Verifying existence of uniform strategies in systems of communicating agents

Natasha Alechina

University of Nottingham

Joint work with Mehdi Dastani and Brian Logan

CLAR 2018 Hangzhou

General area of the talk

- This talk is on specification and verification of multi-agent systems (MAS)
- a MAS is specified in terms of states and joint actions by the agents
- actions can change both the physical properties of the state and the knowledge of agents (e.g. observation and communication actions)
- actions consume and produce resources

General area of the talk continued

- verification is done by model checking (checking whether the system satisfies some properties)
- example properties could be:
  - does agent 1 have a strategy to achieve a state where agent 2 always knows/believes that p is true?
  - do agents 1 and 2 have a strategy to come to know whether p is true, given their resource allocation?
- in general: is there a strategy for a group of agents to achieve/maintain some property, and what kind of resources are required for this (time, energy, communication costs...)

Background 1: temporal logic and model checking

Temporal logic

- temporal logics talk about computational behaviour in state transition systems

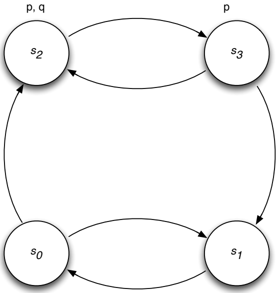

- say things like: 'there is a path (run of the system, computation) where in the next state φ holds', 'always φ', 'φ until ψ'

Temporal logic

- $\bigcirc\varphi$: φ holds in the next state of the path
- $\Box\varphi$: φ holds in every state on the path
- $\varphi\mathcal{U}\psi$: until ψ becomes true, φ holds on the path

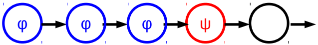

Model checking

- represent a computational system as a state transition system
- express properties of interest in temporal logic (e.g. ‘does the system deadlock?’ $\top \mathcal{U}$ *deadlock*)
- model-checking problem: does a formula φ hold in a state transition system *M*?
- used in verification of hardware and software for a long time

Alternating time temporal logic

- Alternating-Time Temporal Logic (ATL) is a temporal logic which can talk about groups of agents having a strategy to enforce some outcome (temporal property) whatever the other agents in the system do
- ATL is interpreted over concurrent game structures

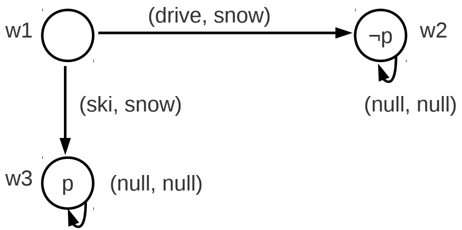

Coalitions, (uniform) strategies

- a strategy is a choice of actions (determined by the current state of the agent or by a finite history = sequence of states)
- a coalition is a group of agents, intuitively with a common goal (such as, discover whether p is true)
- a coalitions's strategy is *uniform* if every agent in the coalition selects actions based on its knowledge (more on knowledge later; for the moment we consider perfect information)

ATL example

- In w_1, agent 1 has a strategy to make sure that in the next state p, and after that p is true forever (choose *ski*, and after that *null*)
- there is only one computational path generated by this strategy, and in the next state p is true
- agent 2 does not have a strategy to enforce p; if it 'chooses' *snow*, agent 1 can perform *ski*, in which case the system is in w_3, or *drive*, in which case it is in w_2, where p is false
- there are two paths, and on one of them $\bigcirc p$ does not hold

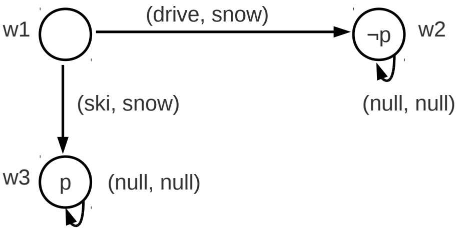

ATL example

- $M, w_1 \models \langle\langle\{1\}\rangle\rangle \bigcirc p$
- $M, w_1 \models \langle\langle\{1\}\rangle\rangle \bigcirc (p \wedge \langle\langle\{1\}\rangle\rangle \Box p)$
- $M, w_1 \not\models \langle\langle\{2\}\rangle\rangle \bigcirc p$

