# Socially Friendly and Group Protecting Coalition Logics

Valentin Goranko[1,2] and Sebastian Enqvist[1]
[1]Stockholm University, Sweden
[2]University of Johannesburg, South Africa (visiting professorship)
valentin.goranko@philosophy.su.se, thesebastianenqvist@gmail.com

## ABSTRACT

We consider extensions of Coalition Logic (CL) which can express statements about inter-related powers of coalitions to achieve their respective goals. In particular, we introduce and study two new extensions of CL. One of them is the "Socially Friendly Coalition Logic" SFCL, which is also a multi-agent extension of the recently introduced "Instantial Neighborhood Logic" INL. SFCL can express the claim that a coalition has a collective strategy to guarantee achieving its explicitly stated goal while acting in a 'socially friendly way', by enabling the remaining agents to achieve other (again, explicitly stated) goals of their choice. The other new extension is the "Group Protecting Coalition Logic" GPCL which enables reasoning about entire coalitional goal assignments, in which every group of agents has its own specified goal. GPCL can express claims to the effect that there is an action profile of the grand coalition such that, by playing it, every sub-coalition of agents can guarantee satisfaction of its own private goal (and thus, protect its own interests) while acting towards achievement of the common goal of the grand coalition. For each of these logics, we discuss its expressiveness, introduce the respective notion of bisimulation and prove bisimulation invariance and Hennessy-Milner property. We then also present sound and complete axiomatic systems and prove decidability for both logics.

## KEYWORDS

coalition logics; multi-agent game models; multi-goal coalitional operators; bisimulations; axiomatic systems; decidability

**ACM Reference Format:**
Valentin Goranko[1,2] and Sebastian Enqvist[1]. 2018. Socially Friendly and Group Protecting Coalition Logics. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 9 pages.

## 1 INTRODUCTION

The Coalition Logic CL introduced and studied in [11, 12] can formalise reasoning about strategic abilities of coalitions of agents to guarantee the achievement of designated objectives regardless of the actions of the remaining agents. More precisely, CL features strategic operators of the type $[C]$, for any group ('coalition') of agents $C$ and, for any formula $\phi$, regarded as expressing the coalitional objective of $C$, $[C]\phi$ intuitively says that the coalition $C$ has a collective action $\sigma_C$ that guarantees the satisfaction of $\phi$ in every outcome state that can occur when the agents in $C$ execute their actions in $\sigma_C$, regardless of the choice of actions of the agents that are not in $C$. Thus, CL can reason about unconditional powers of agents and coalitions to act unilaterally in pursuit of their goals. In this respect CL takes a somewhat one-sided perspective: the agents in some coalition $C$ are viewed as acting *in cooperation* with each other but *in opposition* to all agents outside of the coalition. If one wants to express multi-agent specifications that involve more complex patterns of cooperation vs. opposition, more versatile languages are required. Here we focus on the following two ideas:

**Social friendliness:** Agents can *achieve private goals while leaving room for cooperation with the other agents.*

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

In this paper we propose some extensions of CL with more expressive coalitional operators, expressing abilities of agents and coalitions to guarantee achievement of their group objective while protecting their individual subgroup rights or enabling our agents to achieve their goals, too. Two of these, capturing the two ideas above, we study in detail and provide technical characterisations, incl. complete axiomatisations, for them.

One of these extensions is a variation of CL involving a construction of the type $[C](\phi; \psi_1, \ldots, \psi_k)$ which intuitively says that the coalition $C$ has a collective action that not only guarantees the satisfaction of their objective $\phi$ but also makes it possible that the other agents cooperate to achieve any one of the listed objectives $\psi_1, \ldots, \psi_k$. Thus, this extension of CL enables reasoning that considers the achievement of personal, individual or collective, objectives in a more socially engaged way. The new construction above is also a multi-agent extension and variation of the recently introduced and studied in [15] Instantial Neighbourhood Logic INL. While technically related, the present work takes a different perspective on the meaning and use of that neighbourhood modality.

The other new extension of CL, proposed and studied here, has a different motivation, though in the same vein. It involves new operators of the type $[C_1 \triangleright \phi_1, ..., C_n \triangleright \phi_n]$ intuitively saying that there is a strategy profile $\pi$ for the 'grand coalition' $C = C_1 \cup \ldots \cup C_n$ such that for each $i$, the restriction of $\pi$ to the group of agents $C_i$ is a collective strategy of $C_i$ that enforces the objective $\phi_i$. The intuition is that each group $C_i$ participates in the grand coalition with a collective strategy that, while contributing to the achievement of the common goal, also guarantees the protection of the group interest of $C_i$ expressed by $\phi_i$. The operator above is naturally extended to a full *coalitional goal assignment* that assigns a common goal to *every* coalition (subset of the grand coalition).

Besides the introduction of the two extensions of CL described above, the main technical contributions of this paper are:

- definitions of respective bisimulations and proofs of invariance and Hennessy-Milner property for each of them.
- sound and complete axiomatic systems and proofs of decidability via finite tree-model property for both logics.

*Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.* 

In addition to being directly related to [11], [12] and [15], [14], the present work bears both conceptual and technical connections with the work on ATL with irrevocable strategies [1], [2], strategy contexts [4], coalitional logics of cooperation and propositional control [13, 16], and especially with Strategy logic introduced and studied in [9], [8] and other related works. Indeed, both types of operators introduced here can be translated to Strategy Logic, in a way similar to the standard translation of (even quite expressive) modal logics to first-order logic. However, such translations would result in conceptual, technical and computational overkill. We also point out the strong potential of the logics proposed here for adequate formalisation and treatment of the problems of *rational synthesis* [5] and *rational verification* [17]. These and other connections with already published works will be explored in a follow-up publication.

The paper is organised as follows: after preliminaries on CL and INL in Section 2 we mention several new types of coalitional multi-goal operators in Section 3 and then introduce the Socially Friendly Coalition Logic SFCL and the Group Protecting Coalition Logic GPCL. In Section 4 we define respective bisimulations and prove bisimulation invariance and Hennessy-Milner property for each of them. Then, in Section 5 we provide sound and complete axiomatic systems and prove decidability via finite model property for both logics. We end with brief concluding remarks in Section 6.

# 2 PRELIMINARIES

## 2.1 Multi-agent game models

We fix a finite set of **agents** $\mathrm{Agt} = \{a_1, ..., a_n\}$ and a set of **atomic propositions** AP. Subsets of Agt will also be called **coalitions**.

*Definition 2.1 (Multi-agent game model).* A **game model**[1] for Agt and AP is a tuple

$$\mathcal{M} = (S, \{\Sigma_a\}_{a \in \mathrm{Agt}}, g, V)$$

where $S$ is a non-empty set of **states**; each $\Sigma_a$ is a non-empty set of possible **actions** of agent $a$; $V : \mathrm{AP} \to \mathcal{P}(S)$ is a **valuation** of the atomic propositions from AP in $S$; and $g$ is a **game map** that assigns to each $s \in S$ a strategic game form $g(s) = (\Sigma^s_{a_1}, ....\Sigma^s_{a_n}, o_s)$, where each $\Sigma^s_{a_i} \subseteq \Sigma_{a_i}$ is a non-empty set of actions available to player $a_i$ at $s$, and

$$o_s : \Sigma^s_{a_1} \times ... \times \Sigma^s_{a_n} \to S$$

is a **local outcome function** assigning to any **action profile** $\sigma \in \Sigma^s_{a_1} \times ... \times \Sigma^s_{a_n}$ (we call such $\sigma$ **available at** $s$) the **outcome state** $o_s(\sigma)$ produced by $\sigma$ when applied at $s \in S$.

Now, the **global outcome function** in $\mathcal{M}$ is the partial mapping

$$O : S \times \Sigma_{a_1} \times ... \times \Sigma_{a_n} \dashrightarrow S$$

defined by $O(s, \sigma) = o_s(\sigma)$, whenever $\sigma$ is available at $s$.

Given a coalition $C \subseteq \mathrm{Agt}$, a **joint action** for $C$ in the model $\mathcal{M}$ is a tuple of individual actions $\sigma_C \in \prod_{a \in C} \Sigma_a$. For any such joint action and $s \in S$ such that $\sigma_C$ is available at $s$, we define:

$$O[s, \sigma_C] = \{u \in S \mid \exists \sigma \in \Sigma_{a_1} \times ... \times \Sigma_{a_n} : \sigma|_C = \sigma_C \ \&\ o_s(\sigma) = u\}$$

where $\sigma|_C$ is the restriction of $\sigma$ to $C$.

