ESSAI 2024 course: Logic-based specification and verification of multi-agent systems
Lecture 3:
The temporal logic of coalitional goal assignments

Valentin Goranko
Stockholm University

2nd European Summer School on Artificial Intelligence
ESSAI 2024
Athens, July 15-19, 2024

Introduction: rational behaviour of socially interactive agents

The logic ATL(*) formalises reasoning about *unconditional powers* of agents and coalitions to achieve their goals regardless of the behaviour of all others, which are thus considered as completely adversarial or random.

Real life is much more complex.

Agents in the society act strategically *in their social context*, which may involve both friends and foes, as well as other, unknown agents, who need not be treated a priori as either.

Introduction: cooperation vs protection of interests

Every individual and group of socially interacting agents, while acting and cooperating towards achievement of the collective and societal goals, has their individual and group interests which they want to protect against possible defaults and betrayals by the others.

How can we formalise reasoning about the strategic abilities of agents and coalitions to achieve collective goals while protecting their interests?

This lecture will propose a logical system that addresses and solves that problem.

The main references on this lecture

Precursor:

Sebastian Enqvist and Valentin Goranko:
Socially Friendly and Group Protecting Coalition Logics,

Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'2018), pp 372–380.

The full paper:

Sebastian Enqvist and Valentin Goranko:
The temporal logic of coalitional goal assignments in concurrent multi-player games,

ACM Transactions of Computational Logic, Vol. 23, No. 4, Article 21, 2022.

Available from: `https://dl.acm.org/doi/10.1145/3517128`

The logic TLCGA

Informal introduction and discussion

Coalitional goal assignments and strategic goal operators

Fix a set of agents Agt and a set of goal formulae Φ.

A coalitional goal assignment for Agt in Φ $\gamma : \mathcal{P}(\mathrm{Agt}) \to \Phi$.

We introduce a strategic operator $\langle\!\![\gamma]\!\!\rangle$, informally saying:

There exist a strategy profile Σ for the grand coalition Agt such that for each coalition $C \subseteq \mathrm{Agt}$, the restriction $\Sigma|_C$ of Σ to C is a joint strategy for C forcing the satisfaction of its objective $\gamma(C)$ in all outcome plays enabled by $\Sigma|_C$.

The intuition:

all agents participate in the strategy profile of the grand coalition
with their individual strategies in such a way that,
while contributing to the achievement of the collective goals,
each agent or coalition also guarantees the satisfaction of its own goal
against any possible deviations of all other agents,
thus protecting its individual/coalitional interests.

The temporal logic of coalitional goal assignments TLCGA

TLCGA also involves standard temporal operators of LTL / CTL to express temporalised goals, similarly to ATL.

The language $\mathcal{L}^{\mathsf{TLCGA}}$:

StateFor : $\varphi := p \mid \top \mid \neg\varphi \mid (\varphi \vee \varphi) \mid \langle\!\![\gamma]\!\!\rangle$

PathFor : $\theta := \mathsf{X}\varphi \mid \varphi\mathsf{U}\varphi \mid \mathsf{G}\varphi$

where $p \in \mathsf{AP}$ and $\gamma : \mathcal{P}(\mathsf{Agt}) \to \mathsf{PathFor}$ is a coalitional goal assignment.

Thus, the path formulae are auxiliary, used to express temporalised goals.

$\mathsf{X}\top$ is called a trivial goal and all other goals in PathFor are non-trivial goals.

The family of coalitions $\mathcal{F}$ to which γ assigns non-trivial goals is the support of γ.

Explicit notation for $\langle\!\![\gamma]\!\!\rangle$ with support $\{C_1, ..., C_k\}$:

$$\langle\!\![C_1 \triangleright \phi_1, ..., C_k \triangleright \phi_k]\!\!\rangle$$

Defines the (unique) coalitional goal assignment γ such that

$\gamma(C_1) = \phi_1, ..., C_k = \phi_k$, and $\gamma(C) = \mathsf{X}\top$ for every other $C \in \mathcal{P}(\mathsf{Agt})$.