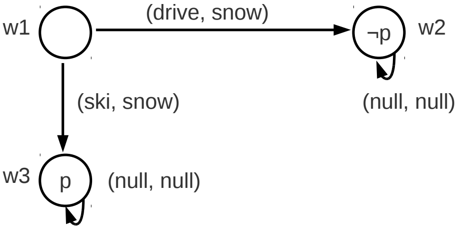

Background 2: epistemic logic

Standard Epistemic Logic

- standard epistemic logic: agent *a* believes/knows ϕ ($B_a\phi$/ $K_a\phi$) iff ϕ is true in all *a*-accessible states
- accessibility relations are the same for all agents in the system; if they are reflexive and transitive, one set of properties for belief/knowledge (S4): all beliefs are consistent, true, and if *a* believes something then it believes that it believes it (positive introspection)
- if accessibility is also symmetric (equivalence relation, often called indistinguishability) then also negative introspection holds (S5)
- what always holds is logical omnisicence: all tautologies are believed, and logical consequences of beliefs are believed

Example: a believes p and is agnostic about q_i

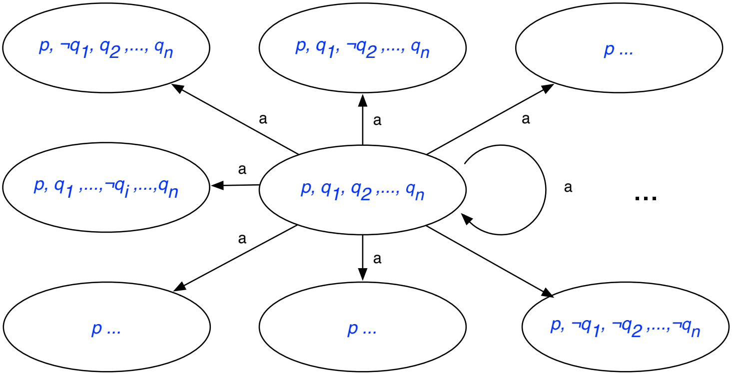

Standard Epistemic Logic

- change of beliefs/knowledge in time:
  - either attach a set of accessible states to each ‘temporal’ state
  - or (for knowledge) make the ‘local state of the agent’ encode the equivalence class for the indistinguishability relation

Syntactic Epistemic Logic

- alternative: in each state, for each agent *a*, represent a set of formulas *a* believes or knows
- instead of a set of accessible states, we have a set of formulas
- this set can change as a result of ontic (actions in the world) and epistemic actions

Advantages of Syntactic Epistemic Logic

- more compact models
- different agents can update their beliefs in different ways:
  - upon receiving a message, can incorporate it or ignore it
  - incorporating the content of the message may involve adding it to the knowledge base and closing it under consequence relation in a different logic (S4, S5, or something much weaker)
  - can avoid logical omniscience

