

Article

Logics for Strategic Reasoning of Socially Interacting Rational Agents: An Overview and Perspectives

Valentin Goranko 

Department of Philosophy, Stockholm University, 10691 Stockholm, Sweden; valentin.goranko@philosophy.su.se

Abstract: This paper is an overview of some recent and ongoing developments of formal logical systems designed for reasoning about systems of rational agents who act in pursuit of their individual and collective goals, explicitly specified in the language as arguments of the strategic operators, in a socially interactive context of collective objectives and attitudes which guide and constrain the agents' behavior.

Keywords: logic; strategic reasoning; rational agents; socially interactive context; concurrent game models

1. Introduction

The concepts of agency and rationality are fundamental for describing and understanding the modern society. Rational agents can be humans, intelligent robots, automated devices, or other autonomous entities that act and interact in pursuit of their individual and collective preferences and goals. In accordance with their various essential attributes, such as beliefs, knowledge, and abilities to observe, learn, remember, reason and communicate, rational agents plan and execute *strategies* (in AI, often also called “policies”) to achieve their goals. The common environment in which agents act and interact is commonly called a *multiagent system*. This generic concept covers a variety of entities, including teams of agents, social groups and networks, organizations, markets, and entire societies. Nowadays, the modeling and analysis of systems of rational agents is a broad interdisciplinary field of extensive research, bringing together ideas, approaches, and methods from humanities, social sciences, artificial intelligence, decision and game theory, mathematics, and computer science.

Agents and groups (also often called “coalitions”) of rational agents in multiagent systems typically have competing, sometimes adversarial to each other, interests and goals. When the goals are purely quantitative, such systems are modeled as *multiplayer games*, studied in game theory, where rational players act so as to optimize their payoffs. Agents may also have *purely qualitative* (“win–lose”) or combined goals. A fundamental aspect of reasoning both *within* and *about* multiagent systems is the reasoning about abilities of agents and coalitions to design and execute *strategies* to achieve their goals, or to prevent others from achieving theirs, hereafter called (*multiagent*) *strategic reasoning*.

Importantly, the behavior of rational agents is usually guided not only by their individual and group goals, but also by those of the others, as well as by common (e.g., societal) goals. Furthermore, their behavior is constrained by various *individual, collective, and societal norms* (such as rights, obligations, and prohibitions). The totality of common goals, attitudes, norms, and other constraints under which agents act and interact I will call *socially interactive context*. That context leads to a complex interplay of *cooperation* and *competition* between the agents in the system.

Since the 1980s, formal logic has emerged as a useful and efficient methodological and technical tool for modeling, analysis, and reasoning about multiagent rational interaction. A rich variety of logical frameworks for strategic reasoning in, and about, multiagent



Citation: Goranko, V. Logics for Strategic Reasoning of Socially Interacting Rational Agents: An Overview and Perspectives. *Logics* **2023**, *1*, 4–35. <https://doi.org/10.3390/logics1010003>

Academic Editor: Fenrong Liu

Received: 26 June 2022

Revised: 25 November 2022

Accepted: 19 January 2023

Published: 6 February 2023



Copyright: © 2023 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

systems have been developed. Early formal logical systems capturing agents' abilities were developed with philosophical motivations and applications in mind, including Parikh's Game Logic [1], where he also suggested perhaps the first applications of logic to multi-agent social scenarios, such as cake-cutting, exchanging secrets (see further Example 10), etc.; Brown's modal logic of ability [2]; and Belnap and Perloff's STIT logics [3]. In the late 1990s–early 2000s, two important logical systems for multiagent strategic reasoning were introduced independently: Pauly's *coalition logic* CL, introduced in [4,5], and Alur, Henzinger, and Kupferman's *alternating time temporal logic* ATL, introduced first in [6], and then, in its final version, in [7] (cf. [8,9] for more recent and detailed discussion on these and related works).

Some remarks¹ are due on conceptual issues arising in regards to the terminology used in this paper, and also generally adopted in the literature on the topic. To begin with, let me emphasize that other, not discussed in the present work important concepts related to rational agency, such as beliefs, desires, intentions, goals, and commitments have been extensively explored in the philosophy and AI literature. Among the most influential such approaches are [10–12], where (especially in the latter) agents' goals and what they imply are explicitly indicated in the language. While being important and influential, these works are not about *strategic reasoning* regarding the agents' goals and their achievement which is the focus of the present paper, so they will not be discussed further here.

Now, let me explain the meaning and use of some key concepts and terms used in this paper, to avoid possible misunderstandings, especially on the background of the alternative approaches mentioned above. First, the use of the word "goal" in the context of these logical frameworks and formal systems should be understood in a precise and possibly somewhat narrow sense. Agents' goals are not some intentional objects permanently or temporarily ascribed to the agents, or intrinsically associated with them. The agents' goals in the present work are not explicitly specified as such in the language, but are determined by the context and identified with properties (respectively, sets) of states in the model—or, more generally, with properties (respectively, sets) of plays in that model—which the agents are assumed to pursue to realize with their strategic behavior. These sets of states, or plays, are represented in the language by formulae which extensionally specify them by means of their formal semantics. Thus, the logical languages considered here cannot express statements such as "the agent *a* has as its goal to realize *G*", nor even "the goal of agent *a* is to realize *G*", but rather "*the agent a has a strategy that guarantees (under the assumed interactive context) the truth of the formula G*", where *G* is regarded as "*a*'s goal" in this context. "Coalitional goals" should be understood likewise. Thus, agents' goals are only contextually specified and can freely change, even within the same formalized statement, but this should be taken only as a terminological choice, rather than a conceptual commitment to what "goal" means. Next, "rationality" of agents here should also be understood in a concrete and precise sense: once a goal (in the sense described above) is assigned to a *rational* agent, that agent is assumed to act in pursuit of that goal, and that assumption naturally constrains the agent's behavior. This is all that "rationality" is assumed to mean here. Lastly, the phrase "socially interactive context" should also be considered merely as a convenient and sufficiently intuitive term, rather than a precise social concept. Though not precisely defined, that term is hopefully well illustrated in the various logical frameworks presented here.

There are other interesting and important approaches to logics for games and strategic behavior, which—mostly for practical reasons—I will only mention here, but will not be discussed further in this paper:

- The *STIT-based approach*, building on the theory of "*Seeing To It That*" (STIT), originating from the seminal work of Belnap and Perloff [3]. For exploring the closer links of the present paper with STIT, the reader is referred to [13] and the overview chapter [14] on using STIT for strategic reasoning, as well as to [15] for temporal extensions of STIT and applications to normative reasoning, to [16] for using STIT for reasoning about social influence, and to [17] for providing semantics of temporal STIT in concurrent game models used in this paper.

- The *dynamic logic* based approach, originating from Parikh's Game Logic [1], and more recently explored further in [18,19] and other related works.

For further references and a broader view on logics for analyzing games, I also refer the reader to [20].

Now, a few words about the main precursors of the logics presented here. The logic CL was introduced in [4,5] with the explicit intention to formalize reasoning about *one-step* (*local*) strategic abilities of agents and coalitions, i.e., strategic abilities to guarantee the achievement of their explicitly specified goals in the immediate outcome of their (individual or collective) action, regardless of the respective actions of all remaining agents. The logic ATL was initially introduced in [6,7] as a logical formalism for formal specification and verification of open (interacting with environment) computer systems where the agents represent concurrently executed processes. Still, it was gradually adopted as one of the most popular and standard logical systems for reasoning about *long-term* strategic abilities of agents and coalitions in extended concurrent multiplayer games. Technically, ATL can be described as an extension of CL with the long-term temporal operators G, F, and U, which were previously adopted, inter alia, in the branching-time temporal logic CTL (which can be regarded as a single-agent fragment of ATL). Both CL and ATL feature a special type of modal operators, denoted by $[C]$ in CL, respectively by $\langle\langle C \rangle\rangle$ in ATL, parameterized with a coalition of agents C , such that, given a coalitional goal of C formalized as (ensuring the truth of) a formula ϕ , the formula $[C]\phi$ 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, in every outcome play (for ATL) that can occur when the agents in C execute their strategies in σ_C , regardless of the choices (strategic or not) of actions of the agents that are not in C .

Thus, both CL and ATL capture reasoning about *unconditional strategic abilities* of agents within coalitions, acting in full cooperation with each other in pursuit of their coalitional goals, aiming to succeed against *any* possible behavior of their opponents, which are thus regarded as their complete adversaries (in the context of CL) or as randomly behaving environment (in the original context of ATL). This, however, is a rather narrow and extreme perspective, as agents in the real world are seldom either complete allies or complete adversaries to each other. In reality, all agents acting in the system (except for the environment, or absolute adversaries) have their own goals and act in pursuit of their fulfillment, rather than just to prevent the proponents from achieving their goals. Thus, the strategic interactions of rational agents usually involve a complex interplay of *cooperation* and *competition*, both driven by the individual and collective goals of all agents, which may be partly or fully allied or adversarial to each other. To adequately model and capture the behavior and interactions of rational agents within a socially interactive context requires much richer and more versatile formal logical frameworks than those provided by the logics CL, ATL and their close variations. Several richer frameworks have been proposed more recently, including *strategy logics* [21–24]; the *socially friendly coalitional logic* SFCL and the *group protecting coalitional logic* GPCL [25], with the latter subsequently extended with temporal operators for long-term goals to the *temporal logic of coalitional goal assignments* TLPGA in [26]; the *logic for conditional strategic reasoning* ConStR (cf. [27,28]); and several recently proposed more special-purpose logics for strategic reasoning, focusing on know-how strategies [29], ethical dilemmas [30], responsibility [31], blameworthiness [32], etc.

This paper aims to provide an accessible and not too technical, yet fairly detailed, overview of the recent developments in the area and to outline some perspectives and challenges for further research. To keep the paper within reasonable scope and size, I will not discuss here the aspects of socially interactive context related to agents' knowledge, nor to normative constraints, but will focus on the strategic reasoning guided only by various patterns of cooperation and competition between agents and coalitions. I will usually assume here complete knowledge of the agents about the system and their full observability (with or without memory) and ability to communicate and coordinate during the "play". These assumptions are sometimes not justified, and then substantial further complications arise, both conceptual and computational, with respect to analyzing the

agents' strategic abilities and designing their strategies in the context of agents' incomplete or imperfect information. For more on these issues, the reader is referred to, e.g., [8,9,33], and the entire volume [34], where the latter is published.

Structure of the paper:

Section 2 provides some technical preliminaries on concurrent game models, including plays and strategies in them, and provides an illustrative example. Then, Section 3 presents a brief informal preview of the main logical systems presented in the next five sections: the basic coalition logic CL in Section 4, its temporal extension ATL in Section 5, the logic of conditional strategic reasoning ConStR in Section 6, the socially friendly coalition logic SFCL in Section 7, and the logic of local coalitional goal assignments LCGA and its temporal extension TLCGA in Section 8. Eventually, Section 9 proposes a uniform logical framework, called *basic strategy logic*, which involves variables over strategies and quantification over them, and thus enables uniform embedding of all other strategic operators mentioned here. I end with brief concluding remarks in Section 10.

2. Technical Preliminaries

The reader is expected to have a basic background in propositional classical and modal logics, within the introductory chapters of, e.g., [35]. In addition, this section provides technical preliminaries needed for the definitions and understanding of the formal semantics of the various strategic operators and logics presented and discussed further. Most of it can be omitted by a reader who is only interested in the informal aspects of the topics, and can be consulted only if, and when, necessary for understanding the further content. For more technical details on these preliminaries, the reader is referred to [8,9], as well as [36] (Chapter 9) on concurrent game models and ATL.

2.1. Concurrent Game Models

Consider a fixed finite nonempty set of **agents** $\text{Agt} = \{a_1, \dots, a_n\}$, also called here **players**, and a fixed countable set of **atomic propositions** Prop . Subsets of Agt will also be called **coalitions**. The **Cartesian product** of a family of sets $\{X_a\}_{a \in \text{Agt}}$ is the set of tuples (x_1, \dots, x_n) where $x_i \in X_{a_i}$ or each $i = 1, \dots, n$, denoted, as usual, by $\prod_{a \in \text{Agt}} X_a$.

Definition 1. Let O be any nonempty set. A **(strategic) game form for the set of agents** Agt **over the set of outcomes** O is a tuple

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

where

- Act is a nonempty set of **actions**;
- $\text{act} : \text{Agt} \rightarrow \mathcal{P}^+(\text{Act})$ is a mapping assigning to each $a \in \text{Agt}$ a nonempty set $\text{act}(a)$ 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 (CGM)** is a tuple

$$\mathcal{M} = (\text{Agt}, S, \text{Act}, g, \text{Prop}, L)$$

where

- S is a nonempty set of **(game) states**;
- Act is a nonempty 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 ;
- Prop is a countable set of **atomic propositions**;

- $L : S \rightarrow \mathcal{P}(\text{Prop})$ is a **labeling function**, assigning to every state in S the set of atomic propositions true at that state.

Thus, 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 . We can now define the **global action function** $\text{act} : \text{Agt} \times S \rightarrow \mathcal{P}^+(\text{Act})$ by setting $\text{act}(a, w) ::= \text{act}_w(a)$. We also define the set $\text{Act}_a ::= \bigcup_{w \in S} \text{act}_w(a)$ of **possible actions** for a .

Given a concurrent game model $\mathcal{M} = (\text{Agt}, S, \text{Act}, g, \text{Prop}, L)$, we define the following.

- An action profile $\zeta \in \prod_{a \in \text{Agt}} \text{Act}_a$ is **available at the state** w if it consists of locally available actions, i.e., if $\zeta \in \prod_{a \in \text{Agt}} \text{act}_w(a)$.
The set of all action profiles that are available at w will be denoted by ActProf_w .
- $\text{out}_{\mathcal{M}}$ is the **global outcome function** assigning to every state w and an action profile $\zeta \in \text{ActProf}_w$ the 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.
- More generally, 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 . The joint action $\zeta|_C$ is **locally available at the state** w iff every individual action in it is locally available for the respective agent in w .
- For any $w \in S$, $C \subseteq \text{Agt}$, and joint action ζ_C that is available at w , we define **the set of possible outcomes from the application of the joint action ζ_C at the state w** :

$$\text{Out}[w, \zeta_C] = \{u \in S \mid \text{there exists } \zeta \in \text{ActProf}_w \text{ such that } \zeta|_C = \zeta_C \text{ and } \text{out}(w, \zeta) = u\}.$$

In the special case when $C = \text{Agt}$, ζ_{Agt} is a full available action profile, so $\text{Out}[w, \zeta_{\text{Agt}}]$ consists of the single outcome $\text{out}(w, \zeta_{\text{Agt}})$.

Remark 3. Note that we only consider deterministic models, where every available action profile produces a single outcome. This is not an essential restriction, because every nondeterministic model can be made deterministic by adding a fictitious additional agent (say, nature) that settles the nondeterminism with its actions which settle the outcome. However, it should be noted that whether nondeterministic models are considered or not will make a difference on the validities of coalition logic CL, as shown in Section 4.4, and in the various other logics where strategic abilities of coalitions can be expressed.

