

Reasoning about cooperation, actions and preferences

Lena Kurzen

Received: 30 November 2008 / Accepted: 6 April 2009 / Published online: 21 April 2009
© The Author(s) 2009. This article is published with open access at Springerlink.com

Abstract In this paper, a logic for reasoning about coalitional power is developed which explicitly represents agents' preferences and the actions by which the agents can achieve certain results. A complete axiomatization is given and its satisfiability problem is shown to be decidable and EXPTIME-hard.

Keywords Modal logic · Multi-agent systems · Collective agency

1 Introduction

Cooperation of agents plays a major role in many fields such as computer science, economics, politics, social sciences and philosophy. Agents can decide to cooperate and to form groups in order to share complementary resources or because as a group they can achieve something better than individually.

When analyzing interactive situations involving multiple agents, we are interested in what results agents can achieve—individually or together as groups. There can be many ways how agents can achieve some result. They can differ significantly, e.g. with respect to their feasibility, costs or side-effects. Hence, it is not only relevant what results groups of agents can achieve but also *how* exactly they can do so. In other words, plans and actions also play a central role if we want to reason about cooperation in an explicit way. However, cooperative ability of agents expressed only in terms of results and actions that lead to these results does not tell us *why* a group of agents would actually decide to achieve a certain result. We also need to take into account the preferences based on which the agents decide what to do. Summarizing, we can say that in interactive situations, the following three questions are of interest:

L. Kurzen (✉)
Universiteit van Amsterdam, Amsterdam, The Netherlands
e-mail: L.M.Kurzen@uva.nl

- **What** results can groups of agents achieve?
- **How** can they achieve something?
- **Why** would they want to achieve a certain result?

The above considerations show that coalitional power, actions/plans and preferences play a major role in interactive situations and are moreover tightly connected. Thus, we argue that a formal theory for reasoning about agents' cooperative abilities in an explicit way should also take into account actions/plans of agents and their preferences.

Modal logics have been used to develop formal models for reasoning about each of these aspects—mostly separately. Coalitional power has mainly been investigated within the frameworks of ATL (Alur et al. 1998), coalition logic (Pauly 2002a) and their extensions.

Recently, there have been some attempts to develop logics for reasoning about coalitional power that also take into account either agents' preferences or actions. One group of such logics looks at cooperation from the perspective of cooperative games (Ågotnes et al. 2007a). In a non-cooperative setting preferences and strategic abilities have been considered in van Otterloo et al. (2004). Another path that has been taken in order to make coalitional power more explicit is to combine cooperation logics with (fragments of) action logics (Sauro et al. 2006; Borgo 2007; Walther et al. 2007).

In this paper, a logic for reasoning about cooperation, actions and preferences (CLA+P) is developed, which is obtained by combining the cooperation logic with actions CLA (Sauro et al. 2006) with a preference logic (van Benthem et al. 2005; van Benthem et al. 2007). Soundness and completeness are shown and the logic's expressivity and computational complexity are investigated.

The remainder of this paper is structured as follows. Section 2 gives an overview of CLA. In Sect. 3, CLA+P is developed, a complete axiomatization is given and its expressivity is discussed. Section 4 gives complexity results and Sect. 5 concludes this work.

2 Cooperation logic with actions

In this section, we briefly present the CLA developed by Sauro et al. (2006), which will be extended in the next section by combining it with a preference logic. The idea of CLA is to make coalitional power explicit by expressing it in terms of the ability to perform actions instead of expressing it directly in terms of the ability to achieve certain outcomes. CLA is a modular modal logic, consisting of an environment module for reasoning about actions and their effects, and an agents module for reasoning about agents' abilities to perform actions. By combining both modules, a framework is obtained in which cooperative ability can be made more explicit.

The environment is modeled as a transition system whose edges are labeled with sets of atomic actions.

Definition 1 (*Environment Model* (Sauro et al. 2006)) An environment model is a set-labelled transition system

$$E = \langle S, Ac, (\rightarrow)_{A \subseteq Ac}, V \rangle.$$

S is a set of states, Ac is a finite set of atomic actions, $\rightarrow_{A \subseteq Ac} S \times S$ and V is a propositional valuation. Each \rightarrow_A is required to be serial.

The intuition being $s \rightarrow_A t$ is that if in s all actions and only the actions in A are performed concurrently, then the next state can be t .

Then a modal language is defined with modalities $[\alpha]$, for α being a propositional formula built from atomic actions. The intended meaning of $[\alpha]\varphi$ is that every transition \rightarrow_A such that $A \models \alpha$ (using the satisfaction relation of propositional logic¹) leads to a φ -state:

$$E, s \models [\alpha]\varphi \text{ iff } \forall A \subseteq Ac, t \in S : \text{if } A \models \alpha \text{ and } s \rightarrow_A t \text{ then } E, t \models \varphi.$$

Due to space restrictions, we cannot go into the underlying philosophy of actions but refer the reader to [Broersen \(2003\)](#) for a detailed discussion of action logics. The restriction to a finite set of actions is reasonable for modelling many concrete situations and also ensures that we have a finite axiomatization.

An environment logic Λ^E is developed, which is sound and complete ([Sauro et al. 2006](#)). It contains seriality axioms and the **K** axiom for each modality $[\alpha]$, for α being consistent. The environment logic can then be used for reasoning about the effects of concurrent actions.

Then an agents module is developed for reasoning about the ability of (groups of) agents to act. Each agent is assigned a set of atomic actions and a group is assigned the actions its members can perform.

Definition 2 (*Agents Model* ([Sauro et al. 2006](#))) An agents model is a triple $\langle Ag, Ac, \text{act} \rangle$, where Ag is a set of agents, Ac is a set of atomic actions and act is a function $\text{act}: Ag \rightarrow \mathcal{P}(Ac)$ such that $\bigcup_{i \in Ag} \text{act}(i) = Ac$. For $G \subseteq Ag$, define $\text{act}(G) := \bigcup_{i \in G} \text{act}(i)$.

We are also interested in agents’ abilities to force more complex actions. A language is developed with expressions $\langle [G] \rangle \alpha$, meaning that the group G can force that a concurrent action is performed that satisfies α . This means that G can perform some set of atomic actions such that no matter what the other agents do, the resulting set of actions satisfies α .

$$\langle Ag, Ac, \text{act} \rangle \models \langle [G] \rangle \alpha \text{ iff } \exists A \subseteq \text{act}(G) : \forall B \subseteq \text{act}(Ag \setminus G) : A \cup B \models \alpha.$$

Then a cooperation logic for actions is developed, which is very much in the style of coalition logic ([Pauly 2002a](#))—the main difference being that it is concerned with the cooperative ability to force *actions*.

Definition 3 (*Coalition Logic for Actions* ([Sauro et al. 2006](#))) The coalition logic for actions Λ^A is defined to be the logic derived from the following set of axioms, with rule of inference modus ponens.

¹ That is, $A \models a$ iff $a \in A$, $A \models \neg\alpha$ iff $A \not\models \alpha$, and $A \models \alpha \wedge \beta$ iff $A \models \alpha$ and $A \models \beta$.

