Positive and negative results on the model-checking problem for ATL with imperfect information

Cătălin Dima, Constantin Enea, Dimitar Guelev and Ferucio Ţiplea

LACL, LIAFA, IMI-BAS Sofia, resp. Univ. Iaşi

Gasics'2010

An example

- Alice and Bob, married, work in the same company.
- When they arrive at work, they are assigned task x or y.
- Task x needs tx time units to be executed.
- Task y needs $ty > tx$ time units to be executed.
- Task y cannot be assigned to both Alice and Bob.
- In the evening Alice and Bob have two objectives:
  1. to pick their child from the nursery,
  2. to do the shopping.
- Supermarket closes early - the one who does y cannot do the shopping.
- Alice and Bob need to exchange information about their assigned task in order to accomplish their two goals.

Observability aspects

- Alice and Bob might not have complete information about the system state.
- So they might need to exchange information in order to achieve their goal.
- Sometimes the goal could be unobservable to both Alice & Bob.
  - Some privacy-related goal: signal anonymously to tax controllers some possible fraud of an unpleasant neighbor...
  - Or any goal which is not in the reach of the two agents, but is of interest to them.
- And what if they get divorced?...
  - Managing changing coalitions?...

Game arenas for n players

Game arena

$\Gamma = (Ag, Q, (C_a)_{a \in Ag}, \delta, Q_0, (\Pi_a)_{a \in Ag}, \lambda)$, with:

- Ag is a finite set of agents.
- Q is a finite set of states.
- C_a is a finite sets of *actions* available to agent a.
- $Q_0 \subseteq Q$ is the set of *initial states*.
- $\delta \subseteq Q \times C \times Q$ is the *transition relation*.
- Π_a ($a \in Ag$) are the atomic propositions observable by agent a.
- $\lambda : Q \to 2^{\Pi}$ is a *state-labeling function* ($\Pi = \bigcup_{a \in Ag} \Pi_a$).

Indistinguishability

- Indistinguishability on states: $q \sim_a q'$ if $\lambda_a(q) = \lambda_a(q')$.
  - Here, $\lambda_a(q) = \lambda(q) \cap \Pi_a$.

Indistinguishability

- Indistinguishability on states: $q \sim_a q'$ if $\lambda_a(q) = \lambda_a(q')$.
  - Here, $\lambda_a(q) = \lambda(q) \cap \Pi_a$.
- Indistinguishability on runs:

Runs ρ and ρ' *indistinguishable (observationally equivalent)* to a, denoted $\rho \sim_a \rho'$, if

1. Equal length: $|\rho| = |\rho'|$,
2. Same actions for a: $act(\rho,i)\big|_a = act(\rho',i)\big|_a$ for all $i < |\rho|$.
3. Same observations for A: $\lambda_a(\rho[i]) = \lambda_a(\rho'[i])$ for all $i \leq |\rho|$.

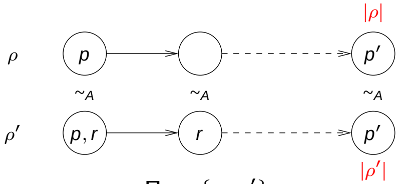

$\Pi_A = \{p, p'\}$

- Synchronous (1) and perfect recall (2+3) observability.
- Other observability variants possible.

Indistinguishability

- Indistinguishability on states: $q \sim_a q'$ if $\lambda_a(q) = \lambda_a(q')$.
  - Here, $\lambda_a(q) = \lambda(q) \cap \Pi_a$.
- Indistinguishability on runs:

Runs ρ and ρ' *indistinguishable (observationally equivalent)* to a, denoted $\rho \sim_a \rho'$, if

1. Equal length: $|\rho| = |\rho'|$,
2. Same actions for a: $act(\rho, i)\big|_a = act(\rho', i)\big|_a$ for all $i < |\rho|$.
3. Same observations for A: $\lambda_a(\rho[i]) = \lambda_a(\rho'[i])$ for all $i \leq |\rho|$.

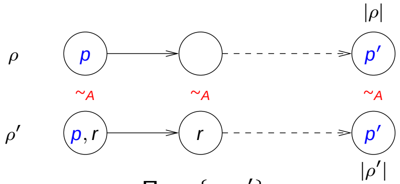

$\Pi_A = \{p, p'\}$

- Synchronous (1) and perfect recall (2+3) observability.
- Other observability variants possible.