It is sometimes more convenient to eliminate the “local view” from the definition of concurrent game models, given by the game map assigning the local strategic game forms, and to redefine these models globally, as follows:

$$\mathcal{M} = (\text{Agt}, S, \text{Act}, \text{act}, \text{Out}, \text{Prop}, L)$$

where the components are defined as above. Clearly, the two definitions are equivalent and will be used interchangeably.

Example 4 (Concurrent game model: Adam and Eve in hotel Life). Figure 1 pictures a concurrent game model formalizing the following story, which the reader is advised to take at face value and with a grain of salt. While it has a clearly allegoric flavor, it is not meant to formalize some meaningful (real or fiction) life story, but just to give an intuitively appealing and sufficiently nontrivial, yet not too complicated, example of a concurrent game model on which to illustrate the components of that concept. Thus, questions about the meaning or idea behind this or that state or transition do not necessarily have good answers.

Two agents, **Adam** and **Eve**, live in hotel “Life”. It has (at least) three rooms: R_1 , R_2 , and R_3 . Initially, both Adam and Eve are in room R_1 . Every day, each of them is able to choose either to stay in the same room (action stay, denoted by s), or to move to another room (action move, denoted by m), with some restrictions described further. Whichever room each of them chooses, they stay there for the night, and then each chooses again to either stay in the same room, or move to another

room, as specified in the model. In some cases, they also have the option to “retreat” (action retreat, denoted by r).

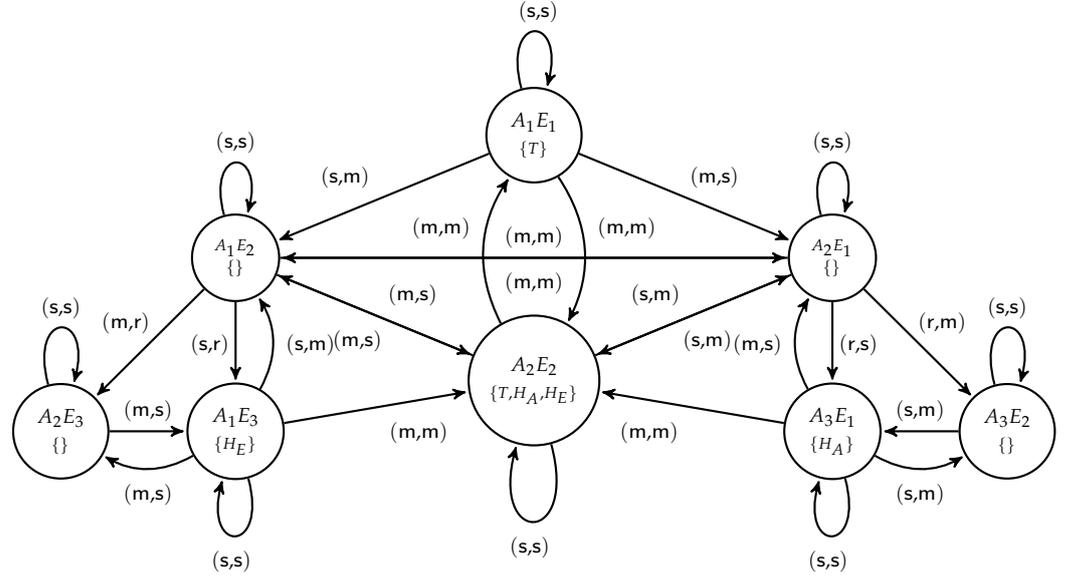


Figure 1. A concurrent game model: Adam and Eve in hotel Life.

Here are the formal components of the model, with some brief explanations:

- The set of agents is $\text{Agt} = \{\text{Adam}, \text{Eve}\}$.
- The set of game states is $\text{St} = \{A_1E_1, A_1E_2, A_2E_1, A_2E_2, A_3E_1, A_1E_3, A_3E_2, A_2E_3\}$ and the names of the states represent the current locations of the agents, e.g., A_1E_2 means that **Adam** is in R1 and **Eve** is in R2, etc. Thus, the names of the states can also be interpreted as pairs of atomic propositions saying “**Adam** is in room X” and “**Eve** is in room Y”. These atomic propositions, however, will not feature in the example.
- There are three atomic propositions of interest for us: $\text{Prop} = \{T, H_A, H_E\}$, where:
 - T , which is true in a state iff **Adam** and **Eve** are in the same room (“together”), i.e., in A_1E_1 and A_2E_2 , as indicated in their labels;
 - H_A , meaning “**Adam** is happy”, true in the states where H_A is listed in the label;
 - H_E , meaning “**Eve** is happy”, true in the states where H_E is listed in the label.
- The global set of actions is $\text{Act} = \{\text{stay}, \text{move}, \text{retreat}\}$, denoted in the figure, respectively, by s, m, r . The transitions caused by the respective ordered pairs of actions are indicated on the figure (the first for **Adam**, and the second for **Eve**). For example, if at state A_1E_1 both **Adam** and **Eve** choose to stay, then the transition, labeled with (s, s) , is looping back to the same state, whereas if **Adam** chooses to stay and **Eve** chooses to move, the transition, labeled with (s, m) , leads to the state A_1E_2 , etc. These transitions also define the global outcome function Out .
- The global action function act (defining the available actions to every agent at every state) can be readily extracted from the figure. For example, $\text{act}(\text{Adam}, A_1E_2) = \{s, m\}$, $\text{act}(\text{Eve}, A_1E_2) = \{s, m, r\}$, etc.
- Lastly, the labeling function $L : \text{St} \rightarrow \mathcal{P}(\text{Prop})$ is explicitly given in the figure, where the label of each state is given in $\{\dots\}$.

2.2. Plays and Strategies in Concurrent Game Models

Given a (globally defined concurrent game model $\mathcal{M} = (\text{Agt}, S, \text{Act}, \text{act}, \text{Out}, \text{Prop}, L)$), a **partial play**, or a **history** in \mathcal{M} is either an element of S or a finite word of the form:

$$h = w_0\zeta_0w_1\dots w_{n-1}\zeta_{n-1}w_n$$

where $w_0, \dots, w_n \in S$ and for each $i < n$, ζ_i is a locally available action profile in ActProf_{w_i} .

Some partial plays in Example 4 are:

A_1E_1 ; $A_1E_1(s, m)A_1E_2$; $A_1E_1(s, m)A_1E_2(m, m)A_2E_1(s, m)A_2E_2$; etc.

The last state in a history h will be denoted by $l(h)$. The set of histories in \mathcal{M} is denoted by $\text{Hist}(\mathcal{M})$. The sequence of states $w_0w_1\dots w_n$ obtained by forgetting the action profiles in the history h is called the **(initial) path induced by h** and denoted $\text{path}(h)$.

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 $\text{Hist}(\mathcal{M})$ an action $\sigma_a(h)$ from $\text{act}(a, w_n)$. A strategy σ_a is **path-based**, if it assigns actions only based on the path generated by the history (not taking into account the intermediate action profiles), i.e., $\sigma_a(h) = \sigma_a(h')$ whenever $\text{path}(h) = \text{path}(h')$. 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 . For every joint strategy Σ_C and a history h , we denote by $\Sigma_C(h)$ the joint action prescribed by Σ_C on h .

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 $w_{n+1} = \text{out}(w_n, \zeta_n)$, and

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

The infinite word $w_0w_1w_2\dots$ obtained by forgetting the action profiles in this infinite play is called the **path induced by Σ at w** , denoted $\text{path}(\Sigma, w)$.

Here are some simple, informally described, positional strategy profiles and the plays induced by them in Example 4:

- Consider the strategy σ_1 which prescribes the action s if Adam and Eve are currently together, or else m , if possible, otherwise, again, s . If both Adam and Eve follow that strategy starting from A_1E_1 , the induced play is
 $\text{play}(A_1E_1, (\sigma_1, \sigma_1)) = A_1E_1(s, s)A_1E_1(s, s)\dots$
 Respectively, the induced play starting from A_1E_2 is
 $\text{play}(A_1E_2, (\sigma_1, \sigma_1)) = A_1E_2(m, m)A_2E_1(m, m)A_1E_2(m, m)\dots$
- Consider the strategy σ_2 which prescribes the action s if the player is currently happy, or else r if possible, otherwise m , if possible, otherwise, s .
 If both Adam and Eve follow σ_2 starting from A_1E_1 , the induced play is
 $\text{play}(A_1E_1, (\sigma_2, \sigma_2)) = A_1E_1(m, m)A_2E_2(s, s)A_2E_2(s, s)\dots$
- If Adam follows σ_1 and Eve follows σ_2 , the induced play starting from A_1E_1 is
 $\text{play}(A_1E_1, (\sigma_1, \sigma_2)) = A_1E_1(s, m)A_1E_2(m, r)A_1E_3(m, s)A_2E_3(m, s)A_1E_3(m, s)\dots$
- Lastly, if Adam follows σ_1 and Eve follows the strategy σ_3 which prescribes the action s only if both players are currently happy or if no other action is available, or else r if possible, otherwise m , then the induced play starting from A_1E_1 is this happy ending:
 $\text{play}(A_1E_1, (\sigma_1, \sigma_3)) = A_1E_1(s, m)A_1E_2(m, r)A_1E_3(m, m)A_2E_2(s, s)A_2E_2(s, s)\dots$

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

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

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 paths obtained from the plays in $\text{Plays}(w, \Sigma, C)$. Since these only depend on the strategies assigned to players in C , I 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 necessarily for all other players.

3. A Variety of Strategic Abilities

This short section is a brief informal preview of the types of strategic abilities and logical operators and systems formalizing them that will be presented further in the paper.

Most of the natural claims about strategic abilities of agents and coalitions (for convenience, hereafter I will often treat individual agents as singleton coalitions) can be stated in terms of existence of a strategy profile Σ (in particular, an action profile ζ) that satisfies certain requirements formalizing these abilities. Here is a natural hierarchy of types of strategic abilities and associated claims about them, which will be discussed further.

1. *Strictly competitive and unconditional*, where all agents, respectively, coalitions, act only in pursuit of their own goals and can be assumed to regard all others either as adversaries, or as behaving randomly. An alternative way of thinking here is that these are abilities of a given agent, respectively, coalition, to achieve goals independently of the actions of all other agents. Both interpretations make good sense in the context of this paper. The typical claim for such kind of *local* (immediate, one step) abilities is:

“The coalition A has a joint action to ensure satisfaction of the coalitional goal of A in every outcome state that may result from that joint action”.

This is the informal semantic reading of the strategic operator $[A]$ in the coalition logic CL [4,5], which will be presented in Section 4.

Respectively, here is a typical *long-term* strategic ability claim:

“The coalition A has a joint strategy to ensure satisfaction of the coalitional goal of A in every outcome play resulting from A following that strategy”.

This is the informal reading of the strategic operator $\langle\langle A \rangle\rangle$ in the alternating time temporal logic ATL and family [7], which will be presented in Section 5.

2. *Competitive, but conditional on the other agents' expected actions*, where coalitions (respectively, agents) still act only in pursuit of their own goals, but, when deciding on their course of action, they take into account the goals and the respective expected actions of the other players, so they are not treated as adversaries (or, behaving randomly), but as rational agents pursuing their own goals. Here is a typical such claim expressing *conditional strategic ability*:

“For some (or, every) joint action of the coalition A that ensures satisfaction of its goal γ_A , the coalition B has a joint action of its own to ensure satisfaction of its goal γ_B .”

There are (at least) two naturally arising readings of such conditional claims, as “*proactive ability*” and as “*reactive ability*”, and two respective versions of local strategic operators formalizing them. These were first introduced in [27] (respectively, as “*de dicto*” and “*de re*” abilities) and further studied in [28]. Conditional abilities will be discussed in more detail in Section 6.

3. *Socially cooperating abilities*, where agents and coalitions still act in pursuit of their own goals, but when deciding on their course of action take into account the goals of other agents in the system and make allowance, if possible, for their satisfaction, too. Thus, agents and coalitions are assumed, i.e., not only rational but also cooperative, whenever possible to reconcile their interests with those of the others. Two natural examples of strategic operators formalizing such abilities are:

- (a) the “*cooperative ability*” operator O_c , again introduced and studied in [27,28], which, when applied to (disjoint) coalitions A and B with respective goals ϕ and ψ , formalizes the statement saying that

“ A has a joint action σ_A which guarantees the satisfaction of ϕ and enables B to also apply a joint action σ_B that guarantees ψ ”.

This will be presented in more detail in Section 6.

- (b) The “socially friendly” coalitional operator SF, introduced and studied in [25], which is a somewhat more general version of O_c , informally says that
- “A has a joint action σ_A that guarantees satisfaction of its goal and also enables the complementary coalition \bar{A} to realize any one of a list of secondary goals by applying a respectively suitable joint action”.

The operator SF and the logic SFCL built on it will be presented in more detail in Section 7.

4. *Abilities for cooperation, protecting the interests of agents and coalitions.* These capture the idea, complementary to that of SFCL, that, while socially responsible rational agents and coalitions contribute with their individual and collective actions to the society of all agents, they wish to perform that in a way that protects their individual and collective interests and goals. Such abilities are expressed by means of the *coalitional goal assignment operator* $\langle \cdot \rangle$, introduced in [25] as a local operator and extended with temporalized long-term goals in [26]. A coalitional goal assignment γ is a mapping which assigns a goal formula to each possible coalition, and the operator $\langle \gamma \rangle$ formalizes the claim that there is a strategy profile σ of Agt, the restriction of which to every coalition C is a joint action σ_C that guarantees the satisfaction of its goal $\gamma(C)$ regardless of any possible behavior of the remaining agents. The coalitional goal assignment operator and the logics LCGA and TCGA built on it will be presented in more detail in Section 8.

4. Coalition Logic and Unconditional Local Strategic Reasoning

The case of *strictly competitive (noncooperative) and unconditional abilities*, presented in this section, is the extreme case of strategic abilities, where agents, respectively, coalitions, act only in pursuit of their own goals and regard all others either as adversaries or as behaving randomly. In this section, I will present briefly the basic multiagent logic CL formalizing these abilities and will illustrate its use with some examples.

4.1. The Basic Logic for Coalitional Strategic Reasoning CL

As noted earlier, a typical claim for such kind of local strategic abilities is of the following type: “The coalition C has a joint action to guarantee satisfaction of the coalitional goal of C in every outcome state that may occur as a result from that joint action”.

This is the informal semantic reading of the strategic operator $[C]$ in the coalition logic CL introduced in [4]; cf. also [5,37]. CL extends the classical propositional logic with *coalitional strategic modal operators* $[C]$, for any coalition of agents C. The formulae of CL are defined by the following formal grammar:

$$\varphi ::= p \mid \top \mid \neg\varphi \mid (\varphi \vee \varphi) \mid [C]\varphi$$

All other standard propositional connectives and constants, including \perp , \wedge , \rightarrow , and \leftrightarrow , are defined as usual. I will write $[i]$ instead of $[\{i\}]$.

4.2. Expressing Claims about Strategic Abilities in CL

Here are some examples of formalizing statements about agents’ strategic abilities in CL, based on the story in Example 4. The reader may wish to check whether they are all intuitively true in the model in Figure 1, say, at state A_1E_1 .

1.

$$[\mathbf{Eve}]H_E \rightarrow \neg[\mathbf{Adam}]\neg H_E$$

If Eve has an action ensuring that she becomes happy, then Adam cannot prevent Eve from reaching a state where she is happy.