- (1) $\langle [G] \top \rangle$, for all $G \subseteq Ag$,
- (2) $\langle [G] \alpha \rightarrow \neg \langle [Ag \setminus G] \neg \alpha \rangle \rangle$,
- (3) $\langle [G] \alpha \rightarrow \langle [G] \beta \rangle$ if $\vdash \alpha \rightarrow \beta$ in propositional logic,
- (4) $\langle [G] a \rightarrow \bigvee_{i \in G} \langle [i] \rangle a \rangle$ for all $G \subseteq Ag$ and atomic $a \in Ac$,
- (5) $\langle ([G_1] \alpha \wedge [G_2] \beta) \rightarrow [G_1 \cup G_2] (\alpha \wedge \beta) \rangle$, for $G_1 \cap G_2 = \emptyset$,
- (6) $\langle ([G] \alpha \wedge [G] \beta) \rightarrow [G] (\alpha \wedge \beta) \rangle$ if α and β have no common atomic actions,
- (7) $\langle [G] \neg a \rightarrow \langle [G] a \rangle$ for atomic $a \in Ac$,
- (8) $\langle [G] \alpha \rightarrow \bigvee \{ \langle [G] \bigwedge \Psi \rangle \mid \Psi \text{ is a set of literals such that } \bigwedge \Psi \rightarrow \alpha \} \rangle$.

Axiom 5 says how groups can join forces. The coalition logic for actions is sound and complete with respect to the class of agents models (Sauro et al. 2006).

Next, agents are introduced as actors into the environment. This is done by combining the environment models with the agents models. Then the agents can perform actions which have the effect of changing the current state of the environment.

Definition 4 (*Multi-agent System* (Sauro et al. 2006)) A multi-agent system (MaS) is a tuple

$$M = \langle S, Ac, (\rightarrow)_{A \subseteq Ac}, V, Ag, \mathbf{act} \rangle,$$

where $\langle S, Ac, (\rightarrow)_{A \subseteq Ac}, V \rangle$ is an environment model and $\langle Ac, Ag, \mathbf{act} \rangle$ an agents model.

Now, we can reason about what states of affairs groups can achieve by performing certain actions. The corresponding language contains all expressions of the previously defined logics and additionally expressions for saying that a group has the power to achieve φ which means that the group can make the system move into a state where φ is true.

Definition 5 (*Language for MaS* (Sauro et al. 2006)) The language for multi-agent systems \mathcal{L}_{cla} is generated by the following grammar:

$$\varphi ::= p \mid \varphi \wedge \varphi \mid \neg \varphi \mid [\alpha] \varphi \mid \langle [G] \rangle \alpha \mid \langle [G] \rangle \varphi$$

for $G \subseteq Ag$ and α being an action expression.

$\langle [G] \rangle \varphi$ means that G can force φ , i.e. G can perform a set of actions such that no matter what the other agents do, the system moves to a φ -state.

$$M, s \models \langle [G] \rangle \varphi \text{ iff } \exists A \subseteq \mathbf{act}(G) \text{ such that } \forall B \subseteq \mathbf{act}(Ag \setminus G), t \in S : \text{ if } s \rightarrow_{A \cup B} t, \text{ then } M, t \models \varphi.$$

A complete axiomatization is obtained by combining the environment logic and the coalition logic for agents by adding two interaction axioms.

Definition 6 (*Cooperation Logic with Actions* (Sauro et al. 2006)) The cooperation logic with actions Λ^{CLA} combines the environment logic Λ^E and the coalition logic for actions Λ^A by adding

- (1) $\langle [G] \rangle \alpha \wedge [\alpha] \varphi \rightarrow \langle [G] \rangle \varphi$,
- (2) $\langle [G] \rangle \varphi \rightarrow \bigvee \{ \langle [G] \rangle \alpha \wedge [\alpha] \varphi \mid \alpha \text{ is the conjunction of a set of atomic actions or their negations} \}$.

CLA provides us with a formal framework for reasoning about what states of affairs groups of agents can achieve and how they can do so. For a detailed discussion of CLA, the reader is referred to [Sauro et al. \(2006\)](#). Now, we proceed by adding preferences to CLA.

3 Cooperation logic with actions and preferences

In this section, a logic for reasoning about cooperation, actions and preferences is developed. This is done by adding a preference logic to CLA. For a more detailed discussion and proofs, see [Kurzen \(2007\)](#).

3.1 Preference logic

There are various ways how preferences can be added to a logic for cooperation and actions. They could range over the actions that he can perform. Alternatively, we can think of each agent having preferences over the set of successor states of the current state.

In the current work, we consider preferences of individual agents ranging over the states of the environment. This is reasonable since by performing actions the agents can change the current state of the environment, and the preferences over those states can be seen as the base of how the agents decide how to act. Such a preference relation can also be lifted in several ways to one over formulas ([van Benthem et al. 2005, 2007](#)).

Definition 7 (*Preference Model* ([van Benthem et al. 2005](#))) A preference model is a tuple

$$M^P = \langle S, Ag, \{ \preceq_i \}_{i \in Ag}, V \rangle,$$

where S is a set of states, Ag is a set of agents, for each $i \in Ag$, $\preceq_i \subseteq S \times S$ is reflexive and transitive, and V is a propositional valuation.

We use a fragment of the preference language developed by [van Benthem et al. \(2007\)](#). It has strict and non-strict preference modalities.

Definition 8 (*Preference Language*) Given a set of propositional variables and a finite set of agents Ag , define the preference language \mathcal{L}_p to be the language generated by the following syntax:

$$\varphi := p \mid \neg \varphi \mid \varphi \vee \varphi \mid \diamond^{\preceq_i} \varphi \mid \diamond^{<_i} \varphi.$$

$\diamond^{\preceq_i} \varphi$ says that there is a state satisfying φ that agent i considers to be at least as good as the current one. The semantics is defined as follows.

$$M^P, s \models \diamond^{\preceq_i} \varphi \quad \text{iff} \quad \exists t : s \preceq_i t \quad \text{and} \quad M^P, t \models \varphi.$$

Analogously for $\diamond^{<_i} \varphi$. The preference relation \preceq is a preorder and $<$ is its largest irreflexive subrelation. Hence, the following axiomatization.

Definition 9 (*Preference Logic Λ^P*) For a given set of agents Ag , let Λ^P be the logic generated by the following axioms for each $i \in Ag$: For \diamond^{\preceq_i} and $\diamond^{<_i}$, we have Duality and **K** and for \diamond^{\preceq_i} also reflexivity and transitivity axioms. Moreover, there are four interaction axioms:

- (1) $\diamond^{<_i} \varphi \rightarrow \diamond^{\preceq_i} \varphi,$
- (2) $\diamond^{\preceq_i} \diamond^{<_i} \varphi \rightarrow \diamond^{<_i} \varphi,$
- (3) $\diamond^{<_i} \diamond^{\preceq_i} \varphi \rightarrow \diamond^{<_i} \varphi,$
- (4) $\varphi \wedge \diamond^{\preceq_i} \psi \rightarrow (\diamond^{<_i} \psi \vee \diamond^{\preceq_i} (\psi \wedge \diamond^{\preceq_i} \varphi)).$

The inference rules are modus ponens, necessitation and substitution.

Transitivity for $\diamond^{<_i}$ follows. We show soundness and completeness using the *bull-dozing* technique (Blackburn et al. 2001) to deal with $<$. For details, we refer to van Benthem et al. (2007).

Theorem 1 Λ^P is sound and complete with respect to the class of preference models.

Proof Follows from Theorem 3.9 in van Benthem et al. (2007). □

The preference logic is able to distinguish between weak and strict preference. This plays a major role in many concepts for reasoning about interaction in multi-agent systems. Due to the logic’s simplicity it still has the modal character and talks about preferences from a local perspective. An additionally \diamond^{\succeq} modality would have increased the expressivity but would have resulted in a global existential modality with respect to all comparable states.

3.2 Environment logic with preferences

As an intermediate step towards a logic for cooperation, actions and preferences, we combine the preference logic and the environment logic. Their models are combined by identifying their states. The agents’ preferences range over the states of the environment. The agents cannot act in the environment, but can rather be seen as observing the environment from the outside while having preferences over its states.

