *S is a set of tuples $(s_1, \ldots, s_n, s_e)$ where $s_e \subseteq \Pi$ and for each $a \in Agt$, $s_a \subseteq \Phi$.*
- *Π is a finite set of propositional variables; $p \in \Pi$ is true in $s \in S$ iff $p \in s_e$.*
- *Act is a non-empty set of actions which includes $idle$, and $d : S \times Agt \to \wp(Act) \setminus \{\emptyset\}$ is a function which assigns to each $s \in S$ a non-empty set of actions available to each agent $a \in Agt$. We assume that for every $s \in S$ and $a \in Agt$, $idle \in d(s, a)$. We denote joint actions by all agents in Agt available at s by $D(s) = d(s, a_1) \times \cdots \times d(s, a_n)$.*
- *for every $s, s' \in S, a \in Agt$, $d(s, a) = d(s', a)$ if $s_a = s'_a$.*
- *$c : Act \times Res \to \mathbb{Z}$ is the function which models consumption and production of resources by actions (a positive integer means consumption, a negative one production). Let $cons_{res}(\alpha) = max(0, c(\alpha, res))$ and $prod_{res}(\alpha) = -min(0, c(\alpha, res))$. We stipulate that $c(idle, res) = 0$ for all $res \in Res$.*
- *$\delta : S \times Act^n \to S$ is a partial function which for every $s \in S$ and joint action $\sigma \in D(s)$ returns the state resulting from executing σ in s.*

We denote by $D_A(s)$ the set of all joint actions by agents in coalition A at s. Let σ be a joint action by agents in A. The set of outcomes of this joint action in s is the set of states reached when A executes σ: $out(s, \sigma) = \{s' \in S \mid \exists \sigma' \in D(s) : \sigma = \sigma'_A \wedge s' = \delta(s, \sigma')\}$ (where σ'_A is the restriction of σ' to A). A *strategy for a coalition* $A \subseteq Agt$ is a mapping $F_A : S^+ \to Act^{|A|}$ (from finite non-empty sequences of states to joint actions by A) such that, for every $\lambda s \in S^+$ (a finite sequence consisting of a sequence λ followed by s), $F_A(\lambda s) \in D_A(s)$. A computation $\lambda \in S^\omega$ is consistent with a strategy F_A iff, for all $i \geq 0$, $\lambda[i+1] \in out(\lambda[i], F_A(\lambda[0, i]))$. Overloading notation, we denote the set of all computations λ consistent with F_A that start from s by $out(s, F_A)$. Given a bound $b \in B$, a computation $\lambda \in out(s, F_A)$ is b-consistent with F_A iff, for every $i \geq 0$, for every $a \in A$,

$$\sum_{j=0}^{j=i-1} tot(F_a(\lambda[0, j])) + b_a \geq cons(F_a(\lambda[0, i]))$$

where $F_a(\lambda[0,j])$ is a's action as part of the joint action returned by F_A for the sequence of states $\lambda[0,j]$; $tot(\sigma) = prod(\sigma) - cons(\sigma)$ is the (vector) difference between the vector $prod(\sigma) = (prod_1(\sigma), \ldots, prod_r(\sigma))$ of resource amounts action σ produces and the vector of resource amounts $cons(\sigma)$ it consumes; b_a is $a's$ resource bound in b. This condition requires that the amount of resources a accumulated on the path so far, plus the original bound, is greater than or equal to the cost of executing the next action by a in the strategy. F_A is a b-strategy if all $\lambda \in out(s, F_A)$ are b-consistent.

In the presence of imperfect information, it makes sense to consider only *uniform* strategies rather than arbitrary ones. A strategy is uniform if after epistemically indistinguishable histories, agents select the same actions. To define uniform strategies, we need the notion of indistinguishable histories. Two states s and t are epistemically indistinguishable by agent a, denoted by $s \sim_a t$, if a has the same local state (knows the same formulas) in s and t: $s \sim_a t$ iff $s_a = t_a$. Two histories $s_1, \ldots, s_k$ and $t_1, \ldots, t_k$ are indistinguishable by a (also denoted by $\sim_a$) if, and only if, for all $j \in [1,k]$, $s_j \sim_a t_j$. An agent's strategy is uniform iff $F_a(\lambda) = F_a(\lambda')$ for all $\lambda \sim_a \lambda'$. A coalition strategy is uniform iff for every agent in the coalition, the strategy is uniform.