TLCGA extends ATL: $\langle\!\langle C \rangle\!\rangle\theta \equiv \langle\!\![C \triangleright \theta]\!\!\rangle$, for any $\theta \in \mathsf{PathFor}$.

An aside: equilibria and co-equilibria

Equilibrium: a strategy profile ensuring that no player can deviate to improve the outcome with respect to his/her private goal.

Good for quantitative goals, but not for qualitative (win/lose) because rational players can still deviate in weak equilibria.

Co-equilibrium: a strategy profile ensuring that every player (and coalition) will achieve their private goal, even if all other players deviate.

Good for qualitative goals, while too strong for quantitative ones.

TLCGA can express both (for qualitative goals),
but it is particularly suitable for expressing co-equilibria.

Expressing co-equilibria with TLCGA:

two examples

Example 1: sheep and wolves, a fragile alliance

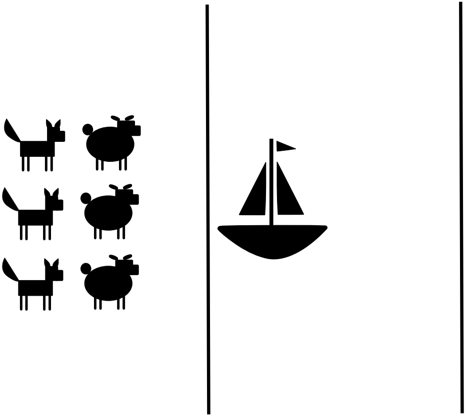

- all animals want to cross the river.
- the boat takes up to 2 animals, no boatman.
- if on either side the wolves outnumber the sheep, they eat them up there.

Do the animals have a 'safe' strategy to cross the river, without any being eaten?

Formalising the Sheep & Wolves puzzle

Specification in ATL:

$$\langle\!\langle \mathsf{Sheep} \cup \mathsf{Wolves} \rangle\!\rangle (\neg e \mathsf{U} c)$$

- c: all animals have crossed the river.
- e: sheep gets eaten.

Not very good: the strategy may be such that the wolves can deviate from it at an opportune moment and eat some of the sheep.

A better specification, in TLCGA:

$$\langle\!\![\mathsf{Sheep} \cup \mathsf{Wolves} \triangleright \mathsf{F} c,\ \mathsf{Sheep} \triangleright (\neg e) \mathsf{U} c]\!\!\rangle$$

Is there a solution satisfying this specification? It depends... (Ask me for a hint.)

Example 2: password protected data sharing

Adapted from (Halpern and Rabin, STOC'1983) and (Parikh, 1985).

Consider a scenario involving two agents, Alice (A) and Bob (B).

Each of them owns a server storing some data, the access to which is protected by a password.

The two agents want to exchange passwords, but neither of them is sure whether to trust the other.

The common goal of the two agents is to cooperate and exchange passwords, but each agent also has the private goal not to give away their password in case the other agent turns out to be untrustworthy and refuses access to their data.

(A side remark: nowadays such deals are arranged by *smart contracts*.)

Use H_A for "*Alice has access to the data on Bob's server*",
and H_B for "*Bob has access to the data on Alice's server*".

So, for instance for Alice, the best possible outcome is H_A
and the worst possible outcome is $\neg H_A \wedge H_B$. Likewise for Bob.

Can the two agents cooperate to safely exchange data?

Data sharing example formalised

An obvious candidate for a formula expressing the common goal:

$$\langle\!\!\lbrack \{A, B\} \triangleright \mathsf{F}\,(H_A \wedge H_B) \rbrack\!\!\rangle$$

Not very good, as that strategy may allow an agent to meanwhile access unilaterally the other's data, and then deviate.

To prevent that, the common goal must be formulated better, as "*eventually reach a state where both agents can access each other's data and until then neither agent should be able to unilaterally access the other's data*," expressed by:

$$\langle\!\!\lbrack \{A, B\} \triangleright (H_A \leftrightarrow H_B)\mathsf{U}(H_A \wedge H_B) \rbrack\!\!\rangle$$

This formula is better, but it does not yet express the stronger requirement that, even if one agent deviates from that strategy profile, the other should still be able to protect her/his interests while still following her/his strategy.