2. The statement above naturally generalizes (with Win being an atomic proposition) to

$$[i]\text{Win} \rightarrow \neg[\text{Agt} \setminus \{i\}]\neg\text{Win}$$

If the agent i has an action to guarantee reaching a “winning” outcome, then the coalition of all other agents cannot prevent i from reaching a “winning” outcome.

It should be intuitively clear that this statement expresses a valid principle of CL. Indeed, we will see shortly that this is the case.

- 3.

$$\neg[\mathbf{Adam}]T \wedge \neg[\mathbf{Eve}]T \wedge [\{\mathbf{Adam}, \mathbf{Eve}\}]T$$

Neither **Adam** nor **Eve** has an action ensuring that they stay together, but they have a joint action ensuring that.

- 4.

$$\neg[\mathbf{Adam}]H_A \wedge \neg[\mathbf{Eve}]H_E \wedge [\{\mathbf{Adam}, \mathbf{Eve}\}](H_A \wedge H_E)$$

Neither **Adam** nor **Eve** has an action ensuring happiness for himself / herself but they have a joint action ensuring happiness for both.

- 5.

$$\neg[\mathbf{Adam}](H_A \wedge \neg[\mathbf{Eve}]H_E)$$

Adam cannot act so as to ensure at the outcome state that both **Adam** is happy and **Eve** does not have an action to ensure reaching her happiness.

- 6.

$$[\{\mathbf{Adam}, \mathbf{Eve}\}](\neg[\mathbf{Adam}]H_A \rightarrow [\mathbf{Eve}]H_E)$$

Adam and **Eve** can act jointly so that at the outcome state **Adam** has an action to ensure reaching his happiness only if **Eve** has an action to ensure reaching her happiness.

- 7.

$$[\mathbf{Adam}][\mathbf{Eve}](H_A \wedge H_E \wedge \neg[\mathbf{Adam}]H_E \wedge \neg[\mathbf{Eve}]H_A)$$

Adam can act to ensure (at the outcome state), that **Eve** can then act to ensure that they are both happy and that each of them can act so as to keep the other happy.

Thus, we see that the language of CL allows expressing meaningful and nontrivial claims about agents’ local (immediate) strategic abilities. However, CL does not allow expressing long-term (temporal) goals and claims, e.g., that **Adam** and **Eve** have a joint strategy to stay together forever, or to eventually reach happiness. For that, a more expressive logic is needed, also involving temporal operators. It is introduced in Section 5.

4.3. Formal Semantics of CL

The formulae of CL are interpreted in concurrent game models. Consider a CGM $\mathcal{M} = (\text{Agt}, S, \text{Act}, \text{act}, \text{Out}, \text{Prop}, L)$. **Truth of a CL-formula ψ at a state s in \mathcal{M}** , denoted $\mathcal{M}, s \models \psi$, as in modal logic, is defined by structural induction on formulae via the clauses:

- $\mathcal{M}, s \models p$ iff $p \in L(s)$, for $p \in \text{Prop}$.
- $\mathcal{M}, s \models \neg\phi$ iff $\mathcal{M}, s \not\models \phi$.
- $\mathcal{M}, s \models (\phi_1 \vee \phi_2)$ iff $\mathcal{M}, s \models \phi_1$ and $\mathcal{M}, s \models \phi_2$.
- $\mathcal{M}, s \models [C]\phi$ iff there exists a joint action σ_C available at s , such that $\mathcal{M}, u \models \phi$ for each $u \in \text{Out}[s, \sigma_C]$.

Note that this definition makes no implicit assumption that the agents in C know the goals of the remaining agents, nor the actions that they may possibly use to achieve these goals.

Here are examples of truth of some of the CL formulae listed in Section 4.2 in the CGM for Example 4, given in Figure 1, which are left for the reader to formally verify.

- $\mathcal{M}, A_1E_1 \models \neg[\mathbf{Adam}]T$; $\mathcal{M}, A_1E_1 \models \neg[\mathbf{Eve}]T$.
- $\mathcal{M}, A_1E_1 \models [\{\mathbf{Adam}, \mathbf{Eve}\}]T$.
- $\mathcal{M}, A_1E_1 \models \neg[\mathbf{Adam}]H_A \wedge \neg[\mathbf{Eve}]H_E \wedge [\{\mathbf{Adam}, \mathbf{Eve}\}](H_A \wedge H_E)$.

- $\mathcal{M}, A_1 E_1 \models [\{\mathbf{Adam}, \mathbf{Eve}\}]([\mathbf{Eve}]H_E \wedge \neg[\mathbf{Adam}]H_A \wedge \neg[\{\mathbf{Adam}, \mathbf{Eve}\}](H_A \wedge H_E))$.

4.4. Axiomatizing the Validities of CL

A CL formula ϕ is called:

- **(Logically) valid** if $\mathcal{M}, s \models \phi$ for every CGM \mathcal{M} and a state $s \in \mathcal{M}$.
- **Satisfiable** if $\mathcal{M}, s \models \phi$ for some CGM \mathcal{M} and a state $s \in \mathcal{M}$.

Pauly [4,5] obtained a complete axiomatization of the valid formulae of CL, by extending the classical propositional logic PL with the following axiom schemes and rule:

- **Agt-maximality:** $\neg[\emptyset] \neg\phi \rightarrow [\mathbf{Agt}] \phi$
This axiom says that the grand coalition Agt can act collectively so as to guarantee any goal that is satisfied in at least one outcome state. Note that the validity of this axiom presupposes that the models are deterministic, implying that the coalition of all agents Agt can enforce any particular possible outcome.
- **Safety:** $\neg[C] \perp$
No coalition has the ability to ensure the falsum will become true.
- **Liveness:** $[C] \top$
Every coalition has the ability to ensure the truth will become true.
- **Superadditivity:** for any $C_1, C_2 \subseteq \mathbf{Agt}$ such that $C_1 \cap C_2 = \emptyset$:

$$([\mathbf{C}_1] \varphi_1 \wedge [\mathbf{C}_2] \varphi_2) \rightarrow [\mathbf{C}_1 \cup \mathbf{C}_2] (\varphi_1 \wedge \varphi_2)$$

This axiom scheme says that two disjoint coalitions, each of which has a joint action to guarantee satisfaction of their own goal, can join forces (simply by each of them following their respective joint action) to guarantee satisfaction of both goals.

- **[C]-monotonicity Rule:**

$$\frac{\varphi_1 \rightarrow \varphi_2}{[C] \varphi_1 \rightarrow [C] \varphi_2}$$

Some important derivable validities include:

- **Outcome Monotonicity:** $[C] (\varphi_1 \wedge \varphi_2) \rightarrow [C] \varphi_1$.
This is derived directly by applying the [C]-monotonicity rule.
- **Coalition Monotonicity:** $[C_1] \varphi \rightarrow [C_1 \cup C_2] \varphi$.
This is derived by applying the superadditivity axiom scheme to the coalitions C_1 and $C_2 \setminus C_1$ with respective goals φ and \top . Note that this validity together with Agt-maximality also implies the validity of $[\mathbf{Agt}] \varphi \vee [\mathbf{Agt}] \neg\varphi$, for any formula φ .
- **Coalition complementarity:** $[C] \varphi \rightarrow \neg[\overline{C}] \neg\varphi$, where $\overline{C} = \mathbf{Agt} \setminus C$.
This is derived by applying the superadditivity axiom scheme to the coalitions C and \overline{C} with respective goals φ and $\neg\varphi$, and then the safety axiom scheme.

The problem of deciding whether a given formula in the logic CL is valid, respectively, satisfiable, is decidable [4,5]. This and other technical results about CL are subsumed by respective results for the more expressive logic ATL, presented in the next section.

5. Logics for Unconditional Long-Term Strategic Reasoning

The logic CL can only capture strategic reasoning about abilities to achieve *local*, i.e., one-step goals. It can be naturally extended with temporal operators to reason about *long-term* unconditional strategic abilities, of the type:

“The coalition C has a joint strategy to guarantee achievement of the coalitional goal of C in every outcome play resulting from C following that strategy”.

This is the informal reading of the strategic operator $\langle\langle C \rangle\rangle$ in the *alternating time temporal logic* ATL and family, introduced (independently from the logic CL) in [7].

I will present and illustrate here the family of ATL-related logics. For more detailed and technically involved surveys on these logics, see [8,9].

5.1. The Alternating-Time Temporal Logic ATL*.

The full alternating-time temporal logic ATL*, introduced in [7], involves long-term temporal operators, the standard repertoire being “next-time” X, “always” G, and “until” U.

The strategic operator in ATL* is $\langle\langle C \rangle\rangle\phi$, denoting the claim that

“The coalition C has a joint strategy that guarantees the satisfaction of the goal ϕ on every outcome play induced by that joint strategy”,

where ϕ is any formula of ATL*. The language of ATL* is formally defined as follows. It involves two sorts of formulae: **state formulae**, that are evaluated at game states, and **path formulae**, that are evaluated on game plays. These are, respectively, defined by the following grammars, where $p \in \text{Prop}$ and $C \subseteq \text{Agt}$:

$$\begin{aligned} \text{State formulae: } \phi &::= p \mid \top \mid \neg\phi \mid (\phi \vee \psi) \mid \langle\langle C \rangle\rangle\gamma. \\ \text{Path formulae: } \gamma &::= \phi \mid \neg\gamma \mid (\gamma \vee \delta) \mid X\gamma \mid G\gamma \mid (\gamma U \delta). \end{aligned}$$

Thus, the path formulae are used to formalize goals.

The language of ATL* is too expressive, and that both creates problems with its intuitive interpretation (see discussions on [38–40]) and raises the complexity cost of computing the truth of an ATL* to 2EXPTIME (cf. [7]). To alleviate these, the fragment ATL was also introduced in [7], which only contains state formulae, defined by the grammar:

$$\phi ::= p \mid \neg\phi \mid (\phi \vee \psi) \mid \langle\langle C \rangle\rangle X\phi \mid \langle\langle C \rangle\rangle G\phi \mid \langle\langle C \rangle\rangle(\phi U \psi),$$

Thus, the coalitional goals in ATL only involve simple patterns with a single temporal operator, which will be called *simple temporal goals*. We define $\langle\langle C \rangle\rangle F\phi$ as $\langle\langle C \rangle\rangle(\top U \phi)$.

The logic CL embeds into ATL as the fragment where the goals can only be of the type $X\phi$, for a state formula ϕ (cf. [41]), i.e., $[C]\phi ::= \langle\langle C \rangle\rangle X\phi$.

Hereafter, in this section, I will focus on the logic ATL.

5.2. Formal Semantics of ATL

Similar to CL, the formulae of ATL* are interpreted in concurrent game models. To keep the presentation less technical, I will only give the formal semantics for the fragment ATL, which consists only of state formulae; for full details of the semantics of ATL*, see, e.g., [8].

Consider a CGM $\mathcal{M} = (\text{Agt}, S, \text{Act}, \text{act}, \text{Out}, \text{Prop}, L)$. **Truth of an ATL-formula ψ at a state s in \mathcal{M}** , denoted $\mathcal{M}, s \models \psi$, is defined by structural induction on formulae. The propositional cases are similar to in CL, and the only new clauses are those for the strategic operators. They all follow the same pattern, which is essentially the semantic clause for $\langle\langle \dots \rangle\rangle$ in ATL*:

$\mathcal{M}, s \models \langle\langle C \rangle\rangle\gamma$ iff there exists a joint strategy Σ_C for C such that γ is true on every play in $\text{Plays}(s, \Sigma_C)$, i.e., every play starting at s and induced by Σ_C .

The specific clauses for the temporal operators defining the goal γ are as follows, where $s_0 = s$:

- $\mathcal{M}, s \models \langle\langle C \rangle\rangle X\phi$ iff there exists a joint strategy Σ_C such that $\mathcal{M}, s_1 \models \phi$ for every play $s_0, s_1, \dots \in \text{Plays}(s, \Sigma_C)$.
- $\mathcal{M}, s \models \langle\langle C \rangle\rangle G\phi$ iff there exists a joint strategy Σ_C such that $\mathcal{M}, s_i \models \phi$ for every play $s_0, s_1, \dots \in \text{Plays}(s, \Sigma_C)$ and $i \geq 0$.
- $\mathcal{M}, s \models \langle\langle C \rangle\rangle(\phi U \psi)$ iff there exists a joint strategy Σ_C such that for every play $s_0, s_1, \dots \in \text{Plays}(s, \Sigma_C)$ there is $i \geq 0$ for which $\mathcal{M}, s_i \models \psi$ and $\mathcal{M}, s_j \models \phi$ for all j such that $0 \leq j < i$.

Satisfiability and validity in ATL are defined similar to in CL.

It can be shown (cf., e.g., [36]) that memoryless strategies suffice for the semantics of ATL formulae, but that is no longer the case even when conjunctions of simple temporal goals are allowed. Indeed, consider the ATL* formula $\phi = \langle\langle a \rangle\rangle(Fp \wedge Fq)$ and a concurrent

game model in Figure 2 for the only agent a . Clearly, there is no positional strategy to satisfy ϕ at s_0 , as it must always prescribe the same action, L or R , there, so either s_1 or s_2 , but not both, will always be visited. On the other hand, there is a simple strategy for a using a bit of memory that can satisfy ϕ , e.g., prescribing the first time when s_0 is visited the action L , and the next time—the action R .

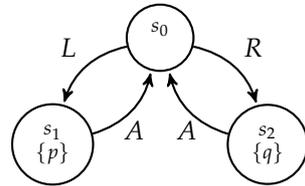


Figure 2. The formula $\langle\langle a \rangle\rangle(F p \wedge F q)$ is true at state s_0 only if the agent a can use some memory.

5.3. Expressing Claims about Strategic Abilities in ATL

Here are some examples of formalizing in ATL statements about agents' long-term strategic abilities.

1.

$$\langle\langle \mathbf{Eve} \rangle\rangle F H_E \rightarrow \neg \langle\langle \mathbf{Adam} \rangle\rangle G \neg H_E$$

If Eve has an action ensuring that she eventually becomes happy, then Adam cannot prevent Eve forever from reaching a state where she is happy.

It should be intuitively clear that this statement expresses a valid principle of ATL. Indeed, this is the case.

2.

$$\langle\langle \mathbf{Yin} \rangle\rangle G \text{ Safe} \wedge \langle\langle \mathbf{Yin} \rangle\rangle F \text{ Goal} \rightarrow \langle\langle \mathbf{Yin} \rangle\rangle (\text{Safe} U \text{ Goal})$$

If Yin has a strategy to keep the system in safe states forever and has a strategy to eventually achieve its goal, then Yin has a strategy to keep the system in safe states until it achieves its goal.

The formula above is not logically valid. Indeed, the strategies of Yin to keep the system in safe states forever and to eventually achieve its goal may be incompatible.

3.

$$\langle\langle \mathbf{Yin} \rangle\rangle G \text{ Safe} \wedge \langle\langle \mathbf{Yang} \rangle\rangle F \text{ Goal} \rightarrow \langle\langle \{\mathbf{Yin}, \mathbf{Yang}\} \rangle\rangle (\text{Safe} U \text{ Goal})$$

If Yin has a strategy to keep the system in safe states forever and Yang has a strategy to eventually reach a goal state, then Yin and Yang together have a strategy to stay in safe states until a goal state is reached.