Individual strategies

- Strategy for a with perfect information: $\sigma : Q^* \to C_a$.
  - Assumes that a knows exactly, at each moment, the system state.
- Individual knowledge for agent a: observation of *sequences* of Π_a values.
- Feasible strategy for agent a (in case of imperfect information):

$$\sigma : \left(2^{\Pi_a}\right)^* \to C_a.$$

  - Also called *de re* strategy.
  - Equivalently, $\sigma : \left[Q/\sim_a\right]* \to C_a$.
- Strategy σ for coalition A is compatible with run $\rho = q_0 \xrightarrow{c_1} q_1 \xrightarrow{c_2} \ldots$ if

$$\forall i \leq |\rho|,\ \sigma(\lambda_A(q_0)\cdots\lambda_A(q_i)) = c_{i+1}\big|_A$$

- $\Sigma_{perf}(a, \Gamma)$: the set of strategies for agent a with perfect information.
- $\Sigma(a, \Gamma)$: the set of feasible strategies for agent a.

Individual strategies

- Strategy for *a* with perfect information: $\sigma : Q^* \to C_a$.
  - Assumes that *a* knows exactly, at each moment, the system state.
- Individual knowledge for agent *a*: observation of *sequences* of Π_a values.
- Feasible strategy for agent *a* (in case of imperfect information):

$$\sigma : \left(2^{\Pi_a}\right)^* \to C_a.$$

  - Also called *de re* strategy.
  - Equivalently, $\sigma : \left[Q/\sim_a\right]* \to C_a$.
- Strategy σ for coalition *A* is compatible with run $\rho = q_0 \xrightarrow{c_1} q_1 \xrightarrow{c_2} \ldots$ if

$$\forall i \leq |\rho|,\ \sigma(\lambda_A(q_0) \cdots \lambda_A(q_i)) = c_{i+1}\big|_A$$

- $\Sigma_{perf}(a, \Gamma)$: the set of strategies for agent *a* with perfect information.
- $\Sigma(a, \Gamma)$: the set of feasible strategies for agent *a*.

Coalitions

- Coalition = subset of agents, meant to achieve some objective.
  - How much information do agents exchange about their current state?
  - How does an agent utilize his information about the state of the other agents in the same coalition to decide what he/she plays as next move?
  - How to avoid revealing too much information to the (other) members of a coalition, as they might leave the coalition?

Coalitions

- Coalition = subset of agents, meant to achieve some objective.
  - How much information do agents exchange about their current state?
  - How does an agent utilize his information about the state of the other agents in the same coalition to decide what he/she plays as next move?
  - How to avoid revealing too much information to the (other) members of a coalition, as they might leave the coalition?
- No communication: agents only agree to follow a common goal.
- Distributed knowledge: a coalition behaves as a single agent, agents exchange all the information about their current state.
- Common knowledge: the choice of the next action represents a common knowledge inside the coalition.

Strategies based on the distributed knowledge

- Actions available to the members of A: $C_A = \prod_{a \in A} C_a$.
- Atomic propositions visible by at least one member of A: $\Pi_A = \prod_{a \in A} \Pi_a$.
- Distributed observability within A: $\lambda_A(q) = \lambda(q) \cap \Pi_A$.
- Distributed knowledge within the coalition A: observation of $\Pi_A = \bigcup_{a \in A} \Pi_a$.
  - Also defines $\sim_A \subseteq Q \times Q$ by $\sim_A = \bigcap_{a \in A} \sim_a$.
- Strategy for coalition A:

$$\sigma : (2^{\Pi_A})^* \to C_A, \qquad \text{or, equivalently, } \sigma : \left[Q/ \sim_A \right]^* \to C_A,$$

- $\Sigma_D(A, \Gamma)$: the set of strategies with distributed knowledge for coalition A.

Strategies based on the distributed knowledge

Possible interpretation:

- The agents send their observation of the system state to a supervisor.
- The supervisor constructs the distributed knowledge of all the agents in A of the system state.
- And, based on that, requests each agent to issue a particular action.

Avoids some unwanted information leak between the participants in a coalition.

Strategies based on the common knowledge

- Actions available to the members of A: $C_A = \prod_{a \in A} C_a$.
- Common knowledge inside a coalition:
  - Let first $\sim_A^{\cup} = \bigcup_{a \in A} \sim_a$.
  - Then the common knowledge relation is:

$$\sim_A^{c} = \left[Q / \sim_a^{\cup} \right]^{\circledast}$$

  where $\circledast$ is the reflexive-transitive closure.
- Strategy (based on common knowledge) for coalition A:

$$\sigma : \left[Q / \sim_A^{\cup} \right]^{*} \to C_A,$$

- $\Sigma_{comm}(A, \Gamma)$: the set of strategies with common knowledge for coalition A.

ATL syntax

$$\phi ::= p \mid \phi \wedge \phi \mid \neg\phi \mid \langle\langle A\rangle\rangle \bigcirc \phi \mid \langle\langle A\rangle\rangle \phi\,\mathcal{U}\,\phi \mid \langle\langle A\rangle\rangle \phi\,\mathcal{W}\,\phi$$

- p ranges over the set Π of atomic propositions,
- A ranges over the set of subsets of Ag.

