Temporal logic of coalitional goal assignments in concurrent multi-player games

(Extended abstract)*

Sebastian Enqvist Valentin Goranko
Department of Philosophy, Stockholm University, Sweden
sebastian.enqvist,valentin.goranko@philosophy.su.se

We introduce and study a natural extension of the logic ATL, called *Temporal Logic of Coalitional Goal Assignments* (TLCGA). It features just one, but quite expressive, coalitional strategic operator, viz. the *coalitional goal assignment* operator $\langle\![\gamma]\!\rangle$, where γ is a mapping assigning to each set of players in the game its coalitional *goal*, formalised by a path formula of the language of TLCGA, i.e. a formula prefixed with a temporal operator X, G, or U, representing the temporalised objective for the respective coalition. Then, the formula $\langle\![\gamma]\!\rangle$ intuitively says that there is a strategy profile Σ for the grand coalition Agt such that for each coalition C, the restriction $\Sigma|_C$ of Σ to C is a collective strategy of C that enforces the satisfaction of its objective $\gamma(C)$ in all outcome plays enabled by $\Sigma|_C$.

Here we establish fixpoint characterizations of the temporal goal assignments in a μ-calculus extension of TLCGA, discuss its expressiveness, construct a sound and complete axiomatic system for TLCGA, and obtain its decidability via finite model property.

1 Introduction

Formalising strategic reasoning has become an increasingly rich and attractive direction of active research and applications of multi-agent modal logics over the past few decades. Early logical systems capturing agents' abilities were developed with philosophical motivations and applications in mind, including Brown's modal logic of ability [6] and Belnap and Perloff's STIT logics [4]. In the late 1990s – early 2000s two seminal works in the area appeared independently: Pauly's Coalition logic CL, introduced in [14, 15], and Alur, Henzinger and Kupferman's Alternating time temporal logic ATL introduced (in its final version) in [3]. The logic CL was introduced with the explicit intention to formalise reasoning about one-step (local) strategic abilities of coalitions of agents to guarantee the achievement of designated objectives in the immediate outcome of their collective action, regardless of the respective actions of the remaining agents. The logic ATL, on the other hand, was introduced as a logical formalism for formal specification and verification of open (interacting with environment) computer systems, where the agents represent concurrently executed processes. However, it was gradually adopted in the research on logics for multi-agent systems as one of the most standard and popular logical systems for reasoning about long-term strategic abilities of agents and coalitions in concurrent multi-player games. The logic ATL can be described as an extension of CL with the long-term temporal operators G and U, adopted in the branching-time temporal logic CTL, which can be regarded as a single-agent fragment of ATL. More precisely, both CL and ATL feature special modal operators[1] $\langle\!\langle C \rangle\!\rangle$, indexed with groups (coalitions) of agents C, such that for any formula ϕ, regarded as expressing the coalitional objective of C, the formula

*This is an extended abstract of a full paper in preparation. All proofs are omitted, but most can be provided on request.
[1] We use here the notation from [3], which was more widely adopted.

Submitted to:
SR 2020

$\langle\!\langle C\rangle\!\rangle\phi$ intuitively says that the coalition C has a collective strategy σ_C that guarantees the satisfaction of ϕ in every outcome (state for CL, respectively, play for ATL) that can occur when the agents in C execute their strategies in σ_C, regardless of the choices (strategic or not) of actions of the agents that are not in C.

Thus, both CL and ATL capture reasoning about *absolute powers* of agents and coalitions to act in pursuit of their goals and succeed unconditionally against *any* possible behaviour of their opponents, which are thus regarded as adversaries (in the context of CL) or as randomly behaving environment (in the context of ATL). This is a somewhat extreme perspective, as strategic interactions of rational agents in the real world usually involve a complex interplay of *cooperation* and *competition*, both driven by the individual and collective objectives of all agents, be they proponents or opponents of the objective in focus. To capture these adequately, richer and more versatile formal logical frameworks are needed. In the recent precursor [8] of the present work we proposed two such extensions of CL with more expressive coalitional operators, respectively implementing the following two ideas relating cooperation and competition in social context:

Social friendliness: Agents can *achieve private goals while allowing for cooperation with the others.*

Group protection: Agents can *cooperate with the others while protecting their private goals.*

The second extension mentioned above, called *Group Protecting Coalition Logic* (GPCL) is the starting point of the present work, which introduces and studies its extension in ATL-like style, called Temporal Logic of Coalitional Goal Assignments (TLCGA). The logic TLCGA features just one, but very expressive, coalitional strategic operator, viz. the *coalitional goal assignment* operator of the type $\langle[\gamma]\rangle$, where γ is a mapping assigning to each coalition (subset) in the family of all agents Agt its coalitional *goals*, which is formalised by a path formula of the language of TLCGA, i.e. a formula prefixed with a temporal operator X, G, or U, representing the temporalised objective for the respective coalition. Then, the formula $\langle[\gamma]\rangle$ intuitively says that there is a strategy profile Σ for the grand coalition Agt such that for each coalition C, the restriction $\Sigma|_C$ of Σ to C is a collective strategy of C that enforces the satisfaction of its objective $\gamma(C)$ in all outcome plays enabled by $\Sigma|_C$. The intuition is that each agent participates in the grand coalition with its individual strategy so that, while contributing to the achievement of the common goal, each coalition also guarantees the protection of its coalitional interest agains any possible deviation of all other agents. The logic TLCGA naturally extends ATL (in particular, CL) by enabling, inter alia, the expression of various important yet nuanced patterns of global strategic interaction, some of which we illustrate with examples in Section 3.4. In particular, the logic TLCGA captures a concept which we call "*co-equilibrium*", which we define and promote here as a new, alternative solution concept, for which we argue that is more natural and applicable than the standard notion of Nash equilibrium in the context of concurrent multi-player games with *qualitative* individual objectives. Existence of a co-equilibrium can be expressed quite simply in TLCGA using the operator $\langle[\gamma]\rangle$.

Besides the introduction of the logic TLCGA the main technical contributions of this paper are:

- Fixpoint characterizations of the temporal goal assignments in a μ-calculus extension of TLCGA.
- bisimulation invariance and Hennessy-Milner property for the logic TLCGA with respect to the GPCL-bisimulation introduced in [8]. (These results are not presented in this extended abstract.)
- sound and complete axiomatic system for TLCGA and decidability via finite model property.