[1] These game models are essentially equivalent to concurrent game models used in [3].

A **rooted game model** is a pair $(\mathcal{M}, s)$ where $\mathcal{M}$ is a game model and $s$ is a state in it, called the **root**.

We will need the following *grafting construction* on rooted game models: Let $\mathfrak{M} = \{(\mathcal{M}_1, s_1), \ldots (\mathcal{M}_m, s_m)\}$ be a (finite) family of (usually pairwise disjoint) rooted game models (which can be assumed to have the same sets of actions), $s$ is a new state that does not belong to any of them, and $g(s) = (\Sigma^s_{a_1}, ....\Sigma^s_{a_n}, o_s)$ is a strategic game form with outcome set $\{s_1, \ldots, s_m\}$. Then the **grafting of** $\mathfrak{M}$ **at** $s$ **via** $g(s)$ is the rooted game model $(\mathcal{G}(\mathfrak{M}, g), s)$ defined with state space being the disjoint union of the state spaces of the models in $\mathfrak{M}$ plus the root state $s$, and sets of actions, game map, and valuation, being the unions of the respective components, where the game map is extended at $s$ by $g(s)$. To save space we leave out the predictable precise technical details.

Later we will be particularly interested in *finite tree-like models*. The class of rooted tree-like models $\mathfrak{T}$ can be defined as the smallest family of rooted models that contains all rooted singleton game models and is closed under the operation of grafting defined above. We omit the formal inductive definition.

## 2.2 Coalition Logic

The formulae of Coalition Logic CL ([11, 12]) are given by the following grammar.

$$\phi ::= p \mid \neg\phi \mid \phi \vee \phi \mid [C]\phi \qquad (p \in \mathrm{AP})$$

Standard definitions of the other propositional connectives apply.

The semantics of CL is defined in terms of **truth of a formula at a state** $s$ **of a game model** $\mathcal{M}$ inductively, by extending the standard clauses for $\neg$ and $\wedge$ with the clause for $[C]$:

- $\mathcal{M}, s \models [C]\phi$ iff there exists a joint action $\sigma_C$ available at $s$, such that $\mathcal{M}, u \models \phi$ for each $u \in O[s, \sigma_C]$.

Now, for every game model $\mathcal{M}$ and a CL formula $\phi$ we can define the **extension of** $\phi$ **in** $\mathcal{M}$ as the set of states $[\![\phi]\!]_{\mathcal{M}}$ in $\mathcal{M}$ where $\phi$ is true. Then the last truth clause above can be re-stated as $[\![[C]\phi]\!]_{\mathcal{M}} = \{s \in S \mid \exists \sigma_C \in \prod_{a \in C} \Sigma_a : O[s, \sigma_C] \subseteq [\![\phi]\!]_{\mathcal{M}}\}$.

A representation theorem for abstract game models and a complete axiomatic system for CL were provided in [11, 12] (later slightly corrected in [6]), and the satisfiability problem was proved decidable and PSPACE-complete. The axiomatic system $Ax_{\mathrm{CL}}$ for CL proposed in [11, 12] extends the classical propositional logic with the following axioms:

(CL1) Agt-**Maximality:** $\neg[\emptyset]\neg\phi \to [\mathrm{Agt}]\phi$
(CL2) **Safety:** $\neg[C]\bot$
(CL3) **Liveness:** $[C]\top$
(CL4) **Superadditivity:** $([C_1]\phi_1 \wedge [C_2]\phi_2) \to [C_1 \cup C_2](\phi_1 \wedge \phi_2)$, for any $C_1, C_2 \subseteq \mathrm{Agt}$ such that $C_1 \cap C_2 = \emptyset$.

and inference rules: **Modus Ponens (MP)** and **Monotonicity**:

$$\frac{\phi \to \psi}{[C]\phi \to [C]\psi}.$$

## 2.3 Instantial Neighbourhood Logic INL

Instantial Neighborhood Logic INL, introduced and investigated in [15], is a modal logic that extends the expressive power of standard neighborhood semantics. In (monotone) **neighborhood semantics**, a model is a structure $(W, R, V)$ where $R \subseteq W \times \mathcal{P}(W)$ is a relation between *points* and *sets of points*. The standard box modality is interpreted in a neighbourhood model $\mathcal{M}$ by the clause:

$\mathcal{M}, w \models \Box\phi$ iff there exists a neighbourhood $Z$ of $w$ such that $\mathcal{M}, u \models \phi$ for each $u \in Z$.

The semantics can be equivalently stated in terms of extensions: $w \in [\![\Box\phi]\!]_{\mathcal{M}}$ iff there exists $Z \subseteq W$ such that $wRZ$ & $Z \subseteq [\![\phi]\!]_{\mathcal{M}}$.

Formulas in the extended logic INL are defined as follows:

$$\phi ::= p \mid \neg\phi \mid \phi \vee \phi \mid \Box(\phi; \psi_1, ..., \psi_n) \qquad (p \in \mathsf{AP})$$

where $(\psi_1, ..., \psi_n)$ is some tuple of formulas, of length $n$. Formally, there is one $n+1$-ary modal operator for each $n \in \omega$, with the case of $n = 0$ reducing to the standard neighbourhood modality $\Box\phi$.

The formulas of INL are evaluated by the following clause for the modal operators:

$w \in [\![\Box(\phi; \psi_1, ..., \psi_n)]\!]_{\mathcal{M}}$ if and only if there exists $Z \subseteq W$ such that $wRZ$, $Z \subseteq [\![\phi]\!]_{\mathcal{M}}$, and $Z \cap [\![\psi_i]\!]_{\mathcal{M}} \neq \emptyset$ for all $i = 1, \ldots, n$.

Thus, while the standard box operator encodes a universal quantification in a neighbourhood of a point in the model, the generalised INL-box operator expresses a combination of both universal and existential claims in such a neighbourhood.

A number of technical results about INL were established in [15], incl. invariance and the Hennessy-Milner property for a suitable notion of bisimulation, strongly complete axiomatisation and PSPACE-completeness of the satisfiability problem.

# 3 COALITION LOGICS WITH MULTI-GOAL OPERATORS

## 3.1 Some new coalitional multi-goal operators

Here is a selection of some natural coalitional multi-goal operators that essentially transcend the expressiveness of CL.

SF **Socially friendly coalitional operator SF**
$[C](\phi; \psi_1, \ldots, \psi_k)$, meaning:
"*C has a collective action $\sigma_C$ that guarantees $\phi$ and enables the complementary coalition $\overline{C}$ to realise any one of $\psi_1, \ldots, \psi_k$ by a suitable collective action*".

SF1 $[C; C_1, \ldots, C_n](\phi; \phi_1, \ldots, \phi_n)$, meaning:
"*C has a collective action $\sigma_C$ that guarantees $\phi$, and is such that, when fixed, each $C_i$ has a collective action that guarantees $\phi_i$*".
This is a refinement of SF, presuming that if $C$ intersects with $C_i$, then the agents in $C \cap C_i$ are already committed to $\sigma_C$.

SF2 $[\langle C_1\rangle\phi_1; \ldots; \langle C_k\rangle\phi_k]$ meaning:
"*$C_1$ has a collective action to guarantee $\phi_1$, and given that action ... $C_k$ has a collective action to guarantee $\phi_k$*".
This is a sequential version of SF1, where the coalitions $C_1, ..., C_k$ are arranged in decreasing priority order.

GIP **Group-interests-protecting coalitional operator GIP**
$\langle\!\langle C_1 \triangleright \phi_1, ..., C_n \triangleright \phi_n \rangle\!\rangle$, meaning:
"*There is an action profile $\sigma$ for the coalition $C_1 \cup \ldots \cup C_n$ such that for each i, the restriction of $\sigma$ to the coalition $C_i$ is an action profile that forces $\phi_i$*".
This operator can be expressed more generally and succinctly as follows. Let us define a **coalitional goal assignment** to be a mapping $\gamma : \mathcal{P}(\mathsf{Agt}) \to \Phi$, where $\Phi$ is the set of formulae of a given 'language of goals' (which may, but need not be, the full logical language under consideration). Thus, for every coalition $C$, $\gamma(C)$ expresses the goal of $C$. In reality, most of the possible coalitions do not form and do not have goals of their own, which can be formalised by assigning the 'truth' $\top$ as a goal to each of them. Now, the coalitional operator defined above naturally generalises simply to $\langle\!\langle\gamma\rangle\!\rangle$.