For that, we need to enrich the goal assignment above with *individual goals*:

$$\langle\!\!\lbrack \{A, B\} \triangleright (H_A \leftrightarrow H_B)\,\mathsf{U}\,(H_A \wedge H_B);\; A \triangleright \mathsf{G}\,(H_B \to H_A);\; B \triangleright \mathsf{G}\,(H_A \to H_B) \rbrack\!\!\rangle$$

The formula above can now be equivalently simplified by replacing the common goal with $\mathsf{F}\,(H_A \wedge H_B)$.

The logic TLCGA: formal introduction

Strategic game forms

A (strategic) game form over a non-empty set of outcomes O is a tuple

$$\mathcal{G} = (\mathsf{Act}, \mathsf{act}, \mathsf{O}, \mathsf{out})$$

where

- Act is a non-empty set of actions,
- $\mathsf{act} : \mathsf{Agt} \to \mathcal{P}^{+}(\mathsf{Act})$ is a mapping assigning to each $\mathsf{a} \in \mathsf{Agt}$ a non-empty set act_a of actions available to the player a,
- $\mathsf{out} : \Pi_{\mathsf{a} \in \mathsf{Agt}}\, \mathsf{act}_\mathsf{a} \to \mathsf{O}$ is a map assigning to every available action profile $\zeta \in \Pi_{\mathsf{a} \in \mathsf{Agt}}\, \mathsf{act}_\mathsf{a}$ a unique outcome in O.

Fix a set of agents Agt and a set of atomic propositions AP.

A concurrent game model for Agt and AP:

$$\mathcal{M} = (\mathsf{S}, \mathsf{Act}, \mathfrak{g}, V)$$

where

- S is a non-empty set of states,
- Act is a non-empty set of actions,
- $\mathfrak{g} : w \mapsto (\mathsf{Act}, \mathsf{act}_w, \mathsf{S}, \mathsf{out}_w)$ is a game map, assigning to each state $w \in \mathsf{S}$ a strategic game form $g(w)$ over the set of outcomes S,
- $V : \mathsf{AP} \to \mathcal{P}(\mathsf{S})$ is a valuation of the atomic propositions in S.

Plays, strategies, strategy profiles

A partial play, or a history in $\mathcal{M}$ is a finite word of the form:

$$w_0 \zeta_0 w_1 ... w_{n-1} \zeta_{n-1} w_n$$

where $w_0, ..., w_n \in \mathsf{S}$ and for each $i < n$, ζ_i is an action profile in $\Pi_{a \in \mathsf{Agt}} \mathsf{act}(a, w_i)$.

A (memory-based) strategy for player a is a map σ_{a} assigning to each history $h = w_0 \zeta_0 ... \zeta_{n-1} w_n$ in Play an action $\sigma_{\mathsf{a}}(h)$ from $\mathsf{act}(\mathsf{a}, w_n)$.

$\mathsf{StratProf}_{\mathcal{M}}(C)$: the set of all joint strategies for a coalition C in $\mathcal{M}$.

$\mathsf{StratProf}_{\mathcal{M}} = \mathsf{StratProf}_{\mathcal{M}}(\mathsf{Agt})$: the set of all strategy profiles in the model $\mathcal{M}$.

Plays and paths induced by joint strategies

Fix a concurrent game model $\mathcal{M} = (\mathsf{S}, \mathsf{Act}, \mathfrak{g}, \mathsf{out}, V)$.