Related work. In addition to the references given at the beginning, the present work bears both conceptual and technical connections with the work on ATL with irrevocable strategies [2], [1], strategy contexts [5], coalitional logics of cooperation and propositional control [10, 16], and especially with Strategy logic introduced and studied in [12], [11] and other related works. Indeed, the operator $\langle[\gamma]\rangle$.

introduced here can be translated to Coalitional Strategy Logic, in a way similar to the standard translation of modal logics to first-order logic. However, such translations would result in a technical overkill, which we have both conceptual and computational reasons to avoid. On the conceptual side, translating TLCGA to Strategy Logic would lose the elegant succinctness and focus of the operator $\langle[\gamma]\rangle$ as the main high-level logical construct of the language and would replace it with its low-level description in Strategy Logic. On the technical side, such translation would map a syntactically simple propositional language with reasonably low computational complexity to a quite more expressive but syntactically much heavier, essentially second-order language, explicitly involving quantification over strategies (being functions from states to actions). Essentially these are the same arguments in favour of preferring modal logic over first-order logic, but amplified by the technical complexity of quantifying over functions rather than individuals.

We also note the relationship of the present work with the logic for local conditional strategic reasoning CSR introduced in [9].

Lastly, we also point out the direct applicability of the logic TLCGA for adequate alternative formalisation of the ideas of *rational synthesis* [7] and *rational verification* [17]. These and other connections with already published works will be explored in a follow-up research.

2 Preliminaries and background

2.1 Concurrent game models, plays, strategies

We fix a finite set of **players/agents** $\mathsf{Agt} = \{\mathsf{a}_1,...,\mathsf{a}_n\}$ and a set of **atomic propositions** AP. Subsets of Agt will also be called **coalitions**. Given a set W, we denote by W^* the set of finite words over W, by W^+ the set of non-empty words from W^*, and by W^ω the set of infinite words over W.

Definition 1. Let O be a non-empty set. A **(strategic) game form over the set of outcomes** O is a tuple $\mathscr{G} = (\mathsf{Act}, \mathsf{act}, \mathsf{O}, \mathsf{out})$, where:

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

We now fix a set of atomic propositions AP.

Definition 2. A **concurrent game model** is a tuple $\mathscr{M} = (\mathsf{S}, \mathsf{Act}, \mathfrak{g}, \mathsf{out}, V)$, where:

- S is a non-empty set of **states**,
- Act is a non-empty set of **actions**,
- $\mathfrak{g}$ is a **game map**, assigning to each state $w \in \mathsf{S}$ a strategic game form $\mathfrak{g}(w) = (\mathsf{Act}, \mathsf{act}_w, \mathsf{S}, \mathsf{out}_w)$ over the set of outcomes S.

  For each $\mathsf{a} \in \mathsf{Agt}$ we will denote by $\mathsf{act}(\mathsf{a}, w)$ the set $\mathsf{act}_w(\mathsf{a})$ of **locally available actions** for a in w. We also define $\mathsf{act}_\mathsf{a} := \bigcup_{w\in\mathsf{S}} \mathsf{act}_w(\mathsf{a})$.

  An **action profile** is any tuple of actions $\zeta \in \Pi_{\mathsf{a}\in\mathsf{Agt}}\, \mathsf{act}_\mathsf{a}$.

  A **locally available action profile at the state** w is any tuple of locally available actions $\zeta \in \Pi_{\mathsf{a}\in\mathsf{Agt}}\, \mathsf{act}_w(\mathsf{a})$. The set of these action profiles will be denoted by $\mathsf{ActProf}_w$.

- out is the **global outcome function** assigning to every state w and a local action profile ζ at w a unique **outcome** $\mathsf{out}(w,\zeta) := \mathsf{out}_w(\zeta)$.
- $V : \mathsf{AP} \to \mathscr{P}(\mathsf{S})$ is a **valuation** of the atomic propositions in S;

Given a coalition $C \subseteq \mathsf{Agt}$, a **joint action** for C in the model $\mathscr{M}$ is a tuple of individual actions $\zeta_C \in \prod_{\mathsf{a} \in C} \mathsf{act}_{\mathsf{a}}$. In particular, for any action profile $\zeta \in \Pi_{\mathsf{a}\in\mathsf{Agt}} \mathsf{act}_{\mathsf{a}}$, $\zeta|_C$ is the joint action obtained by restricting ζ to C. For any $w \in \mathsf{S}$ and joint action ζ_C that is available at w, we define:

$$\mathsf{Out}[w,\zeta_C] = \left\{ u \in S \mid \exists \zeta \in \prod_{\mathsf{a}\in\mathsf{Agt}} \mathsf{act}_w(\mathsf{a}) : \ \zeta|_C = \zeta_C \text{ and } \mathsf{out}(w,\zeta) = u \right\}.$$

A **partial play** is either an element of S or a finite word of the form $w_0\zeta_0w_1...w_{n-1}\zeta_{n-1}w_n$, where $w_0,...,w_n \in \mathsf{S}$ and for each $i < n$, ζ_i is a locally available action profile in $\Pi_{a\in\mathsf{Agt}} \mathsf{act}(a,w_i)$. The set of partial plays is denoted by Play.

A **(memory-based) strategy** for player a is a map σ_{a} assigning to each partial play $h = w_0\zeta_0...\zeta_{n-1}w_n$ in Play an action $\sigma_{\mathsf{a}}(h)$ from $\mathsf{act}(\mathsf{a},w_n)$. Given a coalition $C \subseteq \mathsf{Agt}$, a **joint strategy** for C in the model $\mathscr{M}$ is a tuple Σ_C of individual strategies, one for each player in C. A **(global) strategy profile** Σ is a joint strategy for the grand coalition Agt, i.e. an assignment of a strategy to each player. We denote the set of all strategy profiles in the model $\mathscr{M}$ by $\mathsf{StratProf}_{\mathscr{M}}$, and the set of all joint strategies for a coalition C in $\mathscr{M}$ by $\mathsf{StratProf}_{\mathscr{M}}(C)$. Thus, $\mathsf{StratProf}_{\mathscr{M}} = \mathsf{StratProf}_{\mathscr{M}}(\mathsf{Agt})$.

The **play** induced by a strategy profile Σ at $w \in \mathsf{S}$ is the unique infinite word $\mathsf{play}(w,\Sigma) = w_0\zeta_0w_1\zeta_1w_2\zeta_2...$, such that $w_0 = w$ and, for each $n < \omega$ we have $w_{n+1} = \mathsf{out}(\zeta_n,w_n)$, and $\zeta_{n+1} = \Sigma(w_0\zeta_0...\zeta_nw_{n+1})$.

The infinite word $w_0w_1w_2...$ obtained by simply forgetting the moves of players in this infinite play is called the **computation path** induced by Σ at v, and denoted $\mathsf{path}(\Sigma,v)$.

Note that strategies are memory-based in our semantics: moves of players in a strategy may depend on previous moves of other players, and players have perfect information about previous moves.