*Some observations on relative expressiveness:*

(1) $[C; C_1, \ldots, C_k](\phi; \phi_1, \ldots, \phi_k)$ generalises $[C](\phi; \phi_1, \ldots, \phi_k)$, which is equivalent to $[C; \mathsf{Agt}, \ldots, \mathsf{Agt}](\phi; \phi_1, \ldots, \phi_k)$.
(2) $[\langle C_1\rangle\phi_1; \ldots; \langle C_k\rangle\phi_k]$ is equivalent to $\langle\!\langle C_1 \triangleright \phi_1, C_1 \cup C_2 \triangleright \phi_2, ..., C_1 \cup \ldots \cup C_n \triangleright \phi_k \rangle\!\rangle$.
(3) $[C](\phi; \psi)$ is equivalent to $\langle\!\langle C \triangleright \phi, \mathsf{Agt} \triangleright \psi \rangle\!\rangle$. However, this does not generalise to $[C](\phi; \psi_1, \ldots, \psi_k)$ for $k \geq 2$. Conversely, the GIP operator $\langle\!\langle C_1 \triangleright \phi_1, ..., C_n \triangleright \phi_n \rangle\!\rangle$ cannot be expressed in terms of operators $[C](\phi; \psi_1, \ldots, \psi_k)$.
(These non-expressiveness claims can be verified using the respective notions of bisimulations introduced for each of these operators in Section 4.)

## 3.2 Socially Friendly Coalition Logic SFCL

The formulae of SFCL are given by the following grammar.

$$\phi ::= p \mid \neg\phi \mid \phi \vee \phi \mid [C](\phi; \phi, \ldots, \phi) \qquad (p \in \mathsf{AP})$$

Again, standard definitions of the propositional connectives apply.

Given a finite sequence of formulae $\Psi = \psi_1, \ldots, \psi_k$, we will often write $[C](\phi; \Psi)$. We always assume that $k \geq 1$, but note that the coalition operator $[C]\phi$ from CL, which corresponds to the case $k = 0$, can be expressed as $[C](\phi; \top)$, which we hereafter adopt as an abbreviation. We will denote by $\mathsf{SFCL}_1$ the fragment of SFCL containing only formulae $[C](\phi; \psi)$ where $\psi$ is a single formula. **Modal depth** of the formulae of SFCL is defined as expected, where propositional formulae have depth 0, via the clause $md([C](\phi; \psi_1, \ldots, \psi_k)) = 1 + \max(md(\phi), md(\psi_1), \ldots, md(\psi_k))$.

The formal semantics of SFCL is defined in terms of truth of a formula at a state $s$ of a game model $\mathcal{M}$, just like for CL. The clause for $[C](\phi; \Psi)$, given in terms of its extension in $\mathcal{M}$, is as follows:

$$[\![[C](\phi; \Psi)]\!]_{\mathcal{M}} = \Big\{ s \in S \mid \exists\sigma_C \in \textstyle\prod_{a \in C} \Sigma_a : O[s, \sigma_C] \subseteq [\![\phi]\!]_{\mathcal{M}} \ \& \ O[s, \sigma_C] \cap [\![\psi]\!]_{\mathcal{M}} \neq \emptyset \text{ for each } \psi \in \Psi \Big\}.$$

*Example 3.1.* Consider, for illustration, the family couple Ann and Bill and their son Charlie. Each of them has saved some money. Now, they are sitting in a family meeting and negotiating on how to spend their savings. Ann wishes a complete kitchen renovation, Bill wants a new car, and Charlie dreams of a holiday trip to Disneyland.

Ann considers two options for a kitchen: a cheaper IKEA version, or a designer's luxury version. Bill has in mind three options for a car: a cheap used Ford (black, of course), a more expensive new Nissan, or an even more expensive, vintage pink Cadillac. As for Charlie, he would prefer the whole family to go for an expensive week long family excursion to Disneyland in Paris, but could also settle for a cheaper 2-day car trip to Disneyland Park in California.

The possible actions of every family member or a group are to pay for any option of their wish, that they can afford, and then leave the rest of their savings in the family money pool for the other(s) to use. Let us denote the respective goals by: $K$ (any kitchen), resp. $CK$ (cheap kitchen) and $EK$ (expensive kitchen); $C$ (any car), resp.

$CC$ (cheap car), $AC$ (average car), $EC$ (expensive car); $T$ (any trip), resp. $CT$ (cheap trip), $ET$ (expensive trip).

Calculations have shown that Ann's choices produce the following strategic powers: $[\mathsf{Ann}](K; C, T)$ and $[\mathsf{Ann}](CK; CC \wedge CT, EC, ET)$. On the other hand, it is not the case that $[\mathsf{Ann}](EK; AC \wedge T)$.

Likewise, it turns out that Bill has following strategic powers: $[\mathsf{Bill}](C; K; T)$, and $[\mathsf{Bill}](CC; EK, ET)$. However, $\neg[\mathsf{Bill}](EC; K \wedge T)$.

Lastly, here are some coalitional powers: $[\mathsf{Ann}, \mathsf{Bill}](C \wedge EK; CT)$, $[\mathsf{Ann}, \mathsf{Charlie}](CK \wedge ET; CC)$, $[\mathsf{Bill}, \mathsf{Charlie}](AC \wedge CT; K)$.

Now, one may ask, for instance, whether it can be derived from all of the above whether the whole family can afford buying the pink Cadillac and going with it to Disneyland in California, i.e. whether $[\mathsf{Ann}, \mathsf{Bill}, \mathsf{Charlie}](EC \wedge CT; \top)$, or whether Ann can get her designer's kitchen and still enable Bill to buy a new Nissan or the family to go to Disneyland in Paris, i.e. $[\mathsf{Ann}](EK; AC, ET)$, etc.

### 3.3 Group Protecting Coalition Logic GPCL

The formulae of GPCL are given by the following grammar.

$$\phi ::= p \mid \neg\phi \mid \phi \vee \phi \mid \langle\!\langle \gamma \rangle\!\rangle \qquad (p \in \mathsf{AP})$$

where $p \in \mathsf{AP}$ and $\gamma$ is a coalitional goal assignment. We will use the notation $\langle\!\langle C_1 \triangleright \phi_1, ..., C_k \triangleright \phi_k \rangle\!\rangle$ as an explicit record of $\langle\!\langle \gamma \rangle\!\rangle$ where $\gamma$ is the (unique) coalitional goal assignment defined by $\gamma(C_1) = \phi_1, ..., C_k = \phi_k$, and $\gamma(C) = \top$ for every other $C \in \mathcal{P}(\mathsf{Agt})$.

*Remark.* Alternatively, we could adopt $\langle\!\langle C_1 \triangleright \phi_1, ..., C_k \triangleright \phi_k \rangle\!\rangle$ as the primary constructs in the language, and then regard $\langle\!\langle \gamma \rangle\!\rangle$ as an abbreviation for $\langle\!\langle C_1 \triangleright \gamma(C_1), ..., C_{2^n} \triangleright \gamma(C_{2^n}) \rangle\!\rangle$ where $C_1, ..., C_{2^n}$ is a (fixed) canonical enumeration of $\mathcal{P}(\mathsf{Agt})$. Clearly, both versions are expressively equivalent. Furthermore, while the alternative option seems generally more succinct, that gain vanishes if $\langle\!\langle \gamma \rangle\!\rangle$ is defined as a *partial* goal assignment, assigning only the non-trivial goals.

The semantics of GPCL is defined again in terms of truth of a formula at a state $s$ of a game model $\mathcal{M}$. The clause for $\langle\!\langle \gamma \rangle\!\rangle$, given in terms of extensions, is as follows, where $\Sigma_{\mathsf{Agt}} = \prod_{a \in \mathsf{Agt}} \Sigma_a$:

$$[\![\langle\!\langle \gamma \rangle\!\rangle]\!]_{\mathcal{M}} = \left\{ s \in S \mid \exists \sigma \in \Sigma_{\mathsf{Agt}} : O[s, \sigma_C] \subseteq [\![\gamma(C)]\!]_{\mathcal{M}} \text{ for all } C \subseteq \mathsf{Agt} \right\}$$

Note that the fragment $\mathsf{SFCL}_1$ of SFCL (and therefore CL) embeds into GPCL by defining $[C](\phi; \psi)$ equivalently as $\langle\!\langle C \triangleright \phi, \mathsf{Agt} \triangleright \psi \rangle\!\rangle$.