Assuming that Yin and Yang are distinct agents, the formula above is logically valid, unlike the previous one. Indeed, this is due to the independence of the actions of the two agents, and hence of their abilities to execute their respective strategies.

4.

$$\langle\langle A \rangle\rangle F \langle\langle B \rangle\rangle G \neg \phi$$

The coalition A has a joint strategy to eventually ensure that the coalition B has a strategy to prevent ϕ from ever happening.

This example raises the question of how the semantics works in the case of nested strategic operators. Suppose, the coalitions A and B intersect and a is an agent in both of them. Then, the claim of the external strategic operator $\langle\langle A \rangle\rangle$ requires, inter alia, existence of a strategy for the agent a within a joint strategy σ_A for A that guarantees the eventual satisfaction of the subformula $\langle\langle B \rangle\rangle G \neg \phi$. However, when evaluating that subformula, to justify its truth, one has to identify a respective joint strategy σ_B for $\langle\langle B \rangle\rangle$. Now, the question arises whether the strategy of a within σ_B should not be assumed to be already fixed by σ_A or, conversely, whether the strategy of a within σ_A should not be assumed to be already fixed by σ_B . Note that the standard formal semantics for ATL* (in particular, of ATL) presented here does not impose any such constraints, but, rather, treats these strategies independently. That is, the

standard semantics of ATL^* does not commit the agents in A to the strategies they adopt in order to bring about the truth of the formula $\langle\langle A \rangle\rangle\gamma$. That creates some conceptual issues with the very concept of “strategy”, independently addressed in different ways in [38,39,42,43], where several proposals were made in order to incorporate strategic commitment or uncommitment and persistent strategies in the syntax and semantics of ATL^* . These raise a number of still open technical problems regarding constructing provably complete axiomatizations, proving decidability, and designing decision procedures for the variations of ATL and ATL^* mentioned above.

6. The Logic of Conditional Strategic Reasoning ConStR

Let us now look at the reasoning of (and about) an agent, *conditional on that agent’s knowledge of the goals and choices of actions of the other agents*. Note that, while knowledge is not explicitly included in the syntax of the logics considered here, it is implicit in their reasoning, or in the reasoning of the external observers reasoning about the agents’ abilities, and it is also implicitly assumed in the formal semantics.

6.1. Conditional Strategic Reasoning: An Informal Discussion

I first focus on a simple case: two agents, Alice and Bob, acting independently, and possibly concurrently with other agents. Alice has a goal, γ_A , to achieve. Suppose that Alice has several possible choices of action that would possibly, or certainly, ensure the achievement of her goal. Bob also has a goal, γ_B , of his own and has several possible actions (there may be other agents, besides Alice and Bob, also acting in pursuit of their own goals, but we will ignore them here). Now, based on his knowledge about Alice’s goal and possible choices of actions she may take towards that goal, Bob decides on his own choice of action in pursuit of his goal.

Here is a simple illustrating scenario, discussed in [27]. Alice and Bob are students at Downtown University. Alice is coming to campus today, to meet with her supervisor. Bob wants to meet with Alice somewhere on campus today. Alice may not know that, and they may have no direct communication. In turn, Bob may, or may not, know what Alice is going to do on campus. Now, using his knowledge of what, where, and when Alice intends to do, Bob wants to come up with a plan of where and how to meet her.

This calls for a *conditional strategic reasoning* with statements of the type:

*For some/every action of Alice that guarantees achievement of her goal γ_A ,
Bob has/does not have an action of his own to guarantee achievement of his goal γ_B .*

I will focus here on *local* conditional strategic reasoning, which only refers to the immediate actions of the agents, not to their *long-term global strategies*.

As we will see further, the standard logics for multiagent strategic reasoning, such as CL and ATL , cannot capture this type of conditional reasoning. A more expressive logic for *conditional strategic reasoning* is needed here.

Depending on Bob’s knowledge about Alice’s goal and of her possible or expected choices of action, there can be several possible cases for Bob’s reasoning.

6.1.1. Case 1: Bob Does Not Know Alice’s Goal or Actions

The simplest case is when *Bob does not know Alice’s goal, or her available actions*, and therefore has no a priori expectations about her choice of action. Then, Bob can only ensure that γ_B will occur if he has an action to make γ_B true, *regardless* of how Alice (and all others) act. For instance, in our scenario, suppose the campus has only one entrance. If Bob is standing by the only entrance of the campus, then he is sure to meet Alice when she comes, no matter what she will do there. This can be simply expressed in coalition logic as $[B]\gamma_B$.

6.1.2. Case 2: Bob Knows Alice’s Goal and Possible Actions

Suppose now, that Bob knows Alice’s goal, as well as *all possible actions* of Alice that can ensure the satisfaction of her goal. Thus, Bob knows that Alice will perform one of these actions, but possibly *does not know which one*. For instance, in our scenario, Bob knows

that Alice is coming to campus to meet with her supervisor and she can meet with him either in his office, or in the lecture room, or in the café.

We can express Bob's conditional ability to achieve his goal as follows:

*“Whichever way Alice acts towards achieving her goal γ_A ,
Bob can act so as to ensure achievement of his goal γ_B ”.*

This claim can no longer be expressed in coalition logic, except in the special case when the satisfaction of Alice's goal guarantees satisfaction of Bob's goal, too (thus, Bob need not do anything about that). That can again be simply expressed in CL as $[\emptyset](\gamma_A \rightarrow \gamma_B)$.

6.1.3. Reactive and Proactive Ability

This case above admits two essentially different readings, discussed further: as a *reactive ability*, and as a *proactive ability*. (In [27], these were called *ability de dicto* and *ability de re*.)

Bob's reactive ability. Suppose first, that Bob will know Alice's choice of action when he is to choose his action. Then, Bob's ability to achieve his goal is reactive, meaning that for every action of Alice that ensures her goal γ_A , Bob has an action of his, possibly dependent on Alice's action, that would also ensure his goal, γ_B . For instance, in our scenario, suppose that Alice's supervisor tells Bob where and when he is going to meet with Alice. Then Bob can wait for Alice at the respective place. This claim cannot be expressed in CL, so a new operator is needed for it.

Bob's proactive ability. Suppose now that Bob will not know Alice's action when he is to choose his. For instance, in our scenario, suppose all that Bob knows is that Alice will meet with her supervisor either in his office, or in the lecture room, or in the café, but does not know where. Now, for Bob to ensure that his goal will be achieved, he must have a *uniform* choice of action to make γ_B true, when applied with any action of Alice that ensures the truth of her goal, γ_A . For instance, in our scenario, suppose that all meeting places for Alice and her supervisor are in the same building, so Bob can wait for her at the only entrance of that building. That is, Bob must have a proactive ability to achieve his goal.

This cannot be expressed in CL, either, so a new operator is needed again.

Remark 5. *The notions of proactive and reactive ability, respectively, generalize the notions of α -effectivity and β -effectivity in game theory (cf., e.g., [44]). Indeed, in the context of the Alice and Bob story above, proactive and reactive ability of Bob correspond precisely, respectively, to α -effectivity and β -effectivity when either $\gamma_A \equiv \top$ (so, any action of Alice is equally good for achieving γ_A) or when Bob does not know γ_A and hence can expect any action from Alice. Furthermore, one can still associate proactive and reactive ability with α -effectivity and β -effectivity if one explicitly assumes individual rationality of Alice, implying that she would only choose actions that would ensure γ_A . Thus, one can regard proactive and reactive ability, respectively, as conditional α -effectivity and β -effectivity, terminology which will be used further.*

Lastly, an important point: even though knowledge of the agent (here, Bob) about the other's goals and possible actions is essential, it will not feature in our formal logical language, nor in the formal semantics, but only in the external reasoner's analysis, in which case conditional strategic ability applies.

6.1.4. Case 3: Assuming Alice's Cooperation

Suppose now that Alice also knows Bob's goal and can choose to cooperate with Bob by selecting a suitable action σ_A , that would not only guarantee achievement of her goal γ_A , but will also enable Bob to supplement σ_A with an action σ_B of his, which would then also guarantee achievement of his goal γ_B (we also assume that Alice knows enough about Bob's possible actions). This scenario cannot be formalized in CL, either.

6.2. Modal Operators for Conditional Strategic Reasoning

I will now present three new binary modal operators for conditional strategic reasoning, for any coalitions A and B, with intuitive semantics corresponding to the three

reasoning cases in Section 6.1, i.e., formalizing, respectively, reactive abilities, proactive abilities, and abilities under cooperation.

Below, A and B are coalitions and γ_A and γ_B are their respective goals. In the special case when A and B are singletons, they are assumed to be different agents.

(O_α) $[A]_\alpha(\gamma_A; \langle B \rangle \gamma_B)$ means that the coalition $B \setminus A$ of agents who are in B but not in A has a joint action $\sigma_{B \setminus A}$ such that if A applies any joint action that guarantees the truth of γ_A , then $B \setminus A$ can ensure the truth of γ_B by applying σ_B .

This operator formalizes the notion of a coalition's *proactive ability*, discussed in the special case of single-agent coalitions in Section 6.1.3, respectively, to the game-theoretic notion of *conditional α -effectivity*, hence the notation. Note that the agents in $B \cap A$ (if any) are assumed to act on behalf of A in its pursuit of γ_A .

(O_β) $[A]_\beta(\gamma_A; \langle B \rangle \gamma_B)$ means that for any joint action σ_A of A that guarantees the truth of γ_A , when applied by A there is a joint action $\sigma_{B \setminus A}$ that guarantees γ_B when additionally applied by $B \setminus A$. Note that $[A]_\beta(\perp; \langle B \rangle \gamma_B)$ is vacuously true for any A, B , and γ_B , as then there cannot be such joint actions σ_A that enable satisfying \perp . This may sound odd, but it is no special phenomenon in ConStR, as the same effect occurs in FOL with universal quantification over an empty set of objects.

This operator formalizes a claim of the ability of the coalition B to choose a suitable joint action so as to achieve the goal γ_B assuming that A acts so as to achieve the goal γ_A , if B is to choose their joint action *after* B learns the joint action of A . In this case, the actions of the agents in $B \cap A$ (if any) are assumed to be already fixed by γ_A .

This corresponds to the notion of agents' *reactive ability* discussed in Section 6.1.3, respectively, to the game-theoretic notion of *conditional β -effectivity*, hence the notation.

(O_c) $\langle A \rangle_c(\gamma_A; \langle B \rangle \gamma_B)$ means that A has a joint action σ_A which, when applied, guarantees the truth of γ_A and enables $B \setminus A$ to apply a joint action $\sigma_{B \setminus A}$ that guarantees γ_B when additionally applied by $B \setminus A$.

This operator formalizes Case 3 discussed in Section 6.1, where A knows the goal of B and can choose to cooperate with B by selecting an action among those that ensure satisfaction of γ_A which is also suitable for B .

6.3. The Logic ConStR: Language and Some Definable Operators

We fix a finite nonempty set of agents Agt and a countable set of atomic propositions Prop . The formulae of ConStR, where $p \in \text{Prop}$ and $A, B \subseteq \text{Agt}$ are defined as follows:

$$\phi ::= p \mid \top \mid \neg\phi \mid (\phi \vee \psi) \mid [A]_\alpha(\phi; \langle B \rangle \psi) \mid [A]_\beta(\phi; \langle B \rangle \psi) \mid \langle A \rangle_c(\phi; \langle B \rangle \psi)$$

Here are some definable operators and expressions in ConStR, which can be easily seen from the informal semantics above, and can also be easily verified with the formal semantics introduced further.

- The coalitional operator from CL is definable by means of each of O_c, O_α, O_β as follows:

$$(O_\alpha) [C]\phi ::= [\emptyset]_\alpha(\top; \langle C \rangle \phi).$$

This claims an unconditional ability of C to choose an action that guarantees ϕ .

$$(O_\beta) [C]\phi ::= [\emptyset]_\beta(\top; \langle C \rangle \phi), \text{ or } [C]\phi ::= [\overline{C}]_\beta(\top; \langle C \rangle \phi).$$

The only strategy of the empty coalition guarantees the satisfaction of \top .

$$(O_c) [C]\phi ::= \langle C \rangle_c(\phi; \langle C \rangle \phi) \text{ or } [C]\phi ::= \langle C \rangle_c(\phi; \langle C \rangle \top).$$

- $\langle A \rangle_c(\gamma_A; \langle B \rangle \gamma_B)$ is equivalent to $\langle\langle A \rangle\rangle(\gamma_A \wedge \langle\langle B \setminus A \rangle\rangle \gamma_B)$ in ATL^* .
- The negated O_c operator $\neg\langle A \rangle_c(\phi; \langle B \rangle \psi)$ says that every joint action of A that, when applied, guarantees the truth of ϕ , would prevent B from acting additionally so as to guarantee ψ . This formalizes the conditional reasoning scenario where the goals of A and B are conflicting, so whichever way A acts towards its goal would block B from acting to guarantee achievement of B 's goal.

- $[A]_\beta(\top; \langle B \rangle \psi)$ essentially formalizes the case when the agents in B are not informed about the goal of A, but have to choose their action after learning the action of A.
- On the other hand, $[A](\phi | \psi) ::= [A]_\beta(\phi; \langle \emptyset \rangle \psi)$, also equivalent to $[A]_\alpha(\phi; \langle \emptyset \rangle \psi)$, says that any joint strategy of A that guarantees ϕ to be true also guarantees ψ to be true. That formalizes the claim of an observer who knows both the goal ϕ and the possible joint actions of A, that the outcome from the joint action of A will also satisfy ψ .
- $\langle A \rangle(\phi | \psi) ::= \neg[A](\phi | \neg\psi)$ says that there is a joint strategy of A that ensures the truth of ϕ and also enables the satisfaction of ψ .
 $\langle A \rangle(\phi | \psi)$ is also definable as $\langle A \rangle_c(\phi; \langle \bar{A} \rangle \psi)$, where $\bar{A} = \text{Agt} \setminus A$.
- The coalitional operator $[A]$ from CL is a special case of the above: $[A]\phi ::= \langle A \rangle(\phi | \top)$.

6.4. Formal Semantics of ConStR

Given coalitions $A, B \subseteq \text{Agt}$ and joint actions σ_A for A and σ_B for B, we define the **ordered join of σ_A and σ_B** to be the joint action $\sigma_A \uplus \sigma_B$ for $A \cup B$ which equals σ_A when restricted to A and equals to σ_B when restricted to $B \setminus A$. Thus, in particular, $\sigma_A \uplus \sigma_B = \sigma_A$ for any $B \subseteq A \subseteq \text{Agt}$ and $\sigma_A \uplus \sigma_B = \sigma_A \cup \sigma_B$ when A and B are disjoint.

Now, let $\mathcal{M} = (\text{Agt}, S, \text{Act}, g, \text{Prop}, L)$ be a concurrent game model.