More generally, given a coalition $C \subseteq \mathsf{Agt}$, a state $w \in \mathsf{S}$ and a joint strategy Σ_C for C we define the **set of outcome plays induced by the joint strategy** Σ_C **at** w to be the set of plays

$$\mathsf{Plays}(w,\Sigma_C) = \{\mathsf{play}(w,\Sigma) \mid \Sigma \in \mathsf{StratProf}_{\mathscr{M}} \text{ such that } \Sigma(\mathsf{a}) = \Sigma_C(\mathsf{a}) \text{ for all } \mathsf{a} \in C\}$$

Given a strategy profile Σ we also denote $\mathsf{Plays}(w,\Sigma,C) := \mathsf{Plays}(w,\Sigma|_C)$. We will likewise use the notation $\mathsf{paths}(w,\Sigma,C)$ for the set of computation paths obtained from the plays in $\mathsf{Plays}(w,\Sigma,C)$. Since these only depends on the strategies assigned to players in C, we shall freely use the notation $\mathsf{Plays}(w,\Sigma,C)$ and $\mathsf{paths}(w,\Sigma,C)$ even when Σ is defined for all members of C, but not for all other players.

3 The temporal logic of coalitional goal assignments TLCGA

3.1 Goal assignments, language and syntax of TLCGA

Given a fixed finite set players Agt and a set G of objects, called 'goals', a **(coalitional) goal assignment for** Agt **in** G is a mapping $\gamma : \mathscr{P}(\mathsf{Agt}) \to G$.

We now define the set $\mathsf{StateFor}$ of **state formulae** and the set $\mathsf{PathFor}$ of **path formulae** of TLCGA by mutual induction, using the following BNF:

$\mathsf{StateFor}: \quad \varphi := p \mid \top \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid (\varphi \vee \varphi) \mid \langle\!\![\gamma]\!\!\rangle$

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

where $p \in \mathsf{AP}$ and $\gamma : \mathscr{P}(\mathsf{Agt}) \to \mathsf{PathFor}$ is a goal assignment for Agt in $\mathsf{PathFor}$. The other propositional connectives $\bot$, $\to$ and $\leftrightarrow$, as well as the temporal operator F, are defined as usual. We write: XFor for the set of path formulas of the form $\mathsf{X}\varphi$; UFor for the set of path formulas of the form $\varphi\mathsf{U}\psi$; GFor for the path formulas of the form $\mathsf{G}\varphi$; and UGFor for $\mathsf{UFor} \cup \mathsf{GFor}$.

We denote the language by $\mathscr{L}^{\mathsf{TLCGA}}$, and its next-time fragment (where $\mathsf{PathFor}$ is restricted to XFor) by $\mathscr{L}^{\mathsf{XCGA}}$. The latter is essentially (with some minor notational changes) the language of the logic GPCL introduced in [8].

Intuitively, the path formulae can be regarded as temporal goals. The goal $\mathsf{X}\top$ is called a **trivial goal** and all other goals in $\mathsf{PathFor}$ are **non-trivial goals**. The family of coalitions $\mathscr{F}$ to which the goal assignment γ assigns non-trivial goals is called the **support of** γ, denoted $\mathsf{Support}(\gamma)$, and γ is said to be **supported by** $\mathscr{F}$. Sometimes we will write a goal assignment γ explicitly, like $C_1 \triangleright \theta_1, \ldots, C_n \triangleright \theta_n$, meaning that $\mathsf{Support}(\gamma) = \{C_1, \ldots, C_n\}$ and $\gamma(C_i) = \theta_i$, for $i = 1, \ldots, n$. More notation:

- $\gamma^\top$ 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 mapping each $C' \subseteq C$ to $\gamma(C')$ and mapping all coalitions not contained in C to $\mathsf{X}\top$.

If γ is the unique goal assignment with empty support, we will identify the formula $\langle\![\gamma]\!\rangle$ with $\top$.

3.2 Semantics

The semantics of TLCGA is defined in terms of truth of state formulae at a state, respectively truth of path formulae on a path, in a concurrent game model $\mathscr{M} = (\mathsf{S}, \mathsf{Act}, \mathfrak{g}, \mathsf{out}, V)$. The truth clauses are like in classical logic for the boolean connectives and like in LTL for the temporal operators. The only new clause, for $\langle\![\gamma]\!\rangle$, is as follows, where $s \in \mathsf{S}$:

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

For any state formula $\varphi \in \mathsf{StateFor}$ we define **the extension of** φ **in** $\mathscr{M}$ to be the set of states in $\mathscr{M}$ where φ is true: $[\![\varphi]\!]_{\mathscr{M}} = \{s \in S \mid \mathscr{M}, s \models \varphi\}$. Likewise, we define the extension of any path formula $\theta \in \mathsf{PathFor}$ to be the set of paths in $\mathscr{M}$ where θ is true: $[\![\theta]\!]^p_{\mathscr{M}} = \{\pi \in S \mid \mathscr{M}, \pi \models \theta\}$. The truth clause for $\langle\![\gamma]\!\rangle$ can now be re-stated in terms of formula extensions as follows:

$$[\![\langle\![\gamma]\!\rangle]\!]_{\mathscr{M}} = \big\{ s \in S \mid \exists \Sigma \in \mathsf{StratProf}_{\mathscr{M}} : \mathsf{paths}(s, \Sigma, C) \subseteq [\![\gamma(C)]\!]^p_{\mathscr{M}} \text{ for each } C \subseteq \mathsf{Agt} \big\}.$$

A strategy profile Σ is said to **witness the goal assignment** γ at a state s of a model $\mathscr{M}$, denoted by $\Sigma, s \Vdash \gamma$, if, for every coalition C in the support of γ and every path $\pi \in \mathsf{paths}(s, \Sigma, C)$ in $\mathscr{M}$ we have $\mathscr{M}, \pi \Vdash \gamma(C)$. We then also say that Σ **witnesses the formula** $\langle\![\gamma]\!\rangle$ at the state s in $\mathscr{M}$. Thus, $\mathscr{M}, s \Vdash \langle\![\gamma]\!\rangle$ iff $\langle\![\gamma]\!\rangle$ is witnessed by some strategy profile at s in $\mathscr{M}$.

We note that the formula $\langle\![C \triangleright \mathsf{X}\phi, \mathsf{Agt} \triangleright \mathsf{X}\psi]\!\rangle$ is semantically equivalent to the strategic operator $[C](\phi; \psi)$ in the logic SFCL defined in [8]. Therefore, the corresponding fragment SFCL_1 of SFCL embeds into TLCGA. Note also that the strategic operator $[C]$ from Coalition logic CL is definable as a special case: $[C]\phi := [C](\phi; \top) \equiv \langle\![C \triangleright \mathsf{X}\phi]\!\rangle$.