Definition 10 (*Environment with Preferences*) An environment model with preferences is a tuple

$$E^{\preceq} = \langle S, Ac, (\rightarrow)_{A \subseteq Ac}, \{\preceq_i\}_{i \in Ag}, V \rangle,$$

where $\langle S, Ac, (\rightarrow)_{A \subseteq Ac}, \{\preceq_i\}_{i \in Ag}, V \rangle$ is an environment model and $\langle S, Ag, \{\preceq_i\}_{i \in Ag}, V \rangle$ is a preference model.

We combine the languages for environment and preferences and add expressions saying that “ i (strictly) prefers every α -accessible state”. This will later allow us to express statements saying by which actions groups can(not) achieve an outcome better for (some of) its members.

Convention 1 We will write the symbol \triangleleft in statements that hold for both \preceq and \prec , each uniformly substituted for \triangleleft .

Definition 11 (*Environment Language with Preferences*) The language \mathcal{L}_{ep} contains all expressions of the environment language and the preference language and additionally formulas of the forms $\alpha \subseteq \preceq_i$ and $\alpha \subseteq \prec_i$, for α being an action expression.

Boolean combinations and expressions of previously defined languages are interpreted as usual. For the newly introduced expressions, we have:

$$E^{\preceq}, s \models \alpha \subseteq \triangleleft_i \text{ iff } \forall A \subseteq Ac, t \in S : \text{if } s \rightarrow_A t \text{ and } A \models \alpha \text{ then } s \triangleleft_i t.$$

$\alpha \subseteq \triangleleft_i$ cannot be defined in the preference language and the environment language, because $\alpha \subseteq \preceq_i$ says that each state accessible by an α -transition is also accessible by \preceq . Thus, we would have to be able to refer to particular states. Therefore, we add two inference rules.

$$\text{(PREF-ACT)} \frac{\Box^{\preceq_i} \varphi \rightarrow [\alpha]\varphi}{\alpha \subseteq \preceq_i} \quad \text{(STRICT PREF-ACT)} \frac{\Box^{\prec_i} \varphi \rightarrow [\alpha]\varphi}{\alpha \subseteq \prec_i}$$

In order to obtain a complete axiomatization, two axioms are added which correspond to the converse of the inference rules.

Theorem 2 Let Λ^{EP} be the logic generated by all axioms of the environment logic Λ^E , all axioms of the preference logic Λ^P , and

- (1) $\alpha \subseteq \preceq_i \rightarrow (\Box^{\preceq_i} \varphi \rightarrow [\alpha]\varphi)$,
- (2) $\alpha \subseteq \prec_i \rightarrow (\Box^{\prec_i} \varphi \rightarrow [\alpha]\varphi)$.

The inference rules are modus ponens, substitution, PREF-ACT and STRICT PREF-ACT. Then Λ^{EP} is sound and complete with respect to the class of environment models with preferences.

Proof Soundness is straightforward, and completeness follows from completeness of the sublogics and the closure under the new rules. □

In Λ^{EP} , the performance of concurrent actions changes the current state of the system also with respect to the agents’ “happiness”: transitions can also be a transitions up or down in the agents’ preference orderings.

3.3 Cooperation logic with actions and preferences

Now, we introduce agents as actors by combining the environment models with preferences with agents models. The resulting model is then called a multi-agent system with preferences (henceforth MaSP).

Definition 12 (*Multi-agent System with Preferences*) A multi-agent system with preferences (MaSP) M^{\preceq} is a tuple

$$M^{\preceq} = \langle S, Ac, (\rightarrow)_{A \subseteq Ac}, Ag, \mathbf{act}, \{\preceq_i\}_{i \in Ag}, V \rangle,$$

where $\langle S, Ac, (\rightarrow)_{A \subseteq Ac}, V, Ag, \mathbf{act} \rangle$ is a MaS, $\langle S, Ag, \{\preceq_i\}_{i \in Ag}, V \rangle$ is a preference model and $\langle S, Ac, (\rightarrow)_{A \subseteq Ac}, \{\preceq_i\}_{i \in Ag}, V \rangle$ is an environment with preferences.

Remark 1 Note that given a deterministic MaSP in which each preference relation \preceq_i is total, we can consider each state s as having a strategic game (Osborne and Rubinstein 1994) \mathcal{G}_s attached to it.

$$\begin{aligned} \mathcal{G}_s &= \langle Ag, (\mathcal{P}(\mathbf{act}(i)))_{i \in Ag}, (\lesssim_i)_{i \in Ag} \rangle, \\ \times_{i=1}^n A_i \lesssim_i \times_{i=1}^n A'_i &\text{ iff } t \preceq_i t' \text{ for } s \rightarrow \bigcup_{i \in Ag} A_i t \text{ and } s \rightarrow \bigcup_{i \in Ag} A'_i t'. \end{aligned}$$

Next, we introduce two expressions saying that a group can force the system to move into a φ -state that some agent (strictly) prefers.

Definition 13 (*Language \mathcal{L}_{cla+p}*) The language \mathcal{L}_{cla+p} extends \mathcal{L}_{cla} by formulas of the form

$$\diamond^{\preceq_i} \varphi \mid \diamond^{\prec_i} \varphi \mid \alpha \subseteq \preceq_i \mid \alpha \subseteq \prec_i \mid \langle [G^{\preceq_i}] \rangle \varphi \mid \langle [G^{\prec_i}] \rangle \varphi.$$

The first four expressions are interpreted as in the environment logic with preferences and for the last two we have the following.

$$\begin{aligned} M^{\preceq}, s \models \langle [G^{\prec_i}] \rangle \varphi &\text{ iff } \exists A \subseteq \mathbf{act}(G) : \forall B \subseteq \mathbf{act}(Ag \setminus G), t \in S : \\ &\text{if } s \rightarrow_{A \cup B} t, \text{ then } M^{\preceq}, t \models \varphi \text{ and } s \prec_i t. \end{aligned}$$

Let us now look at how coalitional power to achieve an improvement for an agent is made explicit in CLA+P. We can show that $\langle [G^{\prec_i}] \rangle \varphi$ is equivalent to the existence of an action expression α that G can force such that all transitions of type α lead to a φ -state preferred by i .

Observation 1 *Given a MaSP M^{\preceq} and a state s of its environment, $M^{\preceq}, s \models \langle [G^{\prec_i}] \rangle \varphi$ iff there exists an action expression α such that*

$$M^{\preceq}, s \models \langle [G] \rangle \alpha \wedge [\alpha] \varphi \wedge (\alpha \subseteq \prec_i).$$

Proof Analogous to that of Observation 14 in Sauro et al. (2006). Use the action expression $\bigwedge \Phi(A, G) := \bigwedge (A \cup \{\neg a \mid a \in (\mathbf{act}(G) \setminus A), a \notin \mathbf{act}(Ag \setminus G)\})$ for A being the “witness” of G ’s ability to force φ . □

Now we need axioms establishing a relationship between the newly added formulas and the expressions of the sublogics.

Definition 14 (*Cooperation Logic with Actions and Preferences*) Λ^{CLA+P} is defined to be the smallest logic generated by the axioms of the cooperation logic with actions, the environment logic with preferences and

- (1) $\langle [G] \rangle \alpha \wedge [\alpha] \varphi \wedge (\alpha \subseteq \triangleleft_i) \rightarrow \langle [G^{\triangleleft_i}] \rangle \varphi$,
- (2) $\langle [G^{\preceq_i}] \rangle \varphi \rightarrow \bigvee \{ \langle [G] \rangle \alpha \wedge [\alpha] \varphi \wedge (\alpha \subseteq \preceq_i) \mid \alpha \text{ is a conjunction of action literals} \}$,
- (3) $\langle [G^{\prec_i}] \rangle \varphi \rightarrow \bigvee \{ \langle [G] \rangle \alpha \wedge [\alpha] \varphi \wedge (\alpha \subseteq \prec_i) \mid \alpha \text{ is a conjunction of action literals} \}$.