Example: a believes p and is agnostic about q_i

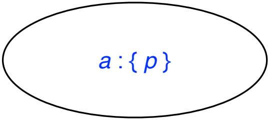

Different possibilities for update: ignore the message

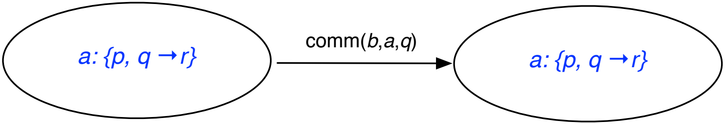

Different possibilities for update: record the fact

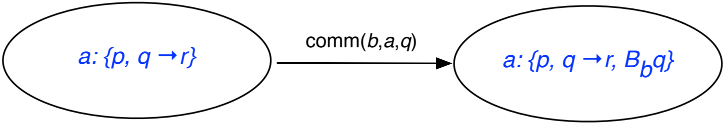

Different possibilities for update: add the content

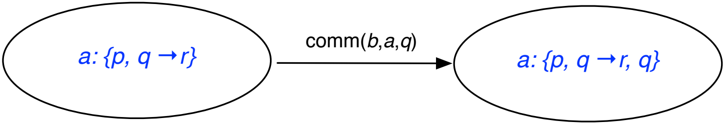

Different possibilities for update: add the content and close under inference

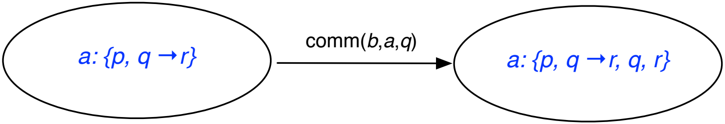

Combination of ATL with resources and epistemics: RB±ATSEL

- Resource-Bounded Alternating Time Syntactic Epistemic Logic (RB±ATSEL) is designed to reason about resource-bounded agents executing both ontic and epistemic actions
- knowledge is modelled syntactically (as a finite set of formulas: the agent's knowledge base):
  - to avoid the problem of logical omniscience
  - to make modelling epistemic actions manageable

What kind of things can RB±ATSEL express

- ‘two robot museum guard robots have a strategy to observe and prevent any attempt approach the artworks in the museum, provided that at least one of them starts fully charged’
- epistemic actions: observing, communicating (anything that changes the agent’s knowledge base without changing the world)
- ontic actions: stopping someone from touching an artwork, charging the battery (changing the world)
- resource allocation: the amount of energy each agent has; there can be multiple resource types: energy, memory, etc.

Concurrent game structure

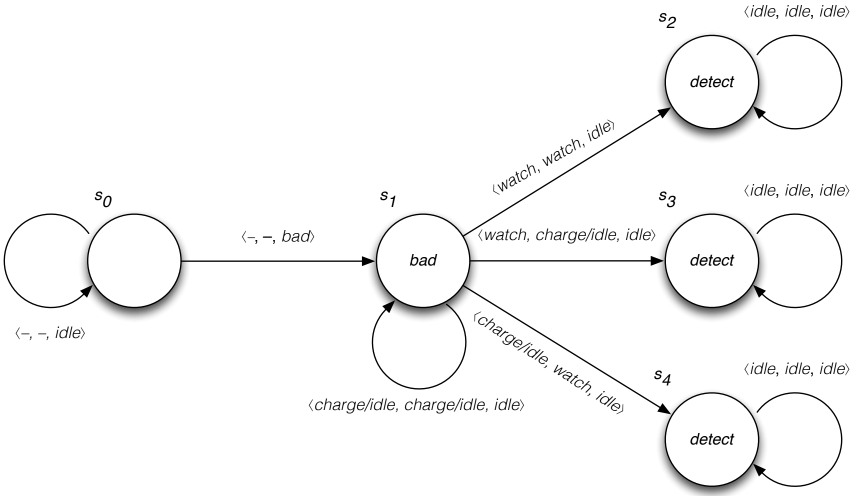

Adding resources (one resource type: energy)