Derived operators:

$$\begin{aligned}
P_A\phi &= \neg K_A \neg\phi & [\![A]\!]\bigcirc\phi &= \neg\langle\langle A\rangle\rangle\bigcirc\neg\phi \\
[\![A]\!]\phi\,\mathcal{U}\,\psi &= \neg\langle\langle A\rangle\rangle(\neg\psi\,\mathcal{W}\,\neg\psi\wedge\neg\varphi) & [\![A]\!]\phi\,\mathcal{W}\,\psi &= \neg\langle\langle A\rangle\rangle(\neg\psi\,\mathcal{U}\,\neg\psi\wedge\neg\varphi) \\
\langle\langle A\rangle\rangle\Diamond\,\phi &= \langle\langle A\rangle\rangle \mathrm{true}\,\mathcal{U}\,\phi & \langle\langle A\rangle\rangle\Box\,\phi &= \langle\langle A\rangle\rangle\phi\,\mathcal{W}\,\mathrm{false} \\
[\![A]\!]\Diamond\,\phi &= [\![A]\!]\mathrm{true}\,\mathcal{U}\,\phi & [\![A]\!]\Box\,\phi &= [\![A]\!]\phi\,\mathcal{W}\,\mathrm{false}
\end{aligned}$$

$\mathcal{W}$ is needed, otherwise not all combinations of coalition and temporal operators can be defined.

Semantics of ATL – perfect information

- $(\Gamma, \rho, i) \models \langle\!\langle A \rangle\!\rangle \bigcirc \phi$ if there exists a tuple of strategies $(\sigma_a)_{a \in A}$ with $\sigma \in \Sigma_{perf}(a, \Gamma)$ such that for any run $\rho' \in \mathrm{Runs}^{\omega}(\Gamma)$ which is compatible with each σ_a we have that $(\Gamma, \rho', i + 1) \models \phi$.
- $(\Gamma, \rho, i) \models \langle\!\langle A \rangle\!\rangle \phi_1 \mathcal{U} \phi_2$ iff there exists a tuple strategies $(\sigma_a)_{a \in A}$ with $\sigma \in \Sigma_{perf}(A, \Gamma)$ such that for every run $\rho' \in \mathrm{Runs}^{\omega}(\Gamma)$ which is compatible with each σ_a there exists $j \geq i$ such that $(\Gamma, \rho', j) \models \phi_2$ and $(\Gamma, \rho', k) \models \phi_1$ for all $k = i, \ldots, j - 1$.
- Similarly for $\mathcal{W}$.

Semantics of ATL – imperfect information, no information exchange

- $(\Gamma, \rho, i) \models \langle\!\langle A \rangle\!\rangle \bigcirc \phi$ if there exists a tuple of strategies $\sigma_a \in \Sigma(a, \Gamma)$ such that for any run $\rho' \in \mathrm{Runs}^{\omega}(\Gamma)$ which is compatible with each σ_a and satisfies $\rho'[0..i] \sim_a \rho[0..i]$ for each $a \in A$ we have that $(\Gamma, \rho', i+1) \models \phi$.
- Similarly for $\mathcal{U}$ and $\mathcal{W}$.

Semantics of ATL – imperfect information, distributed knowledge for coalitions

- $(\Gamma, \rho, i) \models \langle\langle A \rangle\rangle \bigcirc \phi$ if there exists a strategy $\sigma \in \Sigma_D(A, \Gamma)$ such that for any run $\rho' \in \text{Runs}^\omega(\Gamma)$ which is compatible with σ and satisfies $\rho'[0..i] \sim_A \rho[0..i]$ we have that $(\Gamma, \rho', i+1) \models \phi$.
- Strategy $\sigma \in \Sigma_D(A, \Gamma)$: coalition A becomes a single player.

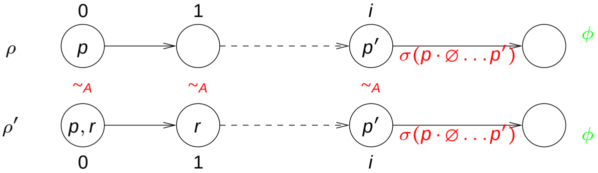

Semantics of ATL – imperfect information, common knowledge for coalitions

- $(\Gamma, \rho, i) \models \langle\!\langle A \rangle\!\rangle \bigcirc \phi$ if there exists a strategy $\sigma \in \Sigma_{comm}(A, \Gamma)$ such that for any run $\rho' \in \mathrm{Runs}^{\omega}(\Gamma)$ which is compatible with σ and satisfies $\rho'[0..i] \sim_A \rho[0..i]$ we have that $(\Gamma, \rho', i+1) \models \phi$.
- Each member of the coalition A plays, at each moment, only those actions that can everybody know he can play at the current moment,
- ... and everybody knows that everybody knows that, etc...
- Thus, by issuing such an action, no member of the coalition gives some advantage (extra information) to the other members of the coalition, which these could use if the “leave” the coalition.