The play induced by a strategy profile Σ at $w \in \mathsf{S}$ in $\mathcal{M}$:

$$\mathsf{play}(w, \Sigma) = w_0 \zeta_0 w_1 \zeta_1 w_2 \zeta_2 \ldots$$

For a coalition $C \subseteq \mathsf{Agt}$ and a joint strategy Σ_C for C, the set of outcome plays induced by the joint strategy Σ_C at w is the set of plays

$$\mathsf{Plays}(w, \Sigma_C) = \left\{ \mathsf{play}(w, \Sigma) \mid \Sigma \in \mathsf{StratProf}_{\mathcal{M}} \text{ s.t. } \Sigma(\mathsf{a}) = \Sigma_C(\mathsf{a}) \text{ for all } \mathsf{a} \in C \right\}$$

Respectively, $\mathsf{paths}(w, \Sigma, C)$ is the set of computation paths obtained from the plays in $\mathsf{Plays}(w, \Sigma|_C)$ by ignoring the action profiles.

Formal semantics of TLCGA

The semantics of TLCGA: defined in terms of truth of state formulae at a state, respectively truth of path formulae on (the path generated by) a play, in a concurrent game model $\mathcal{M} = (\mathsf{S}, \mathsf{Act}, \mathfrak{g}, \mathsf{out}, V)$.

The only essentially new clause, for $\langle\!\![\gamma]\!\!\rangle$:

$\mathcal{M}, s \models \langle\!\![\gamma]\!\!\rangle$ iff there exists a strategy profile $\Sigma \in \mathsf{StratProf}_{\mathcal{M}}$ such that, for each $C \subseteq \mathsf{Agt}$, it holds that $\mathcal{M}, \pi \models \gamma(C)$ for every $\pi \in \mathsf{paths}(s, \Sigma, C)$.

No positional determinacy of TLCGA

Consider the variation $\langle\!\![\cdot]\!\!\rangle_0$ of $\langle\!\![\gamma]\!\!\rangle$, with semantics based on positional strategies.

Consider the model $\mathcal{M} = (\mathsf{S}, \mathsf{Act}, \mathfrak{g}, \mathsf{out}, V)$

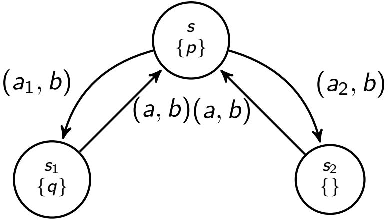

and a goal assignment γ, such that $\gamma(\{\mathsf{a}, \mathsf{b}\}) = p \,\mathsf{U}\, q$ and $\gamma(\{\mathsf{a}\}) = \top \,\mathsf{U}\, \neg(p \vee q)$.

Then, $\mathcal{M}, s \models \langle\!\![\gamma]\!\!\rangle$,
witnessed by any strategy profile Σ such that $\Sigma_{\mathsf{a}}(s) = a_1$ and $\Sigma_{\mathsf{a}}(ss_1s) = a_2$.

However, there is no positional strategy profile witnessing the truth of $\langle\!\![\gamma]\!\!\rangle_0$ at s.

Therefore, memory matters in the semantics of $\langle\!\![\gamma]\!\!\rangle$.

Strategies on paths vs plays

Consider the model $\mathcal{M}$ below, with 3 players: $\{1, 2, 3\}$, where the triples of actions correspond to the order $(1, 2, 3)$ and $*$ denotes any (or, a single) action.

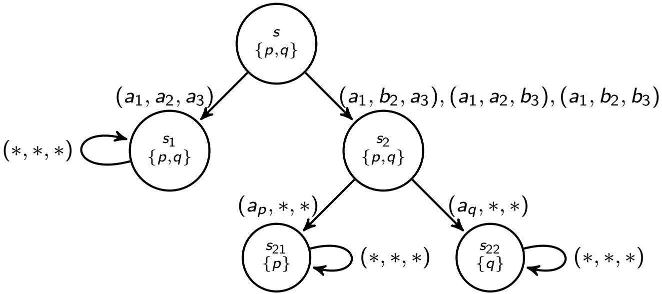

Consider the goal assignment γ, such that:

$\gamma(\{1, 2, 3\}) = \mathsf{G}\,(p \wedge q)$, $\gamma(\{1, 2\}) = \mathsf{G}\,p$ and $\gamma(\{1, 3\}) = \mathsf{G}\,q$. Then:

1. $\mathcal{M}, s \models \langle\!\![\gamma]\!\!\rangle$ in terms of the semantics based on plays-based strategies.
2. $\mathcal{M}, s \not\models \langle\!\![\gamma]\!\!\rangle$ in terms of the semantics based on path-based strategies.