The inference rules are modus ponens, necessitation for action modalities and preference modalities (\Box^{\preceq_i} , \Box^{\prec_i}), substitution of logical equivalents, PREF—ACT and STRICT PREF—ACT.

Theorem 3 *The logic Λ^{CLA+P} is sound and complete with respect to the class of MaSP’s.*

Proof Soundness of the axioms is straightforward and completeness follows from completeness of the sublogics. □

3.4 Expressivity of CLA+P

We now show that in CLA+P, we can express some concepts relevant for reasoning about game-like interaction in multi-agent systems.

3.4.1 Stability

Given a MaSP, $M^{\preceq} = \langle S, Ac, (\rightarrow_A)_{A \subseteq Ac}, Ag, \mathbf{act}, \{ \preceq_i \}_{i \in Ag}, V \rangle$, the following formula characterizes the states that are *individually stable* (*group stable*), i.e. no individual (group) has the power to achieve a strict improvement (for all its members).

$$\psi_{ind. \text{ stable}} := \bigwedge_{i \in Ag} \neg \langle [i]^{\prec_i} \rangle \top.$$

$$\psi_{gr. \text{ stable}} := \bigwedge_{G \subseteq Ag} \bigwedge_{A \subseteq \mathbf{act}(G)} \left(\bigvee_{i \in G} \neg \left(\left(\bigwedge \Phi(A, G) \right) \subseteq \prec_i \right) \right).$$

3.4.2 Dictatorship

We can express that an agent d is a (strong) dictator in the sense that coalitions can only achieve what d (strictly) prefers.

$$\psi_{d= \text{dict.}} := \bigwedge_{G \subseteq Ag} \bigwedge_{A \subseteq \mathbf{act}(G)} \left(\left(\bigwedge \Phi(G, A) \right) \subseteq \triangleleft_d \right).$$

Then, we can also say that there is *no* (strong) dictator:

$$\psi_{no \text{ dict.}} := \bigwedge_{i \in Ag} \neg \left(\bigwedge_{G \subseteq Ag} \bigwedge_{A \subseteq \mathbf{act}(G)} \left(\left(\bigwedge \Phi(G, A) \right) \subseteq \triangleleft_i \right) \right).$$

3.4.3 Enforcing Unanimity

In some situations we might want to impose the condition that groups should only be able to achieve something if they can do so in a way that all members will be better off:

$$\langle [G] \rangle \varphi \rightarrow \left(\bigvee_{A \subseteq \text{act}(G)} \left(\bigwedge_{i \in G} \left(\left(\bigwedge \Phi(A, G) \right) \subseteq \prec_i \right) \wedge \left[\bigwedge \Phi(A, G) \right] \varphi \right) \right)$$

Note that the length of the last four formulas is exponential in the number of atomic actions (and agents).

3.5 CLA+P and coalition logic

Let us now briefly discuss the relation between CLA+P and coalition logic (CL) (Pauly 2002a) in order to illustrate how CLA+P builds upon existing frameworks for reasoning about coalitional power and how the underlying actions that are only implicitly represented in the semantics of CL are made explicit. Given a fixed set of agents Ag , a coalition model $M = \langle (S, E), V \rangle$ with S being a set of states, $E : S \rightarrow (\mathcal{P}(Ag) \rightarrow \mathcal{P}(\mathcal{P}(S)))$ being a playable effectivity function and V being a propositional valuation, we can use Theorem 3.2 of Pauly (2002a) and obtain a corresponding *game frame*, i.e. each state has an associated strategic game form in which each outcome corresponds to some accessible state. Looking back at Remark 1, it is now easy to see how we can construct a corresponding MaS: We take the same set of states and add actions for each of the strategies in the attached games and define the accessibility relation in accordance with the outcome function. Finally, we can define a translation τ of formulas of CL to those of CLA in a straightforward way: formulas $[G]\varphi$ talking about coalitional power in CL are translated into $\langle [G] \rangle \tau(\varphi)$.

If we add preferences to the game forms in the game frame we obtained from the coalition model, then we can transform it into a MaSP in an analogous way.

This shows that the framework of CLA+P is a natural way to make coalitional power as modelled in CL and its extensions more explicit.

4 Complexity of CLA+P

In this section, we analyze the complexity of SAT of CLA+P.

4.1 Decidability of CLA+P

We show that SAT of CLA+P is decidable. The first step is to show that only a restricted class of models of CLA+P needs to be checked.

We start by looking at how we can restrict the class of models with respect to the set of agents. Let $Ag(\varphi)$ denote the set of agents occurring in φ . Now, we ask: Is

every satisfiable φ also satisfiable in a MaSP with set of agents $Ag(\varphi)$? In coalition logic, the answer is negative: the formula $\varphi' = \neg\langle [G] \rangle \{1\} p \wedge \neg\langle \{1\} \rangle q \wedge \langle \{1\} \rangle (p \vee q)$ is only satisfiable in models with at least two agents (Pauly 2002b). However, as in CLA+P the environment models can be nondeterministic, here φ' can indeed be satisfied in a model with only one agent, as the reader can check.

It can be shown that every satisfiable formula $\varphi \in \mathcal{L}_{cla+p}$ is satisfiable in a MaSP with set of agents $Ag(\varphi) \cup \{k\}$, for k being a new agent. k takes the role of all opponents of $Ag(\varphi)$ collapsed into one: k gets the ability to perform exactly the actions that agents not occurring in φ can perform as a group.

Theorem 4 *Every satisfiable formula $\varphi \in \mathcal{L}_{cla+p}$ is satisfiable in the class of MaSP’s with at most $|Ag(\varphi)| + 1$ many agents.*

Proof Assume that $M^{\preceq} = \langle S, Ac, (\rightarrow_A)_{A \subseteq Ac}, Ag, act, \{\preceq_i\}_{i \in Ag}, V \rangle$ satisfies φ . If $Ag \supset Ag(\varphi)$, we construct $M'^{\preceq'} = \langle S, Ac, (\rightarrow)_{A \subseteq Ac}, Ag(\varphi) \cup \{k\}, act', \{\preceq'_i\}_{i \in Ag(\varphi) \cup \{k\}}, V \rangle$, with $act'(k) = \bigcup_{j \in Ag \setminus Ag(\varphi)} act(j)$ and $act'(i) = act(i)$ for $i \neq k$. The preferences are defined as follows: $\preceq'_i = \preceq_i$ for $i \in Ag(\varphi)$ and $\preceq'_k = S \times S$. By induction, we can show that $M^{\preceq}, s \models \varphi$ iff $M'^{\preceq'}, s \models \varphi$. The case where φ is of the form $\langle [G] \rangle \alpha$ follows from the definition of act' . Then the other cases involving coalition modalities follow. □

Next, we want to know how many actions a model needs for satisfying some formula. Consider e.g. $\varphi = \langle [G] \rangle (p \wedge q) \wedge \langle [G] \rangle (\neg p \wedge q) \wedge \langle [G] \rangle (\neg p \wedge \neg q)$. It is only satisfiable in models with $|Ac| \geq 2$. The main task is to find “witnesses” for formulas of the form $\langle [G] \rangle \psi$ in terms of concurrent actions. We can show that every satisfiable φ is satisfiable in a MaSP whose set of atomic actions consists of those in φ , one additional one (a dummy for ensuring that each agent can perform an action), and for every subformula $\langle [G] \rangle \psi$ or $\langle [G^{<i}] \rangle \psi$, one action for each of G ’s members.