3.3 Types of goal assignments

Definition 3. A goal assignment γ supported by a family of coalitions $\mathscr{F}$ is **long-term temporal** if γ maps every coalition in $\mathscr{F}$ either to a U-formula or a G-formula, that is, if $\gamma[\mathscr{F}] \subseteq \mathsf{UGFor}$, where $\gamma[\mathscr{F}] = \{\gamma(C) \mid C \in \mathscr{F}\}$. A goal assignment is called **local**, or **next-time**, if γ maps every coalition in $\mathscr{F}$ to a X-formula, i.e., $\gamma[\mathscr{F}] \subseteq \mathsf{XFor}$. A formula ϕ is said to be in **normal form** if, for every subformula of the form $\langle\!\![\gamma]\!\!\rangle$, the goal assignment γ is either a next-time or a long-term temporal goal assignment.

Definition 4. Let γ be a long-term temporal goal assignment supported by a family $\mathscr{F}$. We say that γ is:

- of **type** U if γ maps at least one element of $\mathscr{F}$ to an U-formula,
- of **type** G if γ maps every element of $\mathscr{F}$ to an G-formula.

We denote the sets of goal assignments of type U and type G respectively by TypeU and TypeG.

3.4 On expressing group objectives in TLCGA

3.4.1 Example 1: Password protected data sharing

This example is adapted from [8], where it was adapted from [13]. Consider the following scenario involving two players, Alice (denoted A) and Bob (denoted B). Each of them owns a server storing some data, the access to which is protected by a password. Alice and Bob want to exchange passwords, but neither of them is sure whether to trust the other. So the common goal of the two players is to cooperate and exchange passwords, but each player also has the private goal not to give away their password in case the other player turns out to be untrustworthy and not provide his/her password. When and how can the two players cooperate to exchange passwords? The answer depends on the kind of actions that Alice and Bob can perform while attempting to achieve their common objective. However, we are more interested now in formalising the problem in TLCGA.

Let us first try to express the common objective by a TLCGA formula. For that, we write 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". Then an obvious candidate for a formula expressing the common goal is the goal assignment formula

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

stating that Alice and Bob have a joint strategy to eventually reach their common objective. However, it is easy to see that this is not good enough. Indeed, while common desired eventual outcome is $H_A \wedge H_B$, but for Alice the worst possible outcome is $\neg H_A \wedge H_B$, whereas the worst possible outcome for Bob is $H_A \wedge \neg H_B$, and each of them would like to avoid their worst possible outcome to happen while trying to achieve the common goal. Thus, the common goal can be formulated better as "*eventually reach a state where both players can access each other's data and until then no player should be able to unilaterally access the other's data*", expressed by the following goal assignment formula:

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

The formula above is ok if both players follow a strategy profile that would realise that goal, but it does not express yet the stronger requirement that even if one of them 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:

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

The common goal can now be reduced to the original one, to produce an equivalent to the above formula:

$$\langle\!\![\{A,B\} \triangleright \mathsf{F}\,(H_A \wedge H_B);\ A \triangleright \mathsf{G}\,(H_B \to H_A);\ B \triangleright \mathsf{G}\,(H_A \to H_B)]\!\!\rangle$$

3.4.2 Example 2: Sheep and wolves: a fragile alliance

This example is a remake with a twist of a well known children's puzzle. A group of 3 wolves and 3 sheep is on the one side of a river and they want to cross the river by boat. There is only one boat that can take 2 animals at a time, but there is no boatman, so one animal has to take the boat back every time, until they all cross the river. The main problem, of course, is that if the wolves ever outnumber the sheep on either side of the river, or on the boat, then the sheep in minority will be promptly eaten up by the wolves. The question is whether, – and if so, how – all animals can cross the river without any sheep being eaten.

Let us formalise the problem in TLCGA. First, some notation. Let Sheep denote the set of all sheep, Wolves denote the set of all wolves, **c** denote the proposition "*all animals have crossed the river*" and **e** denote the proposition "*a sheep gets eaten*". Then the problem seems to be expressed succinctly as the question whether the following formula is true:

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

As in the previous example, this formula is too weak to express the important subtlety that, even if such strategy exists, nothing guarantees that the wolves will not decide to deviate from it and have a gourmet feast with a sheep before (or after) crossing the river. Thus, we need to add an extra goal for all sheep, protecting their interest to stay alive:

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

Now, the common goal can clearly be simplified, while preserving the formula up to equivalence:

$$\langle\!\![\mathsf{Sheep} \cup \mathsf{Wolves} \triangleright \mathsf{F}\, \mathbf{c};\ \mathsf{Sheep} \triangleright \mathsf{G} \neg \mathbf{e}]\!\!\rangle$$

Let us now try to model it as a concurrent game model. We can assume that the river crossing happens instantaneously, so each state of the game is described uniquely (up to re-shuffling of the sheep and of the wolves, which can be considered identical) by the numbers of sheep and wolves on each side of the river, plus the position of the boat (on one or the other side of the river). At each river crossing round, each of the animals has two possible actions: 'stay' or 'go on the boat and cross the river'. The respective transitions are then readily defined, by ensuring that only legitimate transitions can occur, so e.g., if more than two animals decide to jump on the boat at the same time, the state does not change (the transition is a loop). The states satisfying **e** are precisely those where there are more wolves than sheep on any one side of the river, whereas only one state satisfies **c**, where all animals have crossed the river.

And, now, the question: is there a strategy profile satisfying the goal assignment above? The answer, perhaps surprisingly, depends on the specific design of the 'river crossing game'. If it presumes that all animals act simultaneously, then it is easy to see that any joint strategy realising the common goal can be abused by the wolves deviating from it and eating some of the strategy-abiding sheep. The answer to our question in this case is 'No'. However, to level the playing field, the game can be modified so that at every state *first all wolves choose how to act and then all sheep choose how to act*, i.e. formally, every round gets split into two sub-rounds with intermediate states (thus, making it a partly turn-based game). The effect of this change is that now a strategy profile satisfying the goal assignment above could be designed in such a way that the joint strategy of the sheep could involve a suitable joint counter-action to any possible deviation of the wolves that would jeopardise a sheep. We leave to the reader the fun of discovering such joint strategy.

3.4.3 Expressing existence of equilibria and co-equilibria

The fundamental game-theoretic concept of Nash equilibrium can still be applied in the concurrent games that we consider, where the playoff from each play for every player is binary: 1, if that player's goal is satisfied on that play (i.e., the player is a 'winner' in the play), and 0 otherwise (i.e., the player is a 'loser' in the play). However, this notion makes little sense in such qualitative setting, because every strategy profile where no 'loser' can deviate unilaterally to satisfy her objective, is a weak Nash equilibrium. That gives no rational reasons for the losers to adhere to that strategy profile, because any deviation cannot be penalised any further by making their payoff even worse. Thus, we are rather sceptical about the use of Nash equilibria in such games with qualitative objectives as those considered here, and the language of TLCGA is not designed for the purpose of expressing equilibria. Still it enables such expressions, even though in a rather non-succinct way. We illustrate the idea with an example for 2 players, A and B, with respective individual nexttime goals α_A and α_B. First, we can express the existence of an equilibrium satisfying any fixed combination of individual goals. The case when both goals are satisfied by the equilibrium profile is trivial, because no deviations can possibly improve any player's payoff. A more interesting case is when the equilibrium profile satisfies only one goal, say α_A. The following formula expresses existence of such equilibrium:

$$\langle\!\langle [\{A,B\} \triangleright \mathsf{X}\,(\alpha_A \wedge \neg\alpha_B);\ A \triangleright \mathsf{X}\,\neg\alpha_B] \rangle\!\rangle$$