It is known that the model checking problem for ATL with perfect recall and uniform strategies is undecidable. Instead of considering uniform strategies, [1] consider *coalition uniform* strategies. For a coalition A, indistinguishability $s \sim_A s'$ means that A as a whole has the same knowledge in s and s'. Various notions of coalitional knowledge can be used to define $\sim_A$. One of the obvious ones is $s \sim_A t$ iff $\bigcup_{a \in A} s_a = \bigcup_{a \in A} t_a$ (the distributed knowledge of A in s and t is the same). $\sim_A$ can be lifted to histories in the same way as $\sim_a$: $s_1, \ldots, s_k \sim_A t_1, \ldots, t_k$ iff for all $j \in [1,k]$, $s_j \sim_A t_j$.

Definition 2 *A strategy F_A for A is coalition-uniform with respect to $\sim_A$ if for all $\bar{s} \sim_A \bar{t}$, $F_A(\bar{s}) = F_A(\bar{t})$.*

Coalition uniformity is a weaker notion than uniformity. It presupposes that the entire coalition is responsible for choosing actions by individual agents in the coalition, which in turns presupposes some kind of communication between the agents.

The truth definition for RB±ATSEL with coalition-uniform strategies (parameterised by $\sim_A$) is as follows:

- $M, s \models p$ iff $p \in s_e$
- boolean connectives have standard truth definitions
- $M, s \models \langle\langle A^b \rangle\rangle \bigcirc \phi$ iff $\exists$ coalition-uniform b-strategy F_A such that for all $\lambda \in out(s, F_A)$: $M, \lambda[1] \models \phi$
- $M, s \models \langle\langle A^b \rangle\rangle \phi \,\mathcal{U}\, \psi$ iff $\exists$ coalition-uniform b-strategy F_A such that for all $\lambda \in out(s, F_A)$, $\exists i \geq 0$: $M, \lambda[i] \models \psi$ and $M, \lambda[j] \models \phi$ for all $j \in \{0, \ldots, i-1\}$
- $M, s \models \langle\langle A^b \rangle\rangle \Box \phi$ iff $\exists$ coalition-uniform b-strategy F_A such that for all $\lambda \in out(s, F_A)$ and $i \geq 0$: $M, \lambda[i] \models \phi$.
- $M, s \models K_a \phi$ iff $\phi \in s_a$

The following theorem was proved in [1]:

Theorem 1 *The model-checking problem for RB±ATSEL with coalition-uniform strategies, with respect to any decidable notion of $\sim_A$, is decidable.*

3 Uniform Strategies with Explicit Communication

In this section, we first introduce a slightly modified semantics for RB±ATSEL. We introduce an explicit communication step before every action execution step. We also replace coalition uniform strategies by uniform strategies, and change the semantics of the 'next state' operator slightly, to account for the idea that we now have a communication step inserted before each action step. We refer to this semantics $\models_u$ (for 'under uniform strategies'), and call a model belonging to the subclass of RB±ATSEL models with this explicit communication step, a communication model.

Definition 3 *A communication model is an RB±ATSEL model $M = (\Phi, Agt, Res, S, \Pi, Act, d, c, \delta)$ where:*

- *Φ, Agt and Res are as in Definition 1.*
- *$S = S' \cup Q$ where S' and Q are disjoint non-empty sets of tuples $(s_1, \ldots, s_n, s_e)$ where $s_e \subseteq \Pi$ and for each $a \in Agt$, $s_a \subseteq \Phi \cup \{p_Q\}$. S' is the set of 'action states' which are initial states or result from executing a non-communication action; Q is the set of 'communication states' which result from executing*

a communication action. For technical reasons, we require that for each agent $a \in Agt$ and for each $q \in Q$, a distinguished propositional variable p_Q is in each q_a. This distinguishes agents' states in Q from agents' states in S' and allows them to select different actions in states which would otherwise be epistemically indistinguishable. Executing any action in a Q state removes p_Q from the agent's state.

- Π *is a finite set of propositional variables;* $p \in \Pi$ *is true in* $s \in S$ *iff* $p \in s_e$.
- $Act = Act' \cup CA$ *is a union of two disjoint non-empty sets of actions.* Act' *is a set of non-communication actions which contains* $idle$. $CA = \{com(s_a, A) \mid a \in Agt, A \subseteq Agt\}$ *is a set of communication actions. Communication actions for each agent* a *are of the form* $com(s_a, A)$ *which means, send the entire contents of* a*'s state to each agent in* A. *The effect of communication action is adding* s_a *to the local state of each agent in* A.
- $d : S \times Agt \rightarrow \wp(Act) \setminus \{\emptyset\}$ *is a function which assigns to each* $s \in S$ *a non-empty set of actions available to each agent* $a \in Agt$. *For* $s \in S'$, $d(s, a) = CA$. *For* $q \in Q$, $d(q, a) \subseteq Act'$.
- *for every* $s, s' \in S, a \in Agt$, $d(s, a) = d(s', a)$ *if* $s_a = s'_a$.
- $c : Act \times Res \rightarrow \mathbb{Z}$ *is the function which models consumption and production of resources by actions (a positive integer means consumption, a negative one production). Let* $cons_{res}(\alpha) = max(0, c(\alpha, res))$ *and* $prod_{res}(\alpha) = -min(0, c(\alpha, res))$. *We stipulate that* $c(idle, res) = 0$ *for all* $res \in Res$.
- $\delta : S \times Act^n \rightarrow S$ *is a partial function which for every* $s \in S'$ *and joint communication action* $\sigma \in D(s)$ *returns the resulting state in* Q, *and for every* $q \in Q$ *and joint action* $\sigma \in D(s)$ *returns the resulting state in* S'.