The key step in transforming a model satisfying a formula φ into one whose set of actions satisfies the above condition is to appropriately define the action distribution and the accessibility relations. For every α occurring in φ , we have to ensure that two states are related by an α -transition in the new model iff they were in the original one. Additionally, for formulas $\langle [G] \rangle \psi$ and $\langle [G^{<i}] \rangle \psi$, the set of actions introduced for them serves for making explicit how G can force φ .

Theorem 5 *Every satisfiable formula $\varphi \in \mathcal{L}_{cla+p}$ is satisfiable in a MaSP with at most $|Ac(\varphi)| + (\sum_{\langle [G] \rangle \psi \in Sub(\varphi)} |G|) + (\sum_{\langle [G^{\preceq i}] \rangle \psi \in Sub(\varphi)} |G|) + (\sum_{\langle [G^{<i}] \rangle \psi \in Sub(\varphi)} |G|) + 1$ many actions.*

Proof Assume that $M^{\preceq} = \langle S, Ac, (\rightarrow_A)_{A \subseteq Ac}, Ag, act, \{\preceq_i\}_{i \in Ag}, V \rangle$ satisfies φ . We construct a model $M'^{\preceq'} = \langle S, Ac', (\rightarrow')_{A' \subseteq Ac'}, Ag, act', \{\preceq'_i\}_{i \in Ag}, V \rangle$ as follows.

$$Ac' := Ac(\varphi) \cup \bigcup_{\langle [G] \rangle \psi \in Sub(\varphi)} A_G \psi \cup \bigcup_{\langle [G^{\preceq i}] \rangle \psi \in Sub(\varphi)} A_{G^{\preceq i}} \psi \cup \bigcup_{\langle [G^{<i}] \rangle \psi \in Sub(\varphi)} A_{G^{<i}} \psi \cup \{\hat{a}\}.$$

$A_G \psi$ and $A_{G^{<i}} \psi$ consist of newly introduced actions $a_{G\psi j}$, and $a_{G^{<i}\psi j}$, respectively, for each $j \in G$. Action abilities are distributed as follows:

$$\text{act}'(i) := (\text{act}(i) \cap \text{Ac}(\varphi)) \cup \{\hat{a}\} \cup \{a_{Gi} \mid \langle [G] \rangle \psi \in \text{Sub}(\varphi) \text{ or } \langle [G^{\triangleleft_i}] \rangle \psi \in \text{Sub}(\varphi), \text{ for } i \in G\}.$$

For defining the accessibility relation $\rightarrow_{A' \subseteq \text{Ac}'}$, we first define for any state s its set of successors.

$t \in T_{A'}^s$ iff the following conditions are satisfied:

- (1) $\forall [\alpha]\psi \in \text{Sub}(\varphi)$ such that $A' \models \alpha$: If $M^{\preceq}, s \models [\alpha]\psi$, then $M^{\preceq}, t \models \psi$,
- (2) $\forall \alpha \subseteq \triangleleft_i \in \text{Sub}(\varphi)$ such that $A' \models \alpha$: If $M^{\preceq}, s \models \alpha \subseteq \triangleleft_i$, then $s \triangleleft_i t$,
- (3) $\forall \langle [G] \rangle \psi \in \text{Sub}(\varphi)$ such that $A' \models \bigwedge \Phi(A_G \psi, G)$, there is some $\bar{A} \subseteq \text{act}(G)$ such that $s \rightarrow_A t$ for some $A \subseteq \text{Ac}$ such that $A \models \bigwedge \Phi(\bar{A}, G)$, and if $M^{\preceq}, s \models \langle [G] \rangle \psi$ then $M^{\preceq}, s \models [\bigwedge \Phi(\bar{A}, G)]\psi$
- (4) $\forall \langle [G^{\triangleleft_i}] \rangle \psi \in \text{Sub}(\varphi)$ such that $A' \models \bigwedge \Phi(A_{G^{\triangleleft_i} \psi}, G)$, there is some $\bar{A} \subseteq \text{act}(G)$ such that $s \rightarrow_A t$ for some $A \subseteq \text{Ac}$ such that $A \models \bigwedge \Phi(\bar{A}, G)$, and if $M^{\preceq}, s \models \langle [G^{\triangleleft_i}] \rangle \psi$ then $M^{\preceq}, s \models [\bigwedge \Phi(\bar{A}, G)]\psi$ and $M^{\preceq}, s \models (\bigwedge \Phi(\bar{A}, G) \subseteq \triangleleft_i)$.

For any $t \in T_{A'}^s$, we set $s \rightarrow'_{A'} t$. Then we can show by induction on $\psi \in \text{Sub}(\varphi)$ that $M^{\preceq}, s \models \psi$ iff $M'^{\preceq'}, s \models \psi$. □

The next step is to show that every satisfiable formula φ is satisfiable in a model with a certain number of states. Such results are usually obtained by transforming a model into a smaller one using a transformation that preserves the truth of subformulas of φ . In the case of CLA+P, the irreflexivity of the strict preferences and the fact that also $\alpha \subseteq \preceq_i$ is not modally definable in a basic modal language call for a modification of the standard techniques.

We appropriately modify the method of *filtration* (Blackburn et al. 2001) and show that any satisfiable formula $\varphi \in \mathcal{L}_{\text{cla+p}}$ is satisfiable in a model with exponentially many states. The idea of a filtration is to transform a possibly infinite model into a finite one by identifying states that agree on the truth value of each subformula of the considered formula. So, given that we know that φ is satisfied in some MaSP M^{\preceq} with states S , we construct an MaSP \mathcal{M}^{\preceq^f} with set of states $S_{\text{Sub}(\varphi)} = \{|s|_{\text{Sub}(\varphi)} \mid s \in S\}$, where $|s|_{\text{Sub}(\varphi)}$ denotes the equivalence class of the states that in the model M agree with s on the truth values of all $\psi \in \text{Sub}(\varphi)$. The main task is to appropriately define the accessibility relations for actions and preferences in \mathcal{M}^{\preceq^f} such that for $\psi \in \text{Sub}(\varphi)$, we then have that $M^{\preceq}, s \models \psi$ iff $\mathcal{M}^{\preceq^f}, |s| \models \psi$. Here, it is important to note that formulas of the form $\langle [G] \rangle \psi$ and $\langle [G^{\triangleleft_i}] \rangle \psi$ correspond to formulas of the form $\bigvee_{A \subseteq \text{act}(G)} [\bigwedge \Phi(A, G)]\psi$ and $\bigvee_{A \subseteq \text{act}(G)} ([\bigwedge \Phi(A, G)]\psi \wedge (\bigwedge \Phi(A, G) \subseteq \triangleleft_i))$, respectively—for $\bigwedge \Phi(A, G)$ as in the proof of Observation 1. Moreover, the transformation of the model does not change the underlying agents model. Thus, the truth of formulas of the form $\langle [G] \rangle \alpha$ is preserved.