Respectively, the following expresses existence of an equilibrium not satisfying anyone's goal:

$$\langle\!\langle [\{A,B\} \triangleright \mathsf{X}\,(\neg\alpha_A \wedge \neg\alpha_B);\ A \triangleright \mathsf{X}\,\neg\alpha_B;\ B \triangleright \mathsf{X}\,\neg\alpha_A] \rangle\!\rangle$$

Now, the disjunction of all such formulae expresses the existence of any equilibrium. All this generalises to any number of players, though the size of the formula expressing existence of any equilibrium grows exponentially in that number.

However, we wish to define and promote here a new, alternative solution concept, that naturally arises in our framework, vis that of '*co-equilibrium*'. An equilibrium strategy profile means that no player can deviate individually to improve their performance. That concept makes very good sense when players pursue *quantitative* individual objectives which are usually achieved to some degree, leaving room both for possible optimisation and for punishment by the other players when deviating, hence can serve as an effective deterrent from deviation. However, we argue that it does not make very good sense when the individual objectives are *qualitative*, i.e. *win* or *lose*, and losing is the worst possible outcome for the agent, hence it cannot be a deterrent from deviation from a strategy profile where that agent is losing anyway. Furthermore, agents usually participate simultaneously in several coalitions with mutually consistent, yet different objectives, and they try to balance their strategic behaviour so as to serve the collective objectives while protecting their individual interests. In particular, all agents usually have one common, societal objective, say to keep the entire system live and safe, so they enter into a global 'social contract' over that common objective. All these aspects of strategic interaction serve as our motivation to define the notion of **co-equilibrium** in the context of collective and individual qualitative objectives, as a strategy profile that not only ensures satisfaction of the collective objective (the 'social contract') if all players follow it, but moreover also guarantees to every player who adheres to it that even if all other players deviate, that would not affect the satisfaction of his/her individual objective. Thus, a co-equilibrium is a strongly stable solution concept that, we would argue, makes better sense than a Nash equilibrium in the context of games with *qualitative* individual objectives and existence of a co-equilibrium is an important criterium for stability of a society of strategically interacting agents.

Technically, existence of a co-equilibrium can be expressed in TLCGA simply as $\langle\![\gamma]\!\rangle$ where γ is a goal assignment with support consisting of the grand coalition and all singleton sets of agents. That notion can be refined in various ways, e.g., by restricting the players' admissible individual deviations to only those that would still preserve satisfaction of their individual objectives, and possibly also of the collective objective. This notion can also be naturally extended to games with quantitative objectives (payoffs). On the other hand, co-equilibrium is clearly a much more demanding notion and existence of (at least pure) co-equilibria, at least in the strongest sense above is not guaranteed and sufficient conditions for its existence are yet to be investigated. The computational aspects of computing, or even proving existence, of co-equilibria (which can be done by model-checking of TLCGA) are yet to be studied, too. The study of that concept and its natural refinements is left to future work.

4 Fixpoint characterizations of temporal formulae in TLCGA

4.1 The fixpoint property of goal assignments

Definition 5. Given a family of coalitions $\mathscr{F}$ and a goal assignment γ supported by $\mathscr{F}$, we write $\gamma|_{\mathsf{UGFor}}$ for the restriction of γ to the family $\mathscr{F}|_{\mathsf{UGFor}} = \{C \in \mathscr{F} \mid \gamma(C) \in \mathsf{UGFor}\}$. Similarly we write $\gamma|_{\mathsf{XFor}}$ for the restriction of γ to the family $\mathscr{F}|_{\mathsf{XFor}} \subseteq \mathscr{F}$ defined as $\{C \in \mathscr{F} \mid \gamma(C) \in \mathsf{XFor}\}$.

Definition 6. Given a family of coalitions $\mathscr{F}$ and a goal assignment γ supported by $\mathscr{F}$, the **nexttime-extension of** γ is the goal assignment $\Delta\gamma$ defined as follows. First, we define $\sup\Delta\gamma := \left\{\bigcup\mathscr{F}' \mid \emptyset \neq \mathscr{F}' \subseteq \mathscr{F}\right\}$, Then, for each $C \in \sup\Delta\gamma$ we define

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

where as a convention we remove from this formula any conjuncts that reduce to $\top$, which can appear as the result of a conjunction of the empty set (the left conjunct reduces to $\top$) or as $\langle\![\gamma]\!\rangle$ where γ is the empty goal assignment (the right conjunct reduces to $\top$). For all coalitions that are not in $\sup\Delta\gamma$, $\Delta\gamma$ assigns the trivial goal.

Example 1. Consider the example of a goal assignment γ supported by $\mathscr{F} = \{\{a,b\},\{a,c\},\{b,c\}\}$ and defined by the assignment: $\{a,b\} \triangleright p\mathsf{U}q$, $\{a,c\} \triangleright \mathsf{G}r$, $\{b,c\} \triangleright \mathsf{X}s$.

The support of $\Delta\gamma$ will be $\mathscr{F} \cup \{\{a,b,c\}\}$. The action of $\Delta\gamma$ is as follows: $\{a,b\} \triangleright \mathsf{X}\langle\![\{a,b\} \triangleright p\mathsf{U}q]\!\rangle$; $\{a,c\} \triangleright \mathsf{X}\langle\![\{a,c\} \triangleright \mathsf{G}r]\!\rangle$; $\{b,c\} \triangleright \mathsf{X}s$; $\{a,b,c\} \triangleright \mathsf{X}(s \wedge \langle\![\{a,b\} \triangleright p\mathsf{U}q, \{a,c\} \triangleright \mathsf{G}r]\!\rangle)$.

The procedure computing $\Delta\gamma$ is, informally, as follows: for each $\subseteq$-downset $\mathscr{D}$ of coalitions from $\mathscr{F}$, we collect all the formulas φ for which some coalition in $\mathscr{D}$ is mapped to the goal $\mathsf{X}\,\varphi$ into a conjunction, add a conjunct collecting all the longterm goals for coalitions in $\mathscr{D}$ into a single goal assigment, and finally put the resulting conjunction in the scope of an X-operator and assign this goal to the union of $\mathscr{D}$.