We denote by $com(s_A, A)$ the joint action $(com(s_a, A \setminus \{a\}))_{a \in A}$ where each agent $a \in A$ communicates their local state. The result of executing $com(s_A, A)$ in a state $s \in S'$ is a state $q \in Q$ where each agent in A has a local state $q_a = \cup_{a \in A} s_a$ (containing all of the distributed knolwedge of A in s).

The truth definition $\models_u$ for communication models has the following clause modified for $s \in S'$ (coalition uniform replaced by uniform, and $\lambda[1]$ replaced by $\lambda[2]$:

- $M, s \in S' \models \langle\langle A^b \rangle\rangle \bigcirc \phi$ iff $\exists$ uniform b-strategy F_A such that for all $\lambda \in out(s, F_A)$: $M, \lambda[2] \models \phi$

For states $s \in Q$, the truth definition for $\langle\langle A^b \rangle\rangle \bigcirc \phi$ remains standard:

- $M, s \in Q \models \langle\langle A^b \rangle\rangle \bigcirc \phi$ iff $\exists$ uniform b-strategy F_A such that for all $\lambda \in out(s, F_A)$: $M, \lambda[1] \models \phi$

The following clauses are modified (coalition uniform replaced with uniform) for all $s \in S$:

- $M, s \models \langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi$ iff $\exists$ uniform b-strategy F_A such that for all $\lambda \in out(s, F_A)$, $\exists i \geq 0$: $M, \lambda[i] \models \psi$ and $M, \lambda[j] \models \phi$ for all $j \in \{0, \ldots, i-1\}$
- $M, s \models \langle\langle A^b \rangle\rangle \Box \phi$ iff $\exists$ uniform b-strategy F_A such that for all $\lambda \in out(s, F_A)$ and $i \geq 0$: $M, \lambda[i] \models \phi$.

Theorem 2 *The model checking problem for RB±ATSEL over communication models with respect to* $\models_u$ *is decidable.*

The proof of the theorem is by producing a model-checking algorithm for RB±ATSEL over communication models. We obtain this algorithm by modifying the model-checking algorithm for RB±ATSEL for coalition-uniform strategies, for the special case where coalition uniformity is interpreted as selecting the same actions in the states where the distributed knowledge of the coalition (the union of all local states of the agents in the coalition) is the same. The main modification is that the algorithm has an added check for the type of each state that is encountered in the search. If a state results from executing a non-communication action, then the agents in a coalition A execute the communication action $com(s_A, A)$ which results in a state where all agents in A have the same local state, identical to the distributed knolwedge of A in the preceding action state. Note that the choice of $com(s_A, A)$ results in a uniform strategy because each agent in A always communicates the same information to other agents in A when it has the same local state. If the state results from a communication action, the algorithm proceeds as in the algorithm for finding a coalition uniform strategy for RB±ATSEL presented in [1]. However, now the action selection is not just coalition uniform, but uniform, since each agent's local state contains all of the coalition's knowledge.

4 Conclusions and Future Work

In this paper, we propose a model-checking algorithm for the logic RB±ATSEL defined in [1] with *uniform strategies*. Our approach works only for a special kind of models which we call *communication models*. In communication models, agents always perform a communication step before selecting and executing actions. For simplicity, we restrict these communication actions to communicating the entire contents of an agent's state. In the future, we plan to investigate under which conditions the results apply when agents follow a particular communication protocol.

Acknowledgement. This work was supported by the Engineering and Physical Sciences Research Council [grant EP/K033905/1].

References

[1] Natasha Alechina, Mehdi Dastani, and Brian Logan. Verifying existence of resource-bounded coalition uniform strategies. In Subbarao Kambhampati, editor, *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI 2016)*, pages 24–30. IJCAI/AAAI, 2016.

[2] R. Alur, T. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.

[3] C. Dima and F. L. Tiplea. Model-checking ATL under imperfect information and perfect recall semantics is undecidable. *CoRR*, abs/1102.4225, 2011.