Theorem 6 *Every satisfiable $\varphi \in \mathcal{L}_{\text{cla+p}}$ is also satisfiable in a MaSP with $\leq 2^{|\varphi|}$ many states.*

Proof Given that $M^{\preceq}, s \models \varphi$ for some $M^{\preceq} = \langle S, \text{Ac}, (\rightarrow_A)_{A \subseteq \text{Ac}}, \text{Ag}, \text{act}, \{\preceq_i\}_{i \in \text{Ag}}, V \rangle$ and $s \in S$, we obtain $\mathcal{M}^{\preceq^f} = \langle S_{\text{Sub}(\varphi)}, \text{Ac}, (\rightarrow^f)_{A \subseteq \text{Ac}}, \text{Ag}, \text{act}^f, \{\preceq_i^f\}_{i \in \text{Ag}}, V^f \rangle$

by filtrating M^{\preceq} through $Sub(\varphi)$, where the accessibility relations for actions and preferences are defined as follows:

$|s| \rightarrow_A^f |t|$ iff the following conditions are satisfied:

- (1) $\forall [\alpha]\psi \in Sub(\varphi)$ such that $A \models \alpha : \text{if } M^{\preceq}, s \models [\alpha]\psi, \text{ then } M^{\preceq}, t \models \psi,$
- (2) (a) $\forall \alpha \subseteq_{\preceq_i} \in Sub(\varphi)$ such that $A \models \alpha : \text{if } M^{\preceq}, s \models \alpha \subseteq_{\preceq_i}, \text{ then } s \preceq_i t,$
 (b) $\forall \alpha \subseteq_{<_i} \in Sub(\varphi)$ such that $A \models \alpha : \text{if } M^{\preceq}, s \models \alpha \subseteq_{<_i}, \text{ then } s <_i t,$
- (3) $\forall \langle [G] \rangle \psi \in Sub(\varphi)$ such that $A \models \bigwedge \Phi(A', G)$ for some $A' \subseteq \text{act}(G) : \text{if } M^{\preceq}, s \models [\bigwedge \Phi(A', G)]\psi, \text{ then } M^{\preceq}, t \models \psi,$
- (4) (a) $\forall \langle [G^{\preceq_i}] \rangle \psi \in Sub(\varphi)$ such that $A \models \bigwedge \Phi(A', G)$ for some $A' \subseteq \text{act}(G): \text{if } M^{\preceq}, s \models [\bigwedge \Phi(A', G)]\psi \text{ and } M^{\preceq}, s \models (\bigwedge \Phi(A', G) \subseteq_{\preceq_i}), \text{ then } M^{\preceq}, t \models \psi \text{ and } s \preceq_i t.$
 (b) $\forall \langle [G^{<_i}] \rangle \psi \in Sub(\varphi)$ such that $A \models \bigwedge \Phi(A', G)$ for some $A' \subseteq \text{act}(G): \text{if } M^{\preceq}, s \models [\bigwedge \Phi(A', G)]\psi \text{ and } M^{\preceq}, s \models (\bigwedge \Phi(A', G) \subseteq_{<_i}), \text{ then } M^{\preceq}, t \models \psi \text{ and } s <_i t.$

$|s| \preceq_i^f |t|$ iff the following conditions hold:

- (1) (a) $\forall \diamond^{\preceq_i} \psi \in Sub(\varphi): \text{if } M^{\preceq}, t \models \psi \vee \diamond^{\preceq_i} \psi \text{ then } M^{\preceq}, s \models \diamond^{\preceq_i} \psi,$
 (b) If there is some $\diamond^{<_i} \psi \in Sub(\varphi)$, then $s \preceq_i t,$
- (2) If there is some $\alpha \subseteq_{\preceq_i} \in Sub(\varphi)$ or some $\alpha \subseteq_{<_i} \in Sub(\varphi)$, then $s \preceq_i t,$
- (3) If there is some $\langle [G^{\preceq_i}] \rangle \psi \in Sub(\varphi)$ or some $\langle [G^{<_i}] \rangle \psi \in Sub(\varphi)$, then $s \preceq_i t.$

$V^f(p) := \{|s| \mid M, s \models p\}$, for all propositional letters $p \in Sub(\varphi)$. We can show by induction that for all $\psi \in Sub(\varphi)$ and $s \in S$ it holds that $M^{\preceq}, s \models \psi$ iff $M^{\preceq^f}, |s| \models \psi$. This follows from the definitions of $(\rightarrow^f)_{A \subseteq Ac}$ and \preceq^f , and the fact that we do not change the underlying agents model. The interesting cases are those involving strict preferences and those with formulas $\langle [G^{\preceq_i}] \rangle \psi$ and $\alpha \subseteq_{\preceq_i}$. Here, what makes the proof go through is that by conditions 1 b), 2 and 3 of \preceq^f , $|s| \preceq_i |t|$ implies $s \preceq_i t$. Similarly, due to conditions 2 and 4 of \rightarrow_A^f , the truth values of subformulas $\alpha \subseteq_{<_i}$ and $\langle [G^{<_i}] \rangle \psi$ is as in the original model. Moreover, M^{\preceq^f} is a proper MaSP since each \preceq_i^f is reflexive and transitive, and each \rightarrow_A^f is serial. By definition of $S_{Sub(\varphi)}$, $|S_{Sub(\varphi)}| \leq 2^{|\varphi|}$. □

Now, we apply the constructions of the last three proofs successively.

Corollary 1 *Every satisfiable formula $\varphi \in \mathcal{L}_{cla+p}$ is satisfiable in a MaSP of size exponential in $|\varphi|$ satisfying the conditions $|Ag| \leq |Ag(\varphi)| + 1$ and $|Ac| \leq |Ac(\varphi)| + \sum_{\langle [G] \rangle \psi \in Sub(\varphi)} |G| + (\sum_{\langle [G^{\preceq_i}] \rangle \psi \in Sub(\varphi)} |G|) + (\sum_{\langle [G^{<_i}] \rangle \psi \in Sub(\varphi)} |G|) + 1.$*

Having non-deterministically guessed a model of size exponential in $|\varphi|$, we can check in time exponential in $|\varphi|$ whether this model satisfies φ .

Theorem 7 *The satisfiability problem of CLA+P is in NEXPTIME.*

Proof Given φ , we non-deterministically choose a model M^{\preceq} of size exponential in $|\varphi|$ satisfying the conditions $|Ag| \leq |Ag(\varphi)| + 1$ and $|Ac| \leq |Ac(\varphi)| + \sum_{\langle [G] \rangle \psi \in Sub(\varphi)} |G| + (\sum_{\langle [G^{\preceq_i}] \rangle \psi \in Sub(\varphi)} |G|) + (\sum_{\langle [G^{<_i}] \rangle \psi \in Sub(\varphi)} |G|) + 1$. Then, given this model, we can check in time $\mathcal{O}(|\varphi| \cdot |M^{\preceq}|)$, for $|M^{\preceq}|$ being the size of M^{\preceq} , whether M^{\preceq} satisfies

φ . Thus, given a model of size exponential in $|\varphi|$ that also satisfies the conditions on its sets of agents and actions explained earlier, it can be computed in time exponential in $|\varphi|$ whether it satisfies φ . Since it can be checked in time linear in the size of the model whether it is a proper MaSP, we conclude that SAT of CLA+P is in NEXPTIME. \square

This section has shown that SAT of CLA+P is in NEXPTIME. Now we show that the environment logic is already EXPTIME-hard.

4.2 Lower bound

In order to show a lower bound for the complexity of SAT of CLA+P, we show that SAT of the environment logic is EXPTIME-hard. This is done by reduction from the Boolean modal logic $\mathbf{K}_m^{-\cup}$ (Lutz and Sattler 2001; Lutz et al. 2001).

Formulas of $\mathbf{K}_m^{-\cup}$ are interpreted in models $M = \langle W, R_1, \dots, R_m, V \rangle$, where W is a set of states, $R_i \subseteq W \times W$ and V is a valuation.