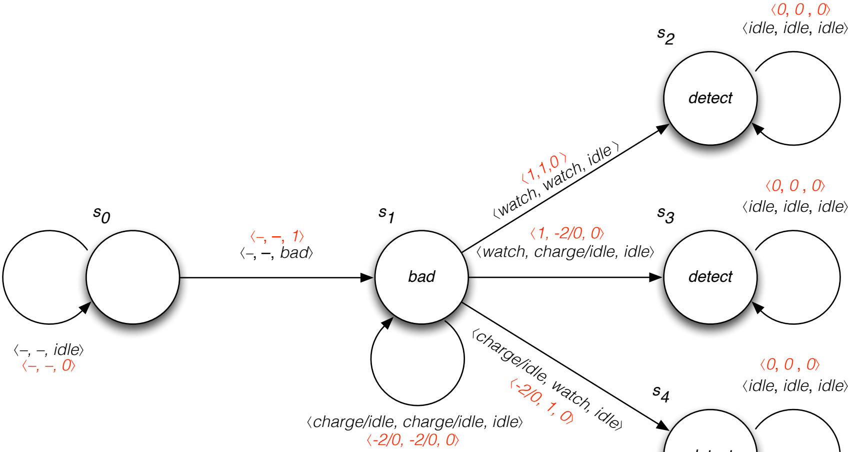

Adding knowledge bases

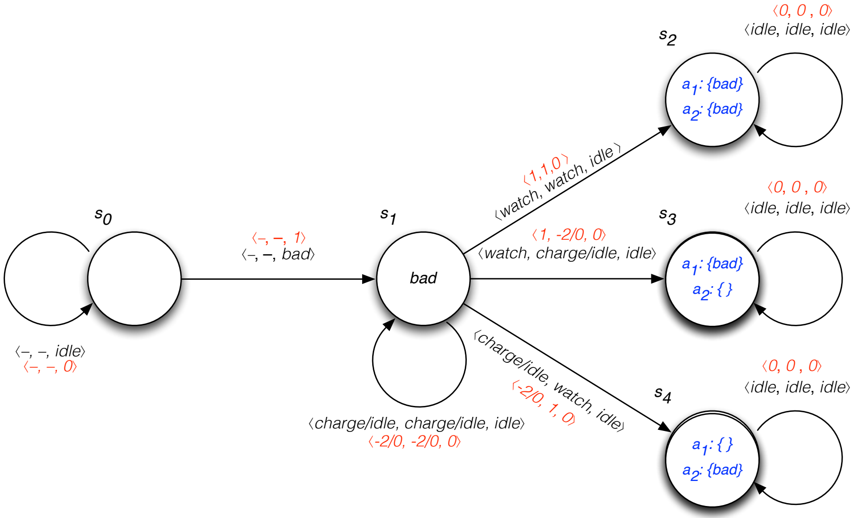

Strategies

- a strategy for coalition A is a mapping from finite sequences of states (histories) to joint actions by agents in A
- if A is the grand coalition (all agents), any strategy of A generates a single run of the system
- otherwise, a strategy corresponds to a tree (each branch of the tree is a run corresponding to a particular choice of actions by A's opponents)
- strategies possible given a particular resource allocation b: a strategy is a b-strategy if for every run generated by this strategy, for each action by A in the strategy, the agents in A will have enough resources to execute it

Language of RB±ATSEL

- In what follows, we assume a set $Agt = \{a_1, \ldots, a_n\}$ of n agents, $Res = \{res_1, \ldots, res_r\}$ a set of r resource types, and a set of propositions Π
- 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 $a \in Agt$.

Meaning of formulas

- $\langle\langle A^b \rangle\rangle \bigcirc \psi$ 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 knowledge definition.

What kind of things can RB±ATSEL express

- if something bad happens (approaching the artwork), one of the guards will know in the next state,provided one of them has one unit of energy:

$$\langle\!\langle \{a_1, a_2\}^{1,0} \rangle\!\rangle \Box (bad \rightarrow \langle\!\langle \{a_1, a_2\}^{0,0} \rangle\!\rangle \bigcirc (K_{a_1} bad \vee K_{a_2} bad))$$

Models of RB±ATSEL

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).
- S is a set of tuples $(s_1, \ldots, s_n, s_e)$ where $s_e \subseteq \Pi$ and for each $a \in Agt$, $s_a \subseteq \Phi$.
- Agt, Res, Π are as before
- 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 $D(s) = d(s, a_1) \times \cdots \times d(s, a_n)$.

Models continued

- 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).
- $\delta : S \times Act^n \to S$ 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.

Costs of strategies

- A strategy is a *b*-strategy (can be carried out under resource bound *b*) if every computation for this strategy can be carried out with initial resource allocation *b* (resources of agents will never drop below 0).