*Example 3.2.* This example is adapted from [10]. Consider a scenario involving two players, Alice ($A$) and Bob ($B$). Each of them owns a server storing some data, to which access is password-protected. The two players want to exchange passwords, but neither player is sure whether to trust the other. So their common goal is to successfully cooperate and exchange passwords, but each player also has the private goal not to give away their password in case the other one turns out to be untrustworthy. Let us 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". So, e.g. the best possible outcome for Alice is $H_A \wedge H_B$ and the worst possible outcome is $\neg H_A \wedge H_B$. When can the two players cooperate to exchange passwords?

Say, for example, that we define a game as follows: each player chooses a password, which may or may not be the correct one, and sends it to the other player. The game ends. In this game, the coalition $\{A, B\}$ can certainly force an outcome where each player has the other's password, i.e. the coalition logic formula

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

holds. But the strategy profile witnessing this does not satisfy the players' individual private goals, since each runs the risk of giving away their password without getting the other's in return.

If, instead, we add a second round to the game, in which each player can either accept the outcome of the first round, or change their password (thus making the shared password useless), then there is a strategy profile in which the players exchange passwords, but which is also safe for each of the two players. This situation is described by the following stronger formula in GPCL:

$$[(\{A, B\} \triangleright H_A \wedge H_B), (A \triangleright H_B \to H_A), (B \triangleright H_A \to H_B)]$$

## 4 BISIMULATIONS AND INVARIANCE

### 4.1 Bisimulations for SFCL

The SFCL-bisimulation is a natural combination of the bisimulations for CL [11] and for INL [15]. We only define SFCL-bisimulation within a game model, which generalises to SFCL-bisimulation *between* game models, by treating both as parts of their disjoint union.

*Definition 4.1 (SFCL-bisimulation).* Let $\mathcal{M} = (S, \{\Sigma_a\}_{a \in \mathsf{Agt}}, g, V)$ be a game model. A binary relation $\beta \subseteq S^2$ is a SFCL-**bisimulation in** $\mathcal{M}$ if it satisfies the following conditions for every pair of states $(s_1, s_2)$ such that $s_1 \beta s_2$ and for every coalition $C$:

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

**Forth:** For any joint action $\sigma^1_C$ of $C$ at $s_1$ there is a joint action $\sigma^2_C$ of $C$ at $s_2$, such that:

- **LocalForth:** For every $u_1 \in O[s_1, \sigma^1_C]$ there exists $u_2 \in O[s_2, \sigma^2_C]$ such that $u_1 \beta u_2$;
- **LocalBack:** For every $u_2 \in O[s_2, \sigma^2_C]$ there exists $u_1 \in O[s_1, \sigma^1_C]$ such that $u_1 \beta u_2$.

**Back:** Conversely, for any joint action $\sigma^2_C$ of $C$ at $s_2$ there is a joint action $\sigma^1_C$ of $C$ at $s_1$, such that the **LocalForth** and **LocalBack** conditions above apply.

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

Proposition 4.2 (SFCL-bisimulation invariance). *Let $\beta$ be a SFCL-bisimulation in a game model $\mathcal{M}$. Then for every SFCL-formula $\theta$ and a pair $s_1, s_2 \in \mathcal{M}$ are such that $s_1 \beta s_2$: $\mathcal{M}, s_1 \models \theta$ iff $\mathcal{M}, s_2 \models \theta$.*

Proof. Routine induction on $\theta$, similar to that in [15]. □

Proposition 4.3 (Hennessy-Milner property). *Let $\beta$ be a SFCL-bisimulation in a finite game model $\mathcal{M}$. Then for any pair $s_1, s_2 \in \mathcal{M}$, $s_1 \beta s_2$ holds iff $s_1$ and $s_2$ are SFCL-equivalent (i.e. satisfy the same SFCL-formulae).*

Proof. (Sketch) Since $\mathcal{M}$ is finite there is a formula $\chi(s)$ for each state $s$ in $\mathcal{M}$ such that $s_1, s_2$ are SFCL-equivalent if and only if $\chi(s_1) = \chi(s_2)$, and such that $\chi(s_1) \wedge \chi(s_2)$ is equivalent to $\bot$ whenever $s_1, s_2$ are not SFCL-equivalent. Our goal is to show that the relation of SFCL-equivalence is itself a SFCL-bisimulation.

The crucial observation is that each state $s$ satisfies the formula:

$$\bigwedge_{C \subseteq \mathsf{Agt}} \bigwedge \Big\{ [C] \left( \chi(v_1) \vee ... \vee \chi(v_m); \chi(v_1), ..., \chi(v_m) \right) \mid \exists \sigma \in \Sigma^s_{a_1} \times ... \times \Sigma^s_{a_n} : O[s, \sigma|_C] = \{v_1, ..., v_m\} \Big\}$$

□

## 4.2 Bisimulations for GPCL

Again, we only define GPCL-bisimulation within a game model, and generalise to bisimulations between models via disjoint unions.

*Definition 4.4 (GPCL-bisimulation).* Let $\mathcal{M} = (S, \{\Sigma_a\}_{a\in\mathrm{Agt}}, g, V)$ be a game model. A binary relation $\beta \subseteq S^2$ is a GPCL-**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 $p \in \mathrm{AP}$: $s_1 \in V(p)$ iff $s_2 \in V(p)$.

**Forth:** For any joint action $\sigma^1$ of Agt at $s_1$ there is a joint action $\sigma^2$ of Agt at $s_2$ such that:
  **LocalBack:** For every coalition $C$ and every $u_2 \in O[s_2, \sigma^2|_C]$, there is some $u_1 \in O[s_1, \sigma^1|_C]$ such that $u_1 \beta u_2$.

**Back:** For any joint action $\sigma^2$ of Agt at $s_2$ there is a joint action $\sigma^1$ of Agt at $s_1$ such that:
  **LocalForth:** For every coalition $C$ and every $u_1 \in O[s_1, \sigma^1|_C]$, there is some $u_2 \in O[s_2, \sigma^2|_C]$ such that $u_1 \beta u_2$.

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

Proposition 4.5 (GPCL-bisimulation invariance). *Let $\beta$ be a GPCL-bisimulation in a game model $\mathcal{M}$. Then for every GPCL-formula $\theta$ and every pair $s_1, s_2 \in \mathcal{M}$ such that $s_1 \beta s_2$:*

$$\mathcal{M}, s_1 \models \theta \text{ iff } \mathcal{M}, s_2 \models \theta$$

Proof. Again, a routine induction on $\theta$. We only consider the case for $\theta = \langle\!\langle C_1 \triangleright \phi_1, ..., C_k \triangleright \phi_k \rangle\!\rangle$: suppose that $\mathcal{M}, s_1 \models \theta$, witnessed by a joint action $\sigma^1$ for Agt at $s_1$. Let $\sigma^2$ be some joint action for Agt at $s_2$ witnessing the Forth condition with respect to $\sigma^1$. We need to show that $O[s_2, \sigma^2_{C_i}] \subseteq [\![\phi_i]\!]$ for each $i \in \{1, ..., k\}$. Suppose $v \in O[s_2, \sigma^2_{C_i}]$. Apply the LocalBack condition to find $v' \in O[s_1, \sigma^1_{C_i}]$ with $v' \beta v$. Since $O[s_1, \sigma^1_{C_i}] \subseteq [\![\phi_i]\!]$ we get $\mathcal{M}, v \models \phi_i$ by the induction hypothesis on $\phi_i$, as required. The converse direction is proved in the same manner. □

Proposition 4.6 (Hennessy-Milner property). *Let $\beta$ be a GPCL-bisimulation in a finite game model $\mathcal{M}$. Then for any pair $s_1, s_2 \in \mathcal{M}$, $s_1 \beta s_2$ holds iff $s_1$ and $s_2$ satisfy the same GPCL-formulae.*

Proof. (Sketch) Since $\mathcal{M}$ is finite, we can define a characteristic formula $\chi(s)$ for each state $s$ in $\mathcal{M}$ as in the proof of 4.3. For a set of states $Z$, let $\chi[Z] = \bigvee\{\chi(v) \mid v \in Z\}$. Our goal is to show that the relation of GPCL-equivalence is itself a GPCL–bisimulation, and the key observation is that each state $s$ satisfies the formula:

$$\bigwedge_{\sigma\in M_1^s\times\ldots M_n^s} \langle\!\langle C_1 \triangleright \chi[O[s, \sigma|_{C_1}], ..., C_k \triangleright \chi[O[s, \sigma|_{C_k}]] \rangle\!\rangle$$

where $C_1, ..., C_k$ is the list of all possible coalitions in $\mathcal{P}(\mathrm{Agt})$. □

# 5 AXIOMATISATIONS AND DECIDABILITY

Here we present axiomatic systems for the logics SFCL and GPCL and prove their completeness with respect to finite tree-like models, thus also implying decidability of each of them. We employ somewhat different proof methods for the two completeness proofs, as each of them has its own merits.