The formal semantics of ConStR extend the semantics of CL to the new operators as follows (recall the definition of $\text{Out}[s, \sigma]$ from Section 2):

- $\mathcal{M}, s \models [A]_\alpha(\phi; \langle B \rangle \psi)$ iff
B has a joint action σ_B such that for every joint action σ_A of A,
if $\mathcal{M}, u \models \phi$ for every $u \in \text{Out}[s, \sigma_A]$, then $\mathcal{M}, u \models \psi$ for every $u \in \text{Out}[s, \sigma_A \uplus \sigma_B]$.
- $\mathcal{M}, s \models [A]_\beta(\phi; \langle B \rangle \psi)$ iff
for every joint action σ_A of A such that $\mathcal{M}, u \models \phi$ for every $u \in \text{Out}[s, \sigma_A]$,
B has a joint action σ_B (generally, dependent on σ_A) such that
 $\mathcal{M}, u \models \psi$ for every $u \in \text{Out}[s, \sigma_A \uplus \sigma_B]$.
- $\mathcal{M}, s \models \langle A \rangle_c(\phi; \langle B \rangle \psi)$ iff
A has a joint action σ_A , such that $\mathcal{M}, u \models \phi$ for every $u \in \text{Out}[s, \sigma_A]$ and
B has a joint action σ_B such that $\mathcal{M}, u \models \psi$ for every $u \in \text{Out}[s, \sigma_A \uplus \sigma_B]$.

Remark 6. The semantics of each of the operators above can be restated to consider joint actions for $B \setminus A$ rather than the whole B. For instance, it can be easily verified for the latter operator, that $\mathcal{M}, s \models [A]_\alpha(\phi; \langle B \rangle \psi)$ iff $B \setminus A$ has a joint action $\sigma_{B \setminus A}$ such that for every joint action σ_A of A, if $\mathcal{M}, u \models \phi$ for every $u \in \text{Out}[s, \sigma_A]$, then $\mathcal{M}, u \models \psi$ for every $u \in \text{Out}[s, \sigma_A \uplus \sigma_{B \setminus A}]$.

As usual, formula ϕ of ConStR is called **valid in ConStR**, denoted $\models_{\text{ConStR}} \phi$, iff $\mathcal{M}, u \models \phi$ for every concurrent game model \mathcal{M} and state $u \in \mathcal{M}$; ϕ is **satisfiable in ConStR**, iff $\mathcal{M}, u \models \phi$ for some concurrent game model \mathcal{M} and state $u \in \mathcal{M}$.

Here are some easy observations on the relationships between the operators in ConStR:

Proposition 7. The following validities and nonvalidities in ConStR hold:

1. $\models_{\text{ConStR}} [a]_\alpha(p; \langle b \rangle q) \rightarrow [a]_\beta(p; \langle b \rangle q)$.
2. $\not\models_{\text{ConStR}} [a]_\beta(p; \langle b \rangle q) \rightarrow [a]_\alpha(p; \langle b \rangle q)$.
3. $\not\models_{\text{ConStR}} \langle a \rangle_c(p; \langle b \rangle q) \rightarrow [a]_\beta(p; \langle b \rangle q)$.
4. $\not\models_{\text{ConStR}} \langle a \rangle_c(p; \langle b \rangle q) \rightarrow [a]_\alpha(p; \langle b \rangle q)$.
5. $\not\models_{\text{ConStR}} [a]_\alpha(p; \langle b \rangle q) \rightarrow \langle a \rangle_c(p; \langle b \rangle q)$.
6. $\not\models_{\text{ConStR}} [a]_\beta(p; \langle b \rangle q) \rightarrow \langle a \rangle_c(p; \langle b \rangle q)$.
7. $\models_{\text{ConStR}} ([A]p \wedge [a]_\beta(p; \langle b \rangle q)) \rightarrow \langle a \rangle_c(p; \langle b \rangle q)$.

Proof.

1. Immediately from the semantic definitions, as $\exists \forall$ implies $\forall \exists$.
2. A counter-model is shown further, in Example 8.

3. A counter-model is shown in Example 8.
 4. Follows from 1 and 3 above.
 5. This holds for the trivial reason that if in a given model the agent a has no action that ensures the truth of p in the given state, then the antecedent is vacuously true, whereas the consequent is false there.
 6. Likewise.
 7. An easy semantic exercise.
-

Example 8. Consider the game model \mathcal{M} in Figure 3, with two players, a and b, where at state s_0 a has 3 actions, $a_1, a_2,$ and a_3 , and b has 2 actions, b_1 and b_2 .

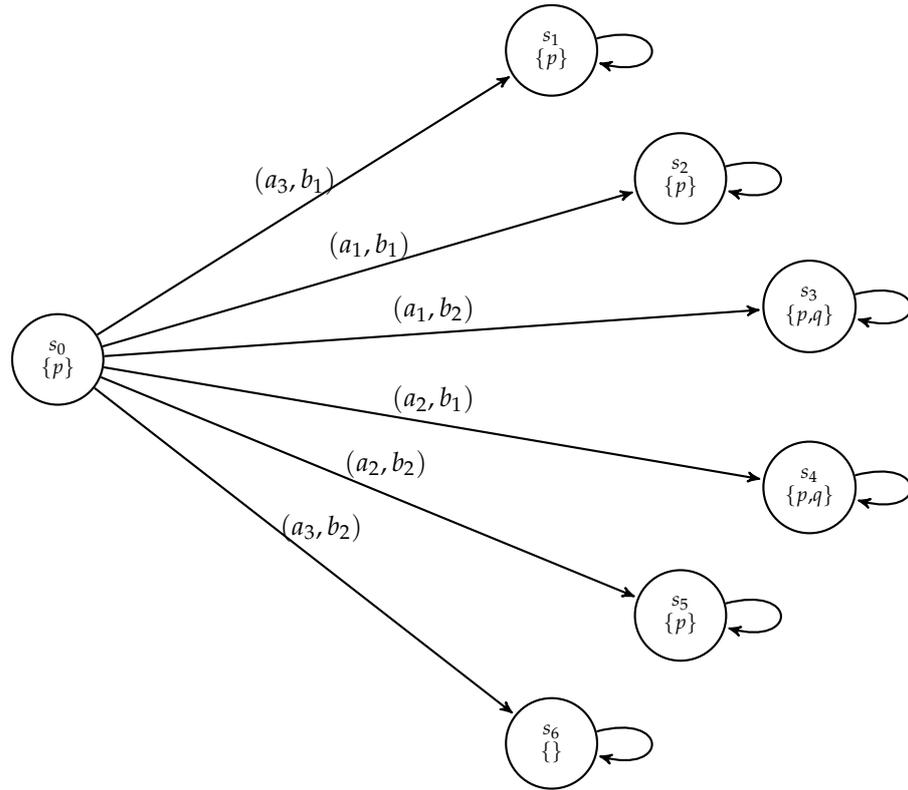


Figure 3. The concurrent game model \mathcal{M} in Example 8.

Note that:

1. $\mathcal{M}, s_0 \not\models [b]q$, whereas $\mathcal{M}, s_0 \models \langle a \rangle_c(p; \langle b \rangle q)$.
Thus, an agent may have only conditional ability to achieve its goal.
2. $\mathcal{M}, s_0 \not\models [a]_\alpha(p; \langle b \rangle q)$. Indeed, neither b_1 nor b_2 ensures q against both choices a_1 and a_2 of a. Thus, b does not have a uniform action to ensure q against any action of a that ensures p .
3. $\mathcal{M}, s_0 \models [a]_\beta(p; \langle b \rangle q)$. Indeed, a has two actions at state s_0 to ensure p : a_1 and a_2 . For each of them, b has an action to ensure q : choose b_2 if a chooses a_1 , and choose b_1 if a chooses a_2 . Therefore, $\not\models_{\text{ConStR}} [a]_\beta(p; \langle b \rangle q) \rightarrow [a]_\alpha(p; \langle b \rangle q)$. Thus, the claim made by the proactive ability operator O_α is stronger than the claim made by the reactive ability operator O_β .
4. On the other hand, if the outcomes of (a_2, b_1) and (a_2, b_2) are swapped, then $[a]_\alpha(p; \langle b \rangle q)$ becomes true at s_0 in the resulting model.
5. Furthermore, clearly, $[a]_\alpha(p; \langle b \rangle q) \rightarrow [a]_\beta(p; \langle b \rangle q)$ is a valid formula in ConStR.
6. If the model is modified by making p true also at s_6 , then in the resulting model \mathcal{M}' we have $\mathcal{M}', s_0 \models \langle a \rangle_c(p; \langle b \rangle q)$ and $\mathcal{M}', s_0 \not\models [a]_\beta(p; \langle b \rangle q)$. Therefore, $\not\models_{\text{ConStR}} \langle a \rangle_c(p; \langle b \rangle q) \rightarrow [a]_\beta(p; \langle b \rangle q)$.
7. Note also that (\mathcal{M}, s_0) does not satisfy the ATL formula $\llbracket a \rrbracket (Xp \rightarrow \langle \langle b \rangle \rangle Xq)$ (where $\llbracket C \rrbracket \phi ::= \neg \langle \langle C \rangle \rangle \neg \phi$); hence, it is not equivalent to $[a]_\beta(p; \langle b \rangle q)$.

For more examples and observations on the mutual nonexpressiveness of the operators in ConStR, see [28]. Moreover, using the notions of bisimulation defined there for each of these operators, one can show that none of the three binary operators in ConStR can be defined in terms of the other two.

6.5. On the Axiomatization and Decidability for ConStR

A system of axioms Ax_{ConStR} for ConStR is presented in [28]. It involves lists of axioms for each of the operators O_c , O_α , and O_β , as well the following two interaction axiom schemes:

$$\text{(ConStR 1)} \quad [A]_\alpha(\phi; \langle B \rangle \psi) \rightarrow [A]_\beta(\phi; \langle B \rangle \psi).$$

$$\text{(ConStR 2)} \quad ([A]\phi \wedge [A]_\beta(\phi; \langle B \rangle \psi) \rightarrow \langle A \rangle_c(\phi; \langle B \rangle \psi).$$

These schemes are generated by uniform substitutions in the respective validities 1 and 7 in Proposition 7. Notably, the only axiom scheme identified there which distinguishes the operators O_α and O_β is the following antimonicity (with respect to A) scheme:

$$[A \cup C]_\alpha(\phi; \langle B \rangle \psi) \rightarrow [A]_\alpha(\phi; \langle B \rangle \psi) \text{ for any } C \subseteq \text{Agt},$$

which is valid for O_α but not O_β . For proof of this claim, and for the other axiom schemes and inference rules for each of the operators of ConStR, see [28].

The only additional (to modus ponens) inference rules are the O-monotonicity rules, for O being each of O_α , O_β , and O_c . For instance, here is the O_c -**monotonicity rule**:

$$\frac{\phi \rightarrow \phi', \psi \rightarrow \psi'}{\langle A \rangle_c(\phi; \langle B \rangle \psi) \rightarrow \langle A \rangle_c(\phi'; \langle B \rangle \psi')}$$

The completeness proof for Ax_{ConStR} is currently still under development, so its completeness is currently still a conjecture, and so is even the completeness of each of the three main subsystems, one for each of the respective operators. For all that I currently know, these seem to be quite challenging problems.

Another conjecture is the decidability of satisfiability in ConStR, even in each of its three main subsystems. Here, I only claim that it can be proved by a model-theoretic argument, based on the *bounded tree-model property*, meaning that every satisfiable formula of SFCL is satisfiable in a finite tree-like model of height and branching factor effectively bounded above in terms of the size of the formula. (Detailed proof of this claim is currently under construction, so the claim is, at present, still a working conjecture.)

7. The Socially Friendly Coalition Logic SFCL

The “socially friendly coalition logic” SFCL was introduced in [25] with the aim to capture and formalize the idea that, while socially responsible (or, “friendly”) rational agents and coalitions act in pursuit of their own goals, when deciding on their course of action they can take into account the goals of other agents in the system and make allowance for cooperation enabling their satisfaction, too, whenever it is possible, to reconcile their goals with those of the others. Formally, the idea is implemented by introducing in the formal language a “socially friendly coalitional operator”, presented here.

7.1. Socially Friendly Coalitional Operators

SF The **socially friendly coalitional operator SF** takes any positive number of formulae $\phi, \psi_1, \dots, \psi_k$ as arguments and places them together into a single formula as follows:

$$[C](\phi; \psi_1, \dots, \psi_k).$$

I will call the formula ϕ above the **primary goal** of the formula (and of the coalition C) and the formulae ψ_1, \dots, ψ_k —the **secondary goals** of the formula.

The intuitive meaning of the formula $[C](\phi; \psi_1, \dots, \psi_k)$ is that

“ C has a joint action σ_C that guarantees satisfaction of ϕ and also enables the complementary coalition \bar{C} to satisfy any one of the goals ψ_1, \dots, ψ_k by applying a respectively suitable joint action”.

The operator **SF** is a multiagent extension of the modal operator $\Box(\psi_1, \dots, \psi_k; \phi)$ (note the different order of the arguments) in the *instantial neighborhood logic* (INL) introduced and studied in [45].

The special case of the “socially friendly coalitional operator” with a single secondary goal $[A](\phi; \psi)$ is equivalent to the operator $\langle A \rangle(\phi | \psi)$ defined in Section 6.

In particular, the strategic operator of CL is also definable here, as $[C]\phi ::= [C](\phi; \top)$.

SF1 A refinement of **SF**: $[C; C_1, \dots, C_n](\phi; \phi_1, \dots, \phi_n)$, meaning:

“ C has a collective action σ_C that guarantees ϕ , and is such that, when fixed, each C_i has a collective action that guarantees ϕ_i ”.

This definition presumes that if C intersects with C_i , then the agents in $C \cap C_i$ are already committed to σ_C . On the other hand, the collective actions claimed to exist for each C_i need not be compatible, similar to in the intuitive semantics of $[C](\phi; \psi_1, \dots, \psi_k)$, which is a special case of SF1, where $C_1 = \dots = C_n = \text{Agt} \setminus C$.

SF2 $[\langle C_1 \rangle \phi_1; \dots; \langle C_k \rangle \phi_k]$ meaning: “ C_1 has a collective action to guarantee ϕ_1 , and given that action ... C_k has a collective action to guarantee ϕ_k ”.

This is a sequential version of SF1 where the coalitions C_1, \dots, C_k are arranged in decreasing priority order.

Hereafter, I will consider **SF** only.

7.2. Syntax and Semantics of the Logic SFCL

The formulae of SFCL are defined by the following grammar:

$$\phi ::= p \mid \top \mid \neg\phi \mid (\phi \vee \phi) \mid [C](\phi; \phi, \dots, \phi) \quad (p \in \text{Prop})$$

The standard definitions of the other propositional connectives apply.

In addition, I define $[C]\phi := [C](\phi; \top)$. Given a finite list of formulae $\Psi = \psi_1, \dots, \psi_k$, I will write $[C](\phi; \Psi)$.

The fragment of SFCL containing only formulae $[C](\phi; \psi)$ where ψ is a single formula will be denoted by SFCL₁.

The formal semantics of SFCL is given in terms of **truth of an SFCL-formula at a state s of a concurrent game model \mathcal{M}** inductively, as in CL, with the following main clause:

$\mathcal{M}, s \models [C](\phi; \Psi)$ iff
 there exists a joint action σ_C of C available at s ,
 such that $\mathcal{M}, u \models \phi$ for each state u in its outcome set $O[s, \sigma_C]$,
 and for each $\psi \in \Psi$ there is a state $v \in O[s, \sigma_C]$ such that $\mathcal{M}, v \models \psi$.