The logic TLCGA:
technical results

Main technical results

- TLCGA-bisimulations and bisimulation invariance of TLCGA formulae.
- Complete axiomatization of the nexttime-fragment XCGA of TLCGA.
- Fixpoint characterization of the long-term temporal goal assignments.
- Complete axiomatization of TLCGA.
- Finite model property and decidability of TLCGA.

Some notation on goal assignments

- $\gamma^{\top}$ is the trivial goal assignment, mapping each coalition to $\mathsf{X}\top$.
- The goal assignment $\gamma[C \triangleright \theta]$ is like γ, but mapping C to θ.
- The goal assignment $\gamma \setminus C$ defined as $\gamma[C \triangleright \mathsf{X}\top]$ is like γ, but excluding C from its support, by replacing its goal with $\mathsf{X}\top$.
- The goal assignment $\gamma|_C$ is defined by restricting γ within C, i.e. mapping each $C' \subseteq C$ to $\gamma(C')$ and mapping all coalitions not contained in C to $\mathsf{X}\top$.

Axioms for the nexttime fragment XCGA

The axiomatic system $\mathsf{Ax}_{\mathsf{XCGA}}$ consists of the following axioms and rules:

(Triv) $\langle\!\![\gamma^{\top}]\!\!\rangle$

(Safe) $\neg\langle\!\![\mathsf{Agt} : \mathsf{X}\bot]\!\!\rangle$

(Mrg) $\langle\!\![C_1 \triangleright \varphi_1]\!\!\rangle \wedge \ldots \wedge \langle\!\![C_n \triangleright \varphi_n]\!\!\rangle \rightarrow \langle\!\![C_1 \triangleright \varphi_1, \ldots, C_n \triangleright \varphi_n]\!\!\rangle$,
where $C_i \cap C_j = \emptyset$ for all $i \neq j$.

(Case) $\langle\!\![\gamma]\!\!\rangle \rightarrow (\langle\!\![\gamma[C \triangleright \mathsf{X}(\varphi \wedge \psi)]]\!\!\rangle \vee \langle\!\![\gamma|_C[\mathsf{Agt} \triangleright \mathsf{X}\neg\psi]]\!\!\rangle)$,
where $\gamma(C) = \mathsf{X}\varphi$

(Con) $\langle\!\![\gamma]\!\!\rangle \rightarrow \langle\!\![\gamma[C \triangleright \mathsf{X}(\varphi \wedge \psi)]]\!\!\rangle$ where:

- $\gamma(C) = \mathsf{X}\varphi$,
- $\gamma(B) = \mathsf{X}\psi$ for some $B \subseteq C$.

Inference rules: Modus Ponens and Goal Monotonicity (G-Mon):

$$\frac{\phi \rightarrow \psi}{\langle\!\![\gamma[C \triangleright \mathsf{X}\,\phi]]\!\!\rangle \rightarrow \langle\!\![\gamma[C \triangleright \mathsf{X}\,\psi]]\!\!\rangle}$$

Completeness and finite model property of XCGA

Theorem (Completeness and FMP of XCGA (AAMAS'2018))

1. *The system* $\mathsf{Ax}_{\mathsf{XCGA}}$ *is sound and complete for the nexttime fragment* XCGA *of* TLCGA.