Example 2. If $\langle\![\gamma]\!\rangle$ is $\langle\![C \triangleright \varphi\mathsf{U}\psi]\!\rangle$, then $\langle\![\Delta\gamma]\!\rangle = \langle\![C \triangleright \mathsf{X}\,\langle\![C \triangleright \varphi\mathsf{U}\psi]\!\rangle]\!\rangle = \langle\![C \triangleright \mathsf{X}\,\langle\![\gamma]\!\rangle]\!\rangle$. So in this special case the nexttime-extension simply pushes the eventuality $\varphi\mathsf{U}\psi$ one step into the future, so to speak. Similarly, if $\langle\![\gamma]\!\rangle$ is $\langle\![C \triangleright \mathsf{G}\varphi]\!\rangle$, then $\langle\![\Delta\gamma]\!\rangle = \langle\![C \triangleright \mathsf{X}\,\langle\![C \triangleright \mathsf{G}\varphi]\!\rangle]\!\rangle = \langle\![C \triangleright \mathsf{X}\,\langle\![\gamma]\!\rangle]\!\rangle$.

Now we define one of the key technical concepts of this work.

Definition 7. The **unfolding formula of a goal assignment** γ supported by $\mathscr{F}$ is defined as follows:

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

where:

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

As before, by convention we remove from this formula all conjuncts that reduce to $\top$ and all disjuncts that reduce to $\bot$. So, for example, if the set $\mathsf{Finish}(\gamma) = \emptyset$, and hence also $\mathsf{UHolds}(\gamma) = \emptyset$, then the formula $\mathsf{unfold}(\gamma)$ reduces to:

$$\bigwedge \mathsf{GHolds}(\gamma) \wedge \langle\![\Delta\gamma]\!\rangle.$$

Definition 8. Let γ be a long-term temporal goal assignment. Then we define the **induction formula for γ on ϕ** as follows

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

after removing redundant conjuncts and disjuncts, as before. So, this formula is like $\mathsf{unfold}(\gamma)$, except that the largest coalition in the support of $\Delta\gamma$ will be mapped to $\mathsf{X}\phi$.

Proposition 1. *For every long-term temporal goal assignment γ we have:* $\mathsf{unfold}(\gamma) = \mathsf{ind}(\gamma, \langle\![\gamma]\!\rangle)$.

Note that if γ is a next-time goal assignment, then $\mathsf{unfold}(\gamma) \equiv \langle\![\gamma]\!\rangle$. For example, suppose γ is supported by $\{\{a\},\{b\}\}$ and maps $\{a\}$ to $\mathsf{X}p$ and $\{b\}$ to $\mathsf{X}q$. Then $\mathsf{unfold}(\gamma)$ is equal to $\langle\![\Delta\gamma]\!\rangle$, which is the following formula:

$$\langle\![\{a\} \rhd \mathsf{X}p, \{b\} \rhd \mathsf{X}q, \{a,b\} \rhd \mathsf{X}(p \wedge q)]\!\rangle$$

which is clearly equivalent to $\langle\![\gamma]\!\rangle = \langle\![\{a\} \rhd \mathsf{X}p, \{b\} \rhd \mathsf{X}q\}]\!\rangle$. In fact, the equivalence always holds; this is by design, and will play a key role for our axiomatization.

Theorem 1 (Fixpoint property)**.** *For any goal assignment γ:* $\langle\![\gamma]\!\rangle \equiv \mathsf{unfold}(\gamma)$.

Therefore, for any long-term temporal goal assignment γ: $\langle\![\gamma]\!\rangle \equiv \mathsf{ind}(\gamma, \langle\![\gamma]\!\rangle)$.

4.2 A μ-calculus of goal assignments

The μ-calculus extension of the language $\mathscr{L}^{\mathsf{TLCGA}}$ of TLCGA will be denoted by $\mathscr{L}_{\mu}^{\mathsf{TLCGA}}$, and the μ-calculus extension of the next-time fragment $\mathscr{L}^{\mathsf{XCGA}}$ – by $\mathscr{L}_{\mu}^{\mathsf{XCGA}}$.

Formally the language $\mathscr{L}_{\mu}^{\mathsf{TLCGA}}$ is given by the following grammar:

$$\mathsf{StateFor}: \quad \varphi := p \mid \top \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid (\varphi \vee \varphi) \mid \langle\![\gamma]\!\rangle \mid \mu z.\psi$$

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

Here, in $\mu x.\psi$ the formula ψ is subject to the usual constraint that every occurrence of the variable x in ψ is positive, in the sense that it is under the scope of an even number (possibly zero) of negations.

A model for $\mathscr{L}_{\mu}^{\mathsf{TLCGA}}$ is just like a model for $\mathscr{L}^{\mathsf{TLCGA}}$, viz a tuple, $\mathscr{M} = (\mathsf{S}, \mathsf{Act}, \mathfrak{g}, \mathsf{out}, V)$, but now the valuation V assigns values to the variable(s) z used in formulae $\mu z.\psi$.

The semantics of $\mathscr{L}_{\mu}^{\mathsf{TLCGA}}$ extends that of $\mathscr{L}^{\mathsf{TLCGA}}$ with the additional clause that the extension $[\![\mu z.\varphi(z)]\!]$ of a least fixpoint-formula in a model $\mathscr{M} = (\mathsf{S}, \mathsf{Act}, \mathfrak{g}, \mathsf{out}, V)$ is given by:

$$[\![\mu z.\varphi(z)]\!] := \bigcap \{Z \subseteq \mathsf{S} \mid [\![\varphi(Z)]\!] \subseteq Z\},$$

where, as expected:

$$[\![\varphi(Z)]\!] := \{w \in \mathsf{S} \mid \mathscr{M}^Z, w \Vdash \varphi(z)\}.$$

Proposition 2 (Fixpoint characterization of TypeU temporal goal assignments)**.** *Suppose that γ is a long-term temporal goal assignment in* TypeU*, and let z be a fresh variable. Then $\langle\![\gamma]\!\rangle \equiv \mu z.\mathsf{ind}(\gamma, z)$.*

Proposition 3 (Fixpoint characterization of TypeG temporal goal assignments)**.** *Suppose that γ is a long-term temporal goal assignment in* TypeG *and let z be a fresh variable. Then $\langle\![\gamma]\!\rangle \equiv \nu z.\mathsf{ind}(\gamma, z)$.*

5 Axiomatization and completeness of TLCGA

5.1 Axiomatic system for TLCGA

Our axiomatic system $\mathsf{Ax}_{\mathsf{TLCGA}}$ consists of the following axiom schemes and rules: (recall notation on goal assignments from Section 3.1).

5.1.1 I. General axiom schemes for goal assignments

(Taut) A complete family of classical tautologies.

(Triv) $\langle\!\![\gamma^{\top}]\!\!\rangle$ (Recall that $\gamma^{\top}$ is the trivial goal assignment, mapping each coalition to $\mathsf{X}\top$)

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

(Merge) $\langle\!\![C_1 \triangleright \theta_1]\!\!\rangle \wedge \ldots \wedge \langle\!\![C_n \triangleright \theta_n]\!\!\rangle \to \langle\!\![C_1 \triangleright \theta_1, \ldots, C_n \triangleright \theta_n]\!\!\rangle$, where $C_1, \ldots, C_n$ are pairwise disjoint.