Uniform strategies

- a strategy is uniform if, after epistemically indistinguishable histories, agents select the same actions
- 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 $s_a = t_a$
- $\sim_a$ can be lifted to sequences of states in an obvious way
- a strategy for A is uniform if it is uniform for every agent in A

Coalition uniform strategies

- for a coalition A, indistinguishability $s \sim_A s'$ means that A as a whole has the same knowledge in the two states
- various notions of coalitional knowledge can be used to define $\sim_A$, for example:
  - $s \sim_A t$ 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)
  - another possible definition of $s \sim_A t$ is $\forall a \in A(s_a = t_a)$
- a strategy for A is coalition uniform with respect to $\sim_A$ if it assigns agents in A the same actions in any two histories indistinguishable in $\sim_A$

Truth definition

- $M, s \models p$ 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 $\phi \in s_a$

Syntactic definition for K_a

- $M, s \models K_a\phi$ iff $\phi \in s_a$: a knows ϕ iff ϕ is in a's state
- more general definition: let alg_a be any algorithmic (terminating) procedure that produces a's knowledge when applied to s_a
- for example, alg_a could be computing the largest subset of some finite set of formulas that is derivable from s_a in a particular logic

Model-checking problem for RB±ATSEL

- given a model M of RB±ATSEL and a RB±ATSEL formula ϕ, return the set of states of M where ϕ is true
- the model-checking problem for ATL with perfect recall and uniform strategies is undecidable (because RB±ATSEL is an extension of ATL with perfect recall)
- The model-checking problem for RB±ATSEL with coalition-uniform strategies, with respect to any decidable notion of $\sim_A$, is decidable [IJCAI 2016].

Adding explicit communication step

- coalition uniformity presupposes that agents can select actions based on the knowledge of other agents in the coalition
- to make this assumption realistic, we add an explicit communication step, with associated costs

Original model (fragment)

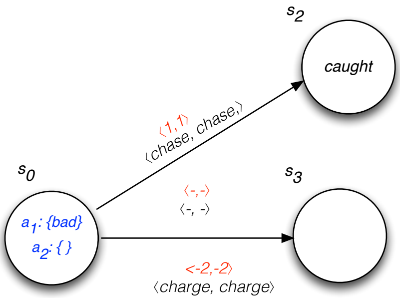

Communication model (fragment)

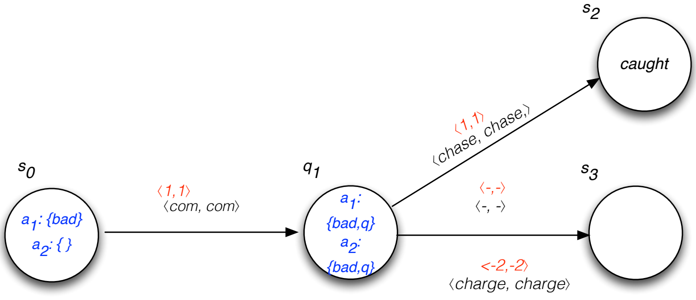

Communication models

- main points:
  - two disjoints sets of states, action states and communication states
  - in action states, only communication actions of the form $com(X, A)$ where $X \subseteq s_a$ (send some contents of state of a to all agents in A) are available
  - the effect of communication action is adding communicated formulas X to the state of every agent in A
  - we changed the truth definition of ‘next’ for communication states (to look two steps ahead)

Model checking for communication models

- Model checking RB±ATSEL over communication models is decidable for perfect recall uniform strategies
- model checking algorithm is obtained by modifying the algorithm for RB±ATSEL for coalition-uniform strategies (for the special case where $\sim_A$ is equivalence of recently communicated formulas)
- the algorithm has an added check for the type of each state that is encountered in the search
- in action states, each agent $a \in A$ executes $com(X, A)$, $X \subseteq s_a$ required by communication protocol ρ, which results in a state where all agents in A have the same recently communicated information
- the choice of $com(X, A)$ results in a uniform strategy because each agent in A always communicates the same information to other agents in A required by the knowledge-based protocol

The cost of communication

- The $com(X, A)$ action can be assigned a cost based e.g., on the number of agents in A and the number of formulas in X

Conclusions

- the model-checking problem for ATL with uniform strategies and perfect recall is undecidable
- however, it is decidable for strategies uniform with respect to e.g., distributed knowledge of the whole coalition
- it is also decidable in models where agents can communicate (following a fixed knowledge-based communication protocol) before action selection

Thank you!
Questions?