2. XCGA *has the finite model property (FMP):*
   every satisfiable formula of XCGA *is satisfiable in a finite model.*

ADDENDUM:
Complete axiomatization of the logic TLCGA

μ-calculus of goal assignments

Language of μCGA:

$$\varphi := p \mid \top \mid \neg\varphi \mid (\varphi \vee \varphi) \mid \langle\!\![\gamma]\!\!\rangle \mid \mu z.\varphi \mid \nu z.\varphi$$

$$\gamma : \mathcal{P}(\mathsf{Agt}) \to \mathsf{Nxt}$$

$$\mathsf{Nxt} : \qquad \psi := \mathsf{X}\varphi$$

Theorem

There exists an effective translation t : TLCGA $\to$ μCGA, that uses only one recursion variable.

Nexttime and long-term temporal goal assignments

A goal assignment is local, or next-time, if γ maps every coalition in its support $\mathcal{F}$ to a X-formula.

A non-trivial goal assignment γ is long-term temporal if γ maps every coalition in its support $\mathcal{F}$ either to an U-formula or a G-formula.

Nexttime-extensions of goal assignments

Given a family of coalitions $\mathcal{F}$ and a goal assignment γ supported by $\mathcal{F}$, the nexttime-extension of γ is the next-time goal assignment $\Delta\gamma$ defined as follows.

First, we define $\sup \Delta\gamma := \left\{ \bigcup \mathcal{F}' \mid \emptyset \neq \mathcal{F}' \subseteq \mathcal{F} \right\}$.

Notation: $\gamma|_{\mathsf{UGFor}}$ is the restriction of γ to $\mathcal{F}|_{\mathsf{UGFor}} = \{ C \in \mathcal{F} \mid \gamma(C) \in \mathsf{UGFor} \}$.

Now, for each $C \in \sup \Delta\gamma$ we define $\Delta\gamma(C) :=$

$$\mathsf{X} \left(\bigwedge \{ \varphi \mid \text{there exists } C' \in \mathcal{F}, C' \subseteq C \text{ such that } \gamma(C') = \mathsf{X}\,\varphi \} \wedge \langle\!\![(\gamma|_C)|_{\mathsf{UGFor}}]\!\!\rangle \right),$$

(with due simplification of empty conjunctions and goals.)

For all coalitions that are not in $\sup \Delta\gamma$, $\Delta\gamma$ assigns the trivial goal.

Intuition: $\Delta\gamma$ describes how the goals in γ are propagated in a single transition step, to the immediate successor states.

Unfolding of goal assignments

$$\mathsf{unfold}(\gamma) := \bigvee \mathsf{Finish}(\gamma) \vee \left(\bigwedge \mathsf{UHolds}(\gamma) \wedge \bigwedge \mathsf{GHolds}(\gamma) \wedge \langle\!\![\Delta\gamma]\!\!\rangle \right),$$

where:

- $\mathsf{Finish}(\gamma) := \left\{ \beta \wedge \langle\!\![\gamma \setminus C]\!\!\rangle \mid \gamma(C) = \alpha \mathsf{U} \beta \right\}$
- $\mathsf{UHolds}(\gamma) := \left\{ \alpha \mid \gamma(C) = \alpha \mathsf{U} \beta, \text{ for some } C, \beta \right\}$
- $\mathsf{GHolds}(\gamma) := \left\{ \chi \mid \gamma(C) = \mathsf{G}\chi, \text{ for some } C \right\}$

Theorem (Fixpoint property, Part I)

For any goal assignment γ:

$$\langle\!\![\gamma]\!\!\rangle \equiv \mathsf{unfold}(\gamma).$$

Unfolding axiom

$$\langle\!\![\gamma]\!\!\rangle \leftrightarrow \mathsf{unfold}(\gamma)$$

Induction formulas

For any long-term temporal goal assignment γ and a formula ϕ, we define the induction formula for γ on ϕ as follows:

$$\mathsf{ind}(\gamma, \phi) := \bigvee \mathsf{Finish}(\gamma) \vee \Big(\bigwedge \mathsf{UHolds}(\gamma) \wedge \bigwedge \mathsf{GHolds}(\gamma) \wedge \langle\!\![\Delta\gamma]\!\!\rangle \big[\bigcup \mathcal{F} \triangleright \mathsf{X}\phi \big] \Big)$$