## 5.1 Axiomatic system for SFCL

The axiomatic system $Ax_{\mathsf{SFCL}}$ for SFCL merges $Ax_{\mathsf{CL}}$ with a multi-agent extension of the axiomatisation of INL from [15] as follows, where $\Theta$ and $\Psi$ are any finite lists of formulae.

(1) **The axioms and rules from** $Ax_{\mathsf{CL}}$**.** (Some of those are subsumed by, or derivable from, the axioms added below.)

(2) **Axioms from INL.**

(INL1) $[C](\phi; \Psi) \to [C](\phi \vee \theta; \Psi)$

(INL2) $[C](\phi; \Psi, \psi) \to [C](\phi; \Psi, \psi \vee \theta)$

(INL3) $[C](\phi; \Psi, \psi) \to [C](\phi; \Psi, \psi \wedge \phi)$

(INL4) $\neg[C](\phi; \bot)$

(INL5) $[C](\phi; \Psi) \to [C](\phi \wedge \neg\theta; \Psi) \vee [C](\phi; \Psi, \theta)$

(INL6) $[C](\phi; \Theta, \theta, \Psi) \to [C](\phi; \Theta, \Psi)$

(INL7) $[C](\phi; \Theta, \theta, \Psi) \to [C](\phi; \Theta, \theta, \Psi, \theta)$

(3) **Additional axioms:**

(SFCL1) $[\mathrm{Agt}](\phi; \psi, \Psi) \leftrightarrow [\mathrm{Agt}](\phi \wedge \psi; \Psi)$.

(SFCL2) $[C](\phi; \psi) \to [C'](\phi; \psi)$, whenever $C \subseteq C' \subseteq \mathrm{Agt}$.

(SFCL3) $[\mathrm{Agt}]\psi \leftrightarrow [C](\top; \psi)$, for any $C \subseteq \mathrm{Agt}$.

(SFCL4) $[C](\phi; \Psi) \wedge [\emptyset]\theta \to [C](\phi \wedge \theta; \Psi)$.

(SFCL5) $[\emptyset](\phi; \psi_1, \ldots, \psi_k) \leftrightarrow [\emptyset]\phi \wedge [\mathrm{Agt}]\psi_1 \wedge \ldots \wedge [\mathrm{Agt}]\psi_k$.

**Rules of inference: (MP)** and $[C]$**-Monotonicity (C-Mon)**:

$$\frac{\phi \to \phi',\ \psi_1 \to \psi_1', \ldots, \psi_k \to \psi_k'}{[C](\phi; \psi_1, \ldots, \psi_k) \to [C](\phi'; \psi_1', \ldots, \psi_k')}.$$

**Some derivable validities:**

The formulae listed below are derivable in $Ax_{\mathsf{SFCL}}$ and will be used in the completeness proof, or are worth noting anyway.

**(D0)** $\neg[C](\bot; \Psi)$. (From (INL3) and (INL4).)

**(D1)** $[C]\phi \wedge \neg[C]\neg\psi \to [C](\phi; \psi)$. (Particular case of (INL5).) Consequently, $[C_1]\phi \wedge [C_2]\psi \to [C_1](\phi; \psi)$, for $C_1 \cap C_2 = \emptyset$

**(D2)** $[C](\phi \wedge \psi_1 \ldots \wedge \psi_k) \to [C](\phi; \psi_1, \ldots, \psi_k)$. (By repeated application of (INL5), using (D0) and (INL1).)

**(D3)** $\neg[C]\neg\psi \to [\mathrm{Agt}]\psi$. (Using (CL1).)

**(D4)** $[C](\phi; \psi_1, \ldots, \psi_k) \to \bigwedge_{i=1}^{k}[\mathrm{Agt}](\phi \wedge \psi_i)$. (Using (SFCL3), (INL5), and (INL6).)

**(D5)** $[\mathrm{Agt}](\phi; \Psi) \leftrightarrow [\mathrm{Agt}](\phi \wedge \bigwedge \Psi)$. (Using (SFCL1).)

**(D6)** If $\vdash \phi, \vdash \psi_1, \ldots \vdash \psi_k$, then $\vdash [C](\phi; \psi_1, \ldots, \psi_k)$. (Using (CL3) and **(C-Mon)**.)

**(D7)** $[\emptyset]\phi \leftrightarrow \neg[\mathrm{Agt}]\neg\phi$ and $[\mathrm{Agt}]\phi \leftrightarrow \neg[\emptyset]\neg\phi$. (Using (CL1), (CL2) and (CL4).)

**(D8)** $[\mathrm{Agt}](\phi \vee \psi) \leftrightarrow ([\mathrm{Agt}]\phi \vee [\mathrm{Agt}]\psi)$. (Using (CL4) and (D7).)

The proof of soundness of $Ax_{\mathsf{SFCL}}$ is routine and left to the reader. For completeness we will prove the following stronger result.

Theorem 5.1 (Finite tree-like model property of SFCL). *Every finite $Ax_{\mathsf{SFCL}}$-consistent set of SFCL formulae $\Gamma$ is satisfied in a finite tree-like game model.*

Proof. (Detailed sketch) We will build finite tree-like rooted game models (i.e. models in $\mathfrak{T}$) satisfying finite $Ax_{\mathsf{SFCL}}$-consistent sets of formulae $\Gamma$, inductively on the greatest modal depth of a formula in $\Gamma$, hereafter denoted $md(\Gamma)$. We restrict the language to only those atomic propositions that occur in formulae in $\Gamma$.

First, $md(\Gamma) = 0$ iff $\Gamma$ is a consistent set of classical propositional formulae. Since $Ax_{\mathsf{SFCL}}$ contains a complete set of propositional tautologies, every such $\Gamma$ is satisfiable in a singleton from $\mathfrak{T}$.

Now, suppose we have constructed satisfying rooted models from $\mathfrak{T}$ for every finite $Ax_{\mathsf{SFCL}}$-consistent set $\Gamma$ such that $md(\Gamma) \leq d$, and now take any such set with $md(\Gamma) = d + 1$. Replace $\Gamma$ with its conjunction, transformed to an equivalent DNF. Then take any $Ax_{\mathsf{SFCL}}$-consistent disjunct from that DNF (there is at least one). It is an elementary conjunction of the type $\delta = \delta_0 \wedge \delta_1 \wedge \ldots \wedge \delta_m$, where $\delta_0$ is a consistent conjunction of propositional literals and each $\delta_i$, for $i > 0$, is a formula of the type $[C](\phi; \psi_1, \ldots, \psi_k)$ or a negation of such a formula, where each $\phi, \psi_1, \ldots, \psi_k$ has modal depth at most $d$. We can assume that the index $k$ is the same for all $\delta_i$, for $i \in \{1, \ldots, m\}$, as the shorter ones can be padded with $\top$s. We can also assume that all coalitions heading the subformulae $\delta_i$ are, unless otherwise specified, proper subsets of Agt, as all subformulae headed with $[\mathrm{Agt}]$ can be equivalently replaced by ones headed with $[\emptyset]$, using Axioms (CL1) and (SFCL1) (and (D5), (D7)). Lastly, we can also assume that all subformulae $\delta_i$ headed by $[\emptyset]$ are of the type $[\emptyset]\psi$, by using axiom (SFCL5) plus the equivalence (D7).

Now, to prove that every consistent elementary conjunction $\delta$ described above is satisfiable we will use the following key technical lemma, which characterises satisfiability of such conjunctions in terms of satisfiability conditions on the arguments occurring in the conjuncts (which are all of lower modal depth).