Results on the model-checking problem

Theorem (Alur & Henzinger & Kupferman)

Model-checking is PTIME-complete for ATL with perfect information

Theorem (Yannakakis, D.& Ţiplea)

Model-checking is undecidable for ATL with imperfect information, perfect recall and non-communicating coalitions.

Theorem (D. & Enea & Guelev)

Model-checking is decidable for ATL with imperfect information, perfect recall and strategies based on distributed knowledge.

Theorem (Diaconu & D.)

Model-checking is undecidable for ATL with imperfect information, perfect recall and strategies based on common knowledge.

Imperfect information, perfect recall and non-communicating coalitions

Proof idea:

- Use two agents.
- Encode configurations of a Turing machine as *levels* in the tree gathering the runs compatible with a winning strategy for coalition $\{1, 2\}$.
- In order to accomplish their objective, the two agents need to construct some weak form of their common knowledge.
  - Agent 1 needs to know what is the current state of agent 2, who needs to know on its turn what does agent 1 know about his current state etc.
- Encoding configurations of a Turing machine on the levels of a tree – an idea utilized in proving that model-checking for LTL with common knowledge is undecidable [van der Meyden].

Imperfect information, perfect recall and strategies based on distributed knowledge

Proof idea:

- Combine (an enhanced version of the) state labeling for ATL with perfect information with state splitting for representing (in a finitary manner) histories with identical observations for the coalition.
- State splitting = *subset construction*, common for transforming games with imperfect information into games with perfect information.
- But the initial games with imperfect information have non-observable objectives (for the agent with imperfect information).
- Fortunately, finite branching ensures that, for any objective of the type $p \mathcal{U} q$, there exists some level n in the three where q must have been accomplished before n on all traces.

Imperfect information, perfect recall and strategies based on common knowledge

Proof idea:

- Adapt the proof of the undecidability of the model-checking problem for LTL with common knowledge.

ATEL: ATL with epistemic operators

- Knowledge operators can be added to this logic:
  - $K_a\phi$: agent *a* knows that ϕ must hold in the current state.
  - $K_A\phi$: the fact that ϕ holds in the current state represents *distributed knowledge* for the coalition A.
  - $C_A\phi$: the fact that ϕ holds in the current state represents *common knowledge* for the coalition A.
- Semantics:

$$
\begin{aligned}
(\Gamma,\rho,i) \models K_a\phi \text{ iff } & (\Gamma,\rho',i) \models \phi, \text{ for all runs } \rho' \in \mathrm{Runs}^\omega(\Gamma) \\
& \text{which satisfy } \rho'[0..i] \sim_a \rho[0..i] \\
(\Gamma,\rho,i) \models K_A\phi \text{ iff } & (\Gamma,\rho',i) \models \phi, \text{ for all runs } \rho' \in \mathrm{Runs}^\omega(\Gamma) \\
& \text{which satisfy } \rho'[0..i] \sim_A \rho[0..i] \\
(\Gamma,\rho,i) \models C_A\phi \text{ iff } & (\Gamma,\rho',i) \models \phi, \text{ for all runs } \rho' \in \mathrm{Runs}^\omega(\Gamma) \\
& \text{which satisfy } \rho'[0..i] \sim_A^c \rho[0..i]
\end{aligned}
$$

ATEL and ATL with imperfect information

- In [Wooldridge et al.] $K_A\langle\langle A\rangle\rangle\phi$ specifies that the coalition A has a distributed knowledge of the fact that they may enforce ϕ.
- But K_A is only applied when initiating the coalition.
- ... and strategies in [Wooldridge et al.] are with perfect information – hence unfeasible.
- If the goal ϕ requires more than one step, then it is unfeasible for the coalition A.
- The goal is feasible only for objectives of the type $\phi = \bigcirc\phi$ with the leading operator in ϕ being a coalition operator.

Model-checking ATEL

Theorem

Model-checking for ATEL with strategies based on distributed knowledge and no common knowledge operators is decidable.

The knowledge operators are handled by the state-splitting technique.

Conclusions and further developments

- Variants of ATL based on the variants of knowledge the agents in a coalition have on the other agents' state.
- If common knowledge must be achieved, then the model-checking problem is undecidable.
- Distributed knowledge makes the model-checking problem decidable.
  - Requires games with partially observable states and partially observable objectives.
  - Practical applications may not require deep nesting of the coalition operators.
  - The existence of ”fair supervisors” in changing coalitions imposes severe restrictions.
- Results on games with incomplete information and non-observable (partially observable) objectives?
  - Determinacy?
  - Memoryless winning without determinacy?