7.3. Example 1: Negotiating the Family Budget

Example 9. This example is an adapted and extended version of one from [25]. Consider 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 for 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 thinks of three options for a car: a cheap used black Ford, a more expensive new green Tesla, 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:

- CK (cheap kitchen), EK (expensive kitchen), $K = CK \vee EK$ (any kitchen);
- CC (cheap car), AC (average car), EC (expensive car), $C = CC \vee AC \vee EC$ (any car);
- CT (cheap trip), ET (expensive trip), $T = CT \vee ET$ (any trip).

Calculations show that Ann's choices and money produce the following strategic powers:

- Ann can afford to pay for an expensive kitchen and then let the others choose some car or some trip, formally: $[Ann](EK; C, T)$.
- Alternatively, Ann can settle with a cheap kitchen, and then let the others choose between some car plus some trip, or an expensive car, or an expensive trip; formally: $[Ann](CK; CC \wedge CT, EC, ET)$.
- However, if Ann opts for an expensive kitchen, then the family cannot afford an average car plus a trip; formally: $\neg[Ann](EK; AC \wedge T)$.

Likewise, it turns out that Bill has the following strategic powers:

- $[Bill](C; K, T)$;
- $[Bill](CC; EK, ET)$.
- However, $\neg[Bill](EC; AK \wedge T)$.

Charlie alone cannot afford to pay in full for any trip, but can only leave his savings in the family pool and try to negotiate some trip with his parents. Formally: $\neg[Charlie](T; \top)$.

Lastly, here are some coalitional powers:

- $[\{Ann, Bill\}](C \wedge EK; CT)$;
- $[\{Ann, Charlie\}](CK \wedge ET; CC)$;
- $[\{Bill, Charlie\}](AC \wedge CT; K)$.

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

7.4. Example 2: Job Applicants

Consider four job candidates, Alice, Bob, Carl, and Diana, who have applied to three companies: Banana, Megasoft, and Fakebook, each advertising one position.

Suppose no candidate has strong preferences between these jobs.

All candidates have been interviewed by each of the companies.

Each company has selected and ranked in a priority order some of them, in which order they are offered the job and have to accept it or not. If an offer is not accepted, it goes to the next-ranked candidate, until the position is filled (if at all).

The rankings have been communicated to the candidates, as follows:

- | | | |
|----------------------------|-------------------|---------------------|
| Banana: | Megasoft: | Fakebook: |
| 1. Alice, 2. Bob, 3. Carl. | 1. Diana, 2. Bob. | 1. Alice, 2. Diana. |

Each of the candidates must now make their strategic choice of which offer to accept, if and when given the option. Suppose the candidates can communicate with each other.

Here are some true statements formalized in SFCL, where Alice, Bob, Carl, and Diana are labeled as A, B, C, D and e_X means "X obtains a job":

- $[A](e_A; e_B \wedge e_D, e_C \wedge e_D, e_B \wedge e_C)$

Indeed, Alice can take the offer from Fakebook and leave it to the others to decide on the other two positions. Note that Bob can then choose either offer and enable either Carl or Diana to obtain a job, but can also decline both offers and thus enable both of them to obtain a job.

- $[A](\top; e_B \wedge e_C \wedge e_D)$
Alternatively, Alice can act selflessly by declining both offers, and thus enable all others to obtain a job (by Bob choosing Fakebook).
- $\neg[C](e_C; \top) \wedge [C](\top; e_C)$
Carl cannot be sure to obtain a job, but all others can be sure of this.
- $\neg[D](e_D; e_B \wedge e_C) \wedge [D](e_D; e_B, e_C)$
Diana cannot be sure to obtain a job and enable the others to see to it that both Bob and Carl obtain a job, but she can ensure that she obtains a job (by accepting the offer from Megasoft) and then the others ensure that either Bob or Carl obtains a job, too.
- $[\{A, D\}](e_A \wedge e_D; e_B, e_C)$
Alice and Diana together can ensure that they both obtain a job and then either of the other two can obtain a job, too, up to their choice.
- $[\{A, D\}](e_A \wedge e_D \wedge \neg e_B \wedge \neg e_C)$
Alice and Diana can also be mean and act together so that only they obtain jobs, by accepting the respective offers from Banana and Megasoft and leaving the Fakebook position unfilled.

7.5. Socially Friendly Coalition Logic SFCL: A System of Axioms

The system of axioms Ax_{SFCL} for SFCL, proposed in [25], combines Ax_{CL} with a multiagent extension of the axiomatization of the instantial neighborhood logic (INL) [45], plus some additional axiom schemes. Here is just one of these additional schemes from Ax_{SFCL} , where Ψ is a finite list of SFCL-formulae and Ψ, θ is the list obtained by appending θ to Ψ :

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

This is a valid scheme. Indeed, for any formula θ , the strategy for C that ensures ϕ and enables each formula in the list Ψ either also ensures $\neg\theta$, in which case it can be added conjunctively to ϕ , or else enables θ , in which case it can be added to the list Ψ .

The only additional (to modus ponens) inference rule is $[C]$ -**monotonicity**:

$$\frac{\phi \rightarrow \phi', \psi_1 \rightarrow \psi'_1, \dots, \psi_k \rightarrow \psi'_k}{[C](\phi; \psi_1, \dots, \psi_k) \rightarrow [C](\phi'; \psi'_1, \dots, \psi'_k)}$$

Note that the rule RE in the axiomatic system for INL in [45] is an admissible rule here, and that can be proved by a routine induction on the structure of ϕ in that rule. For the Boolean cases, that is due to the completeness of the purely propositional fragment (where MP suffices), and the case of $\phi = [C](\dots)$ is handled by using the $[C]$ -monotonicity rule.

A completeness proof for Ax_{SFCL} was presented in [25], but an error was subsequently found in it (on which an erratum has been posted). The correct proof of completeness is currently still under development, and therefore the completeness claim is currently only a working conjecture.

Another conjecture is the decidability of satisfiability in SFCL. Here, I only claim that it can be proved by a model-theoretic argument, similar to the decidability of ConStR, based on the *bounded tree-model property*, meaning that every satisfiable formula of SFCL is satisfiable in a finite tree-like model of height and branching factor effectively bounded above in terms of the size of the formula. A proof of this claim, which is, at present, formally still a working conjecture, will be presented in a subsequent work.

Further open problems and directions for further research are related to various natural generalizations of SF, including SF1 and SF2. To my knowledge, they have not been explored and no technical results are even conjectured for them yet, except for some special cases explored in the context of conditional strategic reasoning in Section 6.

8. The Logic of Local Coalitional Goal Assignments (LCGA)

The logic LCGA was introduced in [25] as the “group protecting coalition logic (GPCL)”, with the aim to formalize the idea, complementary to that of SFCL, that, while socially responsible rational agents and coalitions contribute with their individual and collective actions to the society of all agents, they wish to achieve that in a way that protects their individual and collective interests and goals. Formally, the idea is implemented by introducing in the formal language a “coalitional goal assignment operator” (called “group protecting coalitional operator” in [25]), presented here.

8.1. The Coalitional Goal Assignments Operator

The **coalitional goal assignments operator** takes a list of coalitions C_1, \dots, C_k and a list of formulae representing their goals ϕ_1, \dots, ϕ_k and produces a formula of the type:

$$\langle\langle C_1 \triangleright \phi_1, \dots, C_k \triangleright \phi_k \rangle\rangle.$$

The intuitive meaning of the claim formalized by the formula above is that there is an action profile σ of Agt , such that for each i , the restriction of σ to the coalition C_i is a joint action σ_i that guarantees the satisfaction of ϕ_i , even if some, or all, other agents deviate from the action profile σ . In particular, $[C]\phi$ is equivalent to $\langle\langle C \triangleright \phi \rangle\rangle$.

Now, I will define a more general and succinct representation of that operator. First, a **coalitional goal assignment** is a mapping $\gamma : \mathcal{P}(\text{Agt}) \rightarrow \Gamma$, where Γ is a set of goal formulae (which may, but need not be, the full logical language under consideration). Thus, for every coalition C , $\gamma(C)$ expresses the goal of C .

Usually, most of the possible coalitions do not have coalitional goals and hence do not form. That can be formalized by assigning the “truth” \top as a goal to each of them.

Now, the operator defined above naturally generalizes to one expressed simply as $\langle\langle \gamma \rangle\rangle$, which is a concise representation of $\langle\langle C_1 \triangleright \gamma(C_1), \dots, C_{2^n} \triangleright \gamma(C_{2^n}) \rangle\rangle$, where C_1, \dots, C_{2^n} is a (fixed) canonical enumeration of the set $\mathcal{P}(\text{Agt})$ of all subsets (coalitions) of Agt .

8.2. Syntax and Semantics of the Logic LCGA

The formulae of LCGA are given by the following grammar:

$$\phi ::= p \mid \top \mid \neg\phi \mid (\phi \vee \psi) \mid \langle\langle \gamma \rangle\rangle \quad (p \in \text{Prop})$$

where $p \in \text{Prop}$ and γ is a coalitional goal assignment. All other standard propositional connectives are defined as usual.

I will use the notation $\langle\langle C_1 \triangleright \phi_1, \dots, C_k \triangleright \phi_k \rangle\rangle$ as an explicit record of $\langle\langle \gamma \rangle\rangle$ where γ is the (unique) coalitional goal assignment defined by $\gamma(C_1) = \phi_1, \dots, \gamma(C_k) = \phi_k$, and $\gamma(C) = \top$ for every other $C \in \mathcal{P}(\text{Agt})$. Clearly, both notations for the coalitional goal assignment operator 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 nontrivial goals.

The semantics of LCGA is given in terms of **truth of an LCGA-formula at a state s of a concurrent game model \mathcal{M}** . The only new semantic clause is the one for $\langle\langle \cdot \rangle\rangle$:

$\mathcal{M}, s \models \langle\langle \gamma \rangle\rangle$ iff there exists an action profile $\sigma \in \Sigma_{\text{Agt}}$, available at s and such that $\mathcal{M}, u \models \gamma(C)$ for every $C \subseteq \text{Agt}$ and every $u \in O[s, \sigma|_C]$.

The intuitive reading is spelled out in more detail as follows:

“There is a strategy profile of the grand coalition in which all agents participate with their individual strategies in such a way that, each agent or coalition guarantees the satisfaction of its own goal against any possible deviations of all other agents, thus protecting its individual (respectively, coalitional) interests.”

The clause for $\langle\langle \gamma \rangle\rangle$ can also be given in terms of extensions, where $\Sigma_{\text{Agt}} = \prod_{a \in \text{Agt}} \Sigma_a$:

$$\llbracket \langle\langle \gamma \rangle\rangle \rrbracket_{\mathcal{M}} = \{s \in S \mid \exists \sigma \in \Sigma_{\text{Agt}} : O[s, \sigma|_C] \subseteq \llbracket \gamma(C) \rrbracket_{\mathcal{M}} \text{ for all } C \subseteq \text{Agt} \}$$

8.3. Some Observations and Examples of the Expressiveness of LCGA

1. The operator $\langle A \rangle_c(\phi; \langle B \rangle \psi)$, introduced in Section 6, is definable in terms of $\langle \cdot \rangle$: $\langle A \rangle_c(\phi; \langle B \rangle \psi) = \langle [A \triangleright \phi, A \cup B \triangleright \psi] \rangle$. Nevertheless, it has a different motivation and intuitive interpretation there.
2. The fragment SFCL₁ of SFCL (subsuming CL) embeds into LCGA by defining $[C](\phi; \psi)$ equivalently as $\langle [C \triangleright \phi, \text{Agt} \triangleright \psi] \rangle$.
However, this does not generalize to $[C](\phi; \psi_1, \dots, \psi_k)$ for $k \geq 2$.
Conversely, the operator $\langle [C_1 \triangleright \phi_1, \dots, C_n \triangleright \phi_n] \rangle$ cannot be expressed in terms of the SFCL operators $[C](\phi; \psi_1, \dots, \psi_k)$, either.
These nonexpressiveness claims can be proved by using the respective notions of bisimulations introduced for each of these operators in [25].
3. $\langle [C_1 \triangleright \phi_1, C_1 \cup C_2 \triangleright \phi_2, \dots, C_1 \cup \dots \cup C_k \triangleright \phi_k] \rangle$ is equivalent to the sequential version SF2 of the operator SF, mentioned in Section 7.1 as $\langle [C_1] \phi_1; \dots; [C_k] \phi_k \rangle$, where the coalitions C_1, \dots, C_k and their goals are arranged in decreasing priority order.
4. On the other hand, $\langle [C_1 \triangleright \phi_1, C_1 \triangleright \phi_2, \dots, C_k \triangleright \phi_k] \rangle$ is essentially expressible, up to a natural transformation of the models and semantics, in the group strategic STIT (cf. [13,14]), as $\diamond([C_1 \text{ stit}]X \phi_1 \wedge \dots \wedge [C_k \text{ stit}]X \phi_k)$. This observation (suggested by a reviewer) leads to a new computationally well-behaved fragment of the group strategic STIT, which is known to be not only undecidable, but even nonaxiomatizable; cf. [14].

The following example is from [25], where it was adapted from [1].

Example 10. [password protected data sharing] 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. Thus, 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 and refuses access to their data (a side remark: nowadays, such deals are arranged by smart contracts).

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”. Thus, the best possible outcome for Alice is $H_A \wedge H_B$ and the worst possible outcome is $\neg H_A \wedge H_B$; symmetrically for Bob. Therefore, can the two players cooperate to exchange passwords in a way that satisfies the provisos above? Clearly, we cannot answer that question in any definitive way without having the precise details of the actual setup.

For example, suppose we define the game as follows: each player chooses a password, which may or may not be the correct one, and sends it to the other player, and they perform that simultaneously. Then, 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 LCGA-formula

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

holds true. However, the strategy profile satisfying this goal assignment does not satisfy the players’ individual private goals, since each runs the risk of giving away their password without obtaining the other’s in return, if the other deviates. Therefore, the problem is more correctly described by the following stronger formula of LCGA:

$$\langle \langle \{A, B\} \rangle \triangleright H_A \wedge H_B, A \triangleright H_B \rightarrow H_A, B \triangleright H_A \rightarrow H_B \rangle$$

Clearly, in the simple setup above, such a strategy profile does not exist, but suppose we add a second round to the game, in which each player can first verify the other’s password received in the first round (without being able to access the data yet) and then can either confirm their own password they have sent, if the other’s password is correct, or else change their password (thus making the shared one useless) if the other’s is not correct. Then, there is a strategy profile that satisfies the specification above, and thus enables the players to exchange passwords in a way that protects their individual interests, too.

8.4. Axiomatic System for LCGA

The axiomatic system Ax_{LCGA} for the logic LCGA presented here was first proposed in [25] (for the logic GPCL) and then refined in the present form in [26]. Note, however, that the notation $\langle\langle C \triangleright \phi \rangle\rangle$ here corresponds to the notation used in [25], not the one in [26].

In addition, I will use further the following notation from [26]:

$\gamma[C \triangleright \psi]$, to denote the coalitional goal assignment obtained from γ by reassigning the goal of the coalition C to be ψ , and

$\gamma|_C$, to denote the coalitional goal assignment obtained from γ by “restricting it to C ”, meaning reassigning the trivial goal \top to all coalitions not included in C .

Here are the axiom schemes of Ax_{LCGA} with brief explanations:

(Triv) $\langle\langle \gamma^\top \rangle\rangle$, where γ^\top is the trivial goal assignment, mapping each coalition to \top .