Lemma 5.2. *Let $C_1, \ldots C_m, B_1, \ldots B_{l-1}$ be proper subsets of* Agt *and $B_l = \mathrm{Agt}$; $\delta_0$ be a propositional formula;*
*$\delta_i = [C_i](\phi^i; \psi^i_1, \ldots, \psi^i_k)$, for $i = 1, \ldots, m$;*
*$\delta_{m+j} = \neg[B_j](\eta^j; \theta^j_1, \ldots, \theta^j_k)$, for $j = 1, \ldots, l$,*
*and $\eta^l = \theta^l_1 = \ldots = \theta^l_k = \bot$. Then the following hold.*

***(SAT)*** *The formula*

$$(*) \qquad \delta = \delta_0 \wedge \delta_1 \wedge \ldots \wedge \delta_m \wedge \delta_{m+1} \wedge \ldots \wedge \delta_{m+l}$$

*is satisfiable if and only if:*

1. *$\delta_0$ is propositionally satisfiable,*
2. *Each $\phi^i \wedge \psi^i_j$ is satisfiable, for $i = 1, \ldots, m$ and $j = 1, \ldots, k$.*
3. *For every (possibly empty) set of indices $\{i_1, \ldots, i_e\} \subseteq \{1, \ldots, m\}$ and $j \in \{1, \ldots, l\}$: if the coalitions $C_{i_1}, \ldots C_{i_e}, \overline{B_j}$ are pairwise disjoint (thus, in particular, $C_{i_1} \cup \ldots \cup C_{i_e} \subseteq B_j$), then each of the formulae $\phi^{i_1} \wedge \ldots \wedge \phi^{i_e} \wedge \neg(\eta^j \wedge \theta^j_1, \ldots, \wedge\theta^j_k)$ and $\phi^p \wedge \psi^p_q \wedge \neg(\eta^j \wedge \theta^j_1, \ldots, \wedge\theta^j_k)$, for $p \in \{i_1, \ldots, i_e\}$ and $q \in \{1, \ldots, k\}$, is satisfiable. (When $j = l$ these reduce to satisfiability of $\phi^{i_1} \wedge \ldots \wedge \phi^{i_e}$, as $\phi^p \wedge \psi^p_q$ is covered by (2).)*

***(CON)*** *The 'only if' direction[2] of **(SAT)**, with "satisfiable" replaced throughout by "$Ax_{\mathsf{SFCL}}$-consistent".*

Proof. (Detailed sketch)

**Part I (Necessity)** Whenever we mention in this proof consistency and derivations (stated with $\vdash$), they refer to $Ax_{\mathsf{SFCL}}$.

First, let us check the necessity of the conditions (1)-(3) for the satisfiability, resp. consistency, of $\delta$. We will sketch the deductive reasoning for proving **(CON)**, as the deductive steps can be readily replaced by semantic steps to prove the 'only if' direction of **(SAT)**.

(1) is obviously necessary.

For (2), if some $\phi^i \wedge \psi^i_j$ is inconsistent, then $\delta_i = [C_i](\phi^i; \psi^i_1, \ldots, \psi^i_k)$ is inconsistent, too, hence so is $\delta$. That is because if $\vdash \phi^i \to \neg\psi^i_j$ then $\vdash \neg\delta_i$, by (INL3), (INL4) and (C-Mon).

Lastly, if (3) fails for some such $C_{i_1}, \ldots C_{i_e}, B_j$, then we have

either (**) $\vdash \phi^{i_1} \wedge \ldots \wedge \phi^{i_e} \to (\eta^j \wedge \theta^j_1, \ldots, \wedge\theta^j_k)$.

or (***) $\vdash \phi^p \wedge \psi^p_q \to (\eta^j \wedge \theta^j_1, \ldots, \wedge\theta^j_k)$.

The case of (***) is easier, implying that $\vdash \delta_p \to \neg\delta_{m+j}$, whence the inconsistency of $\delta$. In case of (**) we use $\vdash [C](\phi; \Psi) \to [C]\phi$ (by (C-Mon)) and repeatedly apply the Superadditivity axiom (CL4) to obtain that

$\vdash \delta_{i_1} \wedge \ldots \wedge \delta_{i_e} \to [C_{i_1} \cup \ldots \cup C_{i_e}](\phi^{i_1} \wedge \ldots \wedge \phi^{i_e})$.

From this, by (SFCL2), we obtain that

$\vdash \delta_{i_1} \wedge \ldots \wedge \delta_{i_e} \to [B_j](\phi^{i_1} \wedge \ldots \wedge \phi^{i_e})$.

Then, using (**) and (C-Mon) for $[B_j]$, we obtain that

$\vdash \delta_{i_1} \wedge \ldots \wedge \delta_{i_e} \to [B_j](\eta^j \wedge \theta^j_1, \ldots, \wedge\theta^j_k)$,

whence (using the derivable (D2)

$\vdash \delta_{i_1} \wedge \ldots \wedge \delta_{i_e} \to \neg\delta_{m+j}$, which implies $\vdash \neg\delta$.

**Part II (Sufficiency)** Now, to prove sufficiency of the conditions (1)-(3) we show how to construct a satisfying model for $\delta$ from rooted models for all satisfiable formulae mentioned in (1)-(3). Furthermore, if all latter models are finite tree-like models, then so will be the one for $\delta$, because it will be constructed from them by grafting. Before outlining the construction, let us select these models. The subformula $\delta_0$ will be satisfied at the root of the constructed model, and we can ignore it hereafter. Now, for each $\phi^i \wedge \psi^i_j$ we select a satisfying rooted model $(\mathcal{M}^i_j, s^i_j) \in \mathfrak{T}$. Further, note that for each satisfiable formula $\phi^{i_1} \wedge \ldots \wedge \phi^{i_e} \wedge \neg(\eta^j \wedge \theta^j_1, \ldots, \wedge\theta^j_k)$ referred in (3), either $\nu^j = \phi^{i_1} \wedge \ldots \wedge \phi^{i_e} \wedge \neg\eta^j$ or $\chi^j_r = \phi^{i_1} \wedge \ldots \wedge \phi^{i_e} \wedge \neg\theta^j_r$, for some $r \in \{1, \ldots, k\}$ is satisfiable. Likewise, for each satisfiable formula $\phi^p \wedge \psi^p_q \wedge \neg(\eta^j \wedge \theta^j_1, \ldots, \wedge\theta^j_k)$, either $\nu^j_{p,q} = \phi^p \wedge \psi^p_q \wedge \neg\eta^j$ or $\chi^j_{p,q,r} = \phi^p \wedge \psi^p_q \wedge \neg\theta^j_r$, for some $r \in \{1, \ldots, k\}$ is satisfiable. For each of these, we select $\nu^j$, resp. $\nu^j_{p,q}$, if it is satisfiable, else we select $\chi^j_q$, resp. $\chi^j_{p,q,r}$. In either case, we then select a satisfying rooted model $(\mathcal{M}_j, s_j) \in \mathfrak{T}$. Let $\mathfrak{M}$ be the (finite) family of all these rooted models, plus a singleton model with single action profile with the same state as outcome. We will assume all these models pairwise disjoint. Let $s$ be a state not belonging to any of them.

The rooted model satisfying $\delta$ is defined as the grafting of a set of models in $\mathfrak{M}$ at $s$ by a strategic game form defined as follows. The set of actions for each agent is $\{1, \ldots, m + l\} \times \{1, \ldots, k\}$. Every such action can be regarded as voting in two rounds: the first round in the set $\{1, \ldots, m + l\}$ and the second round in $\{1, \ldots, k\}$. Each action profile determines a successor state of the resulting model, where a satisfiable set of formulae is placed in a label. A satisfying model for the label is rooted at that state, thus completing the grafting.

**Round 1.** The voting in this round is essentially organised as in [7, Lemma 33], from where we borrow the idea, viz: every agent votes for one of the formulae $\delta_1, \ldots, \delta_{m+l}$. For any $i = 1, \ldots, m$, if all agents in $C_i$ vote $i$, we say that these agents *form the coalition* $C_i$ and that the formula $\phi^i$ *is elected* at that round. Then, for each

[2]In fact, both directions hold, but we only need to prove this one here.

set of coalitions $C_{i_1}, \dots C_{i_e}, B_j$ described in condition (3), where the selected satisfiable formula is of type $\nu^j$ or $\nu^j_{p,q}$ that formula is elected iff $j = [(\sum_{i \in I}(v(i) - m)) \mod l] + 1$, where $I$ is the set of indices of all agents $i$ who have voted 'negatively', i.e. in $\{m+1, ..., m+l\}$, and $v(i)$ is $i$'s vote. Clearly, at most one such formula gets elected in this round, and all elected formulae among $\delta_1, ..., \delta_m$ correspond to pairwise disjoint coalitions, which are also disjoint with the coalition $B_j$ of the elected negative formula, if any.

This round of voting ensures that all coalitions $C_1, \dots, C_m$ have respective collective actions to enforce the placement of their universal components $\phi^1, \dots, \phi^k$ in the labels of all successor states generated by all action profiles containing these collective actions. Besides, the elected formula $\nu^j$ or $\nu^j_{p,q}$ (if any) takes care of the respective negative formula $\delta_{m+j}$.