(GrandCoalition) $\langle\!\![\gamma]\!\!\rangle \to (\langle\!\![\gamma[\mathsf{Agt}\triangleright\mathsf{X}(\varphi\wedge\psi)]]\!\!\rangle \vee \langle\!\![\gamma[\mathsf{Agt}\triangleright\mathsf{X}(\varphi\wedge\neg\psi)]]\!\!\rangle)$, where $\gamma(\mathsf{Agt}) = \mathsf{X}\varphi$.

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

(Con) $\langle\!\![\gamma]\!\!\rangle \to \langle\!\![\gamma[C\triangleright\mathsf{X}(\varphi\wedge\psi)]]\!\!\rangle$ where $\gamma(C) = \mathsf{X}\varphi$ and $\gamma(C') = \mathsf{X}\psi$ for some $C' \subseteq C$.

5.1.2 II. General rules of inference:

Modus Ponens and **Goal Monotonicity (G-Mon)**:

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

5.1.3 III. Axioms and rules for the long-term goal assignments

The axioms and rules for the goal assignments of types G and U are given on Figure 1. They are adapted from the respective axioms and rules for least and greatest fixed points in the modal mu-calculus. In the axiom Fix, γ is any goal assignment. In the rule $\mathsf{R\text{-}Ind}$ it is a long-term temporal assignment of type U, and in $\mathsf{R\text{-}CoInd}$ it is a long-term temporal assignment of type G.

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

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

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

Figure 1: Fixpoint axiom and induction rules

Here are some important validities, whose derivations in $\mathsf{Ax}_{\mathsf{TLCGA}}$ are used in the completeness proofs:

Ind $\langle\!\![\gamma]\!\!\rangle \leftrightarrow \mathsf{ind}(\gamma, \langle\!\![\gamma]\!\!\rangle)$ for every long-term temporal goal assignment γ.

(Weakening) $\langle\!\![\gamma]\!\!\rangle \to \langle\!\![C \triangleright \gamma(C)]\!\!\rangle$, for any $C \subseteq \mathsf{Agt}$.

Agt-**Maximality** $\langle\!\![\emptyset \triangleright \mathsf{X}\,\phi]\!\!\rangle \vee \langle\!\![\mathsf{Agt}\triangleright\mathsf{X}\,\neg\phi]\!\!\rangle$.

(Superadditivity) $\langle\!\![C_1 \triangleright \mathsf{X}\phi_1]\!\!\rangle \wedge \langle\!\![C_2 \triangleright \mathsf{X}\phi_2]\!\!\rangle \to \langle\!\![C_1 \cup C_2 \triangleright \mathsf{X}(\phi_1 \wedge \phi_2);\ C_1 \triangleright \mathsf{X}\phi_1;\ C_2 \triangleright \mathsf{X}\phi_2]\!\!\rangle$, if $C_1 \cap C_2 = \emptyset$.

Fix(G) $\langle\!\![C \triangleright \mathsf{G}\chi]\!\!\rangle \to \chi \wedge \langle\!\![C \triangleright \mathsf{X}\,\langle\!\![C \triangleright \mathsf{G}\chi]\!\!\rangle]\!\!\rangle$.

CoInd(G) If $\vdash \phi \to \chi \wedge \langle\!\![C \triangleright \mathsf{X}\phi]\!\!\rangle$ then $\vdash \phi \to \langle\!\![C \triangleright \mathsf{G}\chi]\!\!\rangle$.

FP(G) $\langle\!\![C \triangleright \mathsf{G}\chi]\!\!\rangle \leftrightarrow \chi \wedge \langle\!\![C \triangleright \mathsf{X}\,\langle\!\![C \triangleright \mathsf{G}\chi]\!\!\rangle]\!\!\rangle$.

PreFP(U) $\beta \vee (\alpha \wedge \langle\!\![C \triangleright \mathsf{X}\,\langle\!\![C \triangleright \alpha \mathsf{U} \beta]\!\!\rangle]\!\!\rangle) \to \langle\!\![C \triangleright \alpha \mathsf{U} \beta]\!\!\rangle$.

Ind(U) If $\vdash \beta \vee (\alpha \wedge \langle\!\![C \triangleright \mathsf{X}\phi]\!\!\rangle) \to \phi$ then $\vdash \langle\!\![C \triangleright \alpha \mathsf{U} \beta]\!\!\rangle \to \phi$.

FP(U) $\langle\!\![C \triangleright \alpha \mathsf{U} \beta]\!\!\rangle \leftrightarrow \beta \vee (\alpha \wedge \langle\!\![C \triangleright \mathsf{X}\,\langle\!\![C \triangleright \alpha \mathsf{U} \beta]\!\!\rangle]\!\!\rangle)$

Proposition 4 (Soundness of $\mathsf{Ax}_{\mathsf{TLCGA}}$)**.** *The axiomatic system Ax_{TLCGA} is sound: every derivable formula in Ax_{TLCGA} is valid.*

Recall (cf Section 3.3) that a formula $\phi \in \mathscr{L}^{\mathsf{TLCGA}}$ is in **normal form** if, for every subformula of the form $\langle\!\![\gamma]\!\!\rangle$, the goal assignment γ is either a next-time or a long-term temporal goal assignment.

Proposition 5. *For every formula φ there is a formula ψ which is in normal form, and such that $Ax_{\mathsf{TLCGA}} \vdash \varphi \leftrightarrow \psi$.*

Corollary 1. *For every formula φ there is a semantically equivalent formula ψ which is in normal form.*

5.2 Completeness

Hereafter derivability/provability and consistency refer to the axiomatic system $\mathsf{Ax}_{\mathsf{TLCGA}}$. Given a set of formulae Φ, the maximal consistent subsets of Φ are defined as usual. We also define a suitable notion of "Fischer-Ladner closure" and closed set of formulae, referred further. We omit the details here.

Definition 9. Given a closed set of formulae Φ, a **Φ-atom** is a maximal consistent subset of Φ. We denote by $\mathsf{At}(\Phi)$ the set of all Φ-atoms.

Definition 10. A **nexttime assignment** over a finite set of formulae Φ is a formula of the shape

$$\langle\!\![C_1 \triangleright \mathsf{X}\varphi_1, \ldots, C_k \triangleright \mathsf{X}\varphi_k]\!\!\rangle$$

where each formula φ_i belongs to Φ. A **modal one-step theory** over Φ is a finite set of formulae Γ, such that every formula in Γ is either a nexttime assignment over Φ or the negation of such a formula.