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

Even the grand coalition of all agents cannot ensure the falsum to become true.

(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 scheme generalizes the superadditivity axiom of coalition logic (cf. Section 4.4). It is valid, because, if the coalitions C_1, \dots, C_n are pairwise disjoint, then they can join together their collective strategies for their respective coalitional goals into one strategy profile that ensures satisfaction of all these collective goals.

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

This axiom scheme is valid because any strategy profile for Agt generates a unique successor state. If a state formula ψ is true in it, then it can be added to the coalitional goal of Agt , otherwise its negation can be added to it.

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

This scheme says that 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 goal of C enforced by Σ —or, otherwise, $\neg\psi$ is true in some of these successor states, in which case it can be added to $\gamma|_C$ as the goal of the grand coalition Agt enforced by Σ .

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

Given any coalition C and subcoalition C' , this scheme says that the goal of C' can be added for free to the goal of C . Indeed, if there is any strategy profile Σ that ensures that C and C' can force their respective goals, then Σ also ensures that C can force the conjunction of these goals.

The rules of inference are *modus ponens* (MP) and *goal monotonicity* (G-Mon):

$$\frac{\gamma(C) \rightarrow \psi}{\langle\langle \gamma \rangle\rangle \rightarrow \langle\langle \gamma[C \triangleright \psi] \rangle\rangle}$$

The axiomatic system Ax_{LCGA} was first proved sound and complete in [25] (for the logic GPCL) and in the present form in [26]. Moreover, the satisfiability problem for LCGA is decidable, which follows from completeness results together with the finite model property, as shown in [25,26].

8.5. The Temporal Logic of Coalitional Goal Assignments TCGA

LCGA can only express *local* goals, referring to the immediate outcomes states. Usually, however, agents and coalitions have *long-term* goals, which need explicit temporal operators to be formalized. This was the motivation to introduce and study the temporal extension TCGA of LCGA in [26]. It uses the standard temporal operators X, U, G to express temporalized goals, similarly to ATL (Section 5).

It is convenient to classify the formulae of TCGA in two sorts, **state formulae** and **path formulae**, and to define their sets by mutual induction, as follows:

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

PathFor : $\theta ::= X\varphi \mid (\varphi U \varphi) \mid G\varphi$

where $p \in \text{Prop}$ and $\gamma : \mathcal{P}(\text{Agt}) \rightarrow \text{PathFor}$ is a coalitional goal assignment.

Thus, the path formulae are auxiliary, used to express temporalized goals. As before, I will use $\langle\langle C_1 \triangleright \phi_1, \dots, C_k \triangleright \phi_k \rangle\rangle$ as an explicit notation for $\langle\langle \gamma \rangle\rangle$ with support $\{C_1, \dots, C_k\}$.

Note the change in the notation from LCGA: the goals ϕ_1, \dots, ϕ_k from now on, refer to the current state, not to a successor. Thus, LCGA is embedded in TLCGA as its X-fragment XCGA, where the path formulae are only of the type $X\varphi$.

TLCGA also extends ATL: $\langle\langle C \rangle\rangle \theta \equiv \langle\langle C \triangleright \theta \rangle\rangle$, for any $\theta \in \text{PathFor}$.

8.6. An Aside: Equilibria and Co-Equilibria

The logic TLCGA can be used to express (see [26]) the fundamental game-theoretic concept of *Nash equilibrium*: a strategy profile ensuring that no player can deviate to improve the outcome with respect to his/her private goal. This concept is good for quantitative goals, but not for qualitative (win/lose) ones, because rational players whose goals are not satisfied (the “losers”) can still deviate in weak equilibria. Thus, the arguably more suitable for qualitative goals concept of **co-equilibrium** was proposed in [26], meaning a strategy profile ensuring that *every player (and coalition) achieves their private goal, even if all other agents deviate*. Existence of a co-equilibrium is expressed precisely by the TLCGA formula $\langle\langle a_1 \triangleright \phi_1, \dots, a_k \triangleright \phi_k \rangle\rangle$.

8.7. An Example

The following example, illustrating the use of TLCGA to express natural statements, is adapted from [26].

Example 11. *Three sheep and three wolves are standing on a river bank and all animals want to cross the river (say, because a pride of lions is approaching). There is just one boat, which takes up to two animals. There is no boatman, but the animals can sail the boat across the river. However, if on either side of the river, at any time the wolves outnumber the sheep, then they eat them up there. Of course, the sheep do not want that to happen. Therefore, do the animals have a “safe” strategy to cross the river without any being eaten?*

Let us first look at a simple specification in ATL of the claim that such strategy exists:

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

where

- c is the atomic proposition “all animals have crossed the river”.
- e is the atomic proposition “some sheep are eaten”.

Thus, the formula above says that all animals have a strategy to eventually all cross the river and meanwhile no sheep are eaten. Is this what we want? Not quite, as the strategy may be such that the wolves can deviate from it at an opportune moment and eat some of the sheep. Thus, here is a better specification, in TLCGA:

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

saying “all animals have a collective goal to cross the river, but the coalition of sheep have their coalitional goal that no sheep is ever eaten”, and this is what we really want. Therefore, is there a solution—a strategy profile—satisfying this specification? Well, it depends on a subtle but crucial detail: if all animals act simultaneously, where each animal can either stay on the river bank or board the boat (but if more than two animals board the boat, the boat cannot sail), then it is not difficult to see that such strategy does not exist. However, if at every round the wolves act first, and then the sheep act, then such strategy exists, and I leave the fun to the reader to design such strategy. Note that the strategy of each sheep will now depend also on how the wolves have just acted, so any deviation of the wolves from the agreed-upon “safe” strategy can be counteracted by the strategies of the sheep, and this is the key point.

9. Basic Strategy Logic: A Unifying Formalism for Strategic Reasoning

The semantics of all strategic operators introduced here are defined in terms of quantification in the metalanguage over actions and strategies of the agents and coalitions involved. Thus, the natural idea arises to add *explicit quantification over strategies* in the object language and have all these operators definable there. That idea was realized in the family of **strategy logics**, the first (two-agents) version of which was introduced in [21], and later extended and further studied in [22–24]. The common feature of these strategy logics is that they formally treat strategies as primitive objects, similar to elements of standard structures for first-order logic. I will refer to these generically as **standard strategy logic**, denoted SSL.

Since quantification over strategies is essentially a second-order quantification over functions, the language of SSL is very expressive and it should not be surprising that the standard systems of SSL are not recursively axiomatizable in general. Still, even some restricted fragments of SSL are quite expressive and enable specifying important properties of non-zero-sum games in a simple and natural way. In particular, the “one-alternation” fragment of SSL already subsumes, by expressiveness, the logic ATL^* and is strong enough to express, inter alia, the existence of Nash equilibria and secure equilibria (cf. [21]).

I will define and use here a simple generic version of a strategy logic, which I call **basic strategy logic**, further denoted BSL. It builds on some temporalized language for expressing goals, such as LTL (cf., e.g., [36] (Chapter 6)) that defines a **set of goal formulae** Γ which are evaluated on plays in concurrent game models. BSL extends Γ with standard Boolean connectives, with variables associated with agents and ranging over strategies for them, by means of *strategy assignments*, and with quantification over such variables within the formulae. More precisely, with each agent a , the language of BSL associates a **strategy variable for a** , denoted by x_a . Let $StrVar$ be the set of strategy variables for all agents. Then, for every nonempty coalition of agents $C = \{c_1, \dots, c_m\}$, I will use x_C to denote the tuple of variables $(x_{c_1}, \dots, x_{c_m})$. Furthermore, I will write $\forall x_C$ as a shorthand for the string $\forall x_{c_1} \dots \forall x_{c_m}$, and likewise for $\exists x_C$.

Formally, the language of BSL involves two sets of formulae: a set of **goal formulae** Γ and a set of **state formulae** $StateFor$, defined as follows:

$$StateFor : \quad \varphi ::= p \mid \top \mid \psi \mid \neg\varphi \mid (\varphi \vee \varphi) \mid \exists x_a \varphi,$$

where $p \in Prop$, $\psi \in \Gamma$, and $x_a \in StrVar$. All other standard propositional connectives and constants, including \perp , \wedge , \rightarrow , \leftrightarrow , are defined as usual. In addition, $\forall x_a ::= \neg \exists x_a \neg$.

The semantics of BSL is given in concurrent game models, using, additionally, **strategy assignments**, which are functions defined on $StrVar$ and assigning to every strategy variable x_a a strategy of the respectively specified in the semantics type (e.g., memory-based or positional) for the agent a . Given such strategy assignment \mathbf{v} and a strategy σ_a for agent a , I will denote by $\mathbf{v}[x_a ::= \sigma_a]$ the modified strategy assignment which reassigns the value of x_a to be σ_a . Likewise, $\mathbf{v}[x_C ::= \sigma_C]$ is defined for any coalition C .

Given a state s in a CGM \mathcal{M} , every strategy assignment \mathbf{v} in \mathcal{M} defines the strategy profile $\Sigma(\mathbf{v})$, where $\Sigma(\mathbf{v})_a = \mathbf{v}(x_a)$, and therefore determines the unique play in \mathcal{M} , generated by $\Sigma(\mathbf{v})$, which will be denoted by $play(s, \mathbf{v})$. Then, the truth of any goal formula can be evaluated on that play.

The semantics of BSL can now be defined in terms of truth of a formula at a state s in a given CGM \mathcal{M} , for a given strategy assignment \mathbf{v} , through the clauses:

- $\mathcal{M}, s, \mathbf{v} \models \psi$, for $\psi \in \Gamma$ iff $\mathcal{M}, play(s, \mathbf{v}) \models \psi$
(assuming that truth of the goal formulae in Γ on plays has already been defined).
- $\mathcal{M}, s, \mathbf{v} \models \exists x_a \varphi$ iff $\mathcal{M}, s, \mathbf{v}[x_a ::= \sigma_a] \models \varphi$ for some strategy σ_a for a .

The latter clause extends in a straightforward way to $\exists x_C \varphi$ and $\forall x_C \varphi$, for any coalition C .

Note that, similar to the classical first-order logic, every bound occurrence of a variable x_a is bound by the innermost occurrence of a quantifier binding x_a , so in quantification

patterns such as $Q_1x_a \dots Q_kx_a \dots \phi$, where each Q_i is \exists or \forall , all occurrence of Q_ix_a containing the innermost one in its scope are vacuous.

Now, one can uniformly translate the semantic conditions of the various strategic operators introduced here into formulae in the language of BSL, by using a uniform compositional translation τ , as follows:

CL/ATL*

$$\tau([C]\varphi) ::= \exists x_C \forall x_{\bar{C}} \tau(\varphi),$$

where $\bar{C} = \text{Agt} \setminus C$; likewise, for $\langle\langle C \rangle\rangle \varphi$.

ConStR (O_α)

$$\tau([A]_\alpha(\varphi; \langle B \rangle \psi)) ::= \exists x_{B \setminus A} \forall x_A (\forall x_{\bar{A}} \tau(\varphi) \rightarrow \forall x_{A \cup B} \tau(\psi)).$$

Note that the joint strategy for $B \setminus A$ assigned to $x_{B \setminus A}$ is supposed to ensure, together with the joint strategy for A assigned to x_A , the truth of the (translation of the) goal formula ψ whenever the joint strategy for A assigned to x_A ensures the truth of the (translation of the) goal formula φ against any joint strategy of the agents in \bar{A} , including those in $B \setminus A$.

ConStR (O_β)

$$\tau([A]_\beta(\varphi; \langle B \rangle \psi)) ::= \forall x_A (\forall x_{\bar{A}} \tau(\varphi) \rightarrow \exists x_{B \setminus A} \forall x_{A \cup B} \tau(\psi)).$$

Note the difference with the previous translation, reflecting the difference between the semantics of (O_α) and (O_β): the joint strategy for B now depends on the joint strategy for A , and the strategies of the agents in $B \cap A$ may not differ from those already fixed in the latter.

ConStR (O_c)

$$\tau(\langle A \rangle_c(\varphi; \langle B \rangle \psi)) ::= \exists x_A (\forall x_{\bar{A}} \tau(\varphi) \wedge \exists x_{B \setminus A} \forall x_{A \cup B} \tau(\psi)).$$

This formula captures the semantics of O_c in a straightforward way.

SFCL (SF)

$$\tau([C](\varphi; \psi_1, \dots, \psi_k)) ::= \exists x_C (\forall x_{\bar{C}} \tau(\varphi) \wedge \bigwedge_{m=1}^k \exists x_{\bar{C}} \tau(\psi_m)).$$

Likewise, the formula above formalizes precisely the semantic definition of SF.

SFCL (SF1)

$$\tau([C; C_1, \dots, C_k](\varphi; \psi_1, \dots, \psi_k)) ::= \exists x_C (\forall x_{\bar{C}} \tau(\varphi) \wedge \bigwedge_{m=1}^k \exists x_{C_m \setminus C} \forall x_{\bar{C} \cup C_m} \tau(\psi_m)).$$

This formula says that C has a collective strategy (in particular, a collective action) assigned to x_C that guarantees (the translation of) φ , and is such that, when fixed, each C_m has a collective action (already fixed for the agents in $C_m \cap C$) that guarantees (the translation of) ψ_m against any behavior of the noncommitted agents, i.e., those in $\bar{C} \cup C_m$. This is precisely the semantics of SF1, defined in Section 7.1.

SFCL (SF2) $\tau(\langle\langle C_1 \rangle\rangle \psi_1; \dots; \langle\langle C_k \rangle\rangle \psi_k) ::=$

(for simplicity, assuming that C_1, \dots, C_k are pairwise disjoint)

$$\exists x_{C_1} (\forall x_{\bar{C}_1} \tau(\psi_1) \wedge \exists x_{C_2 \setminus C_1} (\forall x_{\bar{C}_1 \cup C_2} \tau(\psi_2) \wedge \dots \wedge \exists x_{C_k \setminus (C_1 \cup \dots \cup C_{k-1})} \forall x_{\bar{C}_1 \cup \dots \cup C_k} \tau(\psi_k)) \dots).$$

This formula likewise expresses the semantics of the sequential version SF2 of SF1, defined in Section 7.1. To achieve this, every time it claims existence of a joint strategy of C_i it assumes that the strategies of all agents in the already previously mentioned coalitions are already fixed, and that joint strategy must succeed against any behavior of the not-yet-committed agents only.

LCGA

$$\tau(\langle\langle\gamma\rangle\rangle) ::= \exists x_{\text{Agt}} \bigwedge_{C \subseteq \text{Agt}} \forall x_C \tau(\gamma(C)).$$

Again, this formula captures the formal semantics of $\langle\langle\gamma\rangle\rangle$ in a straightforward way.

Thus, the logic BSL is more expressive than any of the other logics introduced here. Still, it is not as expressive as the versions of SSL introduced in [23,24], which decouple the strategy variables and quantification over them from the agents to which they are assigned, and thus enable assigning the same strategy to different agents, which (in my view, is unnatural and) BSL cannot do. Still, the second-order quantification over strategies is fully enabled in BSL, so it is to be expected that BSL is not recursively axiomatizable in general. However, in the natural special case where only finite models over a fixed finite set of agents and with positional strategies are considered, there are clearly only finite such strategies in the model, so model checking of BSL in such models is decidable, whereas satisfiability is recursively enumerable. The complexity of model checking, as well as the questions of axiomatizability and decidability, of the validities of BSL in semantics restricted to such models are left to future work.