**Round 2.** In this round the agents take care of the existential components of the formulae $\delta_1, ..., \delta_{m+l}$, by voting in the set $\{1, ..., k\}$ as follows. We can assume that only agents who did not form coalitions in the first round vote now, as the other votes will be disregarded. Now, if a coalition $C_i$ was formed in the 1st round, while all agents not in $C_i$ did not form coalitions in the 1st round and vote $j$ in the 2nd round, then the formula $\psi^i_j$ is elected now and added to the label of the resulting state for this action profile.

Respectively, for each family of pairwise disjoint coalitions $C_{i_1}, \dots, C_{i_e}, \overline{B_j}$, as described in condition (3), where the selected satisfiable formula is of type $\chi^j_r$ or $\chi^j_{p,q,r}$ (as defined above), if all coalitions $C_{i_1}, \dots, C_{i_e}$ have been formed in round 1, and all agents that are not in any of these coalitions vote j, then the respective formula $\chi^j_r$ or $\chi^j_{p,q,r}$ is elected. Clearly, at most one formula of type $\psi$ or $\chi$ is elected at round 2, and that vote takes respective care of the existential components of positive and negative formulae where the universal components are enforced.

For all action profiles not covered in the cases described in the two rounds, only the truth $\top$ is elected.

That completes the definition of an outcome function $o_s$ at the added root $s$. Note that of every vote (i.e. action profile) the set of elected formulae put in the label of the respective outcome state is satisfiable and there is a satisfying rooted model for it in $\mathfrak{M}$.

Now, each newly created by $o_s$ outcome node is now assigned a respective rooted model from $\mathfrak{M}$ that satisfies the set of elected formulae in its label.

We claim that the rooted model constructed by grafting of $\mathfrak{M}$ at $s$ by the outcome function $o_s$ satisfies $\delta$. The long but routine details are omitted. That completes the proof of the lemma. ◻

We can now prove that every consistent conjunction $\delta$ as in (*) is satisfiable. Indeed, by the claim **(CON)** of Lemma 5.2, the consistency of $\delta$ implies consistency of all formulae referred in conditions (1)-(3), which are all of modal depth lower than $md(\delta)$. By the inductive hypothesis, they are all satisfiable in finite rooted tree-like models. Then, by the 'if' part of **(SAT)**, $\delta$ itself is satisfiable in such a model. That completes the proof of the theorem. ◻

Corollary 5.3 (Completeness of SFCL). *The axiomatic system $Ax_{\mathsf{SFCL}}$ is sound and complete.*

Note that the size of the finite tree-like game model constructed in the proof of Theorem 5.1 is at most exponential in the length of the formula. Still, the number of immediate successors (excl. the roots of trivial singletons) of any node of such a satisfying model can be bounded by a polynomial in the length of the formula (assuming that all agents occurring in the same coalitions of the formula are assigned identical action powers). Therefore, by guessing only the relevant portions of such a satisfying model (current node and its successors) and then checking on the fly, we obtain a decision procedure for SFCL in NPSPACE=PSPACE. The matching lower bound is obtained from the PSPACE-space completeness of the satisfiability problem for CL, cf. [12]. Thus, we obtain the following.

Corollary 5.4 (PSPACE-decidability of SFCL). *Satisfiability in SFCL is decidable and PSPACE-complete.*

## 5.2 Axiomatic system for GPCL

We now present an axiomatic system for the logic GPCL. We begin by introducing some terminology and notation.

**Definition 1.** A **voting profile** is a map $f$ sending each $a_i \in \mathsf{Agt}$ to a goal assignment $\gamma_i = f(a_i)$.

**Definition 2.** Let $f$ be a voting profile. We define the goal assignment merge$(f)$ as follows:

- merge$(f) : C \mapsto \phi$ if $f(a_i) = f(a_j)$ and $f(a_i) : C \mapsto \phi$ for all $a_i, a_j \in C$,
- merge$(f) : C \mapsto \top$ if this holds for no $\phi$.

Some additional terminology and notation:

- The goal assignment $\gamma[C : \phi]$ is like $\gamma$, but mapping $C$ to $\phi$.
- 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 $\top$.
- We let $\gamma^\top$ denote the **trivial goal assignment**, mapping each coalition to $\top$.
- The **modal depth** of a formula is defined in essentially the same way as for SFCL, with $\langle\!\![\cdot]\!\!\rangle$ increasing the depth by 1.

Here is our axiom system $Ax_{\mathsf{GPCL}}$, extending classical propositional logic (all axioms of CL are subsumed here):

(1) **Axioms**.

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

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

(Mrg) $\bigwedge_{a_i \in \mathsf{Agt}} \langle\!\![f(a_i)]\!\!\rangle \to \langle\!\![\mathrm{merge}(f)]\!\!\rangle$

(Case) $\langle\!\![\gamma]\!\!\rangle \to (\langle\!\![\gamma[C : \gamma(C) \wedge \phi]]\!\!\rangle \vee \langle\!\![\gamma|_C[\mathsf{Agt} : \neg\phi]]\!\!\rangle)$

(Con) $\langle\!\![\gamma]\!\!\rangle \to \langle\!\![\gamma[\mathsf{Agt} : \gamma(\mathsf{Agt}) \wedge \gamma(C)]]\!\!\rangle$

**Rules of inference: (MP)** and **Goal Monotonicity (G-Mon)**:

$$\frac{\gamma(C) \to \psi}{\langle\!\![\gamma]\!\!\rangle \to \langle\!\![\gamma[C{:}\psi]]\!\!\rangle}$$

We write $\vdash \psi$ to say that $\psi$ is provable in the system $Ax_{\mathsf{GPCL}}$.

Soundness of $Ax_{\mathsf{GPCL}}$ is routine and left to the reader.

Here are some important validities derivable in $Ax_{\mathsf{GPCL}}$:

**(G1)** Agt-Maximality: $[\emptyset]\varphi \vee [\mathsf{Agt}]\neg\varphi$. (Using (Triv) and (Case).)

**(G2)** Superadditivity: $([C_1]\phi_1 \wedge [C_2]\phi_2) \to [C_1 \cup C_2](\phi_1 \wedge \phi_2)$, if $C_1 \cap C_2 = \emptyset$. (Particular case of (Mrg).)

**(G3)** $\langle\!\![\gamma]\!\!\rangle \to [C]\gamma(C)$, for any $C \subseteq \mathsf{Agt}$. (Using (Triv) and (G-Mon).)

**(G4)** $\langle\!\![\gamma]\!\!\rangle \to (\langle\!\![\gamma[\mathsf{Agt} : \gamma(\mathsf{Agt}) \wedge \phi]]\!\!\rangle \vee \langle\!\![\gamma[\mathsf{Agt} : \gamma(\mathsf{Agt}) \wedge \neg\phi]]\!\!\rangle)$. (From (Case), replacing $C$ with Agt and $\phi$ with $\gamma(\mathsf{Agt}) \to \phi$.)

The main technical result of this section is:

Theorem 5.5 (Finite tree-like model property of GPCL). *Every finite $Ax_{\mathrm{GPCL}}$-consistent set of* GPCL *formulae* $\Gamma$ *is satisfied in a finite tree-like game model.*

From this result we get:

Corollary 5.6 (Completeness of GPCL). *The axiomatic system $Ax_{\mathrm{GPCL}}$ is sound and complete.*

For the rest of the section we focus on the proof of Theorem 5.6.

**Definition 3.** A goal assignment $\gamma$ is said to have depth $\leq k$ if the formula $\langle\!\lbrack\gamma\rbrack\!\rangle$ has depth $k+1$. A goal assignment $\gamma$ is said to *strengthen* the goal assignment $\gamma'$ if $\gamma(C) \vdash \gamma'(C)$ for every coalition $C$. A goal assignment $\gamma$ of depth $\leq k$ is said to be $k$-*maximal* if $\langle\!\lbrack\gamma\rbrack\!\rangle$ is consistent and for every strengthening $\gamma'$ of $\gamma$ of depth $\leq k$ such that $\langle\!\lbrack\gamma'\rbrack\!\rangle$ is consistent, $\gamma$ is also a strengthening of $\gamma'$.