Definition 11. Let Φ be a finite set of formulae. A **consistent game form for** Φ is a game form $(\mathsf{Act}, \mathsf{act}, \mathscr{P}(\Phi), \mathsf{out})$ over the set of outcomes $\mathscr{P}(\Phi)$ such that, for each action profile ζ, $\mathsf{out}(\zeta)$ is a consistent set of formulae. A **maximal consistent game form for** Φ is a game form $(\mathsf{Act}, \mathsf{act}, \mathscr{P}(\Phi), \mathsf{out})$ over outcomes $\mathscr{P}(\Phi)$ such that, for each action profile ζ, $\mathsf{out}(\zeta)$ is a maximal consistent subset of Φ.

Note that if Φ is a closed set of formulae, then a consistent game form for Φ is maximal if and only if for every action profile ζ the set $\mathsf{out}(\zeta)$ is a Φ-atom. Given a strategic game form $G = (\mathsf{Act}, \mathsf{act}, \mathsf{O}, \mathsf{out})$, a coalition C and action profiles ζ', ζ, we write $\zeta' \sim_C \zeta$ to state that $\zeta'|_C = \zeta|_C$.

Theorem 2 (One-step completeness)**.** *Let Γ be a consistent modal one-step theory over Φ, and assume that Φ contains all components of Γ and is closed under conjunctions (up to provable equivalence). Then there exists a maximal consistent game form $\mathscr{M}(\Gamma) = (\mathsf{Act}, \mathsf{act}, \mathscr{P}(\Phi), \mathsf{out})$ for Φ such that, for every goal assignment γ:*

1. *If $\langle\!\![\gamma]\!\!\rangle \in \Gamma$, then there is a profile $\zeta \in \Pi_{a \in \mathsf{Agt}} \mathsf{act}_a$ such that for all C in the support of γ with $\gamma(C) = \mathsf{X}\phi$, and all $\zeta' \sim_C \zeta$, we have $\phi \in \mathsf{out}(\zeta')$.*
2. *If $\neg\langle\!\![\gamma]\!\!\rangle \in \Gamma$, then for every profile $\zeta \in \Pi_{a \in \mathsf{Agt}} \mathsf{act}_a$ there is some C in the support of γ, and some $\zeta' \sim_C \zeta$, for which we have $\overline{\phi} \in \mathsf{out}(\zeta')$ where $\gamma(C) = \mathsf{X}\phi$.*

Theorem 3 (Completeness of $\mathsf{Ax}_{\mathsf{TLCGA}}$)**.** *Let Γ be a finite Ax_{TLCGA}-consistent set of* TLCGA*-formulae. Then Γ is satisfied in some concurrent game model.*

6 Finite model property, decidability and concluding remarks

The completeness proof essentially established a tree-model property and the construction of the satisfying trees can be organised in a way producing regular trees. These regular trees can be folded or modified into finite trees, this eventually proving a finite model property, and hence decidability of the logic TLCGA. An important observation here is that our proof of the one-step completeness theorem produces games in which the set of available actions has a bounded finite size. The precise complexity of satisfiability checking, which is between ExpTime and 2ExpTime, is yet to be determined.

As potential applications of the logic TLCGA we mention the specification and verification of normative systems, as well as rational synthesis and verification.

References

[1] T. Ågotnes, V. Goranko & W. Jamroga (2008): *Strategic Commitment and Release in Logics for Multi-Agent Systems (Extended abstract)*. Technical Report IfI-08-01, Clausthal University of Technology.

[2] Th. Ågotnes, V. Goranko & W. Jamroga (2007): *Alternating-time Temporal Logics with Irrevocable Strategies*. In: *Proc. of TARK XI*, Presses Universitaires de Louvain, Univ. Saint-Louis, Brussels, pp. 15–24.

[3] R. Alur, T. A. Henzinger & O. Kupferman (2002): *Alternating-Time Temporal Logic*. *J. ACM* 49(5), pp. 672–713.

[4] N. Belnap & M. Perloff (1988): *Seeing to it that: a canonical form for agentives*. *Theoria* 54, pp. 175–199.

[5] Thomas Brihaye, Arnaud Da Costa Lopes, François Laroussinie & Nicolas Markey (2009): *ATL with Strategy Contexts and Bounded Memory*. In S. Artëmov & A. Nerode, editors: *Proc. of LFCS'2009*, *LNCS* 5407, Springer, pp. 92–106.

[6] M. A. Brown (1988): *On the Logic of Ability*. *Journal of Philosophical Logic* 17, pp. 1–26.

[7] Dana Fisman, Orna Kupferman & Yoad Lustig (2010): *Rational Synthesis*. In: *Proc. of TACAS 2010*, pp. 190–204. Available at https://doi.org/10.1007/978-3-642-12002-2_16.

[8] Valentin Goranko & Sebastian Enqvist (2018): *Socially Friendly and Group Protecting Coalition Logics*. In: *Proc. of AAMAS 2018*, pp. 372–380. Available at http://dl.acm.org/citation.cfm?id=3237442.

[9] Valentin Goranko & Fengkui Ju (2019): *Towards a Logic for Conditional Local Strategic Reasoning*. In: *Proc. of LORI 2019*, *Lecture Notes in Computer Science* to appear, Springer, p. ?

[10] Wiebe van der Hoek & Michael Wooldridge (2005): *On the logic of cooperation and propositional control*. *Artif. Intell.* 164(1-2), pp. 81–119, doi:10.1016/j.artint.2005.01.003. Available at https://doi.org/10.1016/j.artint.2005.01.003.

[11] Fabio Mogavero, Aniello Murano, Giuseppe Perelli & Moshe Y. Vardi (2017): *Reasoning about Strategies: on the Satisfiability Problem*. *Logical Methods in Computer Science* 13(1).

[12] Fabio Mogavero, Aniello Murano & Moshe Y. Vardi (2010): *Reasoning About Strategies*. In: *Proc. of FSTTCS 2010*, *LIPIcs* 8, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 133–144.

[13] Rohit Parikh (1985): *The logic of games and its applications*. *North-Holland Mathematics Studies* 102, pp. 111–139.

[14] M. Pauly (2001): *Logic for Social Software*. Ph.D. thesis, University of Amsterdam.

[15] M. Pauly (2002): *A Modal Logic for Coalitional Power in Games*. *Journal of Logic and Computation* 12(1), pp. 149–166.

[16] Nicolas Troquard, Wiebe van der Hoek & Michael Wooldridge (2009): *A logic of games and propositional control*. In: *Proc. of AAMAS 2009, Volume 2*, IFAAMAS, pp. 961–968.

[17] Michael Wooldridge, Julian Gutierrez, Paul Harrenstein, Enrico Marchioni, Giuseppe Perelli & Alexis Toumi (2016): *Rational Verification: From Model Checking to Equilibrium Checking*. In: *Proc. of AAAI 2016*, pp. 4184–4191. Available at http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12268.