10. Concluding Remarks: Outlook and Perspectives

Logic-based strategic reasoning in socially interactive context is still in its relatively early stage of development, and the logical systems presented in this paper are just a representative sample of such special-purpose driven developments. While these systems were introduced with different motivations, they share the common purpose of formalizing natural and important patterns of strategic interaction between rational agents. Furthermore, as seen in the previous section, these can all be treated as fragments of the basic strategy logic BSL, so the natural question arises: why, instead of the many bespoke alternatives presented here, is BSL not adopted as a uniform logical language for formalizing strategic reasoning? Of course, that can be achieved, and the study of some fragments of strategy logic have shown that to be a viable approach. However, as I indicated earlier, there are computational reasons in favor of staying within the purely modal framework of the logics introduced here, where actions and strategies are not explicitly referred to and quantified over in the language, but are only present in the semantics. There are many aspects of strategic reasoning, and a logical language that can adequately capture them all would be too heavy to use, so if one only wants to reason about some such specific aspects, then one can naturally look for a minimal language that suffices for that purpose. The standoff between the two approaches—propositional modal logics vs. quantified strategy logics—is quite analogous to the standoff between plain modal logic and first-order logic, and the pros and cons for one over the other approach in both cases are very similar. Still, it would be interesting and useful to compare in deeper detail the expressiveness and the tradeoff between it and the computational complexity of these two approaches in the case of logics for strategic reasoning, to a degree closer to the very-well-explored parallels between modal logic and first-order logic (cf. relevant chapters in [46] for ample details and further references.)

Some major open problems and directions of further exploration of the logic-based strategic reasoning in socially interactive context include:

- Adding agents' *knowledge* in the semantics, and explicitly in the language, by assuming that the agents reason and act under imperfect information.
- Taking into account the *normative aspects and constraints* of the socially interactive context, including obligations, permission, and prohibitions, which socially responsible rational agents must respect in their strategic behavior.
- Analysis of the expressiveness and computational complexity of the basic strategy logic BSL which should eventually determine whether and to what extent BSL may be regarded as a viable alternative to the propositional modal approach behind the logical systems for strategic reasoning presented here.

- In particular, an interesting currently open question is whether there is a finite set of modal strategic operators with semantics that can be translated to BSL which provide expressive completeness, if not for the full language of BSL, then at least for natural and reasonably expressive fragments of it.
- Completeness results for some of the systems of axioms mentioned here, including the three main fragments of ConStR, the entire ConStR, and SFCL. Ultimately, a complete axiomatic system for BSL, if possible, or for substantially rich fragments of it.
- Finite tree-model property and decidability results for the logical systems presented here, as well as other fragments of BSL. In particular, it would be tableaux-based deductive systems and decision methods for the logics studied here, by adapting such systems developed in, e.g., [47,48].

From a more general perspective, this work can be regarded as a step towards developing a unifying and optimally rich, yet computationally feasible, technical framework for logic-based strategic reasoning of rational agents in a socially interactive context.

Funding: This research received no external funding.

Data Availability Statement: Not applicable.

Acknowledgments: I thank the anonymous reviewers for their useful comments and corrections.

Conflicts of Interest: The author declares no conflicts of interest.

Notes

- ¹ These remarks are prompted by an exchange and apparent disagreement with an anonymous reviewer of this paper, who argues that the logical frameworks discussed here do not really deal with goals, rationality, or sociality. In my understanding, the core of that disagreement is in the essentially different perspectives which we apparently have on the topic and on some basic concepts in this paper, as explained further in the text.

References

1. Parikh, R. The logic of games and its applications. *North-Holl. Math. Stud.* **1985**, *102*, 111–139.
2. Brown, M.A. On the Logic of Ability. *J. Philos. Log.* **1988**, *17*, 1–26. [[CrossRef](#)]
3. Belnap, N.; Perloff, M. Seeing to it that: A canonical form for agentives. *Theoria* **1988**, *54*, 175–199. [[CrossRef](#)]
4. Pauly, M. Logic for Social Software. Ph.D. Thesis, University of Amsterdam, Amsterdam, The Netherlands, 2001.
5. Pauly, M. A Modal Logic for Coalitional Power in Games. *J. Log. Comput.* **2002**, *12*, 149–166. [[CrossRef](#)]
6. Alur, R.; Henzinger, T.A.; Kupferman, O. Alternating-Time Temporal Logic. In Proceedings of the 38th IEEE Symposium on Foundations of Computer Science, Miami Beach, FL, USA, 19–22 October 1997; pp. 100–109.
7. Alur, R.; Henzinger, T.A.; Kupferman, O. Alternating-Time Temporal Logic. *J. ACM* **2002**, *49*, 672–713. [[CrossRef](#)]
8. Bulling, N.; Goranko, V.; Jamroga, W. Logics for Reasoning About Strategic Abilities in Multi-player Games. In *Models of Strategic Reasoning: Logics, Games, and Communities*; van Benthem, J., Ghosh, S., Verbrugge, R., Eds.; Springer: Berlin/Heidelberg, Germany, 2015; pp. 93–136.
9. Ågotnes, T.; Goranko, V.; Jamroga, W.; Wooldridge, M. Knowledge and Ability. In *Handbook of Epistemic Logic*; van Ditmarsch, H., Halpern, J., van der Hoek, W., Kooi, B., Eds.; College Publications: London, UK, 2015; pp. 543–589.
10. Bratman, M. *Intentions, Plans, and Practical Reason*; Harvard University Press: Cambridge, MA, USA, 1987.
11. Rao, A.S.; Georgeff, M.P. Modeling Rational Agents within a BDI-Architecture. In Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR'91), Cambridge, MA, USA, 22–25 April 1991; Allen, J.F., Fikes, R., Sandewall, E., Eds.; Morgan Kaufmann: Burlington, MA, USA, 1991; pp. 473–484.
12. Cohen, P.R.; Levesque, H.J. Intention is Choice with Commitment. *Artif. Intell.* **1990**, *42*, 213–261. [[CrossRef](#)]
13. Broersen, J.M. A stit-Logic for Extensive Form Group Strategies. In Proceedings of the 2009 IEEE/WIC/ACM International Conference on Web Intelligence and International Conference on Intelligent Agent Technology—Workshops, Milan, Italy, 15–18 September 2009; IEEE Computer Society: Washington, DC, USA, 2009; pp. 484–487. [[CrossRef](#)]
14. Broersen, J.M.; Herzig, A. Using STIT Theory to Talk About Strategies. In *Models of Strategic Reasoning—Logics, Games, and Communities*; Lecture Notes in Computer Science; van Benthem, J., Ghosh, S., Verbrugge, R., Eds.; Springer: Berlin/Heidelberg, Germany, 2015; Volume 8972, pp. 137–173. [[CrossRef](#)]
15. Lorini, E. Temporal STIT logic and its application to normative reasoning. *J. Appl. Non Class. Logics* **2013**, *23*, 372–399. [[CrossRef](#)]
16. Lorini, E.; Sartor, G. A STIT Logic for Reasoning About Social Influence. *Stud. Log.* **2016**, *104*, 773–812. [[CrossRef](#)]

17. Boudou, J.; Lorini, E. Concurrent Game Structures for Temporal STIT Logic. In Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2018, Stockholm, Sweden, 10–15 July 2018; André, E., Koenig, S., Dastani, M., Sukthankar, G., Eds.; International Foundation for Autonomous Agents and Multiagent Systems: Richland, SC, USA, 2018; pp. 381–389.
18. Ramanujam, R.; Simon, S. Dynamic Logic on Games with Structured Strategies. In Proceedings of the Eleventh International Conference on Principles of Knowledge Representation and Reasoning, Sydney, Australia, 16–19 September 2008; AAAI Press: Palo Alto, CA, USA, 2008; pp. 49–58.
19. van Benthem, J.; Ghosh, S.; Liu, F. Modelling simultaneous games in dynamic logic. *Synthese* **2008**, *165*, 247–268. [[CrossRef](#)]
20. van Benthem, J.; Klein, D. Logics for Analyzing Games. In *The Stanford Encyclopedia of Philosophy*, Summer 2020 ed.; Zalta, E.N., Ed.; Metaphysics Research Lab, Stanford University: Stanford, CA, USA, 2020.
21. Chatterjee, K.; Henzinger, T.A.; Piterman, N. Strategy logic. *Inf. Comput.* **2010**, *208*, 677–693. [[CrossRef](#)]
22. Mogavero, F.; Murano, A.; Vardi, M.Y. Reasoning About Strategies. In *Proceedings of the FSTTCS 2010*; Schloss Dagstuhl—Leibniz-Zentrum fuer Informatik (LIPIcs): Wadern, Germany, 2010; Volume 8; pp. 133–144.
23. Mogavero, F.; Murano, A.; Perelli, G.; Vardi, M.Y. Reasoning About Strategies: On the Model-Checking Problem. *ACM Trans. Comput. Log.* **2014**, *15*, 34:1–34:47. [[CrossRef](#)]
24. Mogavero, F.; Murano, A.; Perelli, G.; Vardi, M.Y. Reasoning about Strategies: on the Satisfiability Problem. *Log. Methods Comput. Sci.* **2017**, *13*.
25. Goranko, V.; Enqvist, S. Socially Friendly and Group Protecting Coalition Logics. In Proceedings of the International Conference on Autonomous Agents and Multiagent Systems AAMAS 2018, Stockholm, Sweden, 10–15 July 2018; International Foundation for Autonomous Agents and Multiagent Systems: Richland, SC, USA, 2018; pp. 372–380.
26. Enqvist, S.; Goranko, V. The temporal logic of coalitional goal assignments in concurrent multi-player games. *ACM Trans. Comput. Log.* **2022**, *23*, 1–58. [[CrossRef](#)]
27. Goranko, V.; Ju, F. Towards a Logic for Conditional Local Strategic Reasoning. Lecture Notes in Computer Science. In Proceedings of the 7th International Workshop LORI 2019, Chongqing, China, 18–21 October 2019; Blackburn, P., Lorini, E., Guo, M., Eds.; Springer: Berlin/Heidelberg, Germany, 2019; Volume 11813, pp. 112–125.
28. Goranko, V.; Ju, F. A Logic for Conditional Local Strategic Reasoning. *J. Log. Lang. Inf.* **2022**, *31*, 167–188. [[CrossRef](#)]
29. Naumov, P.; Yuan, Y. Intelligence in Strategic Games. *J. Artif. Intell. Res.* **2021**, *71*, 521–556. [[CrossRef](#)]
30. Naumov, P.; Yew, R. Ethical Dilemmas in Strategic Games. In Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Virtual, 2–9 February 2021; AAAI Press: Palo Alto, CA, USA, 2021; pp. 11613–11621.
31. Naumov, P.; Tao, J. Two Forms of Responsibility in Strategic Games. In Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event, Montreal, QC, Canada, 19–27 August 2021; Zhou, Z., Ed.; pp. 1989–1995. [[CrossRef](#)]
32. Naumov, P.; Tao, J. Blameworthiness in Security Games. In Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, New York, NY, USA, 7–12 February 2020; AAAI Press: Palo Alto, CA, USA, 2020; pp. 2934–2941. [[CrossRef](#)]
33. van der Hoek, W.; Wooldridge, M. Cooperation, knowledge, and time: Alternating-time temporal epistemic logic and its applications. *Stud. Log.* **2004**, *75*, 125–157. [[CrossRef](#)]
34. van Benthem, J.; Ghosh, S.; Verbrugge, R. (Eds.) *Models of Strategic Reasoning—Logics, Games, and Communities*; LNCS Volume 8972; Springer: Berlin/Heidelberg, Germany, 2015.
35. Blackburn, P.; de Rijke, M.; Venema, Y. *Modal Logic*; Cambridge University Press: Cambridge, UK, 2001.
36. Demri, S.; Goranko, V.; Lange, M. *Temporal Logics in Computer Science*; Cambridge Tracts in Theoretical Computer Science; Cambridge University Press: Cambridge, UK, 2016.
37. van der Hoek, W.; Pauly, M. Modal Logic for Games and Information. In *Handbook of Modal Logic*; Blackburn, P., van Benthem, J., Wolter, F., Eds.; Elsevier: Amsterdam, The Netherlands, 2006; pp. 1077–1148.
38. Ågotnes, T.; Goranko, V.; Jamroga, W. Alternating-time Temporal Logics with Irrevocable Strategies. In Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge, Brussels, Belgium, 25–27 June 2007; Samet, D., Ed.; pp. 15–24.
39. Ågotnes, T.; Goranko, V.; Jamroga, W. *Strategic Commitment and Release in Logics for Multi-Agent Systems (Extended Abstract)*; Technical Report IfI-08-01; Clausthal University of Technology: Clausthal-Zellerfeld, Germany, 2008.
40. Laroussinie, F.; Markey, N.; Oreiby, G. On the Expressiveness and Complexity of ATL. *Log. Methods Comput. Sci.* **2008**, *4*. [[CrossRef](#)]
41. Goranko, V. Coalition Games and Alternating Temporal Logics. In Proceedings of the 8th Conference on Theoretical Aspects of Rationality and Knowledge (TARK VIII), Siena, Italy, 8–10 July 2001; van Benthem, J., Ed.; Morgan Kaufmann: Burlington, MA, USA, 2001; pp. 259–272.
42. Brihaye, T.; Lopes, A.D.C.; Laroussinie, F.; Markey, N. ATL with Strategy Contexts and Bounded Memory. In Proceedings of the Logical Foundations of Computer Science: International Symposium, LICS 2009, Deerfield Beach, FL, USA, 3–6 January 2009; Nerode, A., Ed.; Springer: Berlin/Heidelberg, Germany, 2009; LNCS Volume 5407; pp. 92–106.
43. Walther, D.; van der Hoek, W.; Wooldridge, M. Alternating-time Temporal Logic with Explicit Strategies. In Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge TARK XI, Brussels, Belgium, 25–27 June 2007; Samet, D., Ed.; Presses Universitaires de Louvain: Louvain, Belgium, 2007; pp. 269–278.

44. Abdou, J.; Keiding, H. *Effectivity Functions in Social Choice*; Springer: Berlin/Heidelberg, Germany, 1991.
45. van Benthem, J.; Bezhanishvili, N.; Enqvist, S.; Yu, J. Instantial Neighbourhood Logic. *Rev. Symb. Log.* **2017**, *10*, 116–144. [[CrossRef](#)]
46. Blackburn, P.; van Benthem, J.; Wolter, F. (Eds.) *Handbook of Modal Logic*; Studies in Logic and Practical Reasoning; North-Holland: Amsterdam, The Netherlands, 2007; Volume 3.
47. Goranko, V.; Shkatov, D. Tableau-based decision procedures for logics of strategic ability in multiagent systems. *ACM Trans. Comput. Log.* **2009**, *11*, 1–49. [[CrossRef](#)]
48. Cerrito, S.; David, A.; Goranko, V. Optimal Tableau Method for Constructive Satisfiability Testing and Model Synthesis in the Alternating-Time Temporal Logic ATL^+ . *ACM Trans. Comput. Log.* **2015**, *17*, 4:1–4:34. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.