The following lemma is almost immediate from the definitions together with the Goal Monotonicity rule (G-Mon):

Lemma 5.7. *Let $\gamma$ be a goal assignment of depth $\leq k$. Then $\langle\!\lbrack\gamma\rbrack\!\rangle$ is equivalent to a disjunction of formulas of the form $\langle\!\lbrack\delta\rbrack\!\rangle$ where $\delta$ is a $k$-maximal strengthening of $\gamma$.*

*Definition 5.8.* Let $k \in \omega$. A formula $\alpha$ of modal depth $\leq k$ is said to be a $k$-*atom* if it is consistent and, for all formulas $\psi$ of depth $\leq k$, we have $\vdash \alpha \to \psi$ or $\vdash \alpha \to \neg\psi$.

Note that there are at most finitely many – up to provable equivalence – formulas, hence finitely many atoms of depth $\leq k$ (given that we have fixed a finite set of propositional variables). For each $k$-atom $\alpha$, we may pick one $k$-atom $\alpha^*$ that is provably equivalent to $\alpha$ and such that $\vdash \alpha \leftrightarrow \beta$ iff $\alpha^* = \beta^*$ for all $k$-atoms $\alpha, \beta$. We set:

$$\mathrm{At}_k = \{\alpha^* \mid \alpha \text{ is a } k\text{-atom }\}$$

Further, we will be a bit sloppy with notation and will not distinguish a $k$-atom $\alpha$ from the provably equivalent atom $\alpha^*$. Note that every consistent formula $\phi$ of depth $k$ is equivalent to the disjunction of all formulas $\alpha^*$ where $\alpha$ is a $k$-atom and $\alpha \vdash \phi$.

Lemma 5.9. *Let $\gamma$ be a $k$-maximal goal assignment. Then $\gamma(\mathrm{Agt})$ is a $k$-atom.*

Proof. Repeated use of axiom (G4) together with (Safe). □

We define the (finite) *depth-k canonical model* $\mathcal{M} = (W, g, V)$ as follows (omitting the definition of the global sets of actions). The domain $W$ of the model is the set:

$$\bigcup_{m \leq k} \mathrm{At}_m$$

The valuation is defined by $V(p) = \{\alpha \in W \mid \alpha \vdash p\}$.

The game $g(\alpha)$ associated with a 0-atom is defined arbitrarily.

The non-trivial part of the construction is to define the game form $g(\alpha) = (\Sigma_1, ..., \Sigma_n, o)$ for a given $k+1$-atom $\alpha$. We define this game form as follows: first, let $\{r, p, s\}$ be a given three-element set with a fixed binary relation $W \subseteq \{r, p, s\}^2$ defined by $W = \{(r, s), (s, p), (p, r)\}$. Think of the elements $r, p, s$ as moves "rock, paper, scissors" and the relation $W$ as "wins over": rock wins over scissors, scissors wins over paper and paper wins over rock. Given $a_i \in \mathrm{Agt}$, we define the set of strategies $M_i$ of $a_i$ to be the set:

$$\{r, p, s\}^{\{a_1, \ldots, a_n\}} \times \mathrm{GA}_{\alpha,k} \times \mathrm{cf}(k)$$

where $\mathrm{GA}_{\alpha,k}$ is the set of goal assignments $\gamma$ of depth $\leq k$ such that $\alpha \vdash \langle\!\lbrack\gamma\rbrack\!\rangle$, and $\mathrm{cf}(k)$ is the set of choice functions mapping each depth $\leq k$ goal assignment $\gamma$ to one of its $k$-maximal strengthenings $\delta \in \mathrm{GA}_{\alpha,k}$. Note that $\mathrm{GA}_{\alpha,k} \neq \emptyset$ by (Triv). Note also that, for each $\gamma \in \mathrm{GA}_{\alpha,k}$ there exists a maximal strengthening of $\gamma$ of the required sort, since $\langle\!\lbrack\gamma\rbrack\!\rangle$ is equivalent to the disjunction of formulas $\langle\!\lbrack\delta\rbrack\!\rangle$, where $\delta$ is a $k$-maximal strengthening of $\gamma$, by Lemma 5.7. So the set of actions for each agent is non-empty, as required.

The idea is to let each player vote for a goal assignment, and the votes are then merged to a single goal assignment. The players also play a game of "rock-paper-scissors" to determine who gets to choose the actual outcome of the merged goal assignment, which will generally not be strong enough to uniquely determine an outcome (which should be a $k$-atom).

Given a strategy profile $\pi$ we say that $a_i$ is a *dominant* player if there is no $a_j$ for which $x_j(a_i) W x_i(a_j)$. The *first-round winner* in $\pi$ is $a_i$ for the smallest index $i$ such that $i$ is a dominant player, or is set to $a_0$ if no such player exists.

The *outcome* $o(\pi)$ of a strategy profile $\pi$ is defined to be

$$\sigma_i(\mathsf{merge}(f))(\mathrm{Agt})$$

where:

- $a_i$ is the first-round winner in $\pi$,
- $\sigma_i$ is the choice function chosen by $a_i$, and
- $f$ is the voting profile corresponding to $\pi$, i.e. $f$ sends each $a_j$ to the goal assignment $\gamma_j$ chosen by $a_j$ in the second round.

Note that $o(\pi)$ is a $k$-atom by Lemma 5.9. Moreover, $\alpha \vdash \langle\!\lbrack\mathsf{merge}(f)\rbrack\!\rangle$ (as required for $\sigma_i(\mathsf{merge}(f))$ to be defined), since $\alpha \vdash \langle\!\lbrack f(a_i)\rbrack\!\rangle$ for each $a_i$ and hence $\alpha \vdash \langle\!\lbrack\mathsf{merge}(f)\rbrack\!\rangle$ by the axiom (Mrg).

The following "truth lemma" can now be proved by a reasonably straightforward verification:

Lemma 5.10. *Let $\mathcal{M}$ be the canonical model of some depth $\geq k$, and let $\alpha$ be a $k$-atom. Then $\mathcal{M}, \alpha \Vdash \phi$ iff $\alpha \vdash \phi$, for each formula $\phi$ of depth $\leq k$.*

Theorem 5.5 now follows as a direct consequence, since for every consistent finite set $\Gamma$ of formulas of maximum modal depth $k$ there must be at least one $k$-atom $\alpha$ such that $\alpha \vdash \bigwedge \Gamma$.

Corollary 5.11 (Decidability of GPCL). *The satisfiability problem for* GPCL *is decidable.*

This decidability follows directly from completeness result (Cor. 5.6) together with the finite model property from Theorem 5.5. We believe that the satisfiability problem for GPCL is PSPACE-complete, like SFCL, but leave the proof of this claim for the future.

## 6 CONCLUDING REMARKS

The schemes defining the coalitional goal operators in SFCL and GPCL can be naturally generalised, by allowing more liberal implicit quantification over collective actions of coalitions and associated with them goals. Such generalisation would result in a quite expressive logic in spirit of Strategy Logic, but still in a purely modal style, without explicit mention of strategies in the language. Also, it is natural to extend both logics SFCL and GPCL in the style of the alternating-time temporal logic ATL, by adding the usual temporal operators to reason about interleaved long-term strategic abilities. We leave these extensions to future work.

## Acknowledgements

The work of Valentin Goranko was supported by a research grant 2015-04388 of the Swedish Research Council. The work of Sebastian Enqvist was supported by the research grant 2015-01774 of the Swedish Research Council. We thank the referees for some useful remarks and references.

## REFERENCES

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

[2] 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.

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

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

[5] Dana Fisman, Orna Kupferman, and Yoad Lustig. 2010. Rational Synthesis. In *Proc. of TACAS 2010*. 190–204. https://doi.org/10.1007/978-3-642-12002-2_16

[6] V. Goranko, W. Jamroga, and P. Turrini. 2013. Strategic Games and Truly Playable Effectivity Functions. *Journal of Autonomous Agents and Multi-Agent Systems* 26, 2 (2013), 288–314.

[7] V. Goranko and G. van Drimmelen. 2006. Complete Axiomatization and Decidablity of Alternating-time temporal logic. *Theor. Comp. Sci.* 353 (2006), 93–117.

[8] 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).

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

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

[11] M. Pauly. 2001. *Logic for Social Software*. Ph.D. Dissertation. University of Amsterdam.

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

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

[14] Johan van Benthem, Nick Bezhanishvili, and Sebastian Enqvist. 2017. A Propositional Dynamic Logic for Instantial Neighborhood Models. In *Proc. of LORI 2017 (Lecture Notes in Computer Science)*, Vol. 10455. Springer, 137–150.

[15] Johan van Benthem, Nick Bezhanishvili, Sebastian Enqvist, and Junhua Yu. 2017. Instantial Neighbourhood Logic. *Reviews of Symbolic Logic* 10, 1 (2017), 116–144.

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

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