Theorem (Fixpoint property, Part II)

For any long-term temporal goal assignment γ:

$$\langle\!\![\gamma]\!\!\rangle \equiv \mathsf{ind}(\gamma, \langle\!\![\gamma]\!\!\rangle).$$

Temporal goal assignments of Type U

$$\gamma : F \to \mathsf{PathFor}, \quad \text{for } F = \{C_1, ..., C_n, D_1, ..., D_m\}, \text{ where } n > 0, m \geq 0$$

$$C_i \mapsto \alpha_i \mathsf{U} \beta_i$$

$$D_i \mapsto \mathsf{G}\chi_i$$

For the goal assignments of Type U, $\langle\!\![\gamma]\!\!\rangle$ is a least fixpoint of ind:

Proposition (Fixpoint characterization of Type U goal assignments)

Let γ be a long-term temporal goal assignment of Type U, and let z be a fresh variable not occurring in $\langle\!\![\gamma]\!\!\rangle$. Then

$$\langle\!\![\gamma]\!\!\rangle \equiv \mu z.\mathrm{ind}(\gamma, z).$$

Induction rule

$$\frac{\mathsf{ind}(\gamma, \phi) \to \phi}{\langle\!\![\gamma]\!\!\rangle \to \phi}$$

Temporal goal assignments of Type G

$$\gamma : F \to \mathsf{PathFor}, \quad \text{for } F = \{D_1, ..., D_m\}, \text{ where } m > 0$$

$$D_i \mapsto \mathsf{G}\chi_i$$

For the goal assignments of Type G, $\langle\!\![\gamma]\!\!\rangle$ is a greatest fixpoint of ind:

Proposition (Fixpoint characterization of Type G goal assignments)

Suppose that γ is a long-term temporal goal assignment of Type G. Then

$$\langle\!\![\gamma]\!\!\rangle \equiv \nu z.\mathsf{ind}(\gamma, z).$$

Co-induction rule

$$\frac{\phi \rightarrow \mathsf{ind}(\gamma, \phi)}{\phi \rightarrow \langle\!\![\gamma]\!\!\rangle}$$

Complete axiomatization and FMP of TLCGA

The system $\mathsf{Ax}_{\mathsf{TLCGA}}$ extends $\mathsf{Ax}_{\mathsf{XCGA}}$ with the following axioms and rules:

$$\text{Fix:} \quad \mathsf{unfold}(\gamma) \leftrightarrow \langle\!\![\gamma]\!\!\rangle$$

$$\text{R-Ind:} \ \frac{\mathsf{ind}(\gamma, \phi) \to \phi}{\langle\!\![\gamma]\!\!\rangle \to \phi} \quad (\gamma \in \mathsf{TypeU})$$

$$\text{R-CoInd:} \ \frac{\phi \to \mathsf{ind}(\gamma, \phi)}{\phi \to \langle\!\![\gamma]\!\!\rangle} \quad (\gamma \in \mathsf{TypeG})$$

Theorem (Completeness and FMP of TLCGA)

1. *The system Ax_{TLCGA} is sound and complete for* TLCGA.
2. TLCGA *has the finite model property: every satisfiable formula of* TLCGA *is satisfiable in a finite model.*

Complexity of satisfiability: in ExpSpace, conjectured ExpTime-complete.

Lexture 3: concluding remarks

- TLCGA builds on a natural and well-expressive pattern of logical operators capturing agents' strategic interactions in *social context*, involving both collective and individual, immediate and long-term goals.
- Strategic operators over coalitional goal assignments can express not only the standard notion of (Nash) equilibrium, but also the new solution concept of *co-equlibrium*, especially suitable for games with qualitative goals.
- The logic TLCGA is a quite expressive fragment of Strategy Logic, framed in a purely modal style, without explicit mention of strategies in the language.

  The reward is a neatly axiomatized logic with FMP and decidable validity.
- Essential links with "rational verification" and "rational synthesis".
- Many potential applications.
  For example: specification and verification of smart contracts.

END OF LECTURE 3