

A Logic for Cooperation, Actions and Preferences

Lena Kurzen
Universiteit van Amsterdam
L.M.Kurzen@uva.nl

Abstract

In this paper, a logic for reasoning about cooperation, actions and preferences of agents is developed. It is shown to be sound and complete and the satisfiability problem of its fragment that does not contain strict preferences is shown to be NExpTime-complete.

1. Introduction

When analyzing interactive situations involving multiple agents, we are interested in what results agents can achieve – individually or together as groups. In many cases, agents can have various plans for achieving some result. These plans can differ significantly, e.g. with respect to their feasibility, costs or side-effects. Hence, it is not only relevant which 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 in interactive processes. Cooperative ability of agents expressed only in terms of results and actions that lead to these results does not tell us *why* 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 and moreover tightly connected:

- **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 and preferences play a major role in interactive situations and are moreover tightly connected. Thus, 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.

In logic, these aspects have mostly been addressed separately. Coalitional power has mainly been investigated

within the frameworks of ATL [4], Coalition Logic [12] and their extensions. These logics focus on what results groups can achieve and do not represent explicitly how exactly results can be achieved. 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 [1, 2]. Another path that has been taken in order to make coalitional power more explicit is to combine cooperation logics with action logics [14, 6, 7].

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 [14] with a preference logic [15, 16]. Soundness and completeness are shown and the logics expressivity and computational complexity are investigated.

The remainder of this paper is structured as follows. Section 2 gives a brief overview of the cooperation logic with actions CLA [14]. In Section 3, a cooperation logic with actions and preferences (CLA+P) is developed and soundness and completeness are shown. Also its expressivity is discussed. Section 4 investigates the computational complexity of CLA+P.

2. Cooperation Logic with Actions (CLA)

In this section, we briefly present the cooperation logic with actions (CLA) developed by Sauro et al. [14], 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 results. 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, which at a later step will be populated by agents, is modelled as a labelled transition system whose edges are labeled with sets of atomic actions.

Definition 2.1 [Environment Model [14]] 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, $\rightarrow_{A \subseteq Ac} : S \times S$ for each $A \subseteq Ac$ and V is a propositional valuation. \rightarrow_A is required to be serial for each $A \subseteq Ac$. