

The Temporal Logic of Coalitional Goal Assignments in Concurrent Multiplayer Games

SEBASTIAN ENQVIST, Department of Philosophy, Stockholm University, Sweden

VALENTIN GORANKO, Department of Philosophy, Stockholm University, Sweden and (visiting professorship) Institute for Intelligent Systems, University of Johannesburg, South Africa

We introduce and study a natural extension of the Alternating time temporal logic ATL, called *Temporal Logic of Coalitional Goal Assignments* (TLCGA). It features one new and quite expressive coalitional strategic operator, called 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, U, or G, representing a temporalised objective for the respective coalition, describing the property of the plays on which that objective is satisfied. 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 $\Sigma|_C$.

We establish fixpoint characterizations of the temporal goal assignments in a μ -calculus extension of TLCGA, discuss its expressiveness and illustrate it with some examples, prove bisimulation invariance and Hennessy–Milner property for it with respect to a suitably defined notion of bisimulation, construct a sound and complete axiomatic system for TLCGA, and obtain its decidability via finite model property.

CCS Concepts: • **Theory of computation** → **Modal and temporal logics**; • **Computing methodologies** → **Multi-agent systems**;

Additional Key Words and Phrases: Temporal logic, concurrent multi-player games, coalitional goal assignments

ACM Reference format:

Sebastian Enqvist and Valentin Goranko. 2022. The Temporal Logic of Coalitional Goal Assignments in Concurrent Multiplayer Games. *ACM Trans. Comput. Logic* 23, 4, Article 21 (October 2022), 58 pages.

<https://doi.org/10.1145/3517128>

1 INTRODUCTION

Formalising strategic reasoning has become an increasingly rich and attractive direction of active research and applications of multi-agent modal logic over the past few decades. Early logical systems capturing agents' abilities were developed with philosophical motivations and applications

Sebastian Enqvist was supported by Swedish Research Council grant 2015-01774.

Authors' addresses: S. Enqvist, Department of Philosophy, Stockholm University, Stockholms universitet Filosofiska institutionen SE-106 91, Stockholm, Sweden; email: sebastian.enqvist@philosophy.su.se; V. Goranko, Department of Philosophy, Stockholm University, Universitetsvägen 10 D, Frescati, Stockholm, Sweden and (visiting professorship) Institute for Intelligent Systems, University of Johannesburg, Cnr Kingsway & University Roads, Auckland Park, Johannesburg, South Africa; email: valentin.goranko@philosophy.su.se.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2022 Copyright held by the owner/author(s).

1529-3785/2022/10-ART21

<https://doi.org/10.1145/3517128>

in mind, including Brown’s modal logic of ability [11] and Belnap and Perloff’s STIT logic [8]. In the late 1990s—early 2000s two seminal works in the area appeared independently: Pauly’s **Coalition logic** (CL), introduced in [36, 37], and Alur, Henzinger and Kupferman’s **Alternating time temporal logic** (ATL) introduced (in its final version) in [5], cf also [17].

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 the environment) computer systems, where the agents represent concurrently executed processes. However, it was gradually adopted in the research on logic for multi-agent systems as one of the most standard and popular logical systems for reasoning about the long-term strategic abilities of agents and coalitions in concurrent multiplayer 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 operators¹ $\langle\langle C \rangle\rangle$, indexed with groups (coalitions) of agents C , such that for any formula ϕ , regarded as expressing the coalitional objective of C , the formula $\langle\langle C \rangle\rangle\phi$ 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 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 them proponents or opponents of the objective in focus. To capture these adequately, expressively richer formal logical languages are needed. In the recent precursor [23] of the present work, we proposed two such extensions of CL with additional coalitional operators, respectively implementing the following two ideas relating to cooperation and competition in a social context:

Social friendliness: agents can achieve private goals while leaving room for cooperation with the others and with the rest of the society.

Group protection: agents can cooperate within the society while simultaneously 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 one, very expressive, coalitional strategic operator, viz. the *coalitional goal assignment* operator of the type $\langle\langle \gamma \rangle\rangle$, 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 $\langle\langle \gamma \rangle\rangle$ 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

¹We use here the notation from [5], which was more widely adopted.

$\Sigma|_C$. 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 against any possible deviation of all other agents. The logic of TLCGA naturally extends ATL (in particular, CL) and exhibits richer expressiveness, both purely technical, but also in terms of potential applications (cf. Section 4). Notably, it turns out (cf. Section 3.3) that the semantics of TLCGA is more sophisticated than the semantics of ATL and most of its extensions studied so far, in the sense that it is essential that the class of (memory-based) strategies underlying the semantics of TLCGA includes all *play-based strategies* (taking into account the full history, consisting not only of the sequence of visited states but also including the action profiles causing the transitions between them) rather than only the *path-based strategies* (based on the state histories only), which has been the customary choice both for the logics in the ATL/ATL* family and for the family of Strategy logics (see further) studied so far. More precisely, the two semantics differ for TLCGA, both in terms of truth at a state in a model and in terms of validity (respectively, satisfiability), as shown in Section 3.3, where we also argue that the semantics using play-based strategies is the more natural and faithful one for the intended meaning of the operator $\langle\!\langle\cdot\rangle\!\rangle$. We also analyse there how the two semantics relate technically and show how model checking for the semantics using play-based strategies can be reduced to a model checking for the semantics using path-based strategies.

Besides TLCGA, we also introduce and study, though in less detail, two extensions of this language: TLCGA^+ , in which conjunctions of path formulas are allowed, and a fixpoint language $\mathcal{L}_\mu^{\text{XCGA}}$ in the style of the modal μ -calculus. The first of these two extensions appears naturally in some applications with a connection to game theory, described shortly. The full fixpoint language $\mathcal{L}_\mu^{\text{XCGA}}$ has a theoretical interest on its own, as a very expressive, yet quite well-behaved logic for multiplayer games. Furthermore, since the logics TLCGA and TLCGA^+ embed as fragments of $\mathcal{L}_\mu^{\text{XCGA}}$, the latter can be used as a tool to study these logics and obtain technical results about them. Here, it is used to prove decidability and finite model property of both TLCGA and TLCGA^+ , as well as to establish upper complexity bounds for their satisfiability problems.

As we demonstrate with examples in Section 4, the logic of TLCGA enables the expression of various natural and important nuanced patterns of multiplayer strategic interaction. In particular, the logic TLCGA captures a concept that we call “co-equilibrium”, which we define and promote here as a new, alternative solution concept on the border of non-cooperative and cooperative game theory [33]. We argue (in Section 4.2.2) that is more natural and applicable than the standard notion of Nash equilibrium in the context of concurrent multiplayer games with individual qualitative objectives. The existence of a co-equilibrium can be expressed quite simply in TLCGA using the operator $\langle\!\langle\gamma\rangle\!\rangle$. We also show (in Section 4.2.3) how other naturally defined notions of stable individually and coalitionally stable outcomes can be formalised in TLCGA^+ .

Further motivation for the present work comes from *cooperative game theory* [9, 13, 33]. One natural link is the apparent relationship between concurrent game models and coalitional goal assignments in them studied here, and some classes of cooperative games, such as the so-called *simple games* [38, 46], where the characteristic function defining the game assigns a payoff 0 or 1 to each coalition. An important class of such games are the *voting games* [13]. We only mention here these links with cooperative game theory, but they are left to be explored in a further work. In this article, we only illustrate briefly (in Section 4.2) the expressiveness of TLCGA and TLCGA^+ by showing how natural qualitative analogues of the key notions of *stable strategy profiles* and *core*, defined and studied for a type cooperative games based on concurrent game models in [26], can be expressed in TLCGA^+ .

Main contributions. Besides the introduction of the logic TLCGA and its extensions TLCGA⁺ into $\mathcal{L}_\mu^{\text{XCGA}}$, the main technical contributions of this article are

- Fixpoint characterizations of the main types of long-term goal assignments and translation of TLCGA and TLCGA⁺ into $\mathcal{L}_\mu^{\text{XCGA}}$.
- bisimulation invariance and Hennessy–Milner property for the logic TLCGA and for the full $\mathcal{L}_\mu^{\text{XCGA}}$ with respect to the GPCL-bisimulation introduced in [23].
- sound and complete axiomatic system for TLCGA.
- finite model property and decidability (with triple exponential bound) for the extended logic TLCGA⁺ (hence also of TLCGA).
- double exponential bound on the satisfiability problem for the fixpoint language $\mathcal{L}_\mu^{\text{XCGA}}$.
- complexity bounds for the model checking problems and the satisfiability problems for TLCGA and TLCGA⁺.

Related work and results. In addition to the links and references mentioned so far, the present work bears both conceptual and technical connections with several previously studied logics for strategic reasoning and multiplayer games, including: the logic ATL with irrevocable strategies [3, 4], ATL with strategy contexts [10], coalitional logics of cooperation and propositional control [45, 47], cooperative concurrent games [26], and especially with the family of Strategy Logics, originally introduced in [14] and further extended and studied in [2, 6, 22, 28, 29, 31], and so on. Indeed, the operator $\langle\!\langle\gamma\rangle\!\rangle$ in TLCGA⁺ with the “path-based semantics” mentioned earlier (cf. Section 3.3) can be translated to the Conjunctive Goals fragment² SL [CG] of Strategy Logic [22, 30], in a way similar to the standard translation of modal logics to first-order logic. As shown in Section 3.4, that translation can be used for applying model checking algorithms for SL [CG] to TLCGA⁺ (in particular, TLCGA) and for obtaining the 2ExpTime complexity upper bound for the model checking problem for TLCGA⁺, for both versions of its semantics. Likewise, that translation can be used for applying algorithms and obtaining PSpace complexity upper bound for the satisfiability problem in flat fragment of TLCGA⁺ with the path-based semantics by reduction to the satisfiability problem in flat fragment FSL [CG] of SL [CG], shown in [2] to be decidable and PSpace-complete. On the other hand, a direct argument in Section 9, using embedding of TLCGA and TLCGA⁺ to the μ -calculus extension $\mathcal{L}_\mu^{\text{XCGA}}$ shows that the satisfiability problem in the whole TLCGA⁺ (in particular, TLCGA), with its standard play-based semantics, is in 2ExpTime.

Notwithstanding these important technical benefits, we note that translating TLCGA to (a fragment of) Strategy Logic could generally result in an unwanted surplus of expressiveness, or even in a technical overkill, for which we have both conceptual and computational reasons to avoid, whenever possible. On the conceptual side, translating TLCGA to Strategy Logic would lose of the elegant succinctness and clear focus of the operator $\langle\!\langle\gamma\rangle\!\rangle$ as the main high-level logical construct of the language and would replace it with its low-level description in Strategy Logic.³ On the technical side, such translation would map a syntactically simple propositional language to a generally quite more expressive and syntactically heavier, essentially second-order language, explicitly involving quantification over strategies (being functions from finite sequences of 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. Thus, we eventually adopt and advocate the pragmatic approach of adhering as much as possible to the propositional logic framework of TLCGA for the purposes of expressing and reasoning about properties of strategic interactions of the type mentioned above, while resorting

²We thank an anonymous reviewer for suggesting to consider the translation of TLCGA into that fragment.

³Admittedly, this is a subjective argument and mostly a matter of taste, which we only state but do not try to impose here.

to using translation to fragments of Strategy Logic only when necessary or practically expedient, e.g., for the purpose of using already developed tools for a model or satisfiability checking in these fragments of Strategy Logic.

Our work is also essentially connected with *coalgebraic modal logic* [16, 32, 35], which is an abstract framework for modal logics of state-based evolving systems. Together with the fixpoint characterization of TLCGA, this makes TLCGA in essence a fragment of a *coalgebraic fixpoint logic* [15, 21, 48]. This connection is used to establish decidability and finite model property for our logic. Beyond that, however, our presentation is mostly self-contained, and will not require familiarity with coalgebra.

Still, we want to emphasize that the connection with coalgebraic logic, and coalgebraic fixpoint logics, in particular, is implicitly present throughout the article. In particular, the notion of *one-step completeness*, and the idea of lifting one-step completeness to completeness for the full language, is at the heart of our completeness proof. The notion of one-step completeness is inherently coalgebraic and has been studied in depth by a number of authors [35, 41, 42].

Furthermore, the fact that our translation into fixpoint logic requires only a single recursion variable means that TLCGA is a fragment of a *flat* fixpoint logic [18, 40, 43], and completeness of flat fixpoint logics can be obtained by simpler techniques than the full μ -calculus. There are two main reasons why our completeness proof is not explicitly formulated in coalgebraic terms: first, we note that TLCGA is not a flat μ -calculus per se, but rather embeds into such a logic via a fairly intricate translation. So the results in [43] do not apply directly here, as far as we can see. Second, and more importantly, we want the proof to be as self-contained and accessible without prior knowledge in coalgebraic μ -calculus as possible. It is possible that one could “transfer” completeness of flat coalgebraic μ -calculi to obtain completeness for TLCGA via our translation, but we believe a direct completeness proof is more transparent and provides a better understanding of TLCGA and its semantics.

Lastly, we also note the relationship of the present work with the logic for local conditional strategic reasoning CSR introduced in [24]. Furthermore, we point out the direct applicability of the logic TLCGA for the adequate alternative formalisation of the ideas of *rational synthesis* [20] and *rational verification* [49]. These connections and possible applications [49] are left to future work.

Structure of the article. After some preliminaries in Section 2 on concurrent game models, plays and strategies in them, in Section 3 we introduce and study the formal syntax and semantics of the logic TLCGA, and illustrate its expressiveness with some examples. In Section 5 we obtain fixpoint characterizations of the long-term goal assignments expressed in a suitable μ -calculus extension of TLCGA. We then discuss the connection with coalgebraic modal logic. In Section 6 we introduce the relevant notion of bisimulation for TLCGA and prove bisimulation invariance and the Hennessy–Milner property for it. In Section 7 we provide an axiomatic system for TLCGA for which we prove soundness and completeness. In Section 9 we show the decidability of TLCGA via finite model property. We then end with a brief concluding remarks in Section 10.

2 PRELIMINARIES AND BACKGROUND

2.1 Concurrent Game Models, Plays, Strategies

We fix a finite set of **players/agents** $\text{Agt} = \{a_1, \dots, a_n\}$ 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 any non-empty set. A **(strategic) game form over the set of outcomes** O is a tuple

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

where

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

Definition 2. A **concurrent game model** is a tuple

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

where

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

For every concurrent game model $\mathcal{M} = (S, \text{Act}, g, V)$ we define the following.

- For each $a \in \text{Agt}$ and $w \in S$, the set $\text{act}_w(a)$ consists of the **locally available actions** for a in w . It will also be denoted by $\text{act}(a, w)$. We also define the set $\text{act}_a := \bigcup_{w \in S} \text{act}_w(a)$ of **globally available actions** for a .
- An **action profile** is a tuple of actions $\zeta \in \prod_{a \in \text{Agt}} \text{act}_a$. A **locally available action profile at state** w is any tuple of locally available actions $\zeta \in \prod_{a \in \text{Agt}} \text{act}_w(a)$. The set of these action profiles will be denoted by ActProf_w .
- $\text{out}_{\mathcal{M}}$ is the **global outcome function** assigning to every state w and a local action profile ζ at w a unique **outcome** $\text{out}_{\mathcal{M}}(w, \zeta) := \text{out}_w(\zeta)$. When \mathcal{M} is fixed by the context, it will be omitted from the subscript.
- Given a coalition $C \subseteq \text{Agt}$, a **joint action** for C in \mathcal{M} is a tuple of individual actions $\zeta_C \in \prod_{a \in C} \text{act}_a$. In particular, for any action profile $\zeta \in \prod_{a \in \text{Agt}} \text{act}_a$, $\zeta|_C$ is the joint action obtained by restricting ζ to C .
- For any $w \in S$, $C \subseteq \text{Agt}$, and joint action ζ_C that is available at w , we define:

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

A **partial play**, or a **history** in \mathcal{M} is either an element of S or a finite word of the form:

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

where $w_0, \dots, w_n \in S$ and for each $i < n$, ζ_i is a locally available action profile in $\prod_{a \in \text{Agt}} \text{act}(a, w_i)$. The last state in a history h will be denoted by $l(h)$. The set of histories in \mathcal{M} is denoted by $\text{Hist}(\mathcal{M})$.

A **(memory-based) strategy for player a** is a map σ_a assigning to each history $h = w_0 \zeta_0 \dots \zeta_{n-1} w_n$ in Play an action $\sigma_a(h)$ from $\text{act}(a, w_n)$. Note that strategies are defined here in terms of *histories*, i.e., *partial plays*, not just sequences of states, as it is customary for ATL^* and in particular ATL [5], cf also [12, 17]. This distinction will turn out to be essential for the semantics of

the logic introduced here. A strategy σ_a is **memoryless**, or **positional**, if it assigns actions only based on the current (last) state, i.e., $\sigma_a(h) = \sigma_a(h')$ whenever $l(h) = l(h')$.

Given a coalition $C \subseteq \text{Agt}$, a **joint strategy** for C in the model \mathcal{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 \mathcal{M} by $\text{StratProf}_{\mathcal{M}}$, and the set of all joint strategies for a coalition C in \mathcal{M} by $\text{StratProf}_{\mathcal{M}}(C)$. Thus, $\text{StratProf}_{\mathcal{M}} = \text{StratProf}_{\mathcal{M}}(\text{Agt})$.

Given a strategy profile Σ , the **play** induced by Σ at $w \in S$ is the unique infinite word

$$\text{play}(w, \Sigma) = w_0\zeta_0w_1\zeta_1w_2\zeta_2\dots$$

such that $w_0 = w$ and, for each $n < \omega$ we have $w_{n+1} = \text{out}(\zeta_n, w_n)$, and

$$\zeta_{n+1} = \Sigma(w_0\zeta_0\dots\zeta_nw_{n+1}).$$

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

More generally, given a coalition $C \subseteq \text{Agt}$, a state $w \in S$ 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

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

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

The strategies in our semantics will be memory-based: moves of players in a strategy may depend on previous moves of other players, and players have perfect information and recall about previous moves. Indeed, as we will show in Section 3.2, just like for ATL^+ and ATL^* , but unlike ATL , the restriction to positional strategies generates different semantics for the logic TLCGA which we introduce here.

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 $\gamma : \mathcal{P}(\text{Agt}) \rightarrow G$.

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

$$\text{StateFor} : \quad \varphi := p \mid \top \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid (\varphi \vee \varphi) \mid \langle \gamma \rangle,$$

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

where $p \in \text{AP}$ and $\gamma : \mathcal{P}(\text{Agt}) \rightarrow \text{PathFor}$ is a goal assignment for Agt in PathFor . The other propositional connectives \perp , \rightarrow and \leftrightarrow , as well as the temporal operator F , are defined as usual. We write: $X\text{For}$ for the set of path formulas of the form $X\varphi$; $U\text{For}$ for the set of path formulas of the form $\varphi U \psi$; $G\text{For}$ for the path formulas of the form $G\varphi$; and $UG\text{For}$ for $U\text{For} \cup G\text{For}$.

We denote the language by $\mathcal{L}^{\text{TLCGA}}$, and its next time fragment (where PathFor is restricted to $X\text{For}$) by $\mathcal{L}^{\text{XCGA}}$. The latter is essentially (with some minor notational changes) the language of the logic GPCL introduced in [23].

Intuitively, the path formulae can be regarded as temporal goals. The goal $X\top$ is called a **trivial goal** and all other goals in PathFor are **non-trivial goals**. The family of coalitions \mathcal{F} to which the

goal assignment γ assigns non-trivial goals is called the **support of γ** , denoted $\text{Support}(\gamma)$, and γ is said to be **supported by \mathcal{F}** .

Sometimes we will write a goal assignment γ explicitly, like

$$C_1 \triangleright \theta_1, \dots, C_n \triangleright \theta_n,$$

meaning that $\text{Support}(\gamma) = \{C_1, \dots, C_n\}$ and $\gamma(C_i) = \theta_i$, for $i = 1, \dots, n$.

More notation:

- γ^\top is the **trivial goal assignment**, mapping each coalition to $X\top$.
- The goal assignment $\gamma[C \triangleright \theta]$ is like γ , but mapping C to θ .
- The goal assignment $\gamma \setminus C$ defined as $\gamma[C \triangleright X\top]$ is like γ , but excluding C from its support, by replacing its goal with $X\top$.
- 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 $X\top$.

As a convention, if γ is the unique goal assignment with empty support, we will identify the formula $\langle\!\langle \gamma \rangle\!\rangle$ with \top .

We will also consider a slight extension of the logic TLCGA which allows forming conjunctions of path formulas. This extension, which we call TLCGA^+ , has the same definition of state formulae, whereas the path formulae of TLCGA^+ are defined as follows:

$$\text{PathFor} : \quad \theta := X\varphi \mid \varphi U \varphi \mid G\varphi \mid \theta \wedge \theta.$$

In this article, we will focus mainly on TLCGA, but most of the technical results extend likewise to TLCGA^+ , and we will note that explicitly for some of them.

3.2 Semantics of TLCGA

The semantics of TLCGA is defined in terms of the truth of state formulae at a state, respectively truth of path formulae on (the path generated by) a play, in a concurrent game model $\mathcal{M} = (S, \text{Act}, g, \text{out}, V)$. 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 $\langle\!\langle \gamma \rangle\!\rangle$, is as follows, where $s \in S$:

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

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

$$\llbracket \langle\!\langle \gamma \rangle\!\rangle \rrbracket_{\mathcal{M}} = \left\{ s \in S \mid \exists \Sigma \in \text{StratProf}_{\mathcal{M}} : \text{paths}(s, \Sigma, C) \subseteq \llbracket \gamma(C) \rrbracket_{\mathcal{M}}^p \text{ for each } C \subseteq \text{Agt} \right\}.$$

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

A TLCGA formula ϕ is **valid**, denoted $\Vdash \phi$, if $\mathcal{M}, s \Vdash \phi$ for every concurrent game model \mathcal{M} and a state s in it; respectively ϕ is **satisfiable** if $\mathcal{M}, s \Vdash \phi$ for some concurrent game model \mathcal{M} and some state s in it. Likewise, (local) logical consequence and logical (semantic) equivalence in TLCGA are defined as expected.




Fig. 1. Example showing that memory is needed.

We note that the formula $\llbracket C \triangleright X \phi, \text{Agt} \triangleright X \psi \rrbracket$ is semantically equivalent to the strategic operator $[C](\phi; \psi)$ in the logic SFCL defined in [23]. Therefore, the corresponding fragment SFCL₁ of SFCL embeds into TLCGA. Note also that the strategic operator $[C]$ from Coalition logic CL is definable as a special case: $[C]\phi := [C](\phi; \top) \equiv \llbracket C \triangleright X \phi \rrbracket$.

Furthermore, the strategic operator $\llbracket C \rrbracket$ from ATL is also a special case: $\llbracket C \rrbracket \theta \equiv \llbracket C \triangleright \theta \rrbracket$. Thus, the logic ATL is embedded as a simple fragment of TLCGA.

Also, we note that a natural (downward) monotonicity condition on goal assignments can be imposed, viz. that $\Vdash \gamma(C) \rightarrow \gamma(C')$ for every coalitions C, C' , such that $C' \subseteq C$. This condition can be imposed semantically, up to equivalence, by replacing each $\gamma(C)$ with $\bigwedge_{C' \subseteq C} \gamma(C')$, though the resulting goals cannot be expressed in TLCGA, but in TLCGA⁺. A special case of monotone goal assignments, worth noting, are the “individualistic” goal assignments, where $\gamma(C) \equiv \bigwedge_{a \in C} \gamma(a)$ for every coalition C .

3.3 Variations of the Semantics of TLCGA

3.3.1 Memory-based and the Memoryless Semantics. Let us introduce ad hoc the variation $\llbracket \cdot \rrbracket_0$ of $\llbracket \gamma \rrbracket$, with semantics restricted to positional strategies, i.e., $\mathcal{M}, s \Vdash \llbracket \gamma \rrbracket_0$ iff there exists a positional strategy profile $\Sigma \in \text{StratProf}_{\mathcal{M}}$ such that, for each $C \subseteq \text{Agt}$, it holds that $\mathcal{M}, \pi \Vdash \gamma(C)$ for every $\pi \in \text{paths}(s, \Sigma, C)$.

PROPOSITION 1 (NO POSITIONAL DETERMINACY OF TLCGA). *Let $\text{Agt} = \{a, b\}$. There exist concurrent game model $\mathcal{M} = (S, \text{Act}, g, \text{out}, V)$, state $s \in S$ and a coalitional goal assignment γ , such that $\mathcal{M}, s \Vdash \llbracket \gamma \rrbracket$, but $\mathcal{M}, s \not\Vdash \llbracket \gamma \rrbracket_0$.*

Consequently, the memory-based and the memoryless semantics of $\llbracket \cdot \rrbracket$ are not equivalent.

PROOF. Consider the model $\mathcal{M} = (S, \text{Act}, g, \text{out}, V)$ on Figure 1 and a goal assignment γ , such that $\gamma(\{a, b\}) = p \cup q$ and $\gamma(\{a\}) = \top \cup \neg(p \vee q)$.


Then, $\mathcal{M}, s \Vdash \llbracket \gamma \rrbracket$, witnessed by any strategy profile Σ such that $\Sigma_a(s) = a_1$ and $\Sigma_a(ss_1s) = a_2$.

However, there is no positional strategy profile witnessing the truth of $\llbracket \gamma \rrbracket_0$ at s because any positional strategy for a would have to assign a unique action to any history ending at s , hence not enabling both the satisfaction of $p \cup q$ and of $\top \cup \neg(p \vee q)$ there. \square

Hereafter, we will only work with the memory-based semantics.

3.3.2 Semantics with Path-based Strategies vs. Play-based Strategies. Furthermore, note that (memory-based) strategies are defined here in terms of *plays*, not just paths, as it is customary for ATL^{*} and in particular ATL [5] (cf also [12, 17]), as well as for Strategy Logic [14, 28, 31]. Indeed, the two versions of strategy types affect essentially the semantics, as shown by the following example.

Example 1 (Path-based Strategies vs. Play-based Strategies). 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.

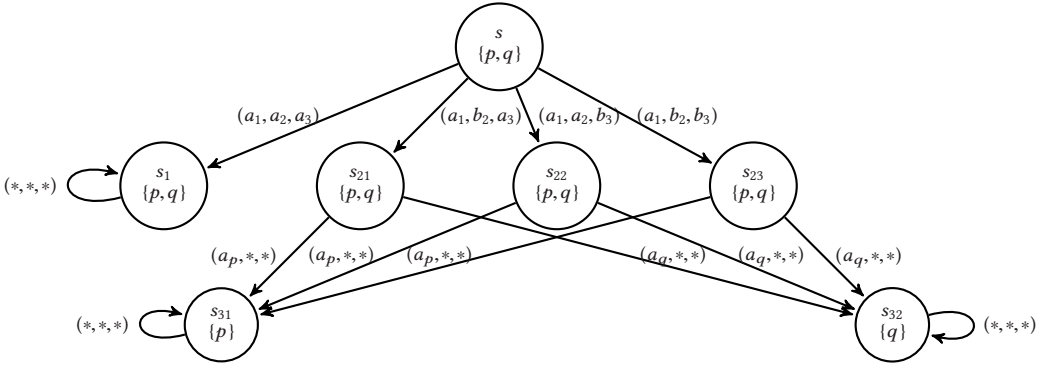


Consider the goal assignment γ , such that $\gamma(\{1, 2\}) = Gp$ and $\gamma(\{1, 3\}) = Gq$. The following hold:

- (1) $\mathcal{M}, s \models \llbracket \gamma \rrbracket$ in terms of the semantics with plays-based strategies adopted here.
Indeed, the strategy profile Σ prescribing the following action profiles: (a_1, a_2, a_3) on the play s ; $(a_p, *, *)$ on the play $s(a_1, a_2, b_3)s_2$; $(a_q, *, *)$ on the plays $s(a_1, b_2, a_3)s_2$ and $s(a_1, b_2, b_3)s_2$; and $(*, *, *)$ on any play ending at s_1, s_{31} , and s_{32} , would ensure the truth of $\llbracket \gamma \rrbracket$ at s .
- (2) $\mathcal{M}, s \not\models \llbracket \gamma \rrbracket$ in terms of the semantics with path-based strategies.
This is because player 1 does not have such a strategy for which both:
 - (a) the coalition $\{1, 2\}$ ensures the satisfaction of the goal Gp by the transition from s_2 to s_{31} , if 3 acts b_3 at s and the game goes to s_2 , and
 - (b) the coalition $\{1, 3\}$ ensures the satisfaction of the goal Gq by the transition from s_2 to s_{32} , if 2 acts b_2 at s and the game goes to s_2 .

Thus, two different semantics for TLCGA emerge. Let us call them respectively **play-based semantics**, hereafter indicated by \models_{play} , and **path-based semantics**, hereafter indicated by \models_{path} . Now, a natural question arises: which is the better/more correct one? We argue that, as the example above indicates, this is the semantics based on play-based strategies, adopted here, because only play-based strategies can detect agents' deviations from the adopted strategy profile of the grand coalition, so as to ensure that the execution of these strategies by the non-deviating agents will still guarantee the fulfilment of their individual and collective goals.

3.3.3 Relating the Path-based and the Play-based Semantics. Still, the two semantics can be reconciled in the example above by splitting off the node s_2 into 3 copies, each being the successor node of s_1 for exactly one action profile, as in Figure 3.3.3. (In fact, for this example, 2 copies suffice, just to split the outcomes from (a_1, b_2, a_3) and (a_1, a_2, b_3) .) That way, different action profiles applied at any given state lead to different successor nodes, so different plays correspond to different paths, hence the two semantics coincide in the resulting model.



This construction generalises to any concurrent game model to produce concurrent game models \mathcal{M} having the following property: for every state $s \in \mathcal{M}$ and action profiles ζ_1, ζ_2 available at s , if $\zeta_1 \neq \zeta_2$ then $\text{out}_{\mathcal{M}}(s, \zeta_1) \neq \text{out}_{\mathcal{M}}(s, \zeta_2)$. We call such models **injective**. Every non-injective concurrent game model \mathcal{M} can be transformed into an injective one $\iota(\mathcal{M})$ by generalising the construction in the example above. That can be done in various ways, but a most economical one is to multiply every state w into as many copies as the maximal number of incoming to w transitional arrows leading from any other given state u labelled with different action profiles applied at u . The labels of these copies, as well as all available action profiles at each of them and their outcomes, are copied from those at w . Thereafter, all multiply-labelled transitions from any state s to w are respectively re-directed to the different copies of w . We leave out the routine technical details of this construction. We will call this construction **state-copying and outcome-splitting (SCOS)**.

Note that plays and paths in any injective model are in a trivial 1–1 correspondence, hence play-based and path-based strategies in injective models coincide. It is then straightforward to show the following, by induction on the formulae in TLCGA^+ .

PROPOSITION 2. *In every injective concurrent game model \mathcal{M} , \Vdash_{play} and \Vdash_{path} coincide, i.e., for every state $s \in \mathcal{M}$ and a state formula $\varphi \in \text{TLCGA}^+$: $\mathcal{M}, s \Vdash_{\text{path}} \varphi$ iff $\mathcal{M}, s \Vdash_{\text{play}} \varphi$.*

We further note that the resulting model from applying the SCOS construction is TLCGA -bisimilar to the original one, in terms of the notion of TLCGA -bisimulation defined in Section 6. Consequently, due to the bisimulation invariance Theorem 3 established there, the SCOS construction preserves, inter alia, the truth of all TLCGA^+ formulae with respect to the play-based semantics, as follows.

PROPOSITION 3. *Given any concurrent game model \mathcal{M} , state $s \in \mathcal{M}$, and a copy s^i of s in the injective model $\iota(\mathcal{M})$ obtained by applying SCOS to \mathcal{M} , for every state formula $\varphi \in \text{TLCGA}^+$: $\mathcal{M}, s \Vdash_{\text{play}} \varphi$ iff $\iota(\mathcal{M}), s^i \Vdash_{\text{play}} \varphi$.*

Combining the two observations above, we obtain the following.

COROLLARY 1. *The **model-checking problem (MC)** for formulae of TLCGA^+ with the play-based semantics is reducible to the model checking problem for TLCGA^+ with the path-based semantics, at the cost of most quadratic blow-up of the size of the model.*

Observe also that in any concurrent game model \mathcal{M} , every path-based strategy is also a play-based strategy (prescribing the same action on every two plays generating the same path). Therefore, for every formula $\langle \gamma \rangle \in \text{TLCGA}^+$, in any concurrent game model \mathcal{M} and a state $s \in \mathcal{M}$: if $\mathcal{M}, s \Vdash_{\text{path}} \langle \gamma \rangle$ then $\mathcal{M}, s \Vdash_{\text{play}} \langle \gamma \rangle$.

We now show that the two semantics also differ with respect to the respective validities (hence, also with respect to the satisfiable formulae) of TLCGA. For that, we will use the idea of Example 1. Consider the goal assignment γ defined there, as well as the goal assignment γ' with support $\{\{1, 2\}, \{1, 3\}, \{1, 2, 3\}\}$, where:

$$\gamma'(\{1, 2\}) = X \{ \{1\} \triangleright G p \}, \gamma'(\{1, 3\}) = X \{ \{1\} \triangleright G q \}, \text{ and } \gamma'(\{1, 2, 3\}) = G (p \wedge q).$$

Then the following hold:

- (1) $\Vdash_{\text{play}} \langle \gamma' \rangle \rightarrow \langle \gamma \rangle$. Indeed, in any given model, a strategy profile satisfying $\langle \gamma \rangle$ can be produced from a strategy profile satisfying $\langle \gamma' \rangle$ by amending the strategy for player 1 in the former strategy profile with those strategies for 1 claimed to exist in the respective successor outcomes from the joint strategies for $\{1, 2\}$ for $\{1, 3\}$. These amendments are done as follows. Consider any successor state w of the current state s , obtained as the outcome from applying at s an action profile ζ obtained, for instance, by extending the joint action of $\{1, 2\}$ prescribed at s by the strategy profile Σ' witnessing the truth of $\langle \gamma' \rangle$ at s with any action of 3 different from the one prescribed by its strategy in Σ' . Then, the strategy for 1 is re-defined on any play of the type $\pi s \zeta w \pi'$ to prescribe the action which a strategy for 1 that is claimed by $\gamma'(\{1, 2\})$ to exist at w prescribes at the play $w \pi'$ in order to ensure the truth of $G p$ on any resulting play. The strategy for 1 on all plays passing through $\zeta' w'$ when ζ' comes from a joint action of $\{1, 3\}$ at s is re-defined likewise. Lastly, when all 3 players are following the strategy profile Σ' witnessing the truth of $\langle \gamma' \rangle$ at s , in the resulting play both $G p$ and $G q$ hold, so no strategy amendment is needed. Since the strategies that we consider are play-based, the two cases of strategy amendments are independent of each other, as they apply to disjoint sets of plays, and are therefore unproblematic to combine (which is not the case for path-based strategies, as Example 1 shows). It is now straightforward to show that the strategy profile Σ , resulting from the replacement of the strategy for 1 in Σ' with the amended strategy described above, witnesses the truth of $\langle \gamma \rangle$.
- (2) $\not\Vdash_{\text{path}} \langle \gamma' \rangle \rightarrow \langle \gamma \rangle$. Indeed, that formula fails in the model displayed in Example 1, because it is straightforward to show that $\mathcal{M}, s \Vdash_{\text{path}} \langle \gamma' \rangle$.

3.4 Standard Translation of the Path-based Semantics of TLCGA⁺ Into Fragments of Strategy Logic and the Complexity of Model Checking TLCGA

We refer here to a standard version SL of Strategy logic as defined e.g., in [28, 31], involving variables ranging over strategies that can be associated with any agents within the formulae by means of strategy assignments, and quantification over such variables. Here we mostly follow the notation for strategy assignments from [2, 30], which is slightly differ from the one in [22].

First, we show that the logic TLCGA⁺ *with the path-based semantics* can be translated to SL in a way very similar to the standard translation of modal logic to first-order logic. The key clause is the translation of the operator $\langle \gamma \rangle$, defined as follows⁴ (using the notation from [2]), for $\text{Agt} = \{a_1, \dots, a_n\}$:

$$\text{tr}(\langle \gamma \rangle) = \exists x_1 \dots \exists x_n \forall y_1 \dots \forall y_n \bigwedge_{C \subseteq \text{Agt}} b_C \gamma(C),$$

where the binding prefix b_C assigns to each agent a_i in Agt the value of the strategy variable x_i if $a_i \in C$, and the value of the strategy variable y_i , otherwise. To shorten the formula up to equivalence, the big conjunction above can be restricted to range only over those coalitions $C \subseteq \text{Agt}$ for which $\gamma(C) \neq X \top$.

⁴This form of the translation was suggested by an anonymous reviewer.

Since the translation above places only conjunctions of temporal goals in the scope of all quantifiers over strategies, TLCGA^+ translates in the Conjunctive-Goal fragment SL [CG]. Using this, we obtain the following:

PROPOSITION 4. *For each of the path-based semantics and the play-based semantics of TLCGA^+ , the respective model checking problem is PTIME-complete in the size of the model and in 2ExpTime in the size of the formula.*

PROOF. For MC in the path-based semantics, the claim follows from the translation \mathbf{tr} above and the respective results in [30, Theorem IV.2] and in [22], for the upper bounds. The lower bound in terms of the size of the model follows from the PTIME-complete complexity of MC for ATL [5], which is embedded into TLCGA. The precise complexities of MC for TLCGA and TLCGA^+ in the size of the formula are currently still open.

For the MC in the play-based semantics, the complexity bounds are obtained from those above, by using the reduction to MC in the path-based semantics provided by Corollary 1. \square

Note further that, when $\langle \gamma \rangle$ is in the **flat fragment** FTLCGA^+ of TLCGA^+ , i.e., all goals $\gamma(C)$ are conjunctions of purely temporal goals (not containing nested γ operators), the translation formula above is in the flat sub-fragment FSL [CG] of SL [CG] [2]. Consequently, that translation can also be used for applying algorithms and obtaining PSpace complexity bound for the satisfiability problem in FTLCGA^+ with the path-based semantics by reduction to the satisfiability problem in FSL [CG], shown in [2] to be decidable and PSpace-complete (notably, lower than the 2ExpTime complexity of model checking for FSL [CG], as shown in [2]). On the other hand, as we show in Section 9, using embedding of TLCGA^+ to a μ -calculus extension $\mathcal{L}_\mu^{\text{XCGA}}$ of the next time fragment $\mathcal{L}^{\text{XCGA}}$ of TLCGA, defined in Section 5.3, the satisfiability problem in $\mathcal{L}_\mu^{\text{XCGA}}$ is in 2ExpTime, hence the satisfiability problem in the whole TLCGA^+ , with its standard play-based semantics, is in 3ExpTime.

4 SOME APPLICATIONS OF TLCGA AND TLCGA^+

4.1 Examples of Expressing and Reasoning About Group Objectives with TLCGA

4.1.1 Example 1: Password Protected Data Sharing. This example is adapted from [23], where it was adapted from [34]. 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 does 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

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

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 the 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:

$$\langle\langle\{A, B\} \triangleright (H_A \leftrightarrow H_B) \cup (H_A \wedge H_B)\rangle\rangle.$$

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:

$$\langle\langle\{A, B\} \triangleright (H_A \leftrightarrow H_B) \cup (H_A \wedge H_B); A \triangleright G(H_B \rightarrow H_A); B \triangleright G(H_A \rightarrow H_B)\rangle\rangle.$$

Note that the common goal can now be simplified to the original one, to produce an equivalent to the above formula:

$$\langle\langle\{A, B\} \triangleright F(H_A \wedge H_B); A \triangleright G(H_B \rightarrow H_A); B \triangleright G(H_A \rightarrow H_B)\rangle\rangle.$$

4.1.2 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. (Spoiler alert: the answer will be gradually revealed further, so the reader may wish to pause here and think on the puzzle before reading further.)

Let us formalise the problem in TLPGA. 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 of whether the following formula is true:

$$\langle\langle\text{Sheep} \cup \text{Wolves} \triangleright (\neg e) \cup c\rangle\rangle.$$

As in the previous example, this formula is too weak to express the important subtlety that, even if such a 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:

$$\langle\langle\text{Sheep} \cup \text{Wolves} \triangleright (\neg e) \cup c; \text{Sheep} \triangleright G\neg e\rangle\rangle.$$

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

$$\langle\langle\text{Sheep} \cup \text{Wolves} \triangleright Fc; \text{Sheep} \triangleright G\neg e\rangle\rangle.$$

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, 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. For example, consider the joint strategy resulting in the play shown below:

S S S W W W	B		
S S W W		B	S W
S S S W W	B		W
S S S		B	W W W
S S S W	B		W W
S W		B	S S W W
S S W W	B		S W
W W		B	S S S W
W W W	B		S S S
W		B	S S S W W
W W	B		S S S W
		B	S S S W W W

At the very first round of this play, a sheep and a wolf cross the river together. If the wolf deviates from this action and stays instead, then two sheep are left to fend against three wolves on one side of the river. We leave it to the reader to convince themselves that any joint strategy that achieves the common goal must encounter a similar situation.

So, 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. Indeed, the joint strategy shown previously can now easily be modified to make it sheep-friendly.

4.2 Some Applications of TLCGA and TLCGA⁺ to Non-cooperative and Cooperative Game Theory

4.2.1 Expressing Nash Equilibria. The fundamental game-theoretic concept of Nash equilibrium can be applied in the concurrent games that we consider, where, given a goal assignment γ , the payoff from each play for every player is binary: 1, if that player’s goal defined by γ 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 a qualitative setting, because every strategy profile where no “loser” can deviate unilaterally to satisfy her objective is a weak Nash equilibrium. That gives no individually rational reasons for the losers to adhere to that strategy profile, because a deviation cannot be penalised any further by making their payoff even worse than it already is. Thus, we are rather sceptical about the use of Nash equilibria in such games with qualitative objectives as those considered here, and we argue further for an alternative solution concept in games with such objectives. Accordingly, even though the language of TLCGA is not designed with the explicit purpose of expressing equilibria by means of formulae, that can be done in terms of characterising the equilibria profiles as those witnessing suitable goal assignments

defined by means of the original goal assignment γ . In this way, TLCSGA also enables expressing *existence of equilibria* by means of formulae.

First, we add some terminology and notation. Let us fix a concurrent game model \mathcal{M} with initial state w . Then, any pair (Σ, γ) of a strategy profile Σ and a goal assignment γ , the pair (Σ, γ) determines a partition of the family of possible coalitions of agents $\mathcal{P}(\text{Agt})$ in two disjoint subsets: the set $W(\Sigma, \gamma)$ of **winning coalitions**, whose goals assigned by γ are satisfied in the play $\text{play}(w, \Sigma)$ in \mathcal{M} starting from w and induced by Σ , and the set $L(\Sigma, \gamma)$ of **losing coalitions**, whose goals assigned by γ are not satisfied in $\text{play}(w, \Sigma)$. In particular, the pair (Σ, γ) determines a partition of Agt in two disjoint subsets, respectively the set $W_0(\Sigma, \gamma)$ of (individual) winners and the set $L_0(\Sigma, \gamma)$ of (individual) losers from Σ with respect to γ .

We now illustrate the idea of expressing equilibria in TLCSGA in the case of a next time goal assignment α with an example for 2 players, A and B , with respective individual next time goals α_A and α_B . First, we can express an equilibrium satisfying any fixed combination of individual goals. In the case when both goals are satisfied by the equilibrium profile, no deviations can possibly improve any player's payoff, so Σ is such equilibrium profile iff Σ witnesses the goal assignment $\{A, B\} \triangleright X(\alpha_A \wedge \alpha_B)$, hence the existence of such equilibrium profile is simply expressed by the formula

$$\langle\langle \{A, B\} \triangleright X(\alpha_A \wedge \alpha_B) \rangle\rangle.$$

A more interesting case is when the equilibrium profile Σ satisfies only one goal, say α_A . This is the case precisely when Σ witnesses the goal assignment $\{A, B\} \triangleright X(\alpha_A \wedge \neg\alpha_B)$; $A \triangleright X \neg\alpha_B$. The goal $X \neg\alpha_B$ of A encodes the claim that if A follows the equilibrium strategy then B cannot deviate to satisfy its goal, thus B has no unilateral beneficial deviation from the equilibrium strategy profile. Thus, the following formula the expresses existence of such equilibrium:

$$\langle\langle \{A, B\} \triangleright X(\alpha_A \wedge \neg\alpha_B); A \triangleright X \neg\alpha_B \rangle\rangle.$$

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

$$\langle\langle \{A, B\} \triangleright X(\neg\alpha_A \wedge \neg\alpha_B); A \triangleright X \neg\alpha_B; B \triangleright X \neg\alpha_A \rangle\rangle.$$

Thus, the language of TLCSGA allows for expressing more refined descriptions of equilibria and the disjunction of all such formulae expresses the existence of any equilibrium with respect to α .

In the more general case of any set of agents Agt and pair (Σ, α) of a strategy profile Σ and a next time goal assignment α , the strategy profile Σ is a Nash equilibrium with respect to α iff Σ witnesses the goal assignment α_Σ^0 defined as follows: $\alpha_\Sigma^0(\text{Agt}) = X \bigwedge_{a \in W_0(\Sigma, \alpha)} \alpha(a)$, $\alpha_\Sigma^0(\text{Agt} \setminus \{a\}) = \neg\alpha(a)$ for each $a \in L_0(\Sigma, \alpha)$, and $\alpha_\Sigma^0(C) = \top$ in all other cases.

Then, the disjunction of all such formulae over all subsets $W_0(\Sigma, \alpha)$ of Agt expresses the existence of any equilibrium, though the size of that formula grows exponentially in the number of agents.

Lastly, for the most general case of any goal assignment γ , the idea is the same, but using TLCSGA⁺ (which allows conjunctions of path formulae of TLCSGA as goals) in order to define the goal $\gamma_\Sigma^0(\text{Agt})$. Note that in the case when all individual goals are of the type $G\psi$, the language of TLCSGA again suffices, because G distributes over conjunctions.

Thus, computing the sets $L_0(\Sigma, \gamma)$ and answering the question of whether the game has any Nash equilibria can be solved by reducing to model checking in TLCSGA⁺; in the special case when all goals are next time formulae, or all are G -type formulae, model checking in TLCSGA suffices.

4.2.2 Co-equilibria. Here we define and promote a new, alternative solution concept, that naturally arises in our framework, viz that of a “*co-equilibrium*”, which is also one of the main motivations for the introduction of the operator $\langle\langle \cdot \rangle\rangle$. Recall that an equilibrium strategy profile means

that no player can deviate individually to improve their performance, and 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. As we argued above, it does not make very good sense when the individual objectives are *qualitative*, i.e., *win* or *lose*, as losing is the worst possible outcome for the player, hence there can be no deterrent from deviation from a strategy profile where that player is losing anyway. Furthermore, players usually participate simultaneously in several coalitions with mutually consistent, yet different objectives. Assuming that they are first of all individually rational and only then collectively rational, players try to adjust their strategic behaviour so as to serve the collective objectives as much as possible while first protecting and ensuring their individual interests. In particular, all players 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, but only on the condition that pursuing it would not compromise the achievement of their individual objectives. These aspects of strategic interaction of individually rational agents serve as our motivation to define the somewhat dual to equilibrium 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.⁵ Thus, a co-equilibrium is a strongly stable solution concept that, we argue, makes better sense than a Nash equilibrium in games with *qualitative* individual objectives and existence of a co-equilibrium is an important criterion for the stability of a society of strategically interacting individually rational agents. Formally, a strategy profile Σ is a **co-equilibrium** with respect to a goal assignment γ iff Σ witnesses the goal assignment γ^* , which is the restriction of γ with support consisting of the grand coalition Agt and all singleton sets of agents. Respectively, the existence of a co-equilibrium with respect to γ can be expressed in TLCA simply as $\langle\langle\gamma^*\rangle\rangle$.

4.2.3 Expressing Other Stable Outcomes. The notion of co-equilibrium is one of a family of stable strategy profiles that can be defined by varying the notion of stability with respect to non-existence of various *beneficial deviations* of players or coalitions. In the case of co-equilibrium, no player or coalition is interested to deviate simply because they are all satisfied by the co-equilibrium strategy profile, so there are no beneficial deviations. This, of course, is the ideal case, which is often not possible in reality, so we will look at some natural relaxations of the notion of a stable outcome.

Let us fix a concurrent game model \mathcal{M} with initial state w , a coalitional goal assignment γ , and a strategy profile Σ . In Section 4.2.1, we defined winning players and coalitions in the outcome $\text{play}(w, \Sigma)$ with respect to γ , the sets $W(\Sigma, \gamma)$, $L(\Sigma, \gamma)$ and respectively $W_0(\Sigma, \gamma)$, $L_0(\Sigma, \gamma)$. In addition, we say that a coalition C is **individually winning in $\text{play}(w, \Sigma)$ with respect to γ** if $C \subseteq W_0(\Sigma, \gamma)$; respectively, C is **individually losing in $\text{play}(w, \Sigma)$ with respect to γ** if $C \subseteq L_0(\Sigma, \gamma)$.

Now, with reference to the fixed coalitional goal assignment γ , we say that a strategy profile Σ is

- (1) **individually stable at w** , if no losing player in $\text{play}(w, \Sigma)$ can deviate from Σ to become a winning player in the resulting strategy profile Σ' .

This is precisely equivalent to the standard notion of Nash equilibrium.

⁵The notion of co-equilibrium, when applied to possibly quantitative objectives, is essentially equivalent to the special case of “ t -immune strategy profile”, introduced in [1], when $t = n - 1$, where n is the number of players.

- (2) **strongly individually stable at w** , if no group of losing players in $\text{play}(w, \Sigma)$ can collectively deviate from Σ to become a group of winning players in the resulting strategy profile Σ' .

This corresponds to Aumann's notion of *strong equilibrium* [7] and is similarly expressible in TLCGA (for next time goal assignments) or TLCGA⁺ (for any goal assignments):

Σ is strongly individually stable with respect to γ iff Σ witnesses the goal assignment γ_Σ^s defined as follows:

$$\begin{aligned} \gamma_\Sigma^s(\text{Agt}) &= \bigwedge_{a \in W_0(\Sigma, \gamma)} \gamma(a), \\ \gamma_\Sigma^s(C) &= \neg \bigwedge_{a \in \bar{C}} \gamma(a) \text{ for each } C \text{ such that } \bar{C} \subseteq L_0(\Sigma, \gamma), \\ &\text{and } \gamma_\Sigma^s(C) = \top \text{ in all other cases.} \end{aligned}$$

- (3) **coalitionally stable at w** , if no losing coalition in $\text{play}(w, \Sigma)$ can collectively deviate from Σ to become a winning coalition in the resulting strategy profile Σ' .

This is similarly expressible in TLCGA (for next time goal assignments) or TLCGA⁺ (for any goal assignments):

Σ is coalitionally stable with respect to γ iff Σ witnesses the goal assignment γ_Σ defined as follows:

$$\begin{aligned} \gamma_\Sigma(\text{Agt}) &= \bigwedge_{C \in W(\Sigma, \gamma)} \gamma(C), \\ \gamma_\Sigma(C) &= \neg \gamma(\bar{C}) \text{ for each } C \text{ such that } \bar{C} \in L(\Sigma, \gamma), \\ &\text{and } \gamma_\Sigma(C) = \top \text{ in all other cases.} \end{aligned}$$

4.2.4 Cooperative Games and TLCGA. The technical ideas outlined above can also be applied to express in TLCGA (or TLCGA⁺) some key notions of the theory of *cooperative games (with transferable utility)* [9, 13, 38], including *beneficial deviations*, *stable outcomes*, and *core* of a cooperative game. We leave the exploration of these to a further work, and here we only outline one example showing how that can be done for the kind of *cooperative concurrent games* studied in [26]. These games are played in concurrent game structures, where each player has a goal expressed by a set of plays starting from some fixed initial state, regarded as the “winning plays” for that player. In particular, cooperative concurrent games with goals expressible in the linear time logic LTL (cf e.g., [17]) are studied in [26], for which it is shown that a suitably defined notion of a core of such a game can be logically characterised using the logic ATL* and the computational complexities of certain decision problems associated with that core have been established.

Let us fix a state w in a concurrent game model \mathcal{M} and consider the cooperative concurrent game $G = G(\mathcal{M}, w)$ generated at w in \mathcal{M} and a strategy profile Σ in G . We note that the terminology from [26], differs somewhat from ours, as they call a winning (respectively, losing) coalition⁶ what we call in Section 4.2.3 *individually winning* (respectively, *individually losing*) coalition. For consistency, we will adhere to our terminology.

To make the parameter w referring to the current initial state explicit, we will denote the set of all losing players in $\text{play}(w, \Sigma)$ with respect to γ by $L_0(w, \Sigma, \gamma)$. Thus, $C \subseteq L_0(w, \Sigma, \gamma)$, i.e., the coalition C is individually losing in $\text{play}(w, \Sigma)$, iff Σ witnesses the TLCGA⁺ goal assignment $\text{Agt} \triangleright \bigwedge_{i \in C} \neg \gamma(i)$. In the case when all $\gamma(i)$ are next time formulae $X\psi(i)$, that goal assignment can be replaced by the TLCGA goal assignment $\text{Agt} \triangleright X \bigwedge_{i \in C} \neg \psi(i)$.

⁶Only individual, but no coalitional goals are considered in [26].

Now, if $C \subseteq L_0(w, \Sigma, \gamma)$, then an *individually beneficial deviation*⁷ for C is any joint strategy Σ_C of C that guarantees for C to be individually winning on any outcome play from w induced by Σ_C . The **core of the game** G , denoted $core(G)$, is defined as the set of **stable strategy profiles** in G , viz. those that admit no beneficial deviations (by any losing coalition). Thus, C has a beneficial deviation at the state w iff $\langle \gamma_C \rangle$ is true at w , where γ_C is the goal assignment $C \triangleright \bigwedge_{i \in C} \gamma(i)$. Therefore, the existence of a coalition that has a beneficial deviation with respect to γ (and Σ) can be expressed as the disjunction of all formulae $\langle \gamma_C \rangle$, overall individually losing coalitions $C \subseteq L_0(w, \Sigma, \gamma)$. Then, Σ is in $core(G)$ iff it does not satisfy that formula at w :

$$\Sigma \in core(G) \quad \text{iff} \quad \mathcal{M}, w \Vdash \bigwedge_{C \subseteq L_0(w, \Sigma, \gamma)} \neg \langle \gamma_C \rangle.$$

Thus, like in the case of Nash equilibria and the other stable outcomes defined in Section 4.2.3, computing the sets $L_0(w, \Sigma, \gamma)$ and answering the question of whether the core of such a game is non-empty can be solved by reducing to model checking in TLCGA⁺; in the special case when all goals are next time formulae, or all are G-type formulae, model checking in TLCGA suffices.

5 FIXPOINT CHARACTERIZATIONS OF TEMPORAL FORMULAE IN TLCGA

In this section, we will show how to embed the logic TLCGA into a suitable fixpoint logic. In fact, the embedding can be extended likewise to the logic TLCGA⁺, and we sketch along the way how this is done. The results in this section will be stated and proved for TLCGA, but their extensions to TLCGA⁺ are quite routine.

5.1 Types of Goal Assignments

Definition 3. A goal assignment γ supported by a family of coalitions \mathcal{F} will be called **long-term temporal** if γ maps every coalition in \mathcal{F} either to a U-formula or a G-formula, that is, if $\gamma[\mathcal{F}] \subseteq \text{UGFor}$, where $\gamma[\mathcal{F}] = \{\gamma(C) \mid C \in \mathcal{F}\}$.

A goal assignment is called **local**, or **next time**, if γ maps every coalition in \mathcal{F} to a X-formula, i.e., $\gamma[\mathcal{F}] \subseteq \text{XFor}$.

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.

To extend this definition to TLCGA⁺, we say that a goal assignment γ in the extended language is long-term temporal if for each coalition C , each conjunct of $\gamma(C)$ is either a U-formula or a G-formula. Clearly, this reduces to the previous definition for the special case of goal assignments in TLCGA.

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

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

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

Again, this can be extended to TLCGA⁺: we say that a long-term temporal goal assignment γ is of type U if there is some coalition C in the support of γ , such that at least one conjunct of $\gamma(C)$ is an U-formula. Otherwise, γ is of type G.

⁷Note that this notion, defined in [26] as “beneficial deviation”, does not actually depend on Σ in any other way but that C is losing in $\text{play}(w, \Sigma)$.

5.2 The Fixpoint Property of Goal Assignments

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

To extend this notion to TLCGA^+ , we define $\gamma|_{\text{UGFor}}(\gamma|_{\text{XFor}})$ by setting, for each coalition C , $\gamma|_{\text{UGFor}}(C)$ to be the conjunction of all formulas $\alpha \cup \beta$ or $G\chi$ that appear as conjuncts of $\gamma(C)$, provided that $\gamma(C)$ has at least one such conjunct, and $\gamma|_{\text{UGFor}}(C) = X\top$ otherwise. The goal assignment $\gamma|_{\text{XFor}}$ is defined similarly.

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

$$\Delta\gamma(C) := X \left(\bigwedge \{ \varphi \mid \text{there exists } C' \in \mathcal{F}, C' \subseteq C \text{ such that } \gamma(C') = X\varphi \} \wedge \llbracket (\gamma|_C)|_{\text{UGFor}} \rrbracket \right),$$

whereas 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 $\llbracket \gamma \rrbracket$ where γ is the empty goal assignment (the right conjunct reduces to \top). For all coalitions that are not in $\text{sup } \Delta\gamma$, $\Delta\gamma$ assigns the trivial goal. Given any formula ϕ , we will sometimes abbreviate the formula $\Delta\gamma[\cup \mathcal{F} \triangleright X\phi]$ by $\Delta\gamma\{\phi\}$.

The definition above may look a bit opaque, but what it does is quite simple. Intuitively, the goal assignment $\Delta\gamma$ describes the conditions that each coalition must ensure to hold for the *next state*, with respect to a given strategy profile Σ , in order for that strategy profile to fulfill the goal assignment γ at the *current state*. To make things more concrete, let us consider two examples.

Example 2. If $\llbracket \gamma \rrbracket$ is $\llbracket C \triangleright \varphi \cup \psi \rrbracket$, then $\llbracket \Delta\gamma \rrbracket = \llbracket C \triangleright X \llbracket C \triangleright \varphi \cup \psi \rrbracket \rrbracket = \llbracket C \triangleright X \llbracket \gamma \rrbracket \rrbracket$. So in this special case, the next time-extension simply pushes the eventuality $\varphi \cup \psi$ one step into the future, so to speak. Similarly, if $\llbracket \gamma \rrbracket$ is $\llbracket C \triangleright G\varphi \rrbracket$, then $\llbracket \Delta\gamma \rrbracket = \llbracket C \triangleright X \llbracket C \triangleright G\varphi \rrbracket \rrbracket = \llbracket C \triangleright X \llbracket \gamma \rrbracket \rrbracket$.

Example 3. Consider the example of a goal assignment γ supported by $\mathcal{F} = \{\{a, b\}, \{c\}, \{b, c\}\}$ and defined by the assignment:

$$\{a, b\} \triangleright pUq, \quad \{c\} \triangleright Gr, \quad \{b, c\} \triangleright Xs.$$

The support of $\Delta\gamma$ will be $\mathcal{F} \cup \{\{a, b, c\}\}$. The action of $\Delta\gamma$ is shown below:

- $\{a, b\} \triangleright X \llbracket \{a, b\} \triangleright pUq \rrbracket$.
- $\{c\} \triangleright X \llbracket \{c\} \triangleright Gr \rrbracket$.
- $\{b, c\} \triangleright X(s \wedge \llbracket \{c\} \triangleright Gr \rrbracket)$.
- $\{a, b, c\} \triangleright X(s \wedge \llbracket \{a, b\} \triangleright pUq, \{c\} \triangleright Gr \rrbracket)$.

The procedure for computing $\Delta\gamma$ goes, informally, as follows: for each \subseteq -downset⁸ \mathcal{D} of coalitions from \mathcal{F} , we collect all the formulas φ for which some coalition in \mathcal{D} is mapped to the goal $X\varphi$ into a conjunction, add a conjunct collecting all the long-term goals for coalitions in \mathcal{D} into a single goal assignment, and finally put the resulting conjunction in the scope of an X -operator and assign this goal to the union of \mathcal{D} .

We are now ready to define one of the key concepts of the article.

⁸I.e., set \mathcal{D} such that for all $Z \in \mathcal{D}$ and all $Z' \in \mathcal{F}$, if $Z' \subseteq Z$ then $Z' \in \mathcal{D}$.

Definition 7. Let γ be a goal assignment, supported by \mathcal{F} . Then we define the following formula:

$$\text{unfold}(\gamma) := \bigvee \text{Finish}(\gamma) \vee \left(\bigwedge \text{UHold}(\gamma) \wedge \bigwedge \text{GHold}(\gamma) \wedge \langle\langle \Delta\gamma \rangle\rangle \right),$$

where

- $\text{Finish}(\gamma) := \{ \beta \wedge \langle\langle \gamma \setminus C \rangle\rangle \mid \gamma(C) = \alpha \cup \beta \}$.
- $\text{ugam} := \{ \alpha \mid \gamma(C) = \alpha \cup \beta, \text{ for some } C, \beta \}$.
- $\text{GHold}(\gamma) := \{ \chi \mid \gamma(C) = G\chi, \text{ for some } C \}$.

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

$$\bigwedge \text{GHold}(\gamma) \wedge \langle\langle \Delta\gamma \rangle\rangle.$$

We call $\text{unfold}(\gamma)$ **the unfolding formula** of γ .

The formula $\text{unfold}(\gamma)$ may require some additional explanation. We shall see that $\text{unfold}(\gamma)$ is in fact equivalent to $\langle\langle \gamma \rangle\rangle$, and the formula can be seen as an analysis of the different possibilities for *how* a given strategy profile may fulfil the goal assignment γ . The first disjunction $\text{Finish}(\gamma)$ describes those cases where one of the coalitions C has an eventuality $\alpha \cup \beta$ as its goal formula; if the formula β happens to be true at the current state then the coalition C trivially attains its goal regardless of its actions. So, for a strategy profile to fulfil the goal assignment γ , it suffices that it fulfils the goals of all coalitions besides C , i.e., the goal assignment $\gamma \setminus C$. In the remaining case, the conditions that a strategy profile must satisfy in order to fulfil the goal assignment γ are divided into two parts: “local” conditions that must be true at the *current state*, which are outside the agents’ control, and those conditions that each coalition must ensure for the *next state*. The latter conditions are described by the formula $\langle\langle \Delta\gamma \rangle\rangle$, as explained earlier. The local conditions are derived from the interpretation of temporal (path) formulas: if the goal of coalition C is an eventuality $\alpha \cup \beta$, and this goal is not trivially fulfilled because β happens to be true, then the goal can only be attained if α is true at the current state. Similarly, a goal $G\chi$ can only be fulfilled if χ is currently true. These conditions are captured by the conjuncts $\text{UHold}(\gamma)$ and $\text{GHold}(\gamma)$, respectively.

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

$$\text{ind}(\gamma, \phi) := \bigvee \text{Finish}(\gamma) \vee \left(\bigwedge \text{UHold}(\gamma) \wedge \bigwedge \text{GHold}(\gamma) \wedge \langle\langle \Delta\gamma \{ \phi \} \rangle\rangle \right),$$

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

PROPOSITION 5. *For every long-term temporal goal assignment γ we have:*

$$\text{unfold}(\gamma) = \text{ind}(\gamma, \langle\langle \gamma \rangle\rangle).$$

PROOF. If γ is long-term temporal and supported by \mathcal{F} then we get:

$$(\gamma \upharpoonright_{\text{UGFor}}) \upharpoonright_{\cup \mathcal{F}} = \gamma.$$

Thus, since there are no next time formulas to consider, $\Delta\gamma$ will map $\cup \mathcal{F}$ to $X\langle\langle \gamma \rangle\rangle$, hence $\Delta\gamma[\cup \mathcal{F} \triangleright X\langle\langle \gamma \rangle\rangle] = \Delta\gamma$. \square

It is not hard to see that, if γ is a next time goal assignment, then $\text{unfold}(\gamma) \equiv \llbracket \gamma \rrbracket$. For example, suppose γ is supported by $\{\{a\}, \{b\}\}$ and maps $\{a\}$ to Xp and $\{b\}$ to Xq . Then $\text{unfold}(\gamma)$ is equal to $\llbracket \Delta\gamma \rrbracket$, which is the following formula:

$$\llbracket \{a\} \triangleright Xp, \{b\} \triangleright Xq, \{a, b\} \triangleright X(p \wedge q) \rrbracket,$$

which is clearly equivalent to $\llbracket \gamma \rrbracket = \llbracket \{a\} \triangleright Xp, \{b\} \triangleright Xq \rrbracket$. In fact, the equivalence always holds; this is by design and will play a key role in our axiomatization.

THEOREM 1 (FIXPOINT PROPERTY). *For any goal assignment γ :*

$$\llbracket \gamma \rrbracket \equiv \text{unfold}(\gamma),$$

and hence for any long-term temporal goal assignment γ :

$$\llbracket \gamma \rrbracket \equiv \text{ind}(\gamma, \llbracket \gamma \rrbracket).$$

PROOF. We prove each implication separately.

Left to right: suppose that $\mathcal{M}, s \Vdash \llbracket \gamma \rrbracket$, where $\mathcal{M} = (S, \text{Act}, g, \text{out}, V)$, and let Σ be some profile witnessing γ at s . Assuming that $\mathcal{M}, s \not\Vdash \bigvee \text{Finish}(\gamma)$, we show that

$$\mathcal{M}, s \Vdash \bigwedge \text{UHold}(\gamma) \wedge \bigwedge \text{GHold}(\gamma) \wedge \llbracket \Delta\gamma \rrbracket.$$

We treat these conjuncts separately. First, note that if $\text{UHold}(\gamma) = \emptyset$ or $\text{GHold}(\gamma) = \emptyset$ then these conjunctions reduce to \top and hence are trivially satisfied.

Suppose $\alpha \in \text{UHold}(\gamma)$. Then there is some coalition C and some β for which $\gamma(C) = \alpha \cup \beta$. The set $\text{paths}(s, \Sigma, C)$ are always non-empty, so consider an arbitrary member π and recall that its first element is s . Since $\Sigma, s \Vdash \gamma$ it follows that $\pi \Vdash \alpha \cup \beta$. Since we assumed that $\mathcal{M}, s \not\Vdash \bigvee \text{Finish}(\gamma)$, we cannot have $\mathcal{M}, s \Vdash \beta$ since this would give $\mathcal{M}, s \Vdash \beta \wedge \llbracket \gamma \rrbracket$ which entails $\mathcal{M}, s \Vdash \beta \wedge \llbracket \gamma|_C \rrbracket$, and $\beta \wedge \llbracket \gamma|_C \rrbracket$ is a member of $\text{Finish}(\gamma)$. Hence we have $\mathcal{M}, s \Vdash \alpha$, as required. The proof that each conjunct from $\text{GHold}(\gamma)$ is satisfied is similar (but simpler).

We now show that $\Sigma, s \Vdash \Delta\gamma$. Pick an arbitrary coalition C in the support of $\Delta\gamma$ and an arbitrary path $\pi \in \text{paths}(s, \Sigma, C)$. We need to show that $\pi \Vdash \Delta\gamma(C)$. Suppose that π is the path generated by a play in $\text{Plays}(w, \Sigma, C)$ of the form:

$$w_0 \zeta_0 w_1 \zeta_1 w_2 \dots$$

where $w_0 = s$. We need to show that

- (1) For each $C' \subseteq C$ in the support of γ , if $\gamma(C') = X\psi$ then $\mathcal{M}, w_1 \Vdash \psi$,
- (2) $\mathcal{M}, w_1 \Vdash \llbracket (\gamma|_{\text{UGFor}})|_C \rrbracket$.

The first item is straightforward, so we focus on item (2). We need to come up with a strategy profile Σ' that witnesses $(\gamma|_{\text{UGFor}})|_C$ at w_1 . We define Σ' as follows. Given a history of the form:

$$v_0 \zeta'_0 v_1 \dots v_{n-1} \zeta'_{n-1} v_n,$$

where $v_0 = w_1$, and a player $a \in C$, we set:

$$\Sigma'(a, v_0 \zeta'_0 v_1 \dots v_{n-1} \zeta'_{n-1} v_n) := \Sigma(a, w_0 \zeta_0 v_0 \zeta'_0 v_1 \dots v_{n-1} \zeta'_{n-1} v_n).$$

Now let $C' \subseteq C$ be a coalition for which $\gamma(C') = \alpha \cup \beta$, and let $\pi' \in \text{paths}(w_1, \Sigma', C')$. Then $s\pi' \in \text{paths}(s, \Sigma, C')$ by the construction of Σ' , hence $s\pi' \Vdash \alpha \cup \beta$. Since $\mathcal{M}, s \not\Vdash \beta$, we get $\pi' \Vdash \alpha \cup \beta$ as well. Similarly we can show that if $C' \subseteq C$ is a coalition for which $\gamma(C') = G\alpha$, and $\pi' \in \text{paths}(w_1, \Sigma', C')$, then $\pi' \Vdash G\alpha$. This shows that $\Sigma', w \Vdash (\gamma|_{\text{UGFor}})|_C$, as required.

Right to left: Suppose $\mathcal{M}, s \Vdash \langle\langle \text{unfold}(\gamma) \rangle\rangle$. We show that $\mathcal{M}, s \Vdash \langle\langle \gamma \rangle\rangle$.

There are two cases to consider: either some formula in $\text{Finish}(\gamma)$ holds at s , or:

$$\mathcal{M}, s \Vdash \bigwedge \text{U Holds}(\gamma) \wedge \bigwedge \text{G Holds}(\gamma) \wedge \langle\langle \Delta\gamma \rangle\rangle.$$

In the first case, there is some C in the support \mathcal{F} of γ for which $\gamma(C) = \alpha\cup\beta$ and:

$$\mathcal{M}, s \Vdash \beta \wedge \langle\langle \gamma \setminus C \rangle\rangle.$$

But then $\alpha\cup\beta$ holds at any path beginning from s , and it follows that $\mathcal{M}, s \Vdash \langle\langle (\gamma \setminus C)[C \triangleright \alpha\cup\beta] \rangle\rangle$. This formula is the same as $\langle\langle \gamma \rangle\rangle$, so we are done.

Suppose the second, more challenging case, and fix a strategy profile Σ for Agt witnessing $\langle\langle \Delta\gamma \rangle\rangle$ at s . Given any locally available action profile $\zeta \in \text{ActProf}_s$, we define the set $\text{fol}(\zeta)$ of **followers of Σ relative to ζ** :

$$\text{fol}(\zeta) := \bigcup \{C \in \mathcal{F} \mid \zeta(a) = \Sigma(a, s) \text{ for all } a \in C\}.$$

Note that $\text{fol}(\zeta)$ belongs to the support of $\Delta\gamma$. By unfolding the definition of $\langle\langle \Delta\gamma \rangle\rangle$, we see that the following conditions hold for each locally available action profile $\zeta \in \text{ActProf}_s$ and any $C \in \mathcal{F}$ such that $C \subseteq \text{fol}(\zeta)$:

- (1) $\text{out}(\zeta, s) \in \llbracket \langle\langle (\gamma|_{\text{UGFor}})|_C \rangle\rangle \rrbracket_{\mathcal{M}}$.
- (2) If $\gamma(C) = X\varphi$ then $\text{out}(\zeta, s) \in \llbracket \langle\langle \varphi \rangle\rangle \rrbracket_{\mathcal{M}}$.

This motivates the following definition: for each locally available action profile $\zeta \in \text{ActProf}_s$ we pick a strategy profile $\underline{\zeta}$ defined for all players in $\bigcup \mathcal{F}$, such that for all $C \subseteq \text{fol}(\zeta)$:

$$\underline{\zeta}, \text{out}(\zeta, s) \Vdash (\gamma|_{\text{UGFor}})|_C.$$

Now, we will build a strategy profile Ω using the strategy profile Σ and all strategy profiles $\underline{\zeta}$ for each possible locally available action profile ζ at s :

Given a player $a \in \bigcup \mathcal{F}$ and $w \in S$, let $\Omega(a, w) = \Sigma(a, w)$ if $w = s$, and some arbitrary available move otherwise. For a history of the form $w_0\zeta_0 \dots \zeta_n w_{n+1}$, if $w_0 \neq s$ then we can again define the move of player a arbitrarily. Otherwise, we set

$$\Omega(a, w_0\zeta_0 w_1 \dots \zeta_n w_{n+1}) := \underline{\zeta}_0(a, w_1\zeta_1 \dots \zeta_n w_{n+1}).$$

We will show that the strategy profile Ω witnesses $\langle\langle \gamma \rangle\rangle$ at s . Let $C \in \mathcal{F}$, and let $\pi \in \text{paths}(s, \Omega, C)$. We need to show that $\pi \Vdash \gamma(C)$. The case where $\gamma(C) = X\varphi$ is immediate, by definition of $\Delta\gamma$ and of Ω at s . We focus on the case where $\gamma(C)$ is of the form $\alpha\cup\beta$. We assume that π is the path generated by a play in $\text{Plays}(s, \Omega, C)$ of the form:

$$w_0\zeta_0 w_1\zeta_1 w_2 \dots$$

so that π equals $w_0 w_1 w_2 \dots$, $w_0 = s$ and $w_1 = \text{out}(s, \zeta_0)$. But then $C \subseteq \text{fol}(\zeta_0)$, and by the construction of the strategy profile Ω , the play $w_1\zeta_1 w_2\zeta_2 w_3 \dots$ belongs to $\text{Plays}(w_1, \underline{\zeta}_0, C)$. So the path $w_1 w_2 w_3 \dots$ satisfies $\alpha\cup\beta$ since, by definition, $\underline{\zeta}_0, w_1 \Vdash (\gamma|_{\text{UGFor}})|_C$. Since $\alpha \in \text{U Holds}(\gamma)$, we have $\mathcal{M}, s \Vdash \alpha$. Since $s = w_0$ it follows that $\pi \Vdash \alpha\cup\beta$ is required. Lastly, the case where $\gamma(C)$ is of the form $G\chi$ is analogous. \square

Finally, we show how to extend these definitions to the logic TLCGA^+ . In this setting, goal assignments are of a more general kind as they map coalitions to conjunctions of X -formulas, U -formulas and/or G -formulas, so we have to account for this slight complication. Given a goal assignment γ , and a coalition C in its support, we can think of $\gamma(C)$ as the set of its conjuncts of the form $X\alpha$, $\alpha\cup\beta$ or $G\chi$. So we abuse notation slightly and write $\theta \in \gamma(C)$ to say that θ has one of these three forms, and is a conjunct of $\gamma(C)$. With this notation in place, we extend the definition

of the next time-extension $\Delta\gamma$ of a goal assignment γ as follows. The support of $\Delta\gamma$ is defined as before, and for each $C \in \text{supp } \Delta\gamma$ we define

$$\Delta\gamma(C) := X \left(\bigwedge \{ \varphi \mid \text{there exists } C' \in \mathcal{F}, C' \subseteq C \text{ such that } X\varphi \in \gamma(C') \} \wedge \llbracket (\gamma|_C) \upharpoonright_{\text{UGFor}} \rrbracket \right),$$

As before we abbreviate the formula $\Delta\gamma[\bigcup \mathcal{F} \triangleright X\phi]$ by $\Delta\gamma\{\phi\}$.

Next, we define the unfolding of a goal assignment: Let γ be a goal assignment, supported by \mathcal{F} . Given a coalition C and a conjunct θ of $\gamma(C)$, we write $\gamma \setminus (C \triangleright \theta)$ for the goal assignment that is like γ , except:

$$(\gamma \setminus (C \triangleright \theta))(C) = \bigwedge \{ \theta' \mid \theta' \in \gamma(C) \wedge \theta' \neq \theta \}.$$

We now define:

$$\text{unfold}(\gamma) := \bigvee \text{Finish}(\gamma) \vee \left(\bigwedge \text{U Holds}(\gamma) \wedge \bigwedge \text{G Holds}(\gamma) \wedge \llbracket \Delta\gamma \rrbracket \right),$$

where

- $\text{Finish}(\gamma) := \{ \beta \wedge \llbracket \gamma \setminus (C \triangleright \alpha \text{U} \beta) \rrbracket \mid \gamma(C) = \alpha \text{U} \beta \}$.
- $\text{U Holds}(\gamma) := \{ \alpha \mid \alpha \text{U} \beta \in \gamma(C), \text{ for some } C, \beta \}$.
- $\text{G Holds}(\gamma) := \{ \chi \mid \text{G} \chi \in \gamma(C), \text{ for some } C \}$.

The induction formula for $\text{ind}(\gamma, \phi)$ is defined as before, by

$$\text{ind}(\gamma, \phi) := \bigvee \text{Finish}(\gamma) \vee \left(\bigwedge \text{U Holds}(\gamma) \wedge \bigwedge \text{G Holds}(\gamma) \wedge \llbracket \Delta\gamma\{\phi\} \rrbracket \right),$$

where $\Delta\gamma\{\phi\}$ is $\Delta\gamma[\bigcup \mathcal{F} \triangleright X\phi]$. With these extended definitions, the proofs of Proposition 5 and Theorem 1 go through as before.

5.3 A μ -calculus of Goal Assignments

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

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

$$\text{StateFor} : \quad \varphi := p \mid \top \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid (\varphi \vee \varphi) \mid \llbracket \gamma \rrbracket \mid \mu x. \varphi,$$

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

Here, in $\mu x. \varphi$ the formula φ is subject to the usual constraint that every occurrence of the variable x in φ is positive, in the sense that it is under the scope of an even number (possibly zero) of negations. We usually denote bound variables $x, y, z \dots$ rather than p, q, r, \dots , but formally we do not introduce a separate supply of fixpoint variables. In the formula $\mu x. \varphi$ the variable x is simply a propositional variable. We define the greatest fixpoint operator as usual:

$$\nu x. \varphi := \neg \mu x. \neg \varphi[\neg x/x],$$

where $\varphi[\neg x/x]$ is the result of the uniform substitution of $\neg x$ for x in φ .

A model for $\mathcal{L}_\mu^{\text{TLCGA}}$ is just like a model for $\mathcal{L}^{\text{TLCGA}}$, viz a tuple, $\mathcal{M} = (S, \text{Act}, \text{g}, \text{out}, V)$, but now the valuation V assigns values to the variable(s) z used in formulae $\mu z. \psi$. Now, for each $Z \subseteq S$, we define the amended valuation $V^Z := V[z \mapsto Z]$, which is like V except that it maps z to Z . We will denote $\mathcal{M}^Z := (S, \text{Act}, \text{g}, \text{out}, V^Z)$ and for any formula $\psi(z)$ in which the variable z may occur free, we will write $\mathcal{M}, s \Vdash \psi(Z)$ to state that $\mathcal{M}^Z, s \Vdash \psi(z)$.

The semantics of $\mathcal{L}_\mu^{\text{TLCGA}}$ extends that of $\mathcal{L}^{\text{TLCGA}}$ with the additional clause that the extension $\llbracket \mu z. \varphi(z) \rrbracket$ of a least fixpoint-formula in a model $\mathcal{M} = (S, \text{Act}, g, \text{out}, V)$ is given by

$$\llbracket \mu z. \varphi(z) \rrbracket := \bigcap \{ Z \subseteq S \mid \llbracket \varphi(Z) \rrbracket \subseteq Z \},$$

where, as expected:

$$\llbracket \varphi(Z) \rrbracket := \{ w \in S \mid \mathcal{M}^Z, w \Vdash \varphi(z) \}.$$

Given a model $\mathcal{M} = (S, \text{Act}, g, \text{out}, V)$, let f_γ be the monotone map on $\mathcal{P}(S)$ induced by $\text{ind}(\gamma, z)$ in the usual way, i.e., for each $Z \subseteq S$, $f_\gamma(Z)$ is the set of states satisfying $\text{ind}(\gamma, z)$ with respect to the amended valuation $V^Z := V[z \mapsto Z]$.

The proofs of the following two propositions are in the appendix.

PROPOSITION 6 (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 not occurring in $\langle \gamma \rangle$. Then $\langle \gamma \rangle \equiv \mu z. \text{ind}(\gamma, z)$.*

PROPOSITION 7 (FIXPOINT CHARACTERIZATION OF TypeG TEMPORAL GOAL ASSIGNMENTS). *Suppose that γ is a long-term temporal goal assignment in TypeG. Then $\langle \gamma \rangle \equiv \nu z. \text{ind}(\gamma, z)$.*

Recall that (memory-based) strategies are defined here in terms of *plays*, not just paths as for ATL^* in [5] (cf also [12, 17]), and we showed in Example 1 that the two versions affect essentially the semantics. In fact, the model in that example also shows that $\nu z. \text{ind}(\gamma, z) \rightarrow \langle \gamma \rangle$ is not valid in the semantics with path-based strategies.

Using the fixpoint characterizations of long-term temporal goal assignments we have established here, we can define an explicit translation $t : \mathcal{L}^{\text{TLCGA}} \rightarrow \mathcal{L}_\mu^{\text{XCGA}}$, preserving the semantics of each $\mathcal{L}^{\text{TLCGA}}$ -formula. To make this precise, the translation has to be defined by induction on a certain well-founded order $<$ over $\mathcal{L}^{\text{TLCGA}}$ -formulas, defined as the smallest transitive relation closed under the following rules:

- If φ is a proper subformula of ψ then $\varphi < \psi$.
- If γ is a goal assignment which is neither next time nor long-term temporal, i.e., the support of γ contains both some coalitions mapped to next time formulas and some coalitions mapped to long-term temporal goals, then $\text{unfold}(\gamma) < \langle \gamma \rangle$.
- If γ is a TypeU or TypeG-goal assignment, and z is a propositional variable not occurring in $\langle \gamma \rangle$, then $\text{ind}(\gamma, z) < \langle \gamma \rangle$.

To see that this order is well-founded, consider for example the clause where γ is in TypeU. Then the formula $\text{ind}(\gamma, z)$ can be viewed as being built up using Boolean connectives from proper subformulas of $\langle \gamma \rangle$ and formulas of the form $\langle \gamma' \rangle$ where γ' is either $\Delta\gamma\{z\}$ or $\gamma|_C$ for some coalition C in the support of γ . Note that the goal assignment $\gamma|_C$ has a smaller support than γ . Furthermore, the goal assignment $\Delta\gamma\{z\}$ is a next time goal assignment, the coalition $\bigcup \mathcal{F}$ is mapped to Xz (where \mathcal{F} is the support of γ) and each coalition $C \neq \bigcup \mathcal{F}$ in the support of $\Delta\gamma\{z\}$ is mapped to $X \langle \gamma|_C \rangle$. But each goal assignment $\gamma|_C$ for $C \neq \bigcup \mathcal{F}$ has smaller support than γ , since if C a proper subset of $\bigcup \mathcal{F}$ then there is some $C' \in \mathcal{F}$ which is not contained in C . So in $\gamma|_C$, the coalition C' will be assigned the trivial goal $X \top$.

With the well-founded order $<$ in place, the translation can be defined by $t(\langle \gamma \rangle) = \mu z. t(\text{ind}(\gamma, z))$ if γ is a type U goal assignment, $t(\langle \gamma \rangle) = \nu z. t(\text{ind}(\gamma, z))$ if γ is of type G and $t(\langle \gamma \rangle) = t(\text{unfold}(\gamma))$ if γ contains both U-formulas and G-formulas as goals. The induction steps for Booleans and next time goal assignments are handled in the standard manner, simply letting the translation commute with these connectives.

As noted at the beginning of the section, it is a relatively routine task to extend our translation to cover the richer language $\mathcal{L}^{\text{TLCGA}^+}$ underlying the logic TLCGA^+ . Again, the translation is by induction on a well-founded ordering over formulas, defined as above but with the extended definitions of $\text{unfold}(y)$ and $\text{ind}(y, z)$. The proofs of Propositions 6 and 7 go through essentially as before.

5.4 A Note on Coalgebras

One reason why the translation of $\mathcal{L}^{\text{TLCGA}}$ into $\mathcal{L}^{\text{XCGA}}$ we have presented is useful is because it establishes a connection with *coalgebraic modal logics*. We note that the next time fragment $\mathcal{L}^{\text{XCGA}}$ of $\mathcal{L}^{\text{TLCGA}}$ is, in fact, an instance of the general framework of coalgebraic modal logics, making the language $\mathcal{L}_\mu^{\text{XCGA}}$ a *coalgebraic fixpoint logic*, as studied in [15, 21, 48]. This connection helps to clarify the place of the logic TLCGA in the landscape of modal fixpoint logics for various kinds of state-based evolving systems. The theory of *universal coalgebra* appears in computer science as a generic framework for modelling a wide range of state-based evolving systems in a uniform manner [39]. It is formulated using the language of basic category theory (functors and natural transformations), see [27] for a standard reference. The key idea of universal coalgebra is to pack all the information about the type of transitions that a system can make (deterministic, non-deterministic, probabilistic etc.) into a functor on the category of sets, which then can be considered as a variable parameter featuring in abstract definitions and results. A coalgebra for a functor F is then just a set X together with a map $f : X \rightarrow FX$, which intuitively represents the evolution of the state-based system. Concurrent game models are in fact coalgebras for a certain functor (together with a valuation of propositional variables), as has been observed several times, e.g., [42]. Furthermore, it can be checked that the semantics of modalities $\langle\!\langle Y \rangle\!\rangle$ in which each goal formula is of the form $X\theta$ for some θ can be phrased in terms of *predicate liftings* for this functor, which is the currently most usual way of interpreting modalities in coalgebras. We omit the details here.

The coalgebraic representation of the logic $\mathcal{L}_\mu^{\text{XCGA}}$, together with the translation from TLCGA into $\mathcal{L}_\mu^{\text{XCGA}}$, gives us access to a wealth of known general results on coalgebraic fixpoint logics. For example, we shall use it to derive decidability and finite model property for the logic TLCGA , as well as a complexity bound on the satisfiability problem for the language $\mathcal{L}_\mu^{\text{XCGA}}$. However, we emphasize that some caution is required here. Universal coalgebra and coalgebraic logic are valuable frameworks, unifying a large class of systems and associated logics, just like universal algebra provides a common language that puts many different algebraic structures under one roof. But, of course, universal algebra does not tell us everything we want to know about specific classes of algebras, like groups or Heyting algebras. Generic results are helpful, but a detailed study of concrete special cases is usually required. The same is true here. In particular, our main technical result in this article on completeness for an axiomatization of TLCGA makes heavy use of ideas from the literature on coalgebraic logic, in particular, the notion of “one-step completeness” is essential [15, 42]. But the proof of one-step completeness is not a trivial consequence of generic results from coalgebra, it requires a careful study of the semantics of next time goal assignments.

Once we have one-step completeness in place, the next step will be to handle the least and greatest fixpoints. We have shown that TLCGA can be translated into $\mathcal{L}_\mu^{\text{XCGA}}$ using a single recursion variable, in effect embedding TLCGA as a fragment of a “flat coalgebraic fixpoint logic” in the sense of Schröder and Venema [44], who prove a generic completeness result for such logics. It is possible that we could obtain completeness for our axiom system by “transferring” Schröder’s and Venema’s completeness result via our translation, but we have chosen to present a direct proof instead since we believe this will be more instructive. That said, we do consider our approach as being very much coalgebraic in spirit, and the connection with coalgebraic fixpoint logic is conceptually important. In particular, the idea that one-step completeness of a (multi-)modal logic can

be “lifted” to give a complete axiomatization of a fixpoint extension [15, 19, 25, 44] is at the heart of our proof.

6 BISIMULATIONS AND BISIMULATION INVARIANCE FOR TLCGA

As noted earlier, the logic GPCL introduced in [23] is essentially the next time fragment $\mathcal{L}^{\text{XCGA}}$ of TLCGA. Therefore, the notion of GPCL-bisimulation (ibid.) also applies to TLCGA. For the reader’s convenience, we introduce it again here, now called TLCGA-**bisimulation** and extend the bisimulation invariance result from [23] to the full logic TLCGA. This notion of bisimulation corresponds to the play-based semantics which is of importance for us. A similar notion can be defined for the path-based semantics and bisimulation invariance of TLCGA formulae with respect to it can be proved likewise, which we leave out.

We only define TLCGA-bisimulation within a single concurrent game model, and generalise to bisimulations between game models via disjoint unions.

Definition 9 (TLCGA-bisimulation). Let

$$\mathcal{M} = (\text{S}, \text{Act}, \text{g}, \text{out}, V),$$

be a game model for the set of agents Agt . A binary relation $\beta \subseteq S^2$ is a TLCGA-**bisimulation** in \mathcal{M} if it satisfies the following conditions for every pair of states (s_1, s_2) such that $s_1\beta s_2$:

Atom equivalence: For every $p \in \text{AP}$: $s_1 \in V(p)$ iff $s_2 \in V(p)$.

Forth: For any action profile ζ^1 of Agt at s_1 there is an action profile ζ^2 of Agt at s_2 such that

LocalBack: For every coalition C and every $u_2 \in \text{Out}[s_2, \zeta^2|_C]$, there is some $u_1 \in \text{Out}[s_1, \zeta^1|_C]$ such that $u_1\beta u_2$.

Back: For any joint action ζ^2 of Agt at s_2 there is a joint action ζ^1 of Agt at s_1 such that

LocalForth: For every coalition C and every $u_1 \in \text{Out}[s_1, \zeta^1|_C]$, there is some $u_2 \in \text{Out}[s_2, \zeta^2|_C]$ such that $u_1\beta u_2$.

States $s_1, s_2 \in \mathcal{M}$ are TLCGA-**bisimulation equivalent**, or just TLCGA-**bisimilar**, if there is a bisimulation β in \mathcal{M} such that $s_1\beta s_2$.

The proof of the following theorem is relatively routine, and can be found in the appendix.

THEOREM 2 (TLCGA-BISIMULATION INVARIANCE). *Let β be a TLCGA-bisimulation in a game model \mathcal{M} . Then for every TLCGA-formula φ and every pair $s_1, s_2 \in \mathcal{M}$ such that $s_1\beta s_2$:*

$$\mathcal{M}, s_1 \Vdash \varphi \text{ iff } \mathcal{M}, s_2 \Vdash \varphi.$$

In fact, the proof of Theorem 2 essentially amounts to a proof of bisimulation invariance for the whole language $\mathcal{L}_\mu^{\text{XCGA}}$, and therefore also for the fragment TLCGA⁺.

THEOREM 3. *Every formula of $\mathcal{L}_\mu^{\text{XCGA}}$ is bisimulation invariant.*

Furthermore, we also have the following result from [23], showing that TLCGA already suffices to capture bisimulation invariance in finite models.

PROPOSITION 8 (HENNESSY–MILNER PROPERTY, cf [23], PROPOSITION 4.6). *Let β be a TLCGA-bisimulation in a finite game model $\mathcal{M} = (\text{S}, \text{Act}, \text{g}, \text{out}, V)$. Then for any pair $s_1, s_2 \in \text{S}$, $s_1\beta s_2$ holds iff s_1 and s_2 satisfy the same TLCGA-formulae.*

PROOF. For the non-trivial direction, we will use only formulae from the fragment $\mathcal{L}^{\text{XCGA}}$. Since \mathcal{M} is finite, we can define, by a standard construction, a “characteristic formula” $\text{char}(s)$ for each state s in \mathcal{M} , such that s_1, s_2 are $\mathcal{L}^{\text{XCGA}}$ -equivalent if and only if $\text{char}(s_1) = \text{char}(s_2)$, and that $\text{char}(s_1) \wedge \text{char}(s_2) \equiv \perp$ whenever s_1, s_2 are not $\mathcal{L}^{\text{XCGA}}$ -equivalent. For a set of states Z , let

$\text{char}[Z] = \bigvee \{\text{char}(v) \mid v \in Z\}$. Our goal is to show that the relation of TLCGA-equivalence is itself a TLCGA-bisimulation, and the key observation is that each state s satisfies the $\mathcal{L}^{\text{XCGA}}$ -formula:

$$\bigwedge_{\zeta \in \text{ActProf}_s} \langle \langle C_1 \triangleright \text{X char}[\text{Out}[s, \zeta|_{C_1}], \dots, C_k \triangleright \text{X char}[\text{Out}[s, \zeta|_{C_k}]] \rangle \rangle,$$

where we list the set $\mathcal{P}(\text{Agt})$ of all possible coalitions as C_1, \dots, C_k . \square

7 AXIOMATIZATION AND ONE-STEP COMPLETENESS OF TLCGA

In this section, we focus exclusively on TLCGA and leave the extension of the axiomatic system presented here to TLCGA⁺, as well as the respective axiomatizations for the path-based semantics (which may turn out to be more problematic) for future work.

7.1 Axiomatic System for TLCGA

Definition 10. Let \mathcal{F} be a set of coalitions. A **voting profile** for \mathcal{F} is a mapping f assigning to each $a_i \in \text{Agt}$ a goal assignment $f(a_i)$. If $f(a_i)(C)$ is a next time formula for each i and $C \in \mathcal{F}$, we say that f is a **one-step voting profile** for \mathcal{F} .

The notion of merging a voting profile, defined next, will be used in some proofs later and we will need some derivable formulae that use it, listed further.

Definition 11. Let f be a voting profile. We define the goal assignment $\text{merge}(f)$ as follows:

- $\text{merge}(f)(C) := \theta$, if $C \neq \emptyset$ and $f(a_i)(C) = \theta$ for each $a_i \in C$,
- $\text{merge}(f)(C) := \text{X}\top$, if $C = \emptyset$ or the above holds for no θ .

Our axioms are as follows (recall notation on goal assignments from Section 3.1).

7.1.1 I. General Axiom Schemes for Goal Assignments.

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

(Safe) $\neg \langle \langle \text{Agt} \triangleright \text{X}\perp \rangle \rangle$

(Merge) $\langle \langle C_1 \triangleright \theta_1 \rangle \rangle \wedge \dots \wedge \langle \langle C_n \triangleright \theta_n \rangle \rangle \rightarrow \langle \langle C_1 \triangleright \theta_1, \dots, C_n \triangleright \theta_n \rangle \rangle$, where C_1, \dots, C_n are pairwise disjoint.

This axiom generalises the Superadditivity axiom of Coalition Logic. The idea is simple: if the coalitions C_1, \dots, C_n are pairwise disjoint, then they can join their collective strategies for their respective coalitional goals into one strategy profile that ensures achievement of all these collective goals.

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

Any strategy profile generates a unique successor state, on which any state formula ψ is either true or its negation is true, so either ψ or $\neg\psi$ can be added to the next time goal of the grand coalition Agt.

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

For any coalition C , state formula ψ , and a strategy profile Σ , either its projection Σ_C to C ensures the truth of ψ in all successor states enabled by Σ_C —in which case ψ can be added to the next time goal of C enforced by Σ —or else $\neg\psi$ is true in some of these successor states, in which case it can be added to $\gamma|_C$ as the next time goal of the grand coalition Agt enforced by Σ .

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

Given any coalition C and sub-coalition C' , the next time goal of C' can be added to the next time goal of C , in the sense that if there is any strategy profile Σ that ensures that C

$$\begin{array}{c}
 \text{Fix: } \text{unfold}(\gamma) \leftrightarrow \langle\!\langle \gamma \rangle\!\rangle \\
 \text{R-Ind: } \frac{\text{ind}(\gamma, \phi) \rightarrow \phi}{\langle\!\langle \gamma \rangle\!\rangle \rightarrow \phi} \quad (\gamma \in \text{TypeU}) \\
 \text{R-CoInd: } \frac{\phi \rightarrow \text{ind}(\gamma, \phi)}{\phi \rightarrow \langle\!\langle \gamma \rangle\!\rangle} \quad (\gamma \in \text{TypeG})
 \end{array}$$

Fig. 2. Fixpoint axiom and induction rules.

and C' can force their respective next time goals $X\phi$ and $X\psi$, then Σ also ensures that C can force the conjunction of these goals.

7.1.2 II. General Rules of Inference:

Modus Ponens and Goal Monotonicity (G-Mon):

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

The meaning of the rule is clear: if ϕ implies ψ , then any coalition C that can ensure the next time goal ϕ within the context of some strategic goal assignment, can also ensure ψ with the same context.

7.1.3 III. Axioms and Rules for the Long-term Goal Assignments. The axioms and rules for the goal assignments of Types 1 and 2, involving long-term temporal operators are given in Figure 2. 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 R-Ind it is a long-term temporal assignment of type U, and in R-CoInd it is a long-term temporal assignment of type G.

We denote the axiomatic system above by Ax_{TLCGA} and will denote derivability in it by $\text{Ax}_{\text{TLCGA}} \vdash$, but will often write just \vdash . Here are some important validities that are derivable in Ax_{TLCGA} , some of which will be used in the proofs further:

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

(Immediately from (Fix), due to Proposition 5).

(Weakening) $\langle\!\langle \gamma \rangle\!\rangle \rightarrow \langle\!\langle C \triangleright \gamma(C) \rangle\!\rangle$, for any $C \subseteq \text{Agt}$. (Using (Triv) and (G-Mon).)

Agt-Maximality $\langle\!\langle \emptyset \triangleright X \phi \rangle\!\rangle \vee \langle\!\langle \text{Agt} \triangleright X \neg \phi \rangle\!\rangle$. (Using (Triv) and (Case).)

(Superadditivity) $\langle\!\langle C_1 \triangleright X \phi_1 \rangle\!\rangle \wedge \langle\!\langle C_2 \triangleright X \phi_2 \rangle\!\rangle \rightarrow \langle\!\langle C_1 \cup C_2 \triangleright X (\phi_1 \wedge \phi_2) \rangle\!\rangle$; $C_1 \triangleright X \phi_1$; $C_2 \triangleright X \phi_2$, if $C_1 \cap C_2 = \emptyset$.

This subsumes the Superadditivity axiom for Coalition Logic CL. It is derivable from (Merge) using twice (Con) to add $X(\phi_1 \wedge \phi_2)$ as a goal assignment to $C_1 \cup C_2$.

(Merge') $\bigwedge_{a_i \in \text{Agt}} \langle\!\langle f(a_i) \rangle\!\rangle \rightarrow \langle\!\langle \text{merge}(f) \rangle\!\rangle$, where f is any voting profile.

This is an essentially equivalent formulation of (Merge). Indeed, (Merge) is a particular case of (Merge'), whereas (Merge') is derivable from (Merge) by first using (Weakening) to detach each $\langle\!\langle C_j \triangleright \theta_j \rangle\!\rangle$ from every $f(a_i)$, for $a_i \in C_j$, if $C_j \neq \emptyset$ and $f(a_i)(C_j) = \theta_j$ for all such a_i .

Fix(G) $\langle\!\langle C \triangleright G \chi \rangle\!\rangle \rightarrow \chi \wedge \langle\!\langle C \triangleright X \langle\!\langle C \triangleright G \chi \rangle\!\rangle \rangle\!\rangle$.

This is a special case of Fix.

CoInd(G) If $\vdash \phi \rightarrow \chi \wedge \langle\!\langle C \triangleright X \phi \rangle\!\rangle$ then $\vdash \phi \rightarrow \langle\!\langle C \triangleright G \chi \rangle\!\rangle$.

This is a special case of the rule CoInd.

In particular, by using Fix(G) and applying G-Mon we obtain $\vdash (\chi \wedge \langle\!\langle C \triangleright X \langle\!\langle C \triangleright G \chi \rangle\!\rangle \rangle\!\rangle) \rightarrow \chi \wedge \langle\!\langle C \triangleright X (\chi \wedge \langle\!\langle C \triangleright G \chi \rangle\!\rangle) \rangle\!\rangle$

Now, by applying $\text{CoInd}(G)$ for $\phi = \chi \wedge \langle C \triangleright X \langle C \triangleright G\chi \rangle \rangle$, we derive

$$\vdash \chi \wedge \langle C \triangleright X \langle C \triangleright G\chi \rangle \rangle \rightarrow \langle C \triangleright G\chi \rangle.$$

Thus, we have derived the fixpoint equivalence for G :

$$\mathbf{FP}(G) \quad \langle C \triangleright G\chi \rangle \leftrightarrow \chi \wedge \langle C \triangleright X \langle C \triangleright G\chi \rangle \rangle.$$

$$\mathbf{PreFP}(U) \quad \beta \vee (\alpha \wedge \langle C \triangleright X \langle C \triangleright \alpha \cup \beta \rangle \rangle) \rightarrow \langle C \triangleright \alpha \cup \beta \rangle.$$

This is a special case of Fix .

$$\mathbf{Ind}(U) \quad \text{If } \vdash \beta \vee (\alpha \wedge \langle C \triangleright X \phi \rangle) \rightarrow \phi \text{ then } \vdash \langle C \triangleright \alpha \cup \beta \rangle \rightarrow \phi.$$

This is a special case of the rule Ind .

In particular, by applying the rule $G\text{-Mon}$ to $\text{PreFP}(U)$ we derive

$$\langle C \triangleright X (\beta \vee (\alpha \wedge \langle C \triangleright X \langle C \triangleright \alpha \cup \beta \rangle \rangle)) \rangle \rightarrow \langle C \triangleright X \langle C \triangleright \alpha \cup \beta \rangle \rangle.$$

Then, by simple propositional inference we derive

$$(\beta \vee (\alpha \wedge \langle C \triangleright X (\beta \vee (\alpha \wedge \langle C \triangleright X \langle C \triangleright \alpha \cup \beta \rangle \rangle)) \rangle)) \rightarrow (\beta \vee (\alpha \wedge \langle C \triangleright X \langle C \triangleright \alpha \cup \beta \rangle \rangle)).$$

Now, by applying $\text{Ind}(U)$ for $\phi = \beta \vee (\alpha \wedge \langle C \triangleright X \langle C \triangleright \alpha \cup \beta \rangle \rangle)$, we derive

$$\vdash \langle C \triangleright \alpha \cup \beta \rangle \rightarrow \beta \vee (\alpha \wedge \langle C \triangleright X \langle C \triangleright \alpha \cup \beta \rangle \rangle).$$

Thus, we have derived the fixpoint equivalence for U :

$$\mathbf{FP}(U) \quad \langle C \triangleright \alpha \cup \beta \rangle \leftrightarrow \beta \vee (\alpha \wedge \langle C \triangleright X \langle C \triangleright \alpha \cup \beta \rangle \rangle).$$

PROPOSITION 9 (SOUNDNESS OF Ax_{TLCGA}). *The axiomatic system Ax_{TLCGA} is sound: every derivable formula in Ax_{TLCGA} is valid.*

PROOF. We show that every axiom is valid and all rules of inference preserve validity.

Checking the validity of the general axiom schemes is fairly routine. Most of these, as well as the preservation of validity by the general rules II , follow from the soundness of the logic GPCL in [23].

The validity of the axiom scheme Fix follows from Theorem 1.

The preservation of validity by the special rule R-Ind can be shown as follows. Suppose $\text{ind}(\gamma, \phi) \rightarrow \phi$ is valid. Take any concurrent game model \mathcal{M} . Then $\mathcal{M} \Vdash \text{ind}(\gamma, \phi) \rightarrow \phi$, hence $\llbracket \phi \rrbracket_{\mathcal{M}}$ is a pre-fixed point of the set operator induced by the formula $\text{ind}(\gamma, z)$ in \mathcal{M} . By Proposition 6, $\llbracket \gamma \rrbracket$ is semantically equivalent to the least fixed point $\mu z.\text{ind}(\gamma, z)$, which is also the least pre-fixed point of $\text{ind}(\gamma, z)$. Therefore, $\mathcal{M} \Vdash \llbracket \gamma \rrbracket \rightarrow \phi$. Thus, $\llbracket \gamma \rrbracket \rightarrow \phi$ is valid.

The preservation of validity by the special rule R-CoInd is proved analogously, using Proposition 7 and the fact that the greatest fixed point $\nu z.\text{ind}(\gamma, z)$, is also its greatest post-fixed point. \square

Recall (cf Section 5.1) that a formula $\phi \in \mathcal{L}^{\text{TLCGA}}$ is in **normal form** if, for every subformula of the form $\llbracket \gamma \rrbracket$, the goal assignment γ is either a next time or a long-term temporal goal assignment.

PROPOSITION 10. *For every formula ϕ there is a formula ψ which is in normal form, and such that $Ax_{\text{TLCGA}} \vdash \phi \leftrightarrow \psi$.*

PROOF. By induction on the structure of formulas, using the axiom Fix for the crucial steps. By design, the unfolding $\text{unfold}(\gamma)$ of any goal assignment γ is a next time goal assignment, and all new goal assignments appearing in the scope of next time operators in the codomain of $\text{unfold}(\gamma)$ will be long-term temporal. So, all mixing of next time and long-term temporal path formulas in $\llbracket \text{unfold}(\gamma) \rrbracket$ will appear in proper subformulas of $\llbracket \gamma \rrbracket$, where the inductive hypothesis is applied. \square

By the soundness, the proposition above implies the following corollary.

Conjunctive formula	Components	Disjunctive formula	Components
$\neg\neg\varphi$	φ	$\varphi \vee \psi$	φ, ψ
$\varphi \wedge \psi$	φ, ψ	$\neg(\varphi \wedge \psi)$	$\neg\varphi, \neg\psi$
$\neg(\varphi \vee \psi)$	$\neg\varphi, \neg\psi$		

Local goal formulae	Components
$\langle\{\gamma\}\rangle$ (positive)	$\{\psi \mid \gamma(C) = X\psi, C \subseteq \text{Agt}\}$
$\neg\langle\{\gamma\}\rangle$ (negative)	$\{\neg\psi \mid \gamma(C) = X\psi, C \subseteq \text{Agt}\}$

Temporal goal formulae	Components
$\langle\{\gamma\}\rangle$ (positive)	$\text{ind}(\gamma, \langle\{\gamma\}\rangle)$
$\neg\langle\{\gamma\}\rangle$ (negative)	$\neg\text{ind}(\gamma, \langle\{\gamma\}\rangle)$

Fig. 3. Types of formulae and their components.

COROLLARY 2. *For every formula φ there is a semantically equivalent formula ψ which is in normal form.*

7.2 Formula Types, Components, and Extended FL-closure of TLPGA Formulae

We use some generic notions and terminology from the literature on tableaux-based satisfiability decision methods (see e.g., [17, Ch.13]). Formulae of TLPGA in normal form can be classified as **literals**: $\top, \neg\top, p, \neg p$, where $p \in \text{AP}$, **conjunctive formulae**, of the type $(\phi \wedge \psi)$ and $\neg(\phi \vee \psi)$; **disjunctive formulae**, of the type $(\phi \vee \psi)$ and $\neg(\phi \wedge \psi)$; **successor formulae**: $\langle\{\gamma\}\rangle$ and $\neg\langle\{\gamma\}\rangle$, where γ is a local (next time) goal assignment; and **long-term temporal formulae**, of the type $\langle\{\gamma\}\rangle$ and $\neg\langle\{\gamma\}\rangle$, where γ is a long-term goal assignment. The formulae in the last four classes have respective **components** that are given in Figure 3. Clearly, every conjunctive (respectively, disjunctive) formula in the table is equivalent to the conjunction (respectively, disjunction) of its components.

Given a formula φ , we define $\bar{\varphi} := \psi$ if φ is of the form $\neg\psi$, and $\bar{\varphi} := \neg\varphi$ otherwise.

Definition 12. The **extended (Fischer-Ladner) closure** of a TLPGA formula in normal form φ is the least set of formulae $\text{ecl}(\varphi)$ such that

- (1) $\varphi \in \text{ecl}(\varphi)$,
- (2) $\text{ecl}(\varphi)$ is closed under taking all components of formulae in $\text{ecl}(\varphi)$.
- (3) $\bar{\psi} \in \text{ecl}(\varphi)$ whenever $\psi \in \text{ecl}(\varphi)$.

For any set of formulae Φ we define $\text{ecl}(\Phi) := \bigcup\{\text{ecl}(\varphi) \mid \varphi \in \Phi\}$.

A set Φ of TLPGA formulae in normal form is said to be **(Fischer-Ladner) closed** iff $\text{ecl}(\Phi) = \Phi$.

The proof of the following proposition can be found in the appendix:

PROPOSITION 11. *The extended closure of any finite set Φ of TLPGA formulae in normal form is finite.*

7.3 One-step Completeness

Hereafter derivability/provability and consistency refer to the axiomatic system Ax_{TLPGA} . Given a set of formulae Φ , the maximal consistent subsets of Φ are defined as usual.

Definition 13. Given a closed set of formulae Φ , a **Φ -atom** is a maximal consistent subset of Φ . We denote by $\text{At}(\Phi)$ the set of all Φ -atoms.

Definition 14. Let Φ be a finite set of formulae. A **next time assignment** over Φ is a formula of the shape

$$\langle\langle C_1 \triangleright X\varphi_1, \dots, C_k \triangleright X\varphi_k \rangle\rangle,$$

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 next time assignment over Φ or the negation of such a formula.

Definition 15. Let Φ be a finite set of formulae. A **consistent game form for Φ** is a game form $(\text{Act}, \text{act}, \mathcal{P}(\Phi), \text{out})$ over the set of outcomes $\mathcal{P}(\Phi)$ such that, for each action profile ζ , $\text{out}(\zeta)$ is a consistent set of formulae. A **maximal consistent game form for Φ** is a game form $(\text{Act}, \text{act}, \mathcal{P}(\Phi), \text{out})$ over outcomes $\mathcal{P}(\Phi)$ such that, for each action profile ζ , $\text{out}(\zeta)$ 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 $\text{out}(\zeta)$ is a Φ -atom.

Given a strategic game form $G = (\text{Act}, \text{act}, \text{O}, \text{out})$, a coalition C and action profiles ζ', ζ , we write $\zeta' \sim_C \zeta$ to state that $\zeta'|_C = \zeta|_C$.

THEOREM 4 (ONE-STEP COMPLETENESS). *Let Γ be a consistent modal one-step theory over a finite set of formulas Φ and assume that Φ contains all components of Γ , also contains $\bar{\psi}$ whenever $\psi \in \Phi$, and is closed under conjunctions (up to provable equivalence). Then there exists a maximal consistent game form $\mathcal{M}(\Gamma) = (\text{Act}, \text{act}, \mathcal{P}(\Phi), \text{out})$ for Φ such that, for every goal assignment γ :*

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

Furthermore, the size of Act is at most exponential in $|\Phi|$.

PROOF. We may assume without loss of generality that, for every next time goal assignment γ over Φ , the set Γ contains either $\langle\langle \gamma \rangle\rangle$ or $\neg\langle\langle \gamma \rangle\rangle$, since otherwise we can extend Γ , using Lindenbaum's lemma, to a consistent (and still finite) set satisfying this assumption.

We consider next time goal assignments over the set of conjunctions of subsets of Φ . We say that such a goal assignment γ is **deterministic** if $\gamma(\text{Agt})$ is (provably equivalent to) the conjunction of a maximal consistent subset of Φ . We then say that a goal assignment γ' is a **strengthening** of γ if, for all C with $\gamma(C) = X\phi$ and $\gamma'(C) = X\phi'$, the formula $\phi' \rightarrow \phi$ is provable. Note that every formula $\langle\langle \gamma \rangle\rangle$ provably implies the disjunction of all formulas $\langle\langle \gamma' \rangle\rangle$ where γ' is a deterministic strengthening of γ over Φ . This follows by repeated applications of the axiom (GrandCoalition).

If γ is deterministic we let $\text{next}(\gamma)$ denote the maximal consistent set Ψ for which $\gamma(\text{Agt}) = X(\wedge \Psi)$. Note that for any deterministic strengthening γ' of a goal assignment γ over Ψ and any C for which $\gamma(C) = X\phi$, $\text{next}(\gamma')$ must contain ϕ as a conjunct. This is a consequence of axioms (Con), (Safe) and the assumption that $\text{next}(\gamma')$ is maximal consistent.

For technical convenience, in this proof, we fix the enumeration of Agt to be a_0, \dots, a_{K-1} , where $K = |\text{Agt}|$. Given an agent $a \in \text{Agt}$, we define act_a^Γ to be the set of all triples (γ, f, k) such that γ is a goal assignment with $\langle\langle \gamma \rangle\rangle \in \Gamma$, $0 \leq k < K$, and f is a function mapping each goal assignment $\gamma' : \mathcal{P}(\text{Agt}) \rightarrow \Phi$ to one of its deterministic strengthenings. To count the number of actions, this is the number of goal assignments γ with $\langle\langle \gamma \rangle\rangle \in \Gamma$, times the number of functions f mapping goal assignments over Φ to deterministic strengthenings, times K . The number of goal assignments

over Φ is $|\Phi|^{2^K}$, and more generally the number of goal assignments over conjunctions of subsets of Φ is $(2^{|\Phi|})^{2^K} = 2^{2^K \cdot |\Phi|}$. So the number of functions f that can appear in an action is

$$(2^{2^K \cdot |\Phi|})^{|\Phi|^{2^K}} = 2^{2^K \cdot |\Phi|^{2^K+1}}.$$

Since the number of agents is fixed, 2^K is a constant and the term $2^K \cdot |\Phi|^{2^K+1}$ is a polynomial in $|\Phi|$. Hence the number of actions is exponential in $|\Phi|$.

Note that $\text{act}_a^\Gamma \neq \emptyset$ for all $a \in \text{Agt}$ since there is at least one goal assignment with $\langle \gamma \rangle \in \Gamma$ by the axiom (Triv), and $\langle \gamma \rangle$ is equivalent to the disjunction of its deterministic strengthenings, so one of these must also be in Γ . Note also that the goal assigned to Agt by any $\langle \gamma \rangle \in \Gamma$ must be consistent by the axiom (Safe). We set $\text{Act}^\Gamma = \bigcup_{a \in \text{Agt}} \text{act}_a^\Gamma$. Given an action profile ζ and $a \in \text{Agt}$, if $\zeta_a = (\gamma, k, f)$ we write $\text{vote}(a, \zeta) = \gamma$, $\text{bet}(a, \zeta) = k$ and $\text{choice}(a, \zeta, \gamma') = f(\gamma')$ for every goal assignment γ' with $\langle \gamma' \rangle \in \Gamma$. We write $\text{vote}(\zeta)$ for the voting profile mapping each $a \in \text{Agt}$ to $\text{vote}(a, \zeta)$. We define the *voting winner* $\text{win}(\zeta)$ to be player a_i where i is determined as follows:

$$i := \left(\sum_{a \in \text{Agt}} \text{bet}(a, \zeta) \right) \bmod K.$$

Finally, we define the outcome of a given action profile ζ as follows:

$$\text{out}(\zeta) := \text{next}(\text{choice}(\text{win}(\zeta), \zeta, \text{merge}(\text{vote}(\zeta)))).$$

We will show that the game form $\mathcal{M}(\Gamma) = (\text{Act}^\Gamma, \text{act}^\Gamma, \mathcal{P}(\Phi), \text{out})$ we have constructed satisfies the criteria listed in the statement of the theorem. First, we shall prove a rather technical auxiliary claim. Before going through its proof, the reader may want to skip ahead to see how the claim is used in the main argument.

CLAIM 1. *Let φ be any formula in Φ , C any coalition, and let ζ be an action profile in $\mathcal{M}(\Gamma)$ such that for every action profile $\zeta' \sim_C \zeta$, we have $\varphi \in \text{out}(\zeta')$. Let γ be any deterministic strengthening of $\text{merge}(\text{vote}(\zeta))$ such that $\langle \gamma \rangle \in \Gamma$ and $\gamma(\text{Agt}) = X(\bigwedge \text{out}(\zeta))$, and let $\gamma(C) = X\psi$. Then there exists a deterministic strengthening γ' of γ such that $\langle \gamma' \rangle \in \Gamma$, and:*

$$\gamma'(C) = X(\psi \wedge \varphi).$$

PROOF OF CLAIM 1. We need to distinguish two cases, for $C = \text{Agt}$ and $C \neq \text{Agt}$. We begin with the easier case where $C = \text{Agt}$. In this case, there is only one action profile $\zeta' \sim_C \zeta$, namely ζ itself. Our assumption thus gives $\varphi \in \text{out}(\zeta)$. Furthermore, we have $\gamma(C) = \gamma(\text{Agt}) = X(\bigwedge \text{out}(\zeta))$ by assumption. But, since $\psi \in \text{out}(\zeta)$ the formula $\bigwedge \text{out}(\zeta)$ is equal (up to provable equivalence) to $\psi \wedge \bigwedge \text{out}(\zeta)$, which is provably equivalent to $\psi \wedge \varphi$ (because γ' is a deterministic strengthening). Hence we can set $\gamma' = \gamma$.

In the case, where $C \neq \text{Agt}$ is more involved. We assumed that $\langle \gamma \rangle \in \Gamma$. By (Case), we have:

$$\langle \gamma[C \triangleright X(\psi \wedge \varphi)] \rangle \vee \langle \gamma[C \triangleright X \neg \varphi] \rangle \in \Gamma.$$

We first show that $\langle \gamma[C \triangleright X \neg \varphi] \rangle \notin \Gamma$. Suppose the contrary, that $\langle \gamma[C \triangleright X \neg \varphi] \rangle \in \Gamma$. Let γ^* be an arbitrary deterministic strengthening of $\gamma[C \triangleright X \neg \varphi]$ such that $\langle \gamma^* \rangle \in \Gamma$. Such strengthening must exist, since we recall that every formula of the form $\langle \delta \rangle$ provably implies the disjunction of all its deterministic strengthenings over Φ . Now we define a new action profile ζ' as follows. First, pick an arbitrary $c \notin C$, which exists since $C \neq \text{Agt}$. For each $a \in C$ set $\zeta'_a = \zeta_a$. For each $a \notin C$ and $a \neq c$, set $\zeta'_a = (\gamma^\top, f, 0)$ where f is arbitrary (recall that γ^\top is the trivial goal assignment with empty support). For c , set $\zeta'_c = (\gamma^\top, f, h)$ where the choice function f chooses γ^* whenever possible, and the bet h is chosen so that the index of the player c is equal to $\sum_{a \in C} \text{bet}(a, \zeta) + h \pmod K$.

This guarantees that c will be the voting winner in ζ' . Clearly $\zeta' \sim_C \zeta$. Furthermore, since γ^* is a strengthening of $\gamma|_C[\text{Agt} \triangleright X \neg\varphi]$, and γ is a strengthening of $\text{merge}(\text{vote}(\zeta))$ by assumption, it follows that γ^* is a deterministic strengthening of $\text{merge}(\text{vote}(\zeta'))$. This is because the only coalitions not mapped to $X \top$ by $\text{merge}(\text{vote}(\zeta'))$ are the ones contained in C , and for any such coalition D we have $\text{merge}(\text{vote}(\zeta'))(D) = \text{merge}(\text{vote}(\zeta))(D)$. So, we get:

$$\begin{aligned} \text{out}(\zeta') &= \text{next}(\text{choice}(\text{win}(\zeta'), \zeta', \text{merge}(\text{vote}(\zeta')))) \\ &= \text{next}(\text{choice}(c, \zeta', \text{merge}(\text{vote}(\zeta')))) \\ &= \text{next}(\gamma^*). \end{aligned}$$

But $\neg\varphi \in \text{next}(\gamma^*)$ since γ^* is a strengthening of $\gamma|_C[\text{Agt} \triangleright X \neg\varphi]$. By consistency of $\text{next}(\gamma^*)$, we have thus found an action profile ζ' such that $\zeta \sim_C \zeta'$ and $\varphi \notin \text{out}(\zeta')$. This is a contradiction with our assumption on the action profile ζ . Thus, we have proved $\langle\!\langle \gamma|_C[\text{Agt} \triangleright X \varphi] \rangle\!\rangle \notin \Gamma$, as desired. It follows that $\langle\!\langle \gamma[C \triangleright X(\psi \wedge \varphi)] \rangle\!\rangle \in \Gamma$. We then define:

$$\gamma' := \gamma[C \triangleright X(\psi \wedge \varphi)].$$

Thus, we have shown that $\langle\!\langle \gamma' \rangle\!\rangle \in \Gamma$, γ' is clearly a strengthening of γ , and it is deterministic since $\gamma'(\text{Agt}) = \gamma(\text{Agt})$. This concludes the proof of the claim. \square

We now prove that the properties (1) and (2) listed in the theorem hold for the game form $\mathcal{M}(\Gamma) = (\text{Act}^\Gamma, \text{act}^\Gamma, \mathcal{P}(\Phi), \text{out})$.

Item (1). Suppose $\langle\!\langle \gamma \rangle\!\rangle \in \Gamma$. Let ζ be defined by letting all players vote for $(\gamma, f, 0)$ where f is an arbitrary, fixed choice function. Let C be in the support of γ , where $\gamma(C) = X\varphi$ and let $\zeta' \sim_C \zeta$. Since all players in C vote for γ in ζ' , and the outcome $\text{out}(\zeta')$ is $\text{next}(\gamma')$ for a deterministic strengthening of $\text{merge}(\text{vote}(\zeta'))$, it follows that $\varphi \in \text{out}(\zeta')$.

Item (2). Suppose $\neg\langle\!\langle \gamma \rangle\!\rangle \in \Gamma$, and let ζ be an arbitrary action profile. We want to show that there is some coalition C and some action profile $\zeta' \sim_C \zeta$ such that $\bar{\varphi} \in \text{out}(\zeta')$, where $\gamma(C) = \varphi$.

We will prove this by reductio ad absurdum. Suppose that for every coalition C with $\gamma(C) = X\varphi$ and every action profile $\zeta' \sim_C \zeta$, we have $\bar{\varphi} \notin \text{out}(\zeta')$. This means that $\varphi \in \text{out}(\zeta')$ since both $\bar{\varphi}$ and φ are in the closure of $\langle\!\langle \gamma \rangle\!\rangle$ and $\text{out}(\zeta')$ is maximal consistent. Let us list all coalitions in the support of γ as C_1, \dots, C_m . Let γ_0 denote the goal assignment $\text{choice}(\text{win}(\zeta), \zeta, \text{merge}(\text{vote}(\zeta)))$. Then $\gamma_0(\text{Agt}) = X(\bigwedge \text{out}(\zeta))$. Furthermore $\langle\!\langle \gamma_0 \rangle\!\rangle \in \Gamma$ by definition of choice, and γ_0 is a deterministic strengthening of $\text{merge}(\text{vote}(\zeta))$. For each $i \in \{1, \dots, m\}$ let ψ_i be the formula such that $\gamma(C_i) = X\psi_i$, and let ψ_i^0 be the formula such that $\gamma_0(C_i) = X\psi_i^0$. We define, for each $i \in \{0, \dots, m\}$, a deterministic goal assignment γ_i such that

- $\langle\!\langle \gamma_i \rangle\!\rangle \in \Gamma$,
- γ_i is a deterministic strengthening of γ_j for all $j < i$,
- $\gamma_i(C_j) = X(\psi_j^0 \wedge \psi_j)$ for all j with $1 \leq j \leq i$, and $\gamma_i(\text{Agt}) = X(\bigwedge \text{out}(\zeta))$.

The goal assignment γ_0 has already been defined, and we can extend the definition inductively to all i by repeatedly applying Claim 1. Note that the induction hypothesis has been tailored so that Claim 1 applies at each inductive step.

Now, consider the goal assignment γ_m . We have $\langle\!\langle \gamma_m \rangle\!\rangle \in \Gamma$. But by definition $\langle\!\langle \gamma_m \rangle\!\rangle$ is a strengthening of $\langle\!\langle \gamma \rangle\!\rangle$, hence $\langle\!\langle \gamma \rangle\!\rangle \in \Gamma$, by Goal Monotonicity. Since we assumed that $\neg\langle\!\langle \gamma \rangle\!\rangle \in \Gamma$, we have reached a contradiction with the consistency of Γ . This concludes the proof of item (2) and thus the proof of the theorem. \square

8 COMPLETENESS OF TLCGA

8.1 Networks

Throughout the rest of this section, we fix a finite, closed set Φ of TLCGA-formulae in normal form. To prove completeness, we will show how to construct a model for each (consistent) atom in $\text{At}(\Phi)$, using the technique of *networks*. The idea behind this technique is to construct a series of approximations to the satisfying model. At each finite stage of the construction, the current approximating network will generally have a number of *defects*, each of which represents some particular reason that the network cannot yet be regarded as a satisfying model. For example, consider a formula of the form $\langle A \triangleright \alpha \cup \beta, B \triangleright G\chi \rangle$. If this formula belongs to the label associated with some node u in a network, then there must exist some strategy profile Σ such that, when restricted to a joint strategy for the coalition A , it ensures that every play generated by it eventually leads to a state where β is in the label, and all states that are visited meanwhile have α in their labels; likewise, when Σ is restricted to a joint strategy for the coalition B , it ensures that on every play generated by it all states that are visited have χ in their labels. For the respective conditions for α and χ it suffices that they are satisfied *locally*, at every step of the construction of the network. Thus, there are two types of conditions to be satisfied by the network: *local conditions*, that can be ensured on-the-fly, i.e., the defects arising from their violation can be fixed step-by-step in the process of the construction and updates of the network; and *eventualities*, which need to be taken special care of. In the example above, if it is not already the case that every play generated by the joint strategy for B obtained from Σ eventually leads to a state where β is in the label, that creates a defect, which we need to be eventually fixed. To do so we first show that we can “push” the defect towards leaves in the network, in the sense that it suffices to fix the defect at each leaf in order to make sure each occurrence of the defect in the current network disappears. Next, we fix each defect associated with a leaf. To do that, we prove that for every atom that contains the formula $\langle A \triangleright \alpha \cup \beta, B \triangleright G\chi \rangle$, we can find a network whose root is labelled by that atom, and in which the occurrence of the formula at the root is not a defect (though it may appear as a defect elsewhere in the network). This network can then be “plugged in” at appropriate leaves in another network in order to fix a defect there. Once a specific occurrence of a defect is fixed it never reappears, but of course, each round of the construction may introduce new defects, like new heads of a hydra appearing where one was previously cut off. These new defects are then taken care of in the *next* round, and so on. By taking a limit of the approximating series we obtain a *perfect* network, i.e., a network with no defects. A “truth lemma” for perfect networks assures that we can view the network obtained in the limit as a model in which each formula attached to the root is true.

Definition 16. A Φ -**network** is a triple $\mathcal{N} = (T, L, \mathcal{G})$ such that

- T is a rooted, finitely branching directed tree,
- $L : T \rightarrow \text{At}(\Phi)$ is a map that assigns to each node of T an atom from $\text{At}(\Phi)$.
- \mathcal{G} is a map that assigns to each non-leaf node u of T a game form $\mathcal{G}(u) = (\text{Act}^u, \text{act}^u, T, \text{out}^u)$, where out^u is subject to the constraint that its codomain is the set of children nodes of u in T .

Definition 17. A network $\mathcal{N} = (T, L, \mathcal{G})$ is said to be a **sub-network** of $\mathcal{N}' = (T', L', \mathcal{G}')$, written $\mathcal{N} \sqsubseteq \mathcal{N}'$, if

- T is a subgraph of T' and the root of T is also the root of T' ,
- If u is any non-leaf node in T then it has the same children in T as in T' and, furthermore, $\mathcal{G}(u) = \mathcal{G}'(u)$,
- $L = L'|_T$.

Definition 18. Given a Φ -network $\mathcal{N} = (T, L, \mathcal{G})$, a **marking** of \mathcal{N} is a map \mathbf{m} from T to the powerset of Φ such that $\mathbf{m}(v) \subseteq L(v)$ for all $v \in T$. (In particular, note that L itself is a marking of

\mathcal{N} .) Given a marking \mathbf{m} of \mathcal{N} , a next time goal assignment γ such that $\langle\!\langle\gamma\rangle\!\rangle \in \Phi$, and a non-leaf node $u \in T$ with $\mathcal{G}(u) = (\text{Act}, \text{act}, T, \text{out})$, we say that the marking \mathbf{m} **verifies the goal assignment γ at u** if there is a strategy profile $\Sigma \in \Pi_{a \in \text{Agt}}$ such that, for every C in the support of γ such that $\gamma(C) = X\psi$ and for every strategy profile Σ' with $\Sigma' \sim_C \Sigma$, we have $\psi \in \mathbf{m}(\text{out}(\Sigma', u))$. We say that \mathbf{m} **refutes the goal assignment γ at u** if for every strategy profile $\Sigma \in \Pi_{a \in \text{Agt}}$ there is some C in the support of γ with $\gamma(C) = X\psi$ and some strategy profile Σ' with $\Sigma' \sim_C \Sigma$ such that $\bar{\psi} \in \mathbf{m}(\text{out}(\Sigma', u))$.

Definition 19. A Φ -network $\mathcal{N} = (T, L, \mathcal{G})$ is said to be **one-step coherent** if, for every non-leaf node $u \in T$ such that $\mathcal{G}(u) = (\text{Act}, \text{act}, T, \text{out})$, the marking L verifies every next time goal assignment γ such that $\langle\!\langle\gamma\rangle\!\rangle \in L(u)$ and refutes every next time goal assignment γ such that $\neg\langle\!\langle\gamma\rangle\!\rangle \in L(u)$.

8.2 Eventualities and Defects

Definition 20. A **TypeU-eventuality** is a formula $\langle\!\langle\gamma\rangle\!\rangle$ where $\gamma \in \text{TypeU}$. A **TypeG-eventuality** is a formula of the form $\neg\langle\!\langle\gamma\rangle\!\rangle$ where $\gamma \in \text{TypeG}$.

Definition 21. Let $\langle\!\langle\gamma\rangle\!\rangle$ be a TypeU-eventuality, where γ is a goal assignment for the family $\mathcal{F} = \{C_1, \dots, C_n, D_1, \dots, D_m\}$ or $\mathcal{F} = \{C_1, \dots, C_n\}$ defined by

$$\gamma(C_1) = \alpha_1 \cup \beta_1, \dots, \gamma(C_n) = \alpha_n \cup \beta_n,$$

and (if $m > 0$)

$$\gamma(D_1) = G\chi_1, \dots, \gamma(D_m) = G\chi_m.$$

Let $\mathcal{N} = (T, L, \mathcal{G})$ be a one-step coherent network. Given a node $u \in T$, we say that $\langle\!\langle\gamma\rangle\!\rangle$ **is partially fulfilled in 0 steps at u in \mathcal{N}** if there is some $i \in \{1, \dots, n\}$ such that $\beta_i \wedge \langle\!\langle\gamma \setminus C_i\rangle\!\rangle \in L(u)$. For any natural number $k \geq 0$, we say that $\langle\!\langle\gamma\rangle\!\rangle$ **is partially fulfilled in $k + 1$ steps at u** if it is either partially fulfilled in 0 steps, or u is a non-leaf node and the following conditions hold:

- $\alpha_i \in L(u)$ for all $i \in \{1, \dots, n\}$,
- $\chi_j \in L(u)$ for all $j \in \{1, \dots, m\}$,
- there is a marking \mathbf{m} that verifies $\langle\!\langle\Delta\gamma\rangle\!\rangle$ at u and is such that for all $v \in T$ such that v is a child of u , $\langle\!\langle\gamma\rangle\!\rangle$ is partially fulfilled in k steps at v whenever $\langle\!\langle\gamma\rangle\!\rangle \in \mathbf{m}(v)$.

Lastly, we say that $\langle\!\langle\gamma\rangle\!\rangle$ **is partially fulfilled at u** if it is partially fulfilled in k steps at u for some $k \geq 0$.

Definition 22. Let $\neg\langle\!\langle\gamma\rangle\!\rangle$ be an eventuality in TypeG, where γ is the goal assignment for the family $\mathcal{F} = \{D_1, \dots, D_m\}$ defined by

$$\gamma(D_1) = G\chi_1, \dots, \gamma(D_m) = G\chi_m.$$

Let $\mathcal{N} = (T, L, \mathcal{G})$ be a one-step coherent network. Given a node $u \in T$, we say that $\neg\langle\!\langle\gamma\rangle\!\rangle$ **is partially fulfilled in 0 steps at u in \mathcal{N}** if there is some $i \in \{1, \dots, n\}$ such that $\bar{\chi}_i \in L(u)$.

For any natural number $k \geq 0$, we say that $\neg\langle\!\langle\gamma\rangle\!\rangle$ **is partially fulfilled in $k + 1$ steps at u** if it is either partially fulfilled in 0 steps, or u is a non-leaf node and there exists a marking \mathbf{m} that refutes $\langle\!\langle\Delta\gamma\rangle\!\rangle$ at u , and such that for all $v \in T$ such that v is a child of u , $\neg\langle\!\langle\gamma\rangle\!\rangle$ is partially fulfilled in k steps at v whenever $\neg\langle\!\langle\gamma\rangle\!\rangle \in \mathbf{m}(v)$. Lastly, we say that $\neg\langle\!\langle\gamma\rangle\!\rangle$ **is partially fulfilled at u** if it is partially fulfilled in k steps at u for some $k \geq 0$.

Definition 23. A **defect** of a network $\mathcal{N} = (T, L, \mathcal{G})$ is a pair (u, ϕ) such that $u \in T$, $\phi \in L(u)$ is an eventuality which is not partially fulfilled at u .

PROPOSITION 12. *Let $\mathcal{N} \sqsubseteq \mathcal{N}'$ and let (u, φ) be a defect of \mathcal{N}' . If u belongs to \mathcal{N} , then (u, φ) is a defect of \mathcal{N} , as well.*

PROOF. A trivial induction on k shows that, if an eventuality of any of the two types is partially fulfilled in k steps at u in \mathcal{N} , then it is partially fulfilled in k steps at the same node in \mathcal{N}' as well. \square

Definition 24. A network is said to be **perfect** if it is one-step coherent, has no leaves, and no defects.

Definition 25. Given a perfect network $\mathcal{N} = (T, L, \mathcal{G})$ we define a game model $\mathcal{M}(\mathcal{N}) = (S, \text{Act}, \mathfrak{g}, V)$ as follows. We take S to be the set of all nodes in T , and $\mathfrak{g} = \mathcal{G}$. Finally, we set $V(p) = \{v \in T \mid p \in L(v)\}$. We call $\mathcal{M}(\mathcal{N})$ the **induced model of the network \mathcal{N}** .

The following proposition will relate truth sets $\llbracket \varphi \rrbracket_{\mathcal{M}(\mathcal{N})}$ of formulas in the induced model of a network to the set of nodes v with $\varphi \in L(v)$. To clearly distinguish the latter from the former, we introduce the following notation:

$$\llbracket \varphi \rrbracket_{\mathcal{N}} := \{v \in T \mid \varphi \in L(v)\}.$$

For the proof of the following proposition, see the appendix.

PROPOSITION 13. *Every Φ -atom that is the label of some node in a perfect Φ -network $\mathcal{N} = (T, L, \mathcal{G})$ is true at the respective state of the model $\mathcal{M}(\mathcal{N})$ induced by that network.*

8.3 Constructing a Perfect Network

8.3.1 Step 1: Extending Leaves in Coherent Networks.

PROPOSITION 14. *Let \mathcal{N} be any finite, one-step coherent network, and let u be a leaf in \mathcal{N} . Then there exists a finite and one-step coherent network \mathcal{N}' such that $\mathcal{N} \sqsubseteq \mathcal{N}'$ and such that u is not a leaf in \mathcal{N}' .*

PROOF. Let $\Gamma = \{\langle \gamma \rangle \mid \langle \gamma \rangle \in L(u)\} \cup \neg\{\langle \gamma \rangle \mid \neg\langle \gamma \rangle \in L(u)\}$. This is a consistent modal one-step theory, so let $(\text{Act}, \text{act}, \mathcal{P}(\Phi), \text{out})$ be the maximal consistent game form $\mathcal{M}(\Gamma)$ provided by Theorem 4. We construct the network $\mathcal{N}' = (T', L', \mathcal{G}')$ as follows. For each atom Ψ in the image of the function out (which is always non-empty), add a new successor v_Ψ to u , and set $L(v_\Psi) = \Psi$. Let S denote the set of successors of u added in this manner. We construct a new game form $\mathcal{G}'(u) = (\text{Act}, \text{act}, S, \text{out}')$ by setting, for each action profile ζ , $\text{out}'(\zeta) := \{v_\Psi\}$ where $\Psi = \text{out}(\zeta)$. This completes the definition of \mathcal{N}' . It is clear that $\mathcal{N} \sqsubseteq \mathcal{N}'$, and the conditions given in Theorem 4 directly entail (by design) that the network \mathcal{N}' is one-step coherent. Since at least one successor was added to u , this node is no longer a leaf in \mathcal{N}' . \square

8.3.2 Step 2: Pushing Defects Towards Leaves.

PROPOSITION 15. *Let \mathcal{N} be a finite, one-step coherent network and let (u, φ) be a defect of \mathcal{N} . Then there exists a set $\{v_1, \dots, v_k\}$ of leaves in \mathcal{N} such that*

- For each $i \in \{1, \dots, k\}$, the pair (v_i, φ) is a defect of \mathcal{N} , and
- For any one-step coherent network \mathcal{N}' such that $\mathcal{N} \sqsubseteq \mathcal{N}'$, if (u, φ) is still a defect in \mathcal{N}' then there is some $i \in \{1, \dots, k\}$ such that (v_i, φ) is a defect of \mathcal{N}' .

PROOF. We focus on the case of a type U eventuality; the case of type G eventualities is very similar. Let $\langle \gamma \rangle$ be a type U eventuality. We say that a defect $(u, \langle \gamma \rangle)$ **one-step generates** a defect $(v, \langle \gamma \rangle)$ if v is one of the children of u , and $(v, \langle \gamma \rangle)$ is a defect of \mathcal{N} . We then say that a

defect $(u, \langle \gamma \rangle)$ **generates** a defect $(v, \langle \gamma \rangle)$ if $(v, \langle \gamma \rangle)$ is a successor of $(u, \langle \gamma \rangle)$ with respect to the transitive closure of the one-step generation relation.

Now suppose that $(u, \langle \gamma \rangle)$ is a defect of \mathcal{N} . We claim that the set of leaves l such that $(l, \langle \gamma \rangle)$ is a defect generated by the defect $(u, \langle \gamma \rangle)$ satisfies the conditions of the proposition. The first condition holds by definition. To prove this, consider any one-step coherent network \mathcal{N}' such that $\mathcal{N} \sqsubseteq \mathcal{N}'$. We show that for every defect $(w, \langle \gamma \rangle)$ of \mathcal{N} , if $(w, \langle \gamma \rangle)$ is still a defect of \mathcal{N}' , then the same holds for some defect $(v, \langle \gamma \rangle)$ of \mathcal{N} that is one-step generated by $(w, \langle \gamma \rangle)$. By repeatedly applying this claim, starting with the defect $(u, \langle \gamma \rangle)$, we eventually reach a leaf l such that $(l, \langle \gamma \rangle)$ is a defect generated by the defect $(u, \langle \gamma \rangle)$, and is still a defect in \mathcal{N}' .

So, let $(w, \langle \gamma \rangle)$ be a non-leaf defect of \mathcal{N} , such that $(w, \langle \gamma \rangle)$ is still a defect of \mathcal{N}' . Suppose, for a contradiction, that for all the children v of w , $(v, \langle \gamma \rangle)$ is *not* a defect of \mathcal{N}' . This means that for all children v of w in \mathcal{N} , and hence for all children of w in \mathcal{N}' since $\mathcal{N} \sqsubseteq \mathcal{N}'$, there is some k_v for which the eventuality $\langle \gamma \rangle$ is partially fulfilled in k_v

steps at v in \mathcal{N} . Let K be the maximum of these numbers k_v , which exists since the set of successors of w is finite. By one-step coherence of the network \mathcal{N}' it follows that $\langle \gamma \rangle$ is partially fulfilled in $K + 1$ steps at w in \mathcal{N}' , witnessed by the labelling function L of \mathcal{N} regarded as a marking of \mathcal{N}' . Thus, we have reached a contradiction, which completes the proof. \square

8.3.3 Step 3: Removing Defects. We now show how to remove defects from a network:

PROPOSITION 16. *Let (u, φ) be a defect of some finite, one-step coherent network \mathcal{N} . Then there exists a finite, one-step coherent network \mathcal{N}' such that $\mathcal{N} \sqsubseteq \mathcal{N}'$, and such that (u, φ) is not a defect of \mathcal{N}' .*

PROOF. By Proposition 15, we may assume w.l.o.g. that the defect (u, φ) is such that u is a leaf: if we can show how to remove the defect φ at a single leaf, then, clearly, we can repeat the procedure to remove φ at each leaf in the set $\{v_1, \dots, v_k\}$. (Note that our procedure for removing a defect at a single leaf v given below will not affect any other leaves, i.e., each leaf in the original network besides v will still be a leaf in the new network.) Combined with Proposition 15 this proves the result.

So, suppose that (u, φ) is a defect and u is a leaf. It is sufficient to show that there is a finite, one-step coherent network \mathcal{N}'' in which the root has the same label as u in \mathcal{N} , and in which the eventuality φ is partially fulfilled. We can then simply identify the root of the network \mathcal{N}' with the leaf u in \mathcal{N} to form a finite, one-step coherent network \mathcal{N}' such that $\mathcal{N}'' \sqsubseteq \mathcal{N}'$ and $\mathcal{N} \sqsubseteq \mathcal{N}'$. By Proposition 12, the eventuality φ is partially fulfilled at u in \mathcal{N}' .

Consider the Φ -atoms Ψ such that $\varphi \in \Psi$ and there exists a finite, one-step coherent network in which the root is labelled by Ψ and the eventuality φ is partially fulfilled. Let δ be the disjunction of all conjunctions of the form $\bigwedge \Psi$ for all such Φ -atoms Ψ . (This is well-defined since the set of all such conjunctions is finite, as long as we disallow conjunctions with redundant multiple occurrences of the same conjunct.) The result then follows from the following claim, which is proved in the appendix. \square

CLAIM 2. $\vdash \varphi \rightarrow \delta$.

8.3.4 Final Step: Putting Everything Together. The following proposition is proved in the appendix.

PROPOSITION 17. *Let \mathcal{N} be a finite, one-step coherent Φ -network. Then there exists a finite, one-step coherent Φ -network \mathcal{N}' such that*

- (1) $\mathcal{N} \sqsubseteq \mathcal{N}'$,
- (2) no leaf of \mathcal{N} is a leaf of \mathcal{N}' ,
- (3) no defect of \mathcal{N} is a defect of \mathcal{N}' .

PROPOSITION 18. *Every Φ -atom is the label of the root of some perfect Φ -network.*

PROOF. Take any Φ -atom Ψ . We construct an infinite chain of finite one-step coherent Φ -networks $\mathcal{N}_0 \sqsubseteq \mathcal{N}_1 \sqsubseteq \dots$ inductively as follows.

We start with $\mathcal{N}_0 = (T_0, L_0, \mathcal{G}_0)$ where $T_0 = \{u_0\}$ is a singleton, $L_0(u_0) = \Psi$ and \mathcal{G}_0 is empty (there are no non-leaf nodes). This is trivially one-step coherent.

Suppose, we have constructed the finite one-step coherent Φ -networks $\mathcal{N}_0 \sqsubseteq \mathcal{N}_1 \sqsubseteq \dots \sqsubseteq \mathcal{N}_n$. Then we apply Proposition 17 to construct a finite, one-step coherent network \mathcal{N}_{n+1} such that $\mathcal{N}_n \sqsubseteq \mathcal{N}_{n+1}$, no leaves in \mathcal{N}_n are still leaves in \mathcal{N}_{n+1} , and no defects in \mathcal{N}_n are still defects in \mathcal{N}_{n+1} .

Finally, we construct the network \mathcal{N} as the union of the chain $\mathcal{N}_0 \sqsubseteq \mathcal{N}_1 \sqsubseteq \dots$. Clearly, it is still one-step coherent and has no leaves and no defects, i.e., it is perfect. \square

We can now state and prove the completeness theorem.

THEOREM 5 (COMPLETENESS OF Ax_{TLCGA}). *Let Γ be a finite Ax_{TLCGA} -consistent set of TLCGA-formulae. Then Γ is satisfied in some concurrent game model.*

PROOF. Let Φ be the extended Fischer-Ladner closure of Γ and let Γ^* be a Φ -atom containing Γ (which exists, by a standard version of Lindenbaum's lemma). By Proposition 18, Γ^* is the label of the root of some perfect Φ -network. Then, by Proposition 13, Γ^* is true at the respective state of the model $\mathcal{M}(\mathcal{N})$ induced by that network. \square

9 FINITE MODEL PROPERTY AND DECIDABILITY

In this section, we show finite model property and decidability for our logic TLCGA. Since we have a truth-preserving and effective translation of the language $\mathcal{L}^{\text{TLCGA}}$ into the fixpoint logic $\mathcal{L}_\mu^{\text{XCGA}}$, it suffices to prove finite model property and decidability for the latter. Here, we will avail ourselves of some abstract results from the literature on coalgebraic modal fixpoint logic. In particular, a general bounded-size model property for coalgebraic μ -calculi was proved in [21], and since $\mathcal{L}_\mu^{\text{XCGA}}$ is an instance of coalgebraic μ -calculus, we are almost done. There is one subtlety that we need to deal with, concerning the notion of “finiteness” of a model. There are two distinct notions of “finite model” that we may consider:

Definition 26. Let $\mathcal{M} = (S, \text{Act}, g, \text{out}, V)$ be a concurrent game model. We say that \mathcal{M} is *state-finite* if S is a finite set. We say that \mathcal{M} is *action-finite* if Act is a finite set. We say that \mathcal{M} is *finite* if it is both state-finite and action-finite.

We get the following “state-finite model property”, as a direct corollary of the general finite model theorem from [21]:

THEOREM 6. *Any satisfiable formula of $\mathcal{L}_\mu^{\text{XCGA}}$ is satisfiable in a state-finite model.*

However, what we want is a proper finite model property. We can obtain this with a little bit of extra work. First, we obtain the following “action-finite model property”:

LEMMA 1. *Any satisfiable formula of $\mathcal{L}^{\text{TLCGA}}$ is satisfiable in an action-finite model.*

PROOF. Since the size of the set of actions in the game forms constructed in the proof of the one-step completeness theorem (Theorem 4) has an exponential upper bound that depends on the size of a finite consistent modal one-step theory, and since our construction of a model for a consistent $\mathcal{L}^{\text{TLCGA}}$ -formula can easily be seen to provide an action-finite model. \square

THEOREM 7 (FINITE MODEL PROPERTY). *Any satisfiable formula of $\mathcal{L}^{\text{TLCGA}}$ is satisfiable in a finite model.*

PROOF. Let φ be a satisfiable formula of $\mathcal{L}^{\text{TLCGA}}$. By soundness of our proof system φ is consistent, and hence is satisfiable in a model $\mathcal{M} = (S, \text{Act}, g, \text{out}, V)$ where the set Act is finite. Hence the corresponding equivalent formula φ' in $\mathcal{L}_\mu^{\text{XCGA}}$ is satisfiable in this model too.

But the frame $(S, \text{Act}, g, \text{out})$ can be equivalently represented as a coalgebra $f : S \rightarrow G^{\text{Act}}S$ for a functor G^{Act} in which the set of actions Act is explicitly encoded, so the triple (S, f, V) is a G^{Act} -model in which φ' is satisfiable. By the finite model property theorem proved in [21], φ' is satisfiable in a G^{Act} -model (S', f', V') for which S' is finite. This model can equivalently be represented as a concurrent game model $(S', \text{Act}, g', \text{out}')$ in which φ' is satisfiable, hence also φ . Since S' is finite and Act is finite, this is a finite model for φ . \square

Together with our completeness result for TLCGA, this implies:

THEOREM 8. *The satisfiability problem for TLCGA is decidable.*

Note that there is no need for an explicit bound on the size of a satisfying finite model for this result; completeness for TLCGA ensures that we can computably enumerate the valid formulas of $\mathcal{L}^{\text{TLCGA}}$, and the finite model property means that we can computably enumerate the non-valid formulas as well. Hence, the set of valid formulas is a computable set.

Since the logic TLCGA^+ also embeds into $\mathcal{L}_\mu^{\text{XCGA}}$, we get a further corollary:

THEOREM 9. *The logic TLCGA^+ has the finite model property.*

The proof of Theorem 8 took a detour via finite model property for the language $\mathcal{L}_\mu^{\text{XCGA}}$. We shall now look closer at the latter language and obtain an upper bound on the complexity of satisfiability for $\mathcal{L}_\mu^{\text{XCGA}}$, under the assumption that the set Agt of agents is fixed. Decidability of TLCGA^+ will then follow as a corollary. To obtain these results, we will require a closer analysis of the one-step satisfiability problem.

Definition 27. Let V be a fixed set of proposition variables. The set of **positive one-step formulas** over V is generated by the following grammar:

$$\langle \gamma_0 \rangle \mid \neg \langle \gamma_1 \rangle \mid \varphi \wedge \varphi \mid \varphi \vee \varphi,$$

where γ_0 is a goal assignment such that for each coalition C in the support of γ_0 we have $\gamma_0(C) = Xp$ for some $p \in V$, called a **positive goal assignment**, and γ_1 is a goal assignment such that for each coalition C in the support of γ_0 we have $\gamma_1(C) = X\neg p$ for some $p \in V$, called a **negative goal assignment**.

A **positive one-step sequent** over V is a finite set Γ of formulas, each of which is either of the form $\langle \gamma_0 \rangle$ where γ_0 is a positive goal assignment, or of the form $\neg \langle \gamma_1 \rangle$ where γ_1 is a negative goal assignment.

Let Γ be a positive one-step sequent over V . A **redistribution** of Γ is a tuple $(C_1, \gamma_1, \dots, C_n, \gamma_n)$ where C_1, \dots, C_n is a set of pairwise disjoint coalitions and each γ_i is a positive goal assignment with $\langle \gamma_i \rangle \in \Gamma$. The intuition behind this notion is as follows: suppose that $\gamma_1, \dots, \gamma_n$ are positive goal assignments with $\langle \gamma_i \rangle \in \Gamma$ for each $i \in \{1, \dots, n\}$. Then according to the sequent Γ , there are action profiles ζ_1, \dots, ζ_n such that for each ζ_i , every coalition C ensures its goal formula $\gamma_i(C)$ with respect to the action profile ζ_i . But then, if C_1, \dots, C_n are any pairwise disjoint coalitions, the coalition $C_1 \cup \dots \cup C_n$ has a joint action in which each coalition C_i ensures its goal formula $\gamma_i(C_i)$, by playing in accordance with the action profile ζ_i . In other words, the action profiles ζ_1, \dots, ζ_n can be combined into a joint action for $C_1 \cup \dots \cup C_n$ by restricting each action profile ζ_i to the coalition C_i , and then “gluing together” the resulting joint actions into a joint action for $C_1 \cup \dots \cup C_n$. The set of redistributions of Γ can be viewed as a set of notations, for each of the different ways that

action profiles corresponding to the positive goal assignments occurring in Γ can be combined in this manner.

Note that a redistribution $(C_1, \gamma_1, \dots, C_n, \gamma_n)$ of Γ can be uniquely represented by a map:

$$f : \mathcal{P} \text{Agt} \rightarrow \Gamma \cup \{*\},$$

defined by setting $f(C_i) = \langle \gamma_i \rangle$ and $f(B) = *$ for $B \notin \{C_1, \dots, C_n\}$. Hence, there are at most $(|\Gamma| + 1)^k$ redistributions of Γ , where $k = 2^{|\text{Agt}|}$. Recall that the set Agt is fixed, so $2^{|\text{Agt}|}$ is a constant, giving a polynomial bound on the number of redistributions.

A goal assignment γ with $\langle \gamma \rangle \in \Gamma$ or $\neg \langle \gamma \rangle \in \Gamma$ we be called a **relevant goal assignment** for Γ . Given a redistribution $R = (C_1, \gamma_1, \dots, C_n, \gamma_n)$, we let $F(R)$ denote the set:

$$\{p \mid \exists i, B \subseteq C_i : \gamma_i(B) = Xp\}.$$

Intuitively, $F(R)$ is the set of goals forced by coalitions that have formed in the action profile represented by the redistribution R . Given a triple (R, γ', C') consisting of a redistribution $R = (C_1, \gamma_1, \dots, C_n, \gamma_n)$, a negative relevant goal assignment γ' , and a coalition C' such that $\gamma'(C') = X\neg q$, we denote by $F(R, \gamma', C')$ the set:

$$\{p \mid \exists i, B \subseteq C_i \cap C' : \gamma_i(B) = Xp\} \cup \{q\}.$$

Intuitively, the set $F(R, \gamma', C')$ is the minimal set of variables that have to be in the outcome of a strategy profile ζ , in which each coalition $C_i \in \{C_1, \dots, C_n\}$ acts towards its goal with respect to the goal assignment γ_i , and at the same time the players in $\text{Agt} \setminus C'$ act to block the goal of C' according to the goal assignment γ' .

Definition 28. Let V be a given set of proposition variables. A **satisfiability constraint** over V is a subset of $\mathcal{P}(V)$. Given a positive one-step sequent Γ over V and a satisfiability constraint \mathcal{S} over V , we say that Γ is **\mathcal{S} -satisfiable** if there exists a game form $\mathcal{G} = (\text{Act}, \text{act}, \mathcal{P}(V), \text{out})$ such that

- If $\langle \gamma \rangle \in \Gamma$ then there exists an action profile ζ such that, for every coalition C , if $\gamma(C) = Xp$ then $p \in \text{out}(\zeta')$ for every $\zeta' \sim_C \zeta$.
- If $\neg \langle \gamma \rangle \in \Gamma$ then for every action profile ζ , there is some coalition C and some $\zeta' \sim_C \zeta$ such that $p \in \text{out}(\zeta')$, where $\gamma(C) = X\neg p$.
- For every action profile ζ , there is some $V' \in \mathcal{S}$ with $\text{out}(\zeta) \subseteq V'$.
- For every $V' \in \mathcal{S}$ there is some action profile ζ with $\text{out}(\zeta) \subseteq V'$.

Given a game form \mathcal{G} satisfying these conditions, we say that \mathcal{G} **\mathcal{S} -satisfies** Γ .

The following proposition reduces the question of whether a positive one-step sequent Γ is \mathcal{S} -satisfiable to checking of a few relatively simple combinatorial conditions on the set \mathcal{S} . The proof is in the appendix.

PROPOSITION 19. *Given a set V of variables, a satisfiability constraint \mathcal{S} over V and a positive one-step sequent Γ over V , Γ is \mathcal{S} -satisfiable if and only if*

- (1) *For each redistribution R of Γ , there is $Z \in \mathcal{S}$ with $F(R) \subseteq Z$.*
- (2) *For each redistribution R of Γ and each negative relevant goal assignment γ' , there is some $C' \subseteq \text{Agt}$ such that*
 - *either $C' = \text{Agt}$ and $q \in Z$ for every $Z \in \mathcal{S}$, where $\gamma'(\text{Agt}) = X\neg q$, or*
 - *$C' \neq \text{Agt}$ and there is some $Z \in \mathcal{S}$ such that $F(R, \gamma', C') \subseteq Z$.*

To illustrate this proposition, let $\text{Agt} = \{a, b\}$, $V = \{p, q, r\}$ and let:

$$\Gamma = \{\langle \{a\} \triangleright Xp \rangle, \langle \{b\} \triangleright Xq \rangle, \neg \langle \{b\} \triangleright X\neg r \rangle\}.$$

Then Γ is \mathcal{S} -satisfiable for $\mathcal{S} = \{\{p, q\}, \{q, r\}\}$. However, Γ is not \mathcal{S} -satisfiable for $\{\{p, q\}, \{p, r\}\}$. Consider the redistribution R in which a “votes” for $\langle\{a\} \triangleright X p\rangle$ and b votes for $\langle\{b\} \triangleright X q\rangle$. Then a forces p on their own, and b forces q on their own, so there should be a possible outcome containing both p and q —which there is. But b should not be able force $\neg r$ on their own since $\neg\langle\{b\} \triangleright X \neg r\rangle \in \Gamma$, so there should be some outcome containing q , which b is forcing on their own, but also r . Formally, this is captured by $F(R, \gamma, \{b\}) = \{q, r\}$, where $\langle\gamma\rangle = \langle\{b\} \triangleright X \neg r\rangle$.

PROPOSITION 20. *Let V be a fixed set of proposition variables and \mathcal{S} a fixed satisfiability constraint. Then the problem of checking whether a given positive one-step sequent Γ is \mathcal{S} -satisfiable is solvable in polynomial time.*

PROOF. An easy consequence of Proposition 19: we only have to check the conditions (1) and (2) on \mathcal{S} for each redistribution of Γ and each negative relevant goal assignment. Recalling that there are at most polynomially many redistributions of Γ , the result follows. \square

We extend the notion of \mathcal{S} -satisfiability to arbitrary positive one-step formulas over V in the obvious manner. Let V be a set of proposition variables and $\mathcal{S} \subseteq V$ a satisfiability constraint. The **one-step satisfiability problem** for V, \mathcal{S} is given a positive one-step formula φ over V , decide whether φ is \mathcal{S} -satisfiable.

PROPOSITION 21. *The one-step satisfiability problem for given V, \mathcal{S} can be solved in nondeterministic polynomial time (measured in the length of a given formula φ).*

PROOF. The procedure for checking \mathcal{S} -satisfiability of φ can be done as follows: we keep in memory a list of formulas L , beginning with the singleton list $[\varphi]$. If the leftmost formula in the memory L that is not “atomic”, i.e., not of the form $\langle\gamma\rangle$ or $\neg\langle\gamma\rangle$ for some γ , has \vee as main connective, then we guess one of the disjuncts non-deterministically and replace the disjunction by the guessed disjunct. If the main connective of the leftmost non-atomic formula is \wedge then we remove that formula and instead add each conjunct to the list. Clearly, this reduction can only go on for a number of steps that are bounded by the length of the formula φ , and once each formula in the list is atomic, we apply Proposition 20 to check whether the one-step sequent read off from the list is \mathcal{S} -satisfiable. \square

It follows from a general result in [21] that, if the one-step satisfiability problem for fixed V, \mathcal{S} is solvable in exponential time, then the satisfiability problem for the full μ -calculus language is in double exponential time. From Proposition 21 we thus get

THEOREM 10. *The satisfiability problem for $\mathcal{L}_\mu^{\text{XCGA}}$ is in 2ExpTime .*

As a corollary, we get:

THEOREM 11. *The logic TLCGA^+ , (hence, also TLCGA) is decidable in 3ExpTime .*

PROOF. The translation from TLCGA^+ to $\mathcal{L}_\mu^{\text{XCGA}}$ increases the size of a formula by at most a single exponential, so the result follows from Theorem 10. \square

We do not know whether the 2ExpTime bound given by Theorem 10 is tight. Many variants of the μ -calculus share the same ExpTime -complete complexity of the satisfiability problem as the standard modal μ -calculus, as demonstrated in [15]. However, the conditions under which the general ExpTime -bound of [15] holds are not trivial to verify; in fact we believe they may fail for the language $\mathcal{L}_\mu^{\text{XCGA}}$. Furthermore, even if the satisfiability problem for $\mathcal{L}_\mu^{\text{XCGA}}$ happens to be in ExpTime , this bound does not immediately transfer to the language $\mathcal{L}^{\text{TLCGA}}$, since our translation of $\mathcal{L}^{\text{TLCGA}}$ into $\mathcal{L}_\mu^{\text{XCGA}}$ can produce an exponential blowup in the formula size. In summary, the

2ExpTime-bound for satisfiability in $\mathcal{L}_\mu^{\text{XCGA}}$ stated above is the best we can currently claim for sure about the complexity of that satisfiability problem. The problem of determining the precise complexity of satisfiability, both for $\mathcal{L}^{\text{TLCGA}}$ and for $\mathcal{L}_\mu^{\text{XCGA}}$, may very well turn out to be a difficult one. We leave this as a question for future research.

10 CONCLUDING REMARKS

The present work falls in the line of research employing formal logical methods for modelling, expressing, and reasoning about strategic interactions in multi-agent systems, and in particular multiplayer games, initiated with the introduction of logics such as CL, ATL, and Strategy Logic. The coalitional goal assignment operator $\langle \cdot \rangle$ introduced here covers as special cases the modal operators for strategic abilities featuring in CL and ATL. Furthermore, whereas these strategic operators assume purely adversarial behaviour of the opposing agents and coalitions, and express unconditional ability of the proponent coalition to achieve its objective against any such objective, the operator $\langle \cdot \rangle$ expresses a natural combination of cooperative and non-cooperative interactions which is more common and realistic in “real-life” multi-agent scenarios. Whereas this operator is already quite expressive, it formalises but one general and important pattern of such interaction. Other such patterns, formalising variants of *conditional* strategic reasoning, have been proposed and studied in [24], and certainly, more are currently being identified and studied. The patterns of strategic interaction formalised by $\langle \cdot \rangle$ also enable the expression of new versions of socially relevant solution concepts, such as co-equilibrium, discussed in Section 4, opening new perspectives towards a formal analysis of multi-agent strategic interaction.

On the more technical aspects of our work, we note that, while the most important logical questions about TLCGA have been resolved here, there is still substantial technical follow-up work to do, including complete axiomatizations of important variations and extensions, such as TLCGA^+ and possibly the full $\mathcal{L}_\mu^{\text{XCGA}}$, as well as development of practically efficient tableau-based algorithmic methods for deciding satisfiability and well as for solving the relevant model checking problems. In particular, identifying the exact complexities of these problems, for which reasonable upper bounds have been established in Section 9, is yet to be completed. For solving these problems and exploring further research directions, we hope that the interaction with fragments of Strategy logic, initiated in Section 3.4, will prove fruitful both ways. In particular, the distinction between path-based and play-based semantics established in Section 3.3 suggests that the latter type of semantics is worth exploring in the context of Strategy logic, too.

Lastly, we note that in this article we have mainly explored the theoretical foundations and the purely logical and computational aspects of the logics TLCGA, TLCGA^+ , and $\mathcal{L}_\mu^{\text{XCGA}}$, whereas their potential for applications to game-theory and, more generally, to formal modelling and reasoning about multi-agent strategic interaction has only been indicated in Section 4, but is yet to be unfolded in future work.

APPENDIX

A ELECTRONIC APPENDIX

A.1 Proof of Proposition 6

PROOF. We prove each implication of the equivalence separately.

Right to left: By Theorem 1, the truth set of $\langle \gamma \rangle$ is a fixpoint of the operator defined by the formula $\text{ind}(\gamma, z)$. Therefore, $\mu z. \text{ind}(\gamma, z) \rightarrow \langle \gamma \rangle$ is semantically valid.

Left to right: Conversely, suppose that $\mathcal{M}, v \Vdash \langle \gamma \rangle$. Since γ is a long-term temporal goal assignment of type U, we can assume that its support is $\mathcal{F} = \{C_1, \dots, C_n, D_1, \dots, D_m\}$, and that it

maps each C_i to $\alpha_i \cup \beta_i$, and maps each D_j to $G\chi_j$. Then γ is witnessed by some strategy profile Σ , such for each $i \in \{1, \dots, n\}$, the formula $\alpha_i \cup \beta_i$ holds on every computation path in $\text{paths}(v, \Sigma, C_i)$ and for each $j \in \{1, \dots, m\}$, the formula $G\chi_j$ holds on every computation path in $\text{paths}(v, \Sigma, D_j)$. We define a set T of histories as follows: we set $(v_0 \tau_0 v_1 \dots v_{h-1} \tau_{h-1} v_h) \in T$ if and only if

- (i) $v_0 = v$ and the word $v_0 \dots v_h$ is an initial segment of some computation path in $\text{paths}(v, \Sigma, \cup \mathcal{F})$, and
- (ii) there are no indices i, j such that $0 \leq j < h$ and $0 \leq i \leq n$ and $\mathcal{M}, v_j \Vdash \beta_i$.

It is clear that $v \in T$, and that T can be viewed as a tree rooted at v where the successors of a node \vec{w} in T are the elements of T of the form $\vec{w}u$ for some u . Furthermore, since $\Sigma, v \Vdash \gamma$ and $C_i \subseteq \cup \mathcal{F}$ for each $i \in \{1, \dots, n\}$, the formula $\alpha_i \cup \beta_i$ holds on each computation path in $\text{paths}(v, \Sigma, \cup \mathcal{F})$. Hence, the tree T is *Noetherian*, i.e., it has no infinite branches. This means that we may reason by *bar induction*: to show that a property $P(x)$ holds for every element of T , we show that $P(l)$ holds if l is a leaf, and that if $P(\vec{w} \cdot w')$ holds for every child $\vec{w} \cdot w' \in T$ of some non-leaf node \vec{w} of T , then $P(\vec{w})$ holds as well.

We shall prove the following statement by bar induction: for every word \vec{w} in T with last element u , we have $\mathcal{M}, u \Vdash \mu z.\text{ind}(\gamma, z)$. It follows eventually that $\mathcal{M}, v \Vdash \mu z.\text{ind}(\gamma, z)$, as required.

Case 1: \vec{w} is a leaf. Since \vec{w} is a leaf in T , the last element u of \vec{w} must satisfy one of the formulae β_1, \dots, β_n , since otherwise the extension $\vec{w} \cdot \text{out}((\Sigma(a)(\vec{w})_{a \in \text{Ag}_t}), u)$ (where $(\Sigma(a)(\vec{w})_{a \in \text{Ag}_t})$ is the action profile assigned by Σ at u) would be a child of \vec{w} in T . Furthermore, none of these formulae are true in any elements of \vec{w} *except* the last one, since $\vec{w} \in T$. Since, for every computation path $\pi \in \text{paths}(u, \Sigma, \cup \mathcal{F})$, the infinite word $\vec{w} \cdot \pi$ belongs to $\text{paths}(v, \Sigma, \cup \mathcal{F}) \subseteq \text{paths}(v, \Sigma, C_i)$, it follows that for each $i \in \{1, \dots, n\}$ the formula $\alpha_i \cup \beta_i$ holds for all computation paths in $\text{paths}(u, \Sigma, C_i)$. Likewise, for each $j \in \{1, \dots, m\}$ the formula $G\chi_j$ holds for all computation paths in $\text{paths}(u, \Sigma, D_j)$. Let us fix some β_i that holds at u . We get:

$$\mathcal{M}, u \Vdash \beta_i \wedge \langle \gamma \setminus C_i \rangle,$$

where the conjunct on the right is witnessed by the strategy profile Σ . Since this formula is in $\text{Finish}(\gamma)$, we have $\mathcal{M}, u \Vdash \mu z.\text{ind}(\gamma, z)$ as required.

Case 2: \vec{w} is not a leaf, and the **inductive hypothesis (IH)** holds for all of its children in T .

Since \vec{w} is not a leaf in T , it has some child in T . By definition of T this means that the formulae β_1, \dots, β_n are false on every state in \vec{w} , including its last element, which we denote by u . Let us pick any $B \subsetneq \cup \mathcal{F}$. It is not hard to show that the following formula

$$\langle \langle B_1 \triangleright \gamma(B_1), \dots, B_k \triangleright \gamma(B_k) \rangle \rangle,$$

where B_1, \dots, B_k are the coalitions in \mathcal{F} that are contained in B , holds for each $u' \in \text{out}(u, \Sigma, B)$, witnessed by the strategy profile Σ . The key thing to note is that, if $C_i \in \{B_1, \dots, B_k\}$ for some $i \in \{1, \dots, n\}$, then for every $u' \in \text{out}(u, \Sigma, B)$ and every path $\pi \in \text{paths}(u', \Sigma, C_i)$, the path $\vec{w} \cdot \pi$ is in $\text{paths}(v, \Sigma, C_i)$ and, therefore, satisfies the path formula $\alpha_i \cup \beta_i$. But, since β_i was false everywhere in \vec{w} , this means that the path formula $\alpha_i \cup \beta_i$ holds on the computation path π as well.

It now follows that for all $B \subsetneq \cup \mathcal{F}$, the path formula

$$X \langle \langle B_1 \triangleright \gamma(B_1), \dots, B_k \triangleright \gamma(B_k) \rangle \rangle,$$

where B_1, \dots, B_k are as above, holds for each computation path in $\text{paths}(u, \Sigma, B)$. Furthermore, by the bar induction hypothesis the path formula $X\mu z.\text{ind}(\gamma, z)$ holds for all $\pi \in \text{paths}(u, \Sigma, \cup \mathcal{F})$, since for every such path $us_0s_1s_2 \dots$ the word $\vec{w} \cdot s_0$ is a child of \vec{w} in T . Putting all this together, we get:

$$\mathcal{M}, u \Vdash \langle \langle \Delta \gamma \left[\bigcup \mathcal{F} \triangleright X\mu z.\text{ind}(\gamma, z) \right] \rangle \rangle,$$

witnessed by the strategy profile Σ . Besides, each α_i for $i \in \{1, \dots, n\}$ and each χ_j for $j \in \{1, \dots, m\}$ holds at u , since $\vec{w} \in T$, so we get:

$$\mathcal{M}, u \Vdash (\alpha_1 \wedge \dots \wedge \alpha_n \wedge \chi_1 \wedge \dots \wedge \chi_m \wedge \langle\langle \Delta_Y [\bigcup \mathcal{F} \triangleright X\mu z.\text{ind}(y, z)] \rangle\rangle\rangle).$$

This is a disjunct of the unfolding of the fixpoint formula $\mu z.\text{ind}(y, z)$, and so we get $\mathcal{M}, u \Vdash \mu z.\text{ind}(y, z)$, as required. This concludes the proof. \square

A.2 Proof of Proposition 7

PROOF. Given a model $\mathcal{M} = (S, \text{Act}, g, \text{out}, V)$, we prove separately each implication of the equivalence in \mathcal{M} . The proof structure is dually analogous to the proof of Proposition 6, but the argument is different.

Left to right: By Theorem 1, the truth set of $\langle\langle Y \rangle\rangle$ is a fixpoint of the operator defined by the formula $\text{ind}(y, z)$. Therefore, $\langle\langle Y \rangle\rangle \rightarrow v z.\text{ind}(y, z)$ is semantically valid.

Right to left: Let $\theta = v z.\text{ind}(y, z)$ and suppose $\mathcal{M}, v \Vdash \theta$. Consequently, we get $\mathcal{M}, v \Vdash \text{ind}(y, \theta)$, where $\text{ind}(y, \theta) = \chi_1 \wedge \dots \wedge \chi_m \wedge \langle\langle \Delta_Y \{\theta\} \rangle\rangle$, where γ is supported by $\mathcal{F} = \{D_1, \dots, D_m\}$, respectively mapped to $G\chi_1, \dots, G\chi_n$.

Thus, $\mathcal{M}, v \Vdash \chi_i$ for each i and $\mathcal{M}, v \Vdash \langle\langle \Delta_Y \{\theta\} \rangle\rangle$.

So, there exists a strategy profile Σ_0 such that $\Sigma_0, v \Vdash \Delta_Y \{\theta\}$,

hence, for each $D_i \in \mathcal{F}$ and every play $\pi \in \text{Plays}(v, \Sigma_0, D_i)$, we have $\mathcal{M}, \pi \Vdash \Delta_Y \{\theta\}(D_i)$.

Recall that for each $C \subsetneq \bigcup \mathcal{F}$, the respective goal of $\Delta_Y \{\theta\}$ is

$$\Delta_Y \{\theta\}(C) = \Delta_Y(C) = X \left\langle\left\langle (D_j \triangleright G\chi_j)_{D_j \subseteq C} \right\rangle\right\rangle$$

and note that $(D_j \triangleright G\chi_j)_{D_j \subseteq C}$ is precisely $\gamma|_C$. Thus, for each such C and every state $w \in \text{Out}[v, \Sigma_0|_C]$ we have $\mathcal{M}, w \Vdash \langle\langle \gamma|_C \rangle\rangle$. Let $\Sigma^{w, C}$ be a strategy profile witnessing $\gamma|_C$ at w .

Now we will define a strategy profile Σ that witnesses γ at v . Intuitively, for each player a , Σ will combine Σ_0 at v with the strategies $\Sigma^{w, C}$ for $a \in C$, ensured to exist at the respective outcome states w , applied on the respective histories passing through w . For any player $a \in \bigcup \mathcal{F}$ we define $\Sigma(a)$ as follows, on any history h in \mathcal{M} starting at v .

- (1) If h has length 0, i.e., $h = v$, we define $\Sigma(a, h) := \Sigma_0(a, v)$.
- (2) Let $h = v\zeta h'$, where $h' = v_1\zeta_1 \dots v_{n-1}\zeta_{n-1}v_n$. We first compare the action profile ζ with the action profile $\zeta_0 = \Sigma_0(v)$ prescribed by Σ_0 at v . Let C_ζ be the set of players in $\bigcup \mathcal{F}$ whose actions in ζ and ζ_0 are the same. Then $v_1 \in \text{Out}[v, \Sigma_0|_{C_\zeta}]$, so $\mathcal{M}, v_1 \Vdash \langle\langle \gamma|_{C_\zeta} \rangle\rangle$. If $a \notin C_\zeta$, then $\Sigma(a, h)$ is defined arbitrarily. Suppose $a \in C_\zeta$.

If $C_\zeta \subsetneq \bigcup \mathcal{F}$, then we use the strategy σ_a for player a from the strategy profile Σ^{v_1, C_ζ} to define $\Sigma(a, h) := \sigma_a(h')$.

In the case when $C_\zeta = \bigcup \mathcal{F}$, we have that $\Delta_Y \{\theta\}(\bigcup \mathcal{F}) = X\theta$, hence $\mathcal{M}, v_1 \Vdash \theta$. In this case, we define $\Sigma(a, h) := \Sigma_1(a, h')$, where Σ_1 is the strategy recursively defined by the same procedure on paths starting from v_1 . Note that this definition is correct because for every path h , being finite, the procedure will eventually reach the case (1) of a subpath of length 0, where the strategy is defined explicitly.

For all other histories h , the action $\Sigma(a, h)$ is defined arbitrarily.

By virtue of the construction of Σ , for $D_j \in \mathcal{F}$ and every play $\pi \in \text{Plays}(v, \Sigma, D_j)$, we have $\mathcal{M}, \pi \Vdash G\chi_j$. The proof is direct for all $D_j \subsetneq \bigcup \mathcal{F}$, while for $D_j = \bigcup \mathcal{F}$, if applicable, it follows from the fact, proved by induction on n , that $\mathcal{M}, v_n \Vdash \chi_j$, where $v_0 = v$ and $\pi = v_0\zeta_0 v_1\zeta_1 v_2\zeta_2 \dots$

Therefore, $\Sigma, v \Vdash \gamma$, hence $\mathcal{M}, v \Vdash \langle\langle Y \rangle\rangle$. \square

A.3 Proof of Proposition 11

PROOF. The proof can be done by several nested inductions: first, on the number of formulae in Φ , then—in the inductive step—by structural induction on the additional formula φ . The only non-trivial point now is when $\varphi = \langle\!\langle\gamma\rangle\!\rangle$, where γ is a long-term temporal goal assignment, to show that φ adds finitely many formulae to $\text{ecl}(\Phi)$ when recursively taking components. Let \mathcal{F} be the support of γ . We can assume that $\gamma \in \text{TypeU}$, as the case of $\gamma \in \text{TypeG}$ is similar but simpler. We will prove the claim by a 3rd nested induction on the number of coalitions in $\mathcal{F} = \{C_1, \dots, C_n, D_1, \dots, D_m\}$ (we can assume $m > 0$) Let $\gamma(C_1) = \alpha_1 \cup \beta_1, \dots, \gamma(C_n) = \alpha_n \cup \beta_n$ and $\gamma(D_1) = G\chi_1, \dots, \gamma(D_m) = G\chi_m$. Then the component of γ added to $\text{ecl}(\Phi)$ is (by Proposition 5) $\text{ind}(\gamma, \langle\!\langle\gamma\rangle\!\rangle) = \text{unfold}(\gamma) = \bigvee \text{Finish}(\gamma) \vee (\bigwedge \text{UHold}(\gamma) \wedge \bigwedge \text{GHold}(\gamma) \wedge \langle\!\langle\Delta\gamma\rangle\!\rangle)$. By the (nested) inductive hypothesis in the structural induction on φ and the innermost inductive hypothesis, every disjunct $(\beta_i \wedge \langle\!\langle\gamma \setminus C_i\rangle\!\rangle) \in \text{Finish}(\gamma)$ adds only finitely many new components to $\text{ecl}(\Phi)$. Again by the inductive hypothesis on the structure of φ , all subformulae α_i and χ_j , hence all formulae in $\text{UHold}(\gamma)$ and $\text{GHold}(\gamma)$, add only finitely many new components, too. Finally, note that $\Delta\gamma$ is a goal assignment and all goals in it are either X -prefixed goals in γ or $X\gamma$, hence $\langle\!\langle\Delta\gamma\rangle\!\rangle$ only adds finitely many new components, too. The case of negated goal assignment $\varphi = \neg\langle\!\langle\gamma\rangle\!\rangle$, is completely analogous. That completes all inductive steps of the nested inductions, and the proof itself. \square

A.4 Proof of Proposition 13

PROOF. We prove, by structural induction on all formulae $\varphi \in \Phi$, that for every $u \in T$, $\varphi \in L(u)$ iff $\mathcal{M}(\mathcal{N}), u \Vdash \varphi$. The clauses for propositional variables and Boolean connectives are standard, so we omit them. Since all formulae in Φ are assumed to be in normal form, for each formula in Φ of the shape $\langle\!\langle\gamma\rangle\!\rangle$, the goal assignment γ is either next time or long-term temporal. The induction step in the first case follows from the fact that the network \mathcal{N} is assumed to be one-step coherent and the IH. Indeed, since $L(u)$ is a Φ -atom, exactly one of $\langle\!\langle\gamma\rangle\!\rangle$ and $\neg\langle\!\langle\gamma\rangle\!\rangle$ is in $L(u)$. Then, since γ is next time and Φ is closed, for every coalition C , $\gamma(C) = X\psi$ where $\psi \in \Phi$. Now, if $\langle\!\langle\gamma\rangle\!\rangle \in L(u)$ then the marking L verifies γ at u and, by the IH applied to all $\gamma(C)$, we obtain that $\mathcal{M}(\mathcal{N}), u \Vdash \langle\!\langle\gamma\rangle\!\rangle$. Likewise, if $\neg\langle\!\langle\gamma\rangle\!\rangle \in L(u)$, then the marking L refutes γ at u and, by the IH applied to all $\gamma(C)$, we obtain that $\mathcal{M}(\mathcal{N}), u \not\Vdash \langle\!\langle\gamma\rangle\!\rangle$. This completes the case.

Now, we focus on the case where γ is long-term temporal. The claim will be proved by a sub-induction on the size (number of coalitions) of the support \mathcal{F} of γ .

Base case: $\mathcal{F} = \emptyset$. Then γ is the trivial goal assignment, and since a perfect network has no leaves this case is trivial.

Induction step: We now suppose that $\mathcal{F} = \{C_1, \dots, C_n\}$ is of size n , and that the induction hypothesis holds for all γ' with support of size $< n$. We divide the proof that the induction hypothesis holds for γ into two sub-cases, depending on the type of γ . We will first prove the easier sub-case.

Case: γ is in TypeG . Let γ be supported by $\mathcal{F} = \{D_1, \dots, D_m\}$ and defined by: $\gamma(D_1) = G\chi_1, \dots, \gamma(D_m) = G\chi_m$. We claim that the set

$$\llbracket \langle\!\langle\gamma\rangle\!\rangle \rrbracket_{\mathcal{N}} := \{u \in T \mid \langle\!\langle\gamma\rangle\!\rangle \in L(u)\},$$

is the greatest post-fixpoint νf of the monotone map f induced by $\text{ind}(\gamma, z)$ in $\mathcal{M}(\mathcal{N})$, defined by

$$f(Z) = \left\{ u \in T \mid \mathcal{M}(\mathcal{N}), u \Vdash_{[z \mapsto Z]} \text{ind}(\gamma, z) \right\}.$$

By Proposition 7, the greatest post-fixpoint νf defined above is equal to $\llbracket \langle\!\langle\gamma\rangle\!\rangle \rrbracket_{\mathcal{M}(\mathcal{N})}$, whence the case will follow.

For the inclusion $[\langle Y \rangle]_{\mathcal{N}} \subseteq \nu f$, we reason by co-induction. That is, we prove that the set $[\langle Y \rangle]_{\mathcal{N}}$ is a post-fixpoint of f . Recall that $\text{ind}(y, z) = \chi_1 \wedge \dots \wedge \chi_m \wedge \langle \Delta Y \{z\} \rangle$, where $\Delta Y \{\phi\}$ denotes $\Delta Y [\bigcup \mathcal{F} \triangleright X \phi]$. To show that $[\langle Y \rangle]_{\mathcal{N}}$ is a post-fixpoint of f , suppose that $u \in [\langle Y \rangle]_{\mathcal{N}}$, i.e., that $\langle Y \rangle \in L(u)$. Since Φ is closed under taking components, $\text{ind}(y, \langle Y \rangle) \in \Phi$, and hence $\langle \Delta Y \{\langle Y \rangle\} \rangle \in \Phi$ whenever $\langle Y \rangle \in \Phi$. By the post-fixpoint axiom Fix and Proposition 5 we have $\text{ind}(y, \langle Y \rangle) \in L(u)$, as well.

Since $L(u)$ is an atom, it follows that $\chi_i \in L(u)$ for each $i = 1, \dots, n$ and $\langle \Delta Y \{\langle Y \rangle\} \rangle \in L(u)$. Then, by the inductive hypothesis, it follows (note that z does not occur free in χ_i) that $\mathcal{M}(\mathcal{N}, u \Vdash_{[z \mapsto [\langle Y \rangle]_{\mathcal{N}}]} \chi_i$ for each $i = 1, \dots, n$. Furthermore, using one-step coherence of the network \mathcal{N} , and since $\langle \Delta Y \{\langle Y \rangle\} \rangle \in L(u)$, we have $\mathcal{M}(\mathcal{N}, u \Vdash_{[z \mapsto [\langle Y \rangle]_{\mathcal{N}}]} \langle \Delta Y \{z\} \rangle$. To see this, $\Delta Y \{\langle Y \rangle\}$ is a next time goal assignment, so one-step coherence guarantees that there is some $\Sigma \in \Pi_{a \in \text{Agt}}$ such that, for every C in the support of $\Delta Y \{\langle Y \rangle\}$ such that $\Delta Y \{\langle Y \rangle\}(C) = X\psi$ and for every strategy profile Σ' with $\Sigma' \sim_C \Sigma$, we have $\psi \in L(\text{out}(\Sigma'), u)$. In particular, this means that for each $C \subseteq \bigcup \mathcal{F}$ in the support of $\Delta Y \{\langle Y \rangle\}$, each $D_i \subseteq C$, and each $\Sigma' \sim_C \Sigma$, we have $\chi_i \in L(\text{out}(\Sigma', u))$, and so by the induction hypothesis $\mathcal{M}(\mathcal{N}, \text{out}(\Sigma', u) \Vdash_{[z \mapsto [\langle Y \rangle]_{\mathcal{N}}]} \chi_i$. Also, for each $\Sigma' \sim_{\bigcup \mathcal{F}} \Sigma$, we have $\langle Y \rangle \in L(\text{out}(\Sigma', u))$, and hence $\mathcal{M}(\mathcal{N}, \text{out}(\Sigma', u) \Vdash_{[z \mapsto [\langle Y \rangle]_{\mathcal{N}}]} z$ by definition of $[\langle Y \rangle]_{\mathcal{N}}$. Putting these facts together we get $\mathcal{M}(\mathcal{N}, u \Vdash_{[z \mapsto [\langle Y \rangle]_{\mathcal{N}}]} \langle \Delta Y \{z\} \rangle$ as claimed. We now get $\mathcal{M}(\mathcal{N}, u \Vdash_{[z \mapsto [[\langle Y \rangle]_{\mathcal{N}}]} \text{ind}(y, z)$, hence $u \in f([\langle Y \rangle]_{\mathcal{N}})$. Thus, $[\langle Y \rangle]_{\mathcal{N}}$ is a post-fixpoint of f .

For the converse inclusion $\nu f \subseteq [\langle Y \rangle]_{\mathcal{N}}$, we reason contrapositively: suppose that v is some node in T that does not belong to $[\langle Y \rangle]_{\mathcal{N}}$. This means that $\langle Y \rangle \notin L(v)$, hence $\neg \langle Y \rangle \in L(v)$ since $L(v)$ is an atom over Φ , and Φ is a closed set of formulas and hence closed under single negations. We will prove that $v \notin \nu f$. Since νf is the intersection of its approximants $f^\xi(T)$ where ξ ranges over ordinals, it suffices to find a *finite* ordinal $k < \omega$ such that $v \notin f^k(T)$. Here we recall that $f^k(T)$ is defined inductively by $f^0(T) = T$, $f^{i+1} = f(f^i(T))$.

Note that the formula $\neg \langle Y \rangle$ is a TypeG-eventuality. So, since \mathcal{N} was assumed to be a perfect network, there is some $k < \omega$ for which this eventuality is partially fulfilled in k steps at v . Hence it suffices to prove, by induction on k , that for all $w \in [\neg \langle Y \rangle]_{\mathcal{N}}$, if the eventuality $\neg \langle Y \rangle$ is fulfilled in k steps at w then $w \notin f^{k+1}(T)$. We refer to the induction hypothesis on k as the *innermost* induction hypothesis. We refer to the induction hypothesis of the structural induction on complexity of formulas as the *outermost* induction hypothesis.

For the base case of the innermost induction, when $k = 0$, let w be some element of $[\neg \langle Y \rangle]_{\mathcal{N}}$ at which the eventuality $\neg \langle Y \rangle$ is partially fulfilled in 0 steps. If $w \in f^1(T) = f(T)$, then $\mathcal{M}(\mathcal{N}, w \Vdash_{[z \mapsto T]} \text{ind}(y, z)$, hence $\mathcal{M}(\mathcal{N}, w \Vdash \chi_i$ for each χ_i . By the outermost induction hypothesis on χ_i , each χ_i is in $L(w)$. Since the eventuality $\neg \langle Y \rangle$ is partially fulfilled in 0 steps at w and $w \in [\neg \langle Y \rangle]_{\mathcal{N}}$, we have $\overline{\chi}_i \in L(w)$ for some χ_i , which contradicts the consistency of $L(w)$. Hence $w \notin f(T)$.

For the induction step, suppose $k = j + 1$ and the innermost induction hypothesis holds for j . Let w be some element of $[\neg \langle Y \rangle]_{\mathcal{N}}$ at which the eventuality $\neg \langle Y \rangle$ is partially fulfilled in k steps. Suppose, for a contradiction, that $w \in f^{k+1}(T) = f(f^k(T)) = f(f^{j+1}(T))$. Then $\mathcal{M}(\mathcal{N}, w \Vdash_{[z \mapsto f^{j+1}(T)]} \text{ind}(y, z)$. Let F^j be the set of all $x \in [\neg \langle Y \rangle]_{\mathcal{N}}$ such that $\neg \langle Y \rangle$ is partially fulfilled in j steps at x . By the innermost induction hypothesis on j , we get $f^{j+1}(T) \cap F^j = \emptyset$. Since $\neg \langle Y \rangle$ is partially fulfilled in $k = j + 1$ steps at w , either $\overline{\chi}_i \in L(w)$ for some χ_i or there exists a marking \mathbf{m} that refutes $\langle \Delta Y \rangle$ at w , and such that for all $w' \in T$ such that w' is a child of w , $w' \in F^j$ whenever $\neg \langle Y \rangle \in \mathbf{m}(w')$. In the former case we immediately get a contradiction. So we focus on the latter case.

Since $\mathcal{M}(\mathcal{N}, w \Vdash_{[z \mapsto f^{j+1}(T)]} \text{ind}(y, z)$, we have $\mathcal{M}(\mathcal{N}, w \Vdash_{[z \mapsto f^{j+1}(T)]} \langle \Delta Y \{z\} \rangle$. Let Σ be a witnessing strategy profile for $\Delta Y \{z\}$ at w . Then, for every $\Sigma' \sim_{\bigcup \mathcal{F}} \Sigma$ we have

$\mathcal{M}(\mathcal{N}), \text{out}(\Sigma', w) \Vdash_{[z \mapsto f^{j+1}(T)]} z$, i.e., $\text{out}(\Sigma', w) \in f^{j+1}(T)$. Furthermore, for each $C \subsetneq \bigcup \mathcal{F}$ in the support of $\Delta\gamma\{z\}$, each $D_i \subseteq C$ and each $\Sigma' \sim_C \Sigma$ we have $\mathcal{M}(\mathcal{N}), \text{out}(\Sigma', w) \Vdash_{[z \mapsto f^{j+1}(T)]} \chi_i$, and so $\chi_i \in L(\text{out}(\Sigma', w))$ by the outermost induction hypothesis on χ_i . On the other hand, since the marking \mathbf{m} refutes $\langle \Delta\gamma \rangle$ at w there must be some C^* in the support of $\Delta\gamma = \Delta\gamma\{\langle \Delta\gamma \rangle\}$ and some $\Sigma' \sim_{C^*} \Sigma$ such that $\bar{\psi} \in \mathbf{m}(\text{out}(\Sigma', w))$, where $\Delta\gamma\{\langle \Delta\gamma \rangle\}(C^*) = X\bar{\psi}$. If $C^* = \bigcup \mathcal{F}$ then $\bar{\psi} = \langle \Delta\gamma \rangle$ so $\bar{\psi} = \neg\langle \Delta\gamma \rangle$, so $\neg\langle \Delta\gamma \rangle \in \mathbf{m}(\text{out}(\Sigma', w))$. Recall that \mathbf{m} was such that $w' \in F^j$ whenever $\neg\langle \Delta\gamma \rangle \in \mathbf{m}(w')$, so we get $\text{out}(\Sigma', w) \in F^j$. But, since $\Sigma' \sim_{\bigcup \mathcal{F}} \Sigma$, we must also have $\text{out}(\Sigma', w) \in f^{j+1}(T)$, which is a contradiction since $f^{j+1}(T) \cap F^j = \emptyset$. On the other hand, if $C^* \subsetneq \bigcup \mathcal{F}$, then, since C^* is in the support of $\Delta\gamma$, there must be some D_i in the support of γ with $D_i \subseteq C^*$, and so χ_i is a conjunct of $\bar{\psi}$ and therefore inconsistent with $\bar{\psi} \in \mathbf{m}(\text{out}(\Sigma', w)) \subseteq L(\text{out}(\Sigma', w))$. But, since $\Sigma' \sim_{C^*} \Sigma$, we get $\chi_i \in L(\text{out}(\Sigma', w))$, which is a contradiction since $L(\text{out}(\Sigma', w))$ is consistent by assumption. So, in either case we get a contradiction, hence we have shown that $w \notin f^{j+1}(T)$, as desired.

Case: γ is in TypeU. The argument is similar to the above, but somewhat more complicated. Let γ be supported by $\mathcal{F} = \{C_1, \dots, C_n, D_1, \dots, D_m\}$ and defined by

$$\gamma(C_1) = \alpha_1 \cup \beta_1, \dots, \gamma(C_n) = \alpha_n \cup \beta_n,$$

and

$$\gamma(D_1) = G\chi_1, \dots, \gamma(D_m) = G\chi_m.$$

Since γ is in TypeU we have $\{C_1, \dots, C_n\} \neq \emptyset$.

We claim that the set

$$\langle \Delta\gamma \rangle_{\mathcal{N}} := \{u \in T \mid \langle \Delta\gamma \rangle \in L(u)\},$$

is the least pre-fixpoint μf of the monotone map f induced by $\text{ind}(\gamma, z)$ in $\mathcal{M}(\mathcal{N})$, defined by

$$f(Z) = \left\{ u \in T \mid \mathcal{M}(\mathcal{N}), u \Vdash_{[z \mapsto Z]} \text{ind}(\gamma, z) \right\}.$$

By Proposition 7, the least pre-fixpoint μf is equivalent to $\llbracket \langle \Delta\gamma \rangle \rrbracket_{\mathcal{M}(\mathcal{N})}$, whence the case will follow.

For the inclusion $\mu f \subseteq \llbracket \langle \Delta\gamma \rangle \rrbracket_{\mathcal{N}}$, we reason by least fixpoint induction. That is, we prove that the set $\llbracket \langle \Delta\gamma \rangle \rrbracket_{\mathcal{N}}$ is a pre-fixpoint of f . Recall that

$$\text{ind}(\gamma, z) = \bigvee_{1 \leq i \leq n} (\beta_i \wedge \langle \Delta\gamma \setminus C_i \rangle) \vee \left(\bigwedge_{1 \leq i \leq n} \alpha_i \wedge \bigwedge_{1 \leq i \leq m} \chi_i \wedge \langle \Delta\gamma\{z\} \rangle \right),$$

where $\Delta\gamma\{\phi\}$ denotes $\Delta\gamma[\bigcup \mathcal{F} \triangleright X\phi]$. To show that $\llbracket \langle \Delta\gamma \rangle \rrbracket_{\mathcal{N}}$ is a pre-fixpoint of f , suppose that $u \in f(\llbracket \langle \Delta\gamma \rangle \rrbracket_{\mathcal{N}})$. We need to show that $u \in \llbracket \langle \Delta\gamma \rangle \rrbracket_{\mathcal{N}}$, i.e., that $\langle \Delta\gamma \rangle \in L(u)$. By definition of f , either there is some $i \in \{1, \dots, n\}$ for which $\mathcal{M}(\mathcal{N}), u \Vdash_{z \mapsto \llbracket \langle \Delta\gamma \rangle \rrbracket_{\mathcal{N}}} \beta_i \wedge \langle \Delta\gamma \setminus C_i \rangle$, or $\mathcal{M}(\mathcal{N}), u \Vdash_{z \mapsto \llbracket \langle \Delta\gamma \rangle \rrbracket_{\mathcal{N}}} \alpha_i$ for each $i \in \{1, \dots, n\}$, $\mathcal{M}(\mathcal{N}), u \Vdash_{z \mapsto \llbracket \langle \Delta\gamma \rangle \rrbracket_{\mathcal{N}}} \chi_i$ for each $i \in \{1, \dots, m\}$, and $\mathcal{M}(\mathcal{N}), u \Vdash_{z \mapsto \llbracket \langle \Delta\gamma \rangle \rrbracket_{\mathcal{N}}} \langle \Delta\gamma\{z\} \rangle$.

In the former case, the induction hypothesis on β_i gives $\beta_i \in L(u)$, and since the support of $\gamma \setminus C_i$ is smaller than that of γ , the induction hypothesis for the induction on the size of the support gives $\langle \Delta\gamma \setminus C_i \rangle \in L(u)$. Since Φ is a closed set, we have $\beta_i \wedge \langle \Delta\gamma \setminus C_i \rangle \in \Phi$, and since $L(u)$ is an atom we get $\beta_i \wedge \langle \Delta\gamma \setminus C_i \rangle \in L(u)$. It follows, again by closure of Φ and $L(u)$ being an atom, that $\langle \Delta\gamma \rangle \in L(u)$ as required.

In the latter case, the induction hypotheses on $\alpha_1, \dots, \alpha_n$ and χ_1, \dots, χ_m ensure that these formulas are all in $L(u)$. We will show that $\langle \Delta\gamma\{\langle \Delta\gamma \rangle\} \rangle \in L(u)$, from which it will follow using the fact that $L(u)$ is an atom and Φ is closed that $\langle \Delta\gamma \rangle \in L(u)$. To prove this, we use that $\mathcal{M}(\mathcal{N}), u \Vdash_{z \mapsto \llbracket \langle \Delta\gamma \rangle \rrbracket_{\mathcal{N}}} \langle \Delta\gamma\{z\} \rangle$. It suffices to show that the marking L does not refute $\langle \Delta\gamma\{\langle \Delta\gamma \rangle\} \rangle$, because by one-step

coherence it follows that $\neg\langle\Delta\gamma\{\langle\gamma\rangle\}\rangle \notin L(u)$ and hence $\langle\Delta\gamma\{\langle\gamma\rangle\}\rangle \in L(u)$. So suppose, for a contradiction, that L refutes $\langle\Delta\gamma\{\langle\gamma\rangle\}\rangle$. Since $\mathcal{M}(\mathcal{N}), u \Vdash_{z \mapsto [\langle\gamma\rangle]_{\mathcal{N}}} \langle\Delta\gamma\{z\}\rangle$, there is a witnessing strategy profile Σ such that $\mathcal{M}(\mathcal{N}), \text{out}(\Sigma', u) \Vdash_{z \mapsto [\langle\gamma\rangle]_{\mathcal{N}}} z$, and hence $\langle\gamma\rangle \in L(\text{out}(\Sigma', u))$, for each coalition $\Sigma' \sim_{\cup \mathcal{F}} \Sigma$, and $\mathcal{M}(\mathcal{N}), \text{out}(\Sigma', u) \Vdash_{z \mapsto [\langle\gamma\rangle]_{\mathcal{N}}} \psi$ for each $E \subsetneq \cup \mathcal{F}$ in the support of $\Delta\gamma\{z\}$ with $\Delta\gamma\{z\}(E) = X\psi$. On the other hand, since L refutes $\langle\Delta\gamma\{\langle\gamma\rangle\}\rangle$, there must be some E^* in the support of $\langle\Delta\gamma\{\langle\gamma\rangle\}\rangle$ (which is the same as the support of $\Delta\gamma\{z\}$), and some $\Sigma' \sim_{E^*} \Sigma$ such that $\psi \notin \text{out}(\Sigma', u)$ where $\Delta\gamma\{\langle\gamma\rangle\}(E^*) = X\psi$. If $E^* = \cup \mathcal{F}$, then $\psi = \langle\gamma\rangle$, so $\langle\gamma\rangle \notin L(\text{out}(\Sigma', u))$. But we already know that $\langle\gamma\rangle \in L(\text{out}(\Sigma', u))$ for each $\Sigma' \sim_{\cup \mathcal{F}} \Sigma$, so this is a contradiction. On the other hand, suppose that $E^* \subsetneq \cup \mathcal{F}$. Then $\psi = \langle\gamma|_{E^*}\rangle$, so $\langle\gamma|_{E^*}\rangle \notin L(\text{out}(\Sigma', u))$. But the support of $\gamma|_{E^*}$ is smaller than that of γ since $E^* \subsetneq \cup \mathcal{F}$, so the induction hypothesis for the induction on the size of supports applies, and we get $\mathcal{M}(\mathcal{N}), \text{out}(\Sigma', u) \not\Vdash_{z \mapsto [\langle\gamma\rangle]_{\mathcal{N}}} \langle\gamma|_{E^*}\rangle$. But, since $\Delta\gamma\{\langle\gamma\rangle\}(E^*) = X\langle\gamma|_{E^*}\rangle$ and $\Sigma' \sim_{E^*} \Sigma$, we have $\mathcal{M}(\mathcal{N}), \text{out}(\Sigma', u) \Vdash_{z \mapsto [\langle\gamma\rangle]_{\mathcal{N}}} \langle\gamma|_{E^*}\rangle$, so we get a contradiction again. Thus, we have shown that $u \in [\langle\gamma\rangle]_{\mathcal{N}}$, as required. Therefore, $\mu f \subseteq [\langle\gamma\rangle]_{\mathcal{N}}$.

Now for the converse inclusion, $[\langle\gamma\rangle]_{\mathcal{N}} \subseteq \mu f$, suppose $u \in [\langle\gamma\rangle]_{\mathcal{N}}$, i.e., $\langle\gamma\rangle \in L(u)$. Note that $\langle\gamma\rangle$ is a TypeU-eventuality. So, since \mathcal{N} is perfect, there is some $k < \omega$ such that the eventuality is partially fulfilled in k steps at u . Since the least fixpoint μf is the union of its ordinal approximants $f^\xi(\emptyset)$, where ξ ranges over ordinals, it suffices to show by induction on k that for all $w \in \langle\gamma\rangle$, if the eventuality $\langle\gamma\rangle$ is fulfilled in k steps at w then $w \in f^{k+1}(\emptyset)$. We have several nested inductions at this point, so we refer to the induction hypothesis for the structural induction on complexity of formulas as the *outermost induction hypothesis*, that of the induction on the size of supports as the *middle induction hypothesis*, and the induction hypothesis on k as the *innermost induction hypothesis*.

For the base case, where $k = 0$, if $\langle\gamma\rangle$ is fulfilled in 0 steps at w then there is some C_i with $\beta_i \wedge \langle\gamma \setminus C_i\rangle \in L(u)$. The outermost induction hypothesis on β_i and the middle induction hypothesis on $\langle\gamma \setminus C_i\rangle$, together with closure of Φ and $L(w)$ being an atom, immediately give $\mathcal{M}(\mathcal{N}), w \Vdash \beta_i \wedge \langle\gamma \setminus C_i\rangle$, and therefore $\mathcal{M}(\mathcal{N}), w \Vdash_{[z \mapsto \emptyset]} \text{ind}(\gamma, z)$, since the variable z does not appear in β_i or $\langle\gamma \setminus C_i\rangle$. By definition we get $w \in f(\emptyset) = f^1(\emptyset)$.

Now let $w \in [\langle\gamma\rangle]_{\mathcal{N}}$, and suppose $\langle\gamma\rangle$ is partially fulfilled in $k = j + 1$ steps at w where the induction hypothesis holds for j . If $\langle\gamma\rangle$ is partially fulfilled in j steps at w then the innermost induction hypothesis applies and we are done. Otherwise, since $\langle\gamma\rangle$ is partially fulfilled in $k = j + 1$ steps at w , we have $\alpha_i \in L(w)$ for all $i \in \{1, \dots, n\}$, and $\chi_j \in L(w)$ for all $j \in \{1, \dots, m\}$, and there exists a marking \mathbf{m} that verifies $\langle\Delta\gamma\{\langle\gamma\rangle\}\rangle$ at w and is such that for every child w' of w with $\langle\gamma\rangle \in \mathbf{m}(w')$ the eventuality $\langle\gamma\rangle$ is partially fulfilled in j steps at w' . The outermost induction hypothesis gives $\mathcal{M}(\mathcal{N}), w \Vdash \alpha_i$ for all $i \in \{1, \dots, n\}$, and $\mathcal{M}(\mathcal{N}), w \Vdash \chi_j$ for all $j \in \{1, \dots, m\}$. Since z does not appear in any of these formulas we get $\mathcal{M}(\mathcal{N}), w \Vdash_{[z \mapsto f^j(\emptyset)]} \alpha_i$ for all $i \in \{1, \dots, n\}$, and $\mathcal{M}(\mathcal{N}), w \Vdash_{[z \mapsto f^j(\emptyset)]} \chi_j$ for all $j \in \{1, \dots, m\}$.

Let Σ be some strategy profile witnessing that the marking \mathbf{m} verifies $\langle\Delta\gamma\{\langle\gamma\rangle\}\rangle$ at w . Let E be some coalition in the support of $\Delta\gamma\{z\}$, which is the same as the support of $\Delta\gamma\{\langle\gamma\rangle\}$. If $E = \cup \mathcal{F}$, let $\Sigma' \sim_{\cup \mathcal{F}} \Sigma$. Then $\langle\gamma\rangle \in \mathbf{m}(\text{out}(\Sigma', w))$, so $\langle\gamma\rangle$ is partially fulfilled in j steps at w . By the innermost induction hypothesis, $\text{out}(\Sigma', w) \in f^j(\emptyset)$. So $\mathcal{M}(\mathcal{N}), \text{out}(\Sigma', w) \Vdash_{z \mapsto f^j(\emptyset)} z$. On the other hand, if $E \subsetneq \cup \mathcal{F}$, then $\Delta\gamma\{z\}(E)$ is of the form $X\langle\gamma|_E\rangle$. So, if $\Sigma' \sim_E \Sigma$ then $\langle\gamma|_E\rangle \in \mathbf{m}(\text{out}(\Sigma', w)) \subseteq L(\text{out}(\Sigma', w))$. Since the support of $\gamma|_E$ is smaller than that of γ , the middle induction hypothesis applies and we get $\mathcal{M}(\mathcal{N}), \text{out}(\Sigma', w) \Vdash \langle\gamma|_E\rangle$. Since the variable z does not appear in $\langle\gamma|_E\rangle$ we get $\mathcal{M}(\mathcal{N}), \text{out}(\Sigma', w) \Vdash_{[z \mapsto f^j(\emptyset)]} \langle\gamma|_E\rangle$. So, we get $\mathcal{M}(\mathcal{N}), w \Vdash_{z \mapsto f^j(\emptyset)} \langle\Delta\gamma\{z\}\rangle$. Collecting all the facts we have established, we get $\mathcal{M}(\mathcal{N}), w \Vdash_{[z \mapsto f^j(\emptyset)]} \text{ind}(\gamma, z)$, so $w \in f(f^j(\emptyset)) = f^{j+1}(\emptyset) = f^k(\emptyset)$ as required.

We have thus shown that $[\langle\gamma\rangle]_{\mathcal{N}} \subseteq \mu f$, and the proof is completed. \square

A.5 Proof of Theorem 2

PROOF. Structural induction on φ . We prove further, in Corollary 2, that every TLCGA-formula is equivalent to one in normal form, so we can assume that φ is in a normal form. The Boolean cases are routine, so we only consider the case for $\varphi = \langle\!\langle\gamma\rangle\!\rangle$, where γ is either a next time assignment, or a long-term temporal assignment in TypeU or in TypeG.

(1) We first consider the case of next time assignments. Let γ be a next time assignment defined by $\gamma(C_1) = X\phi_1, \dots, \gamma(C_n) = X\phi_n$. Suppose that $\mathcal{M}, s_1 \Vdash \langle\!\langle\gamma\rangle\!\rangle$, witnessed by a joint action ζ^1 for Agt at s_1 . Let ζ^2 be some joint action for Agt at s_2 witnessing the Forth condition with respect to ζ^1 . We need to show that $\text{Out}[s_2, \zeta_{C_i}^2] \subseteq \llbracket \phi_i \rrbracket$ for each $i \in \{1, \dots, k\}$. Suppose $v \in \text{Out}[s_2, \zeta_{C_i}^2]$. Apply the LocalBack condition to find $v' \in \text{Out}[s_1, \zeta_{C_i}^1]$ with $v' \beta v$. Since $\text{Out}[s_1, \zeta_{C_i}^1] \subseteq \llbracket \phi_i \rrbracket$ we get $\mathcal{M}, v \Vdash \phi_i$ by the induction hypothesis on ϕ_i , as required. Thus, $\mathcal{M}, s_2 \Vdash \langle\!\langle\gamma\rangle\!\rangle$.

The converse direction is proved in the same way.

(2) Next, we claim that for any goal assignment γ , if the bisimulation invariance claim holds for $\langle\!\langle\gamma\rangle\!\rangle$ and for all proper subformulae⁹ of $\langle\!\langle\gamma\rangle\!\rangle$ and all pairs of bisimilar states s_1, s_2 in the state space S of \mathcal{M} , then it also holds likewise for $\langle\!\langle\Delta\gamma\rangle\!\rangle$, because this is a special case of a next time assignment, involving only $\langle\!\langle\gamma\rangle\!\rangle$ and subformulae of $\langle\!\langle\gamma\rangle\!\rangle$. This claim we will use further, when proving the inductive steps for long-term assignments.

(3) Let $\gamma \in \text{TypeG}$ and assume that the bisimulation invariance claim holds for all proper subformulae of $\langle\!\langle\gamma\rangle\!\rangle$ and all pairs of bisimilar states s_1, s_2 in the state space S of \mathcal{M} . We will prove the bisimulation invariance of $\langle\!\langle\gamma\rangle\!\rangle$ for all such pairs of states by using the fixpoint characterisation in Proposition 7 and proving the claim for $vz.\text{ind}(\gamma, z)$, instead.

To prove that claim we hereafter treat $\text{ind}(\gamma, Z)$ as a (monotone) operator on sets of states and use the characterisation of greatest fixed points given by Knaster–Tarski theorem, according to which $vz.\text{ind}(\gamma, z) = \bigcap_{\alpha \in \text{Ord}} \text{ind}(\gamma, Z^\alpha)$, where the family $\{Z^\alpha\}_{\alpha \in \text{Ord}}$ of subsets of S is defined by transfinite induction, as usual:

$$Z^0 = S; Z^{\alpha+1} = \text{ind}(\gamma, Z^\alpha); Z^\lambda = \bigcap_{\alpha < \lambda} \text{ind}(\gamma, Z^\alpha) \text{ for limit ordinals } \lambda.$$

For technical convenience we will treat each Z^α both as a set of states and as a formula for which this set is its extension in \mathcal{M} , noting that bisimulation invariance of a formula ϕ in a model \mathcal{M} is equivalent to the closure under bisimulation of its extension $\llbracket \phi \rrbracket_{\mathcal{M}}$.

It suffices to prove bisimulation invariance (respectively, bisimulation closure) of each approximant formula Z^α , as the closure under bisimulation is preserved in the intersection of any family of sets. We prove these closures by a nested transfinite induction on α . The only non-trivial case is that of successor ordinals.

Recall that, for a goal assignment $\gamma \in \text{TypeG}$, we have $\text{ind}(\gamma, \phi) = \bigwedge \text{GHolds}(\gamma) \wedge \langle\!\langle\Delta\gamma\{\phi\}\!\rangle\!\rangle$.

Thus, $Z^{\alpha+1} = \text{ind}(\gamma, Z^\alpha) = \bigwedge \text{GHolds}(\gamma) \wedge \langle\!\langle\Delta\gamma\{Z^\alpha\}\!\rangle\!\rangle$. Since each formula in $\text{GHolds}(\gamma)$ is a proper subformula of $\langle\!\langle\gamma\rangle\!\rangle$, its bisimulation invariance follows from the inductive hypothesis of the main induction, so it only remains to show the bisimulation invariance of $\langle\!\langle\Delta\gamma\{Z^\alpha\}\!\rangle\!\rangle$, assuming the bisimulation invariance of $\Delta\gamma\{Z^\alpha\}$. This claim is a particular case of the claim 2 for next time extensions of goal assignments, proved above.

This completes the inductive step for $\langle\!\langle\gamma\rangle\!\rangle$ with $\gamma \in \text{TypeG}$.

(4) Lastly, the inductive step for $\langle\!\langle\gamma\rangle\!\rangle$ with $\gamma \in \text{TypeU}$ is analogous, with the respective changes:

- Using the fixpoint characterisation in Proposition 6 we prove the claim for $\mu z.\text{ind}(\gamma, z)$, instead.

⁹Note that the proper subformulae of $\langle\!\langle\gamma\rangle\!\rangle$ are all formulae assigned as goals by γ and their subformulae.

- To prove that claim we use the characterisation of least fixed points given by Knaster–Tarski theorem, according to which $\mu z.\text{ind}(\gamma, z) = \bigcup_{\alpha \in \text{Ord}} \text{ind}(\gamma, Z^\alpha)$, where the family $\{Z^\alpha\}_{\alpha \in \text{Ord}}$ is defined by transfinite induction, as usual: $Z^0 = \emptyset$; $Z^{\alpha+1} = \text{ind}(\gamma, Z^\alpha)$; $Z^\lambda = \bigcup_{\alpha < \lambda} \text{ind}(\gamma, Z^\alpha)$ for limit ordinals λ .
- Again, it suffices to prove the bisimulation invariance (respectively, bisimulation closure) of each approximant formula Z^α , which we do by a nested transfinite induction on α , with the only non-trivial case being that of successor ordinals, for which we use the definition of $\text{ind}(\gamma, \phi)$ for $\gamma \in \text{TypeU}$ and the inductive hypotheses for Z^α , and the case of next time extensions of goal assignments. The differences from the previous case of $\gamma \in \text{TypeG}$, coming from the additional subformulae in $\text{Finish}(\gamma)$ and $\text{UHold}(\gamma)$, are inessential.

This completes the structural induction and the whole proof. \square

A.6 Full Proof of Proposition 16

PROOF. By Proposition 15, we may assume w.l.o.g. that the defect (u, φ) is such that u is a leaf: if we can show how to remove the defect φ at a single leaf, then, clearly, we can repeat the procedure to remove φ at each leaf in the set $\{v_1, \dots, v_k\}$. (Note that our procedure for removing a defect at a single leaf v given below will not affect any other leaves, i.e., each leaf in the original network besides v will still be a leaf in the new network.) Combined with Proposition 15 this proves the result.

So, suppose that (u, φ) is a defect and u is a leaf. It is sufficient to show that there is a finite, one-step coherent network \mathcal{N}'' in which the root has the same label as u in \mathcal{N} , and in which the eventuality φ is partially fulfilled. We can then simply identify the root of the network \mathcal{N}' with the leaf u in \mathcal{N} to form a finite, one-step coherent network \mathcal{N}' such that $\mathcal{N}'' \sqsubseteq \mathcal{N}'$ and $\mathcal{N} \sqsubseteq \mathcal{N}'$. By Proposition 12, the eventuality φ is partially fulfilled at u in \mathcal{N}' .

Consider the Φ -atoms Ψ such that $\varphi \in \Psi$ and there exists a finite, one-step coherent network in which the root is labelled by Ψ and the eventuality φ is partially fulfilled. Let δ be the disjunction of all conjunctions of the form $\bigwedge \Psi$ for all such Φ -atoms Ψ . (This is well-defined since the set of all such conjunctions is finite, as long as we disallow conjunctions with multiple occurrences of the same conjunct.) The result then follows from the following claim:

CLAIM 3. $\vdash \varphi \rightarrow \delta$.

To prove the claim, we consider the two cases for the eventuality φ .

Case: φ is a TypeU eventuality. Then φ is of the form $\langle \gamma \rangle$, where γ is a goal assignment supported by a set of coalitions \mathcal{F} for which $\gamma(C)$ is an U-formula for at least one $C \in \mathcal{F}$. Say that $\mathcal{F} = \{C_1, \dots, C_n, D_1, \dots, D_m\}$ and γ is defined by

$$\gamma(C_1) = \alpha_1 \cup \beta_1, \dots, \gamma(C_n) = \alpha_n \cup \beta_n,$$

and (if $m > 0$)

$$\gamma(D_1) = G\chi_1, \dots, \gamma(D_m) = G\chi_m.$$

Our aim is to prove:

$$\vdash \text{ind}(\gamma, \delta) \rightarrow \delta,$$

and thereafter apply the induction rule to conclude the Claim. It suffices to show that, if Ψ is any atom that is consistent with $\text{ind}(\gamma, \delta)$, then $\bigwedge \Psi$ is, in fact, one of the disjuncts of δ . Indeed, suppose that $\not\vdash \text{ind}(\gamma, \delta) \rightarrow \delta$; by Lindenbaum's Lemma there exists a maximal consistent set of formulae Γ containing $\text{ind}(\gamma, \delta)$ but also the negation of δ . Then $\Gamma \cap \Phi$ is a Φ -atom that is consistent with $\text{ind}(\gamma, \delta)$, but $\bigwedge(\Gamma \cap \Phi)$ cannot be a disjunct of δ since that would make Γ inconsistent.

So, suppose Ψ is a Φ -atom for which the set $\Psi \cup \{\text{ind}(\gamma, \delta)\}$ is consistent. Then Ψ is consistent with at least one of the disjuncts of $\text{ind}(\gamma, \delta)$.

If Ψ is consistent with one of the disjuncts $\beta_i \wedge \langle\!\langle \gamma \mid C_i \rangle\!\rangle$ then, in fact, this disjunct must be a member of Ψ , since Ψ is a Φ -atom and each formula $\beta_i \wedge \langle\!\langle \gamma \mid C_i \rangle\!\rangle$ belongs to Φ . In this case, it is trivial to construct a (singleton) network in which $\langle\!\langle \gamma \rangle\!\rangle$ is partially fulfilled in 0 steps.

If Ψ is consistent with $\alpha_1 \wedge \dots \wedge \alpha_n \wedge \chi_1 \wedge \dots \wedge \chi_m \wedge \langle\!\langle \Delta\gamma \mid \delta \rangle\!\rangle$, then $\alpha_1, \dots, \alpha_n, \chi_1, \dots, \chi_m \in \Psi$ and the set $\Psi^+ := \Psi \cup \{\langle\!\langle \Delta\gamma \mid \delta \rangle\!\rangle\}$ is consistent. Let Φ^+ be the extended Fischer-Ladner closure of the set $\Phi \cup \{\delta\}$, and let Θ be the maximal modal one-step theory that is contained in Ψ^+ . Then, since $\Theta \subseteq \Psi^+$, it is consistent by our assumption, and it is a one-step theory over Φ^+ . By one-step completeness (Theorem 4), there exists a finite maximal consistent game form $\mathcal{M}(\Theta) = (\text{Act}, \text{act}, \mathcal{P}(\Phi^+), \text{out})$ for Φ^+ such that, for every goal assignment γ' :

- (1) If $\langle\!\langle \gamma' \rangle\!\rangle \in \Theta$, then there is an action profile $\zeta \in \Pi_{a \in \text{Agt}} \text{act}_a$ such that for all C in the support of γ' , and all $\zeta' \sim_C \zeta$, we have $\psi \in \text{out}(\zeta')$, where $\gamma'(C) = X\psi$.
- (2) If $\neg\langle\!\langle \gamma' \rangle\!\rangle \in \Theta$, then for every profile $\zeta \in \Pi_{a \in \text{Agt}} \text{act}_a$ there is some C in the support of γ' , and some $\zeta' \sim_C \zeta$, for which $\bar{\psi} \in \text{out}(\zeta')$, where $\gamma'(C) = X\psi$.

Since $\langle\!\langle \Delta\gamma \mid \delta \rangle\!\rangle \in \Theta$, the first clause ensures that there is a profile $\rho \in \Pi_{a \in \text{Agt}} \text{act}_a$ such that for all C in the support of $\Delta\gamma[\bigcup \mathcal{F} \triangleright X\delta]$, and all $\zeta' \sim_C \rho$, we have $\psi \in \text{out}(\zeta')$, where $\Delta\gamma\{\delta\}(C) = X\psi$.

In particular, for $C = \bigcup \mathcal{F}$ this entails that $\delta \in \text{out}(\zeta')$ for all $\zeta' \sim_{\bigcup \mathcal{F}} \rho$.

We now construct a network showing that $\bigwedge \Psi$ is a disjunct of δ as follows. For each action profile $\zeta \in \Pi_{a \in \text{Agt}} \text{act}_a$, we pick a Φ -network \mathcal{N}_ζ according to the following rule.

- if $\zeta \sim_{\bigcup \mathcal{F}} \rho$, then let \mathcal{N}_ζ be a network of which the root is labelled $\text{out}(\zeta) \cap \Phi$, and in which the eventuality $\langle\!\langle \gamma \rangle\!\rangle$ is partially fulfilled at the root. Such a network exists since, if $\zeta \sim_{\bigcup \mathcal{F}} \rho$ then $\delta \in \text{out}(\zeta)$, hence $\bigwedge(\text{out}(\zeta) \cap \Phi)$ is a disjunct of δ .
- Otherwise, let \mathcal{N}_ζ be a network consisting of a single node labelled $\text{out}(\zeta) \cap \Phi$.

We form the network $\mathcal{N} = (T, L, \mathcal{G})$ by taking the disjoint union of the networks \mathcal{N}_ζ , for each action profile $\zeta \in \Pi_{a \in \text{Agt}} \text{act}_a$, together with a new root r labelled Ψ , with an edge to the root of each \mathcal{N}_ζ , and letting $\mathcal{G}(r) := (\text{Act}, \text{act}, T, \text{out}')$ where out' maps each action profile ζ to the root of \mathcal{N}_ζ . Note that $\mathcal{N}_\zeta \sqsubseteq \mathcal{N}$ for each ζ , and that for each profile ζ we have $L(\text{out}'(\zeta)) = \text{out}(\zeta) \cap \Phi$.

Since $\Pi_{a \in \text{Agt}} \text{act}_a$ is a finite set, each \mathcal{N}_ζ is one-step coherent, and due to the clauses (1) and (2) of Theorem 4, the network \mathcal{N} is a finite and one-step coherent network. Moreover, to see that the eventuality $\langle\!\langle \gamma \rangle\!\rangle$ is partially fulfilled at the root r , we define a marking \mathbf{m} by setting

$$\mathbf{m}(v) := \left\{ \theta \in \Phi \mid \text{there exists } C \subseteq \bigcup \mathcal{F} \text{ such that } \Delta\gamma(C) = X\theta \text{ and } \zeta \sim_C \rho \right\},$$

if v is the root of the network \mathcal{N}_ζ for some $\zeta \in \Pi_{a \in \text{Agt}} \text{act}_a$, and $\mathbf{m}(v) := \emptyset$ otherwise. By definition, this marking verifies $\Delta\gamma$, witnessed by the action profile ρ . Furthermore, if $\langle\!\langle \gamma \rangle\!\rangle \in \mathbf{m}(v)$ then, since $\Delta\gamma(C) \neq X\langle\!\langle \gamma \rangle\!\rangle$ for all $C \neq \bigcup \mathcal{F}$, v must be the root of some network \mathcal{N}_ζ where $\zeta \sim_{\bigcup \mathcal{F}} \rho$. This means that $\langle\!\langle \gamma \rangle\!\rangle$ is partially fulfilled at v in \mathcal{N}_ζ , hence in \mathcal{N} , since $\mathcal{N}_\zeta \sqsubseteq \mathcal{N}$. It follows that $\langle\!\langle \gamma \rangle\!\rangle$ is partially fulfilled in at most $k + 1$ steps at r in \mathcal{N} , where k is the maximum number such that $\langle\!\langle \gamma \rangle\!\rangle$ is partially fulfilled in at most k steps at the root of one of the finitely many networks \mathcal{N}_ζ .

Case: φ is a TypeG eventuality. Then φ is of the form $\neg\langle\!\langle \gamma \rangle\!\rangle$, where γ is a goal assignment supported by a set of coalitions \mathcal{F} and $\gamma(C)$ is an G-formula for all $C \in \mathcal{F}$. Say that $\mathcal{F} = \{D_1, \dots, D_m\}$ and γ is defined by $\gamma(D_1) = G\chi_1, \dots, \gamma(D_m) = G\chi_m$. Our aim is to prove:

$$\vdash \neg\delta \rightarrow \text{ind}(\gamma, \neg\delta),$$

and thereafter apply the co-induction rule to conclude that $\vdash \neg\delta \rightarrow \langle\!\langle \gamma \rangle\!\rangle$, hence $\vdash \neg\langle\!\langle \gamma \rangle\!\rangle \rightarrow \delta$, as required. It suffices to show that, if Ψ is any atom that is consistent with $\neg\text{ind}(\gamma, \neg\delta)$, then $\bigwedge \Psi$ is,

in fact, one of the disjuncts of δ . Indeed, suppose that $\not\models \neg\delta \rightarrow \text{ind}(\gamma, \neg\delta)$. Then $\not\models \neg\text{ind}(\gamma, \neg\delta) \rightarrow \delta$ so by Lindenbaum's Lemma there exists a maximal consistent set of formulae Γ containing $\neg\text{ind}(\gamma, \neg\delta)$ but also the negation of δ . Then $\Gamma \cap \Phi$ is a Φ -atom that is consistent with $\neg\text{ind}(\gamma, \neg\delta)$, but $\bigwedge(\Gamma \cap \Phi)$ cannot be a disjunct of δ since that would make Γ inconsistent.

We recall that $\text{ind}(\gamma, \neg\delta)$ is the formula $\chi_1 \wedge \dots \wedge \chi_m \wedge \llbracket \Delta\gamma[\bigcup \mathcal{F} \triangleright X \neg\delta] \rrbracket$, so $\neg\text{ind}(\gamma, \neg\delta)$ is provably equivalent to $\overline{\chi}_1 \vee \dots \vee \overline{\chi}_m \vee \neg\llbracket \Delta\gamma[\bigcup \mathcal{F} \triangleright X \neg\delta] \rrbracket$. So, suppose Ψ is a Φ -atom for which the set $\Psi \cup \{\neg\text{ind}(\gamma, \neg\delta)\}$ is consistent. Then Ψ is consistent with at least one of the disjuncts of $\neg\text{ind}(\gamma, \neg\delta)$. If Ψ is consistent with one of the disjuncts $\overline{\chi}_i$ then, in fact, this disjunct must be a member of Ψ , since Ψ was a Φ -atom and each formula $\overline{\chi}_i$ belongs to Φ . In this case, it is trivial to construct a (singleton) network in which $\neg\llbracket \gamma \rrbracket$ is partially fulfilled in 0 steps. If Ψ is consistent with $\neg\llbracket \Delta\gamma[\bigcup \mathcal{F} \triangleright X \neg\delta] \rrbracket$, then the set $\Psi^+ := \Psi \cup \{\neg\llbracket \Delta\gamma[\bigcup \mathcal{F} \triangleright X \neg\delta] \rrbracket\}$ is consistent. Let Φ^+ be the extended Fischer-Ladner closure of the set $\Phi \cup \{\neg\delta\}$, and let Θ be the set of all one-step formulae over Φ^+ belonging to Ψ^+ . Then, since $\Theta \subseteq \Psi^+$, it is consistent by our assumption, and it is a one-step theory over Φ^+ . By one-step completeness (Theorem 4), there exists a finite maximal consistent game form $\mathcal{M}(\Theta) = (\text{Act}, \text{act}, \mathcal{P}(\Phi^+), \text{out})$ such that, for every goal assignment γ' :

- (1) If $\llbracket \gamma' \rrbracket \in \Theta$, then there is a profile $\zeta \in \Pi_{a \in \text{Agt}} \text{act}_a$ such that for all C in the support of γ' , and all $\zeta' \sim_C \zeta$, we have $\psi \in \text{out}(\zeta')$, where $\gamma'(C) = X\psi$.
- (2) If $\neg\llbracket \gamma' \rrbracket \in \Theta$, then for every profile $\zeta \in \Pi_{a \in \text{Agt}} \text{act}_a$ there is some C in the support of γ' , and some $\zeta' \sim_C \zeta$, for which we have $\overline{\psi} \in \text{out}(\zeta')$, where $\gamma'(C) = X\psi$.

Since $\neg\llbracket \Delta\gamma[\bigcup \mathcal{F} \triangleright X \neg\delta] \rrbracket \in \Theta$, the second clause ensures that for every profile $\rho \in \Pi_{a \in \text{Agt}} \text{act}_a$ there is some C in the support of $\Delta\gamma[\bigcup \mathcal{F} \triangleright X \neg\delta]$, and some $\zeta' \sim_C \rho$, for which we have $\overline{\psi} \in \text{out}(\zeta')$, where $\gamma'(C) = X\psi$. We pick a choice function selecting such a pair $(c(\rho), f(\rho))$ for each ρ , where $c(\rho) \in \mathcal{F}$, $\rho \sim_{c(\rho)} f(\rho)$ and $\overline{\psi} \in \text{out}(f(\rho))$, where $\gamma(c(\rho)) = X\psi$.

We now construct a network showing that $\bigwedge \Psi$ is a disjunct of δ as follows. For each action profile $\zeta \in \Pi_{a \in \text{Agt}} \text{act}_a$, we pick a Φ -network \mathcal{N}_ζ according to the following rule.

- if $\zeta = f(\rho)$ for some ρ such that $c(\rho) = \bigcup \mathcal{F}$, then let \mathcal{N}_ζ be a network of which the root is labelled $\text{out}(\zeta) \cap \Phi$, and in which the eventuality $\neg\llbracket \gamma \rrbracket$ is partially fulfilled at the root. Such a network exists since, if $\zeta = f(\rho)$ for some ρ then $\gamma(\bigcup \mathcal{F}) = X \neg\delta$ and so $\overline{\delta} = \delta \in \text{out}(\zeta)$, hence $\bigwedge(\text{out}(\zeta) \cap \Phi)$ is a disjunct of δ .
- Otherwise, let \mathcal{N}_ζ be a network consisting of a single node labelled $\text{out}(\zeta) \cap \Phi$.

We form the network $\mathcal{N} = (T, L, \mathcal{G})$ by taking the disjoint union of the networks \mathcal{N}_ζ , for each action profile $\zeta \in \Pi_{a \in \text{Agt}} \text{act}_a$, together with a new root r labelled Ψ with an edge to the root of each \mathcal{N}_ζ , and letting $\mathcal{G}(r) = (\text{Act}, \text{act}, T, \text{out}')$ where out' maps each action profile ζ to the root of \mathcal{N}_ζ . Note that $\mathcal{N}_\zeta \sqsubseteq \mathcal{N}$ for each ζ , and that for each profile ζ we have $L(\text{out}'(\zeta)) = \text{out}(\zeta) \cap \Phi$. Since $\Pi_{a \in \text{Agt}} \text{act}_a$ is a finite set, each \mathcal{N}_ζ is one-step coherent, and due to the clauses (1) and (2) of Theorem 4, the network \mathcal{N} is a finite and one-step coherent network. Moreover, to see that the eventuality $\neg\llbracket \gamma \rrbracket$ is partially fulfilled at the root r , we define a marking m by setting

$$\mathbf{m}(v) := \left\{ \overline{\theta} \in \Phi \mid \text{there exists } \rho \text{ such that } \Delta\gamma(c(\rho)) = X\theta \text{ and } f(\rho) = \zeta \right\},$$

if v is the root of the network \mathcal{N}_ζ for some $\zeta \in \Pi_{a \in \text{Agt}} \text{act}_a$, and $\mathbf{m}(v) = \emptyset$ otherwise. By definition, this marking refutes $\Delta\gamma$. Furthermore, if $\neg\llbracket \gamma \rrbracket \in \mathbf{m}(v)$ then, since $\Delta\gamma(C) \neq X\llbracket \gamma \rrbracket$ for all $C \neq \bigcup \mathcal{F}$, there must be some ρ such that $c(\rho) = \bigcup \mathcal{F}$, and v is the root of $\mathcal{N}_{f(\rho)}$. This means that $\neg\llbracket \gamma \rrbracket$ is partially fulfilled at v in $\mathcal{N}_{f(\rho)}$, hence in \mathcal{N} , since $\mathcal{N}_{f(\rho)} \sqsubseteq \mathcal{N}$. It follows that $\neg\llbracket \gamma \rrbracket$ is partially fulfilled in at most $k + 1$ steps at r in \mathcal{N} , where k is the maximum number such that $\neg\llbracket \gamma \rrbracket$ is partially fulfilled in at most k steps at the root of one of the finitely many networks \mathcal{N}_ζ . \square

A.7 Proof of Proposition 17

PROOF. Repeated use of Propositions 14 and 16, as follows. First, we fix an enumeration $\mathcal{L} = l_1, l_2, \dots, l_k$ of all leaves in \mathcal{N} and an enumeration $\mathcal{D} = (u_1, \phi_1), \dots, (u_d, \phi_d)$ of all defects in \mathcal{N} . Then, we construct a finite chain of finite one-step coherent networks $\mathcal{N}_0 \sqsubseteq \mathcal{N}_1 \sqsubseteq \dots$, inductively as follows.

We begin with $\mathcal{N}_0 = \mathcal{N}$. Suppose, we have constructed the finite one-step coherent Φ -networks $\mathcal{N}_0 \sqsubseteq \mathcal{N}_1 \sqsubseteq \dots \mathcal{N}_n$ and let \mathcal{L}_n be an enumeration of all leaves of \mathcal{N} that are still in \mathcal{N}_n and \mathcal{D}_n be an enumeration of all defects in \mathcal{N} that are still in \mathcal{N}_n . Then we do the following.

(1) If \mathcal{D}_n is non-empty, we pick the first defect (u, ϕ) listed in it and apply Proposition 16 to construct a finite, one-step coherent network \mathcal{N}_{n+1} such that $\mathcal{N}_n \sqsubseteq \mathcal{N}_{n+1}$ and (u, ϕ) is not a defect in \mathcal{N}_{n+1} . Then we update the list \mathcal{D}_n to \mathcal{D}_{n+1} by removing (u, ϕ) and all other defects in that list that may have been resolved in \mathcal{N}_{n+1} . Note that, by Proposition 12, newly occurring defects in \mathcal{N}_{n+1} may only occur at newly added nodes. We update likewise the list \mathcal{L}_n to \mathcal{L}_{n+1} by removing the leaves in \mathcal{N}_n that are no longer leaves in \mathcal{N}_{n+1} .

(2) If \mathcal{D}_n is empty, but \mathcal{L}_n is non-empty, then we pick the first leaf l currently listed in \mathcal{L}_n and apply Proposition 14 to construct a finite, one-step coherent network \mathcal{N}_{n+1} such that $\mathcal{N}_n \sqsubseteq \mathcal{N}_{n+1}$ and l is not a leaf in \mathcal{N}_{n+1} . Then we update the list \mathcal{L}_n to \mathcal{L}_{n+1} by removing l and all other leaves in \mathcal{N}_n that are no longer leaves in \mathcal{N}_{n+1} . Note that, again by Proposition 12, newly occurring defects in \mathcal{N}_{n+1} may only occur at newly added nodes, so the list of original defects in \mathcal{N} remains empty.

(3) If both \mathcal{D}_n and \mathcal{L}_n are empty, the last constructed network \mathcal{N}_n is the desired \mathcal{N}' .

Note that both lists \mathcal{L} and \mathcal{D} are finite, neither of them gets extended in any step, and at least one of them strictly decreases in each step. Therefore, they will both become empty in finitely many steps, hence the construction is guaranteed to terminate. Clearly, once the procedure terminates, we obtain a network none of whose leaves are leaves of \mathcal{N} and such that none of its defects are defects in \mathcal{N} . \square

A.8 Proof of Proposition 19

PROOF. For the direction left to right, suppose $\mathcal{G} = (\text{Act}, \text{act}, \mathcal{P}(V), \text{out})$ is some game form that \mathcal{S} -satisfies Γ . By assumption, for each action profile ζ there is some $Z \in \mathcal{S}$ with $\text{out}(\zeta) \subseteq Z$. We show that \mathcal{S} satisfies both conditions (1) and (2).

For (1), let $R = (C_1, \gamma_1, \dots, C_n, \gamma_n)$ be a redistribution of Γ . Since $\gamma_1, \dots, \gamma_n$ are positive relevant goal assignments, there must be strategy profiles ζ_1, \dots, ζ_n such that for each ζ_i , coalition C and $\zeta'_i \sim_C \zeta_i$, we have $p \in \text{out}(\zeta'_i)$ where $\gamma_i(C) = Xp$. Define the new strategy profile ζ^* by setting $\zeta_a^* = (\zeta_i)_a$ if $a \in C_i$, and ζ_a^* arbitrarily chosen if $a \notin C_1 \cup \dots \cup C_n$. It is easy to see that $F(R) \subseteq \text{out}(\zeta^*)$, hence $F(R) \subseteq Z$ for some $Z \in \mathcal{S}$.

For (2), let $R = (C_1, \gamma_1, \dots, C_n, \gamma_n)$ be a redistribution of Γ and let γ' be a negative relevant goal assignment. Let ζ_1, \dots, ζ_n and ζ^* be defined as before, so that $F(R) \subseteq \text{out}(\zeta^*)$. Since \mathcal{G} \mathcal{S} -satisfies Γ , there must be some coalition C and some $\zeta'_i \sim_C \zeta^*$ such that $q \in \text{out}(\zeta'_i)$ where $\gamma'(C) = Xq$. If $C = \text{Agt}$ then q must belong to the outcome of every action profile, hence to every member of \mathcal{S} . If $C \neq \text{Agt}$, then we get $F(R, \gamma', C) \subseteq \text{out}(\zeta'_i) \subseteq Z$ for some $Z \in \mathcal{S}$.

For right to left, suppose \mathcal{S} satisfies the constraints stated in the proposition. Enumerate Agt as $a_0, \dots, a_{|\text{Agt}|-1}$. Let $\mathcal{G} = (\text{Act}, \text{act}, \mathcal{P}(V), \text{out})$ be defined as follows. First, given an agent $a \in \text{Agt}$, we define the set of actions act_a to be the set of triples (γ, f, k) where

- γ is either a positive relevant goal assignment, or a distinguished symbol $*$,
- $k \in \{0, \dots, |\text{Agt}| - 1\}$,
- f is a map sending each redistribution R of Γ to some element Z of \mathcal{S} with $F(R) \subseteq Z$.

Note that each player has the same set of actions, and the set of actions is non-empty since, by assumption, for each redistribution R of Γ there is some element Z of \mathcal{S} with $F(R) \subseteq Z$.

Given an action profile ζ we say that **a coalition C has formed behind a positive relevant goal assignment** if that goal assignment was chosen by every member of C in ζ . A given action profile ζ thus partitions Agt into disjoint coalitions $\{C_1, \dots, C_n\}$ where each coalition C_i is a maximal one that has formed behind some positive relevant goal assignment (or the special symbol $*$). Given an action profile ζ with corresponding partition of the form $\{C_1, \dots, C_n\}$, where each coalition C_i has formed behind the goal assignment γ_i , or possibly of the form $\{C_1, \dots, C_n, D\}$, where each coalition C_i has formed behind the goal assignment γ_i and each player in D has played an action with $*$ as first component, we define the redistribution $\text{red}(\zeta)$ induced by ζ to be $R = (C_1, \gamma_1, \dots, C_n, \gamma_n)$. Suppose each $a \in \text{Agt}$ has chosen the action (γ_a, Z_a, k_a) . Then let

$$\text{bet}(\zeta) = \sum_{a \in \text{Agt}} k_a \bmod |\text{Agt}|.$$

We now define the outcome function by setting $\text{out}(\zeta) = f(\text{red}(\zeta))$, where f is the function chosen by player a_i for $i = \text{bet}(\zeta)$. Informally speaking, each agent picks a goal assignment that they want to act towards, and also votes on a function intended to pick the actual outcome of an action profile depending on which coalitions have formed in favor of which goal assignments. To decide which voter gets to determine the outcome of an action profile, each player also has to bet on a number k , and the winner of this bet is the agent indexed by $\sum_{a \in \text{Agt}} k_a \bmod |\text{Agt}|$.

We now show that \mathcal{G} \mathcal{S} -satisfies Γ .

It clearly holds that $\text{out}(\zeta) \in \mathcal{S}$ for every action profile ζ and, conversely, for each $Z \in \mathcal{S}$, Z is the outcome of the profile in which each player chooses the action $(*, Z, 0)$. Now suppose $\langle \gamma \rangle \in \Gamma$ for some positive goal assignment γ . We need to find an action profile ζ such that for every coalition C , and every $\zeta' \sim_C \zeta$, we have $p \in \text{out}(\zeta')$ where $\gamma(C) = Xp$. But this is easy: the action profile ζ defined by $\zeta_a = (\gamma, Z, k)$ for all $a \in \text{Agt}$, with Z, k arbitrarily chosen, clearly does the job.

On the other hand, suppose $\neg \langle \gamma' \rangle \in \Gamma$ for some negative goal assignment γ' . Let ζ be an arbitrary action profile in \mathcal{G} . We want to find some coalition B and some $\zeta' \sim_B \zeta$ such that $q \in \text{out}^{\Gamma}(\zeta')$, where $\gamma'(B) = X\neg q$. Let $R = (C_1, \gamma_1, \dots, C_n, \gamma_n)$ be the redistribution of Γ defined to consist of those C_i, γ_i such that C_i has formed in ζ behind the goal assignment γ_i , and each $C_i \subseteq C'$. By the second condition on \mathcal{S} , there is some coalition C' such that:

- (i) either $C' = \text{Agt}$ and $q \in Z$ for every $Z \in \mathcal{S}$, where $\gamma'(\text{Agt}) = X\neg q$, or
- (ii) $C' \neq \text{Agt}$ and there is some $Z \in \mathcal{S}$ such that $F(R, \gamma', C') \subseteq Z$.

In the first case, we have $\zeta \sim_{\text{Agt}} \zeta$ and $q \in \text{out}(\zeta) \in \mathcal{S}$. In the second case, $C' \neq \text{Agt}$ and there is some $Z \in \mathcal{S}$ such that $F(R, \gamma', C') \subseteq Z$. Let $R' = \{C'_1, \gamma'_1, \dots, C'_m, \gamma'_m\}$ be the redistribution defined by $\{C'_1, \dots, C'_m\} = \{C_i \cap C' \mid i \in \{1, \dots, n\} \ \& \ C_i \subseteq C'\}$, and let γ'_j be γ_i for $C'_j = C_i \cap C'$. This is well-defined since $\{C_1, \dots, C_n\}$ consists of pairwise disjoint sets. Furthermore, note that $F(R') = Z$, since:

$$\begin{aligned} F(R') &= \{p \mid \exists i, B \subseteq C'_i : \gamma'_i(C'_i) = Xp\} \\ &\subseteq \{p \mid \exists i, B \subseteq C_i \cap C' : \gamma_i(B) = Xp\} \\ &\subseteq \{p \mid \exists i, B \subseteq C_i \cap C' : \gamma_i(B) = Xp\} \cup \{q\} \\ &= F(R, \gamma', C') \\ &\subseteq Z. \end{aligned}$$

Define a new action profile ζ' as follows: pick some function f with $f(R') = Z$, and pick some distinguished player $c \in \text{Agt} \setminus C'$. For each $a \in C'$, set $\zeta'_a = \zeta_a$. For $a \in \text{Agt} \setminus C'$, set $\zeta'_a = (*, f, 0)$ if $a \neq c$. Finally set $\zeta'_c = (*, f, k)$, where k is chosen to ensure that $\text{bet}(\zeta')$ is the index of player c . Clearly, $\zeta' \sim_{C'} \zeta$, and it remains only to show $q \in \text{out}(\zeta')$. But $q \in Z$, and since $R' = \text{red}(\zeta')$ we get $\text{out}(\zeta') = f(R') = Z$. Informally, what happens here is that the players in $\text{Agt} \setminus C'$ pick one member among them, c , that will do the actual work towards ensuring that the outcome of ζ' contains q . So that player picks Z containing q as the outcome, and the other players in $\text{Agt} \setminus C'$ lay down their bets by betting on 0 so that c can make sure to win the bet, thus ensuring that the outcome is Z . \square

ACKNOWLEDGMENTS

We thank the referees for the corrections and critical comments which contributed to essential improvements in the content of the article.

REFERENCES

- [1] Ittai Abraham, Danny Dolev, Rica Gonen, and Joseph Y. Halpern. 2006. Distributed computing meets game theory: Robust mechanisms for rational secret sharing and multiparty computation. In *Proceedings of the 25th Annual ACM Symposium on Principles of Distributed Computing*. Eric Ruppert and Dahlia Malkhi (Eds.). ACM, 53–62. DOI: <https://doi.org/10.1145/1146381.1146393>
- [2] Erman Acar, Massimo Benerecetti, and Fabio Mogavero. 2019. Satisfiability in strategy logic can be easier than model checking. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence, AAAI 2019, 31st Innovative Applications of Artificial Intelligence Conference, IAAI 2019 and 9th AAAI Symposium on Educational Advances in Artificial Intelligence*. AAAI Press, 2638–2645. DOI: <https://doi.org/10.1609/aaai.v33i01.33012638>
- [3] Th. Ågotnes, V. Goranko, and W. Jamroga. 2007. Alternating-time temporal logics with irrevocable strategies. In *Proceedings of the TARK XI*. Presses Universitaires de Louvain, Univ. Saint-Louis, Brussels, 15–24.
- [4] T. Ågotnes, V. Goranko, and W. Jamroga. 2008. *Strategic Commitment and Release in Logics for Multi-Agent Systems (Extended Abstract)*. Technical Report IfI-08-01. Clausthal University of Technology.
- [5] R. Alur, T. A. Henzinger, and O. Kupferman. 2002. Alternating-time temporal logic. *Journal of the ACM* 49, 5 (2002), 672–713.
- [6] Benjamin Aminof, Vadim Malvone, Aniello Murano, and Sasha Rubin. 2016. Graded strategy logic: Reasoning about uniqueness of nash equilibria. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*. Catholijn M. Jonker, Stacy Marsella, John Thangarajah, and Karl Tuyls (Eds.). ACM, 698–706. Retrieved from <http://dl.acm.org/citation.cfm?id=2937027>.
- [7] R. J. Aumann. 1959. Acceptable points in general cooperative n -person games. In *Proceedings of the Contributions to the Theory of Games, volume IV*. A. W. Tucker and L. D. Luce (Eds.). Princeton University Press, 287–324.
- [8] N. Belnap and M. Perloff. 1988. Seeing to it that: A canonical form for agentives. *Theoria* 54, 3 (1988), 175–199.
- [9] Rodica Branzei, Dinko Dimitrov, and Stef Tijs. 2008. *Models in Cooperative Game Theory* (2nd ed.). Springer-Verlag, Berlin.
- [10] Thomas Brihaye, Arnaud Da Costa Lopes, François Laroussinie, and Nicolas Markey. 2009. ATL with strategy contexts and bounded memory. In *Proceedings of the LFC'S'09*. S. Artëmov and A. Nerode (Eds.). Springer, Berlin, 92–106.
- [11] M. A. Brown. 1988. On the logic of ability. *Journal of Philosophical Logic* 17, 1 (1988), 1–26.
- [12] N. Bulling, V. Goranko, and W. Jamroga. 2015. Logics for reasoning about strategic abilities in multi-player games. In *Proceedings of the Models of Strategic Reasoning: Logics, Games, and Communities*. J. van Benthem, S. Ghosh, and R. Verbrugge (Eds.). Springer, Berlin, 93–136.
- [13] Georgios Chalkiadakis, Edith Elkind, and Michael J. Wooldridge. 2011. *Computational Aspects of Cooperative Game Theory*. Morgan & Claypool Publishers. DOI: <https://doi.org/10.2200/S00355ED1V01Y201107AIM016>
- [14] Krishnendu Chatterjee, Thomas A. Henzinger, and Nir Piterman. 2010. Strategy logic. *Information and Computation* 208, 6 (2010), 677–693. DOI: <https://doi.org/10.1016/j.ic.2009.07.004>
- [15] C. Cirstea, C. Kupke, and D. Pattinson. 2011. Exptime tableaux for the coalgebraic μ -calculus. *Logical Methods in Computer Science* 7, 3 (2011), 1–33.
- [16] Corina Cirstea and Dirk Pattinson. 2007. Modular construction of complete coalgebraic logics. *Theoretical Computer Science* 388, 1–3 (2007), 83–108.
- [17] Stéphane Demri, Valentin Goranko, and Martin Lange. 2016. *Temporal Logics in Computer Science*. Cambridge University Press, Cambridge.

- [18] Sebastian Enqvist. 2018. Flat modal fixpoint logics with the converse modality. *Journal of Logic and Computation* 28, 6 (2018), 1065–1097.
- [19] Sebastian Enqvist, Fatemeh Seifan, and Yde Venema. 2019. Completeness for μ -calculi: A coalgebraic approach. *Annals of Pure and Applied Logic* 170, 5 (2019), 578–641.
- [20] Dana Fisman, Orna Kupferman, and Yoav Lustig. 2010. Rational synthesis. In *Proceedings of the TACAS 2010*. Javier Esparza and Rupak Majumdar (Eds.). Lecture Notes in Computer Science, Vol. 6015. Springer, Berlin, 190–204.
- [21] Gaëlle Fontaine, Raul Leal, and Yde Venema. 2010. Automata for coalgebras: An approach using predicate liftings. In *Proceedings of the International Colloquium on Automata, Languages, and Programming*. Springer, Berlin, 381–392.
- [22] Patrick Gardy, Patricia Bouyer, and Nicolas Markey. 2020. Dependences in strategy logic. *Theory of Computing Systems* 64, 3 (2020), 467–507. DOI: <https://doi.org/10.1007/s00224-019-09926-y>
- [23] Valentin Goranko and Sebastian Enqvist. 2018. Socially friendly and group protecting coalition logics. In *Proceedings of the AAMAS 2018*. 372–380. Retrieved from <http://dl.acm.org/citation.cfm?id=3237442>.
- [24] Valentin Goranko and Fengkui Ju. 2019. Towards a logic for conditional local strategic reasoning. In *Proceedings of the LORI 2019*. P. Blackburn and E. Lorini (Eds.). Lecture Notes in Computer Science, Vol. 11813. Springer, 112–125.
- [25] V. Goranko and G. van Drimmelen. 2006. Complete axiomatization and decidability of alternating-time temporal logic. *Theory of Computing Systems* 353 (2006), 93–117.
- [26] Julian Gutierrez, Sarit Kraus, and Michael J. Wooldridge. 2019. Cooperative concurrent games. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. Edith Elkind, Manuela Veloso, Noa Agmon, and Matthew E. Taylor (Eds.). International Foundation for Autonomous Agents and Multiagent Systems, 1198–1206. Retrieved from <http://dl.acm.org/citation.cfm?id=3331822>.
- [27] S. MacLane. 1971. *Categories for the Working Mathematician*. Number 5 in Graduate text in Math. Springer-Verlag, Berlin.
- [28] Fabio Mogavero, Aniello Murano, Giuseppe Perelli, and Moshe Y. Vardi. 2014. Reasoning about strategies: On the model-checking problem. *ACM Transactions on Computational Logic* 15, 4 (2014), 34:1–34:47.
- [29] Fabio Mogavero, Aniello Murano, Giuseppe Perelli, and Moshe Y. Vardi. 2017. Reasoning about strategies: On the satisfiability problem. *Logical Methods in Computer Science* 13, 1 (2017), 1–37.
- [30] Fabio Mogavero, Aniello Murano, and Luigi Sauro. 2013. On the boundary of behavioral strategies. In *Proceedings of the 28th Annual ACM/IEEE Symposium on Logic in Computer Science*. IEEE Computer Society, 263–272. DOI: <https://doi.org/10.1109/LICS.2013.32>
- [31] Fabio Mogavero, Aniello Murano, and Moshe Y. Vardi. 2010. Reasoning about strategies. In *Proceedings of the FSTTCS 2010*. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 133–144.
- [32] Lawrence S. Moss. 1999. Coalgebraic logic. *Annals of Pure and Applied Logic* 96, 1–3 (1999), 277–317.
- [33] Martin Osborne and Ariel Rubinstein. 2011. *A Course in Game Theory*. MIT Press, electronic edition.
- [34] Rohit Parikh. 1985. The logic of games and its applications. In *Topics in the Theory of Computation, Selected Papers of the International Conference on 'Foundations of Computation Theory' (FCT'83)*, Marek Karplinski and Jan van Leeuwen (Eds.).
- [35] Dirk Pattinson. 2003. Coalgebraic modal logic: Soundness, completeness and decidability of local consequence. *Theoretical Computer Science* 309, 1–3 (2003), 177–193.
- [36] M. Pauly. 2001. *Logic for Social Software*. Ph.D. Dissertation. University of Amsterdam.
- [37] M. Pauly. 2002. A modal logic for coalitional power in games. *Journal of Logic and Computation* 12, 1 (2002), 149–166.
- [38] K. G. Ramamurthy. 1990. *Coherent Structures and Simple Games*. Theory and Decision Library, Vol. 6. Springer Netherlands.
- [39] Jan J. M. M. Rutten. 2000. Universal coalgebra: A theory of systems. *Theoretical Computer Science* 249, 1 (2000), 3–80.
- [40] Luigi Santocane and Yde Venema. 2010. Completeness for flat modal fixpoint logics. *Annals of Pure and Applied Logic* 162, 1 (2010), 55–82.
- [41] Lutz Schröder. 2008. Expressivity of coalgebraic modal logic: The limits and beyond. *Theoretical Computer Science* 390, 2 (2008), 230.
- [42] Lutz Schröder and Dirk Pattinson. 2009. PSPACE bounds for rank-1 modal logics. *ACM Transactions on Computational Logic* 10, 2 (2009), 1–33.
- [43] Lutz Schröder and Yde Venema. 2010. Flat coalgebraic fixed point logics. In *Proceedings of the CONCUR*, Paul Gastin and François Laroussinie (Eds.). Lecture Notes in Computer Science, Vol. 6269. Springer, Berlin, 524–538.
- [44] Lutz Schröder and Yde Venema. 2010. Flat coalgebraic fixed point logics. In *Proceedings of the International Conference on Concurrency Theory*. Springer, Berlin, 524–538.
- [45] Nicolas Troquard, Wiebe van der Hoek, and Michael Wooldridge. 2009. A logic of games and propositional control. In *Proceedings of the AAMAS 2009, Volume 2*. IFAAMAS, 961–968.
- [46] Ad M. A. van Deemen. 1997. *Coalition Formation and Social Choice*. Theory and Decision Library, Vol. 19. Springer Netherlands.

- [47] Wiebe van der Hoek and Michael Wooldridge. 2005. On the logic of cooperation and propositional control. *Artificial Intelligence* 164, 1–2 (2005), 81–119. DOI : <https://doi.org/10.1016/j.artint.2005.01.003>
- [48] Yde Venema. 2006. Automata and fixed point logic: A coalgebraic perspective. *Information and Computation* 204, 4 (2006), 637–678.
- [49] Michael Wooldridge, Julian Gutierrez, Paul Harrenstein, Enrico Marchioni, Giuseppe Perelli, and Alexis Toumi. 2016. Rational verification: From model checking to equilibrium checking. In *Proceedings of the AAAI 2016*. 4184–4191. Retrieved from <http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12268>.

Received 29 December 2020; revised 17 December 2021; accepted 4 February 2022