Definition 15 Let $\mathcal{R}_1, \dots, \mathcal{R}_m$ be atomic modal parameters. Then the set of modal parameters of $\mathbf{K}_m^{-\cup}$ is the smallest set containing $\mathcal{R}_1, \dots, \mathcal{R}_m$ that is closed under \neg and \cup . The language $\mathcal{L}_m^{-\cup}$ is generated by the following grammar:

$$\varphi ::= p \mid \varphi \wedge \varphi \mid \neg\varphi \mid \langle \mathcal{S} \rangle \varphi \quad \mathcal{S} ::= \mathcal{R}_i \mid \neg\mathcal{S} \mid \mathcal{S}_1 \cup \mathcal{S}_2.$$

The extension $\mathcal{E}(\mathcal{S}) \subseteq W \times W$ of a parameter \mathcal{S} in a model is as follows.

$$\begin{aligned} \mathcal{E}(\mathcal{R}_i) &= R_i \\ \mathcal{E}(\neg\mathcal{S}) &= (W \times W) \setminus \mathcal{E}(\mathcal{S}) \\ \mathcal{E}(\mathcal{S}_1 \cup \mathcal{S}_2) &= \mathcal{E}(\mathcal{S}_1) \cup \mathcal{E}(\mathcal{S}_2) \end{aligned}$$

Formulas of $\mathcal{L}_m^{-\cup}$ are interpreted in a model $M = \langle W, R_1, \dots, R_m, V \rangle$ as follows: Propositional letters and boolean combinations are interpreted in the standard way and for modal formulas we have

$$M, w \models \langle \mathcal{S} \rangle \varphi \text{ iff } \exists w' \in W : (w, w') \in \mathcal{E}(\mathcal{S}) \text{ and } M, w' \models \varphi.$$

We define a translation τ consisting of two components τ_1 for formulas and τ_2 for models. Let us extend the language \mathcal{L}_e by a propositional letter $q \notin \mathcal{L}_m^{-\cup}$. Then τ_1 is defined as follows:

$$\begin{aligned} \tau_1(p) &= p & \tau^S(\mathcal{R}_i) &= a_i \\ \tau_1(\varphi_1 \wedge \varphi_2) &= \tau_1(\varphi_1) \wedge \tau_1(\varphi_2) & \tau^S(\mathcal{S}_1 \cup \mathcal{S}_2) &= \tau^S(\mathcal{S}_1) \vee \tau^S(\mathcal{S}_2) \\ \tau_1(\neg\varphi) &= \neg\tau_1(\varphi) & \tau^S(\neg\mathcal{S}) &= \neg\tau^S(\mathcal{S}) \\ \tau_1(\langle \mathcal{S} \rangle \varphi) &= \neg[\tau^S(\mathcal{S})](q \vee \neg\tau_1(\varphi)) \end{aligned}$$

τ_2 translates a model M of $\mathbf{K}_m^{-\cup}$ into an environment model $\tau_2(M) = \langle W \cup \{u\}, Ac, (\rightarrow)_{A \subseteq Ac}, V' \rangle$ with u being a newly introduced state and

$$w \rightarrow_A w' \text{ iff } A = \{a_i | (w, w') \in R_i\} \text{ or } w' = u.$$

Thus, each \rightarrow_A is serial and $\tau_2(M)$ is an environment model. $V'(q) = \{u\}$, and for all $p \neq q$, $V'(p) = V(p)$. Before showing that for any $\varphi \in \mathcal{L}_m^{-\cup}$ and $M \in \mathbb{M}_m^{-\cup}$ for any state $w \in W : M, w \models \varphi$ iff $\tau_2(M), w \models \tau_1(\varphi)$, we proof a lemma saying that if in M w' is \mathcal{S} -accessible from w , then in $\tau_2(M)$, w' is accessible from w by a transition of type $\tau^S(\mathcal{S})$.

Convention 2 For $M = \langle W, R_1, \dots, R_m, V \rangle \in \mathbb{M}_m^{-\cup}$ and $\tau_2(M) = \langle W \cup \{u\}, Ac, (\rightarrow)_{A \subseteq Ac}, V' \rangle$, define $A_{w,w'} := \{a_i \in Ac | (w, w') \in R_i\}$.

Lemma 1 Let $M = \langle W, R_1, \dots, R_m, V \rangle$ be a model of $\mathbf{K}_m^{-\cup}$. Then for any modal parameter \mathcal{S} and for any states $w, w' \in W$ it holds that

$$(w, w') \in \mathcal{E}(\mathcal{S}) \text{ iff in } \tau_2(M) : \exists A \subseteq Ac : w \rightarrow_A w' \text{ and } A \models \tau^S(\mathcal{S}).$$

Proof Note that by definition of $(\rightarrow_A)_{A \subseteq Ac}$, the righthand side is equivalent to $A_{w,w'} \models \tau^S(\mathcal{S})$. Then the proof goes by induction on \mathcal{S} . □

Theorem 8 For any formula $\varphi \in \mathcal{L}_m^{-\cup}$ and any model M of $\mathbf{K}_m^{-\cup}$, it holds that for any state w in M :

$$M, w \models \varphi \text{ iff } \tau_2(M), w \models \tau_1(\varphi).$$

Proof By induction. Base case and boolean cases are straightforward. Let $\varphi = \langle S \rangle \psi$.

(\Rightarrow) $M, w \models \langle S \rangle \psi \Leftrightarrow \exists w' : (w, w') \in \mathcal{E}(\mathcal{S})$ and $M, w' \models \psi$. By the previous lemma and induction hypothesis, $\tau_2(M), w \models \tau_1(\langle S \rangle \psi)$.

(\Leftarrow) $\tau_2(M), w \models \tau_1(\langle S \rangle \psi) \Leftrightarrow \tau_2(M), w \models \neg[\tau^S(\mathcal{S})](q \rightarrow \neg\tau_1(\psi))$. Then $\exists w' \in W \cup \{u\}, \exists A \in Ac : A \models \tau^S(\mathcal{S}), w \rightarrow_A w'$ and $\tau_2(M), w' \models \neg q \wedge \tau_1(\psi)$. Thus, $w' \neq u$. By induction hypothesis and the previous lemma, $M, w' \models \psi$ and $(w, w') \in \mathcal{E}(\mathcal{S})$. Hence, $M, w \models \langle S \rangle \psi$. □

Theorem 9 SAT of Λ^E is EXPTIME-hard.

Proof We can polynomially transform any $\varphi \in \mathcal{L}_m^{-\cup}$ into $\tau_1(\varphi)$. Now, if $\tau_1(\varphi)$ is satisfiable in an environment model $E = \langle W, Ac, (\rightarrow)_{A \subseteq Ac}, V \rangle$ with $Ac = \{a_1, \dots, a_m\}$, then φ is satisfiable in a model $M = \langle W, R_1, \dots, R_m, V \rangle$, where $(w, w') \in R_i$ iff $\exists A \subseteq Ac$ such that $a_i \in Ac$ and $w \rightarrow_A w'$. This can be shown by induction on φ .

If $\tau_1(\varphi)$ is not satisfiable in an environment model then it cannot be satisfiable in any model M of $\mathbf{K}_m^{-\cup}$ because otherwise by the previous theorem, $\tau_1(\varphi)$ would be satisfied in $\tau_2(M)$. □

Corollary 2 SAT of CLA+P is EXPTIME-hard.

This section has shown that the satisfiability problem of CLA+P is EXPTIME-hard but still decidable. This rather high complexity is due to the environment logic which itself is already EXPTIME-hard.

5 Conclusions and future work

We developed a modular modal logic that allows for reasoning about the coalitional power of agents, actions and their effects, and agents' preferences. The current approach is based on the logic CLA (Sauro et al. 2006) which is combined with a preference logic (van Benthem et al. 2007). The resulting logic CLA+P, which is shown to be sound and complete, allows us to make explicit how groups can achieve certain results. We can also express how a group can achieve that a transition takes place that is an improvement for some agent.

In the framework of CLA+P, it can be expressed how the abilities to perform certain actions are distributed among the agents, what are the effects of the concurrent performance of these actions and what are the agents' preferences over those effects. Moreover, in CLA+P, we can distinguish between different ways how groups can achieve some result—not only with respect to the actions that lead to some result, but also with respect to the preferences. We can for instance express that a group can achieve some result in a way that is 'good' for all its members in the sense that after the achievement all of them are better off. This then also allows us to axiomatize properties that one might want to impose onto a multi-agent system, e.g. the restriction that groups can only achieve the truth of a certain formula if this can be done without making anybody worse off. Thus, CLA+P provides a framework for reasoning about interactive situations in an explicit way that gives us more insights into the cooperative abilities of agents. Comparing CLA+P to CL shows that CLA+P naturally builds on game frames underlying the semantics of CL and makes both the agents' actions and the preferences explicit that are only implicitly represented in the semantics of CL.

The satisfiability problem of CLA+P is shown to be decidable and EXPTIME-hard. Keeping in mind that using CLA+P we can talk about strict preferences, intersections of accessibility relations as well as the property of one relation being a subset of another, EXPTIME-hardness is not a surprising result. Even though the modular models of CLA+P are rather special, its complexity is in accordance with general results concerning the connection between expressive power and complexity of modal logics for reasoning about coalitional power and preferences (Dégrement and Kurzen 2009).

We showed that the satisfiability problem of the underlying environment logic is by itself already EXPTIME-hard. Thus, we identified one cause for the high complexity of CLA+P. It is mostly due to the fact that the accessibility relation of the models can be arbitrary: there does not need to be any relation between \rightarrow_A , \rightarrow_B and $\rightarrow_{A \cap B}$. Whereas this generality allows us to model a lot of dynamic processes, from a computational viewpoint, it seems to be appealing to change the environment logic in order to decrease computational complexity. Also, when comparing our models to the game frames of CL, we can see that restricting ourselves to deterministic environment models can be reasonable. The same holds for considering only total preorders as preference relation. This would also increase the expressive power (Dégrement and Kurzen 2009).

There are several immediate ways to extend the logic developed in this paper. First of all, we can follow the ideas of Ågotnes et al. (2007b) and add a restricted form of

quantification that allows statements of the form $\langle [P^{\preceq_i}] \rangle \psi$ saying that there is some group G that has property P and $\langle [G^{\preceq_i}] \rangle \psi$. In the current work, we chose a rather simple preference logic with unary modalities. In order to increase the expressive power with respect to the agents' preferences over outcomes that can be achieved, it seems very promising to consider more expressive preference logics that also contain binary preference modalities saying e.g. that an agent prefers every φ state over every ψ state (van Benthem et al. 2007).

Moreover, it will be interesting to develop a CLA and preferences based on a logic for reasoning about complex plans such as the one developed by Gerbrandy and Sauro (2007).

Acknowledgements The author thanks Ulle Endriss, Eric Pacuit, Luigi Sauro, Johan van Benthem, Olivier Roy and the anonymous referees for their help and their useful comments. She also thanks the participants of the *Workshop on Logic and Intelligent Interaction* at ESSLLI 2008. Their comments have helped to improve earlier versions. This work is supported by the Dutch Organization for Scientific Research (NWO), TACTICS project, grant number 612.000.525.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

- Ågotnes, T., Dunne, P., van der Hoek, W., & Wooldridge, M. (2007a). Logics for coalitional games. In J. van Benthem, S. Ju, & F. Veltman (Eds.), *A meeting of the minds* (pp. 3–20). London, UK: College Publications.
- Ågotnes, T., van der Hoek, W., & Wooldridge, M. (2007b). Quantified coalition logic. In *IJCAI '07: Proceedings of the Twentieth International Joint Conference on Artificial Intelligence* (pp. 1181–1186). Hyderabad, India.
- Alur, R., Henzinger, T. A., & Kupferman, O. (1998). Alternating-time temporal logic. *Lecture Notes in Computer Science*, 1536, 23–60.
- Blackburn, P., de Rijke, M., & Venema, Y. (2001). *Modal logic, No. 53 in Cambridge tracts in theoretical computer science*. UK: Cambridge University Press.
- Borgo, S. (2007). Coalitions in action logic. In *IJCAI '07: Proceedings of the Twentieth International Joint Conference on Artificial Intelligence* (pp. 1822–1827). Hyderabad, India.
- Broersen, J. (2003). Modal action logics for reasoning about reactive systems. Ph.D. thesis, Faculteit der Exacte Wetenschappen, Vrije Universiteit Amsterdam.
- Dégremont, C., & Kurzen, L. (2009). Modal logics for preferences and cooperation: Expressivity and complexity. In *KRAMAS 2008: Proceedings of the Workshop on Knowledge Representation for Agents and Multi-Agent Systems*. (to appear).
- Gerbrandy, J., & Sauro, L. (2007). Plans in cooperation logic: A modular approach. In *Proceedings of the IJCAI Workshop on Nonmonotonic Reasoning, Action and Change (NRAC 2007)*. Hyderabad, India.
- Kurzen, L. (2007). Logics for cooperation, actions and preferences. Master's thesis, Universiteit van Amsterdam, the Netherlands.
- Lutz, C., & Sattler, U. (2001). The complexity of reasoning with boolean modal logics. In F. Wolter, H. Wansing, M. de Rijke & M. Zakharyashev (Eds.), *Advances in modal logics* (Vol. 3). Stanford: CSLI Publications.
- Lutz, C., Sattler, U., & Wolter, F. (2001). Modal logics and the two-variable fragment. In *Annual Conference of the European Association for Computer Science Logic CSL'01*. Paris, France: Springer Verlag.
- Osborne, M. J., & Rubinstein, A. (1994). *A course in game theory*. Cambridge, MA: MIT Press.
- Pauly, M. (2002a). A modal logic for coalitional power in games. *Journal of Logic and Computation*, 12(1), 149–166.

- Pauly, M. (2002b). On the complexity of coalitional reasoning. *International Game Theory Review*, 4, 237–254.
- Sauro, L., Gerbrandy, J., van der Hoek, W., & Wooldridge, M. (2006). Reasoning about action and cooperation. In *AAMAS '06: Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multi-agent Systems* (pp. 185–192). Hakodate, Japan.
- van Benthem, J., Roy, O., & Girard, P. (2007). Everything else being equal: A modal logic approach to ceteris paribus preferences. *ILLC Prepublication Series, PP-2007-09*.
- van Benthem, J., van Otterloo, S., & Roy, O. (2005). Preference logic, conditionals and solution concepts in games. In H. Lagerlund, S. Lindström, & R. Sliwinski (Eds.), *Modality matters: Twenty-five essays in honour of Krister Segerberg, Uppsala Philosophical Studies 53*. University of Uppsala.
- van Otterloo, S., van der Hoek, W., & Wooldridge, M. (2004). Preferences in game logics. In *AAMAS '04: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems* (pp. 152–159). IEEE Computer Society.
- Walther, D., van der Hoek, W., & Wooldridge, M. (2007). Alternating-time temporal logic with explicit strategies. In *TARK '07: Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge* (pp. 269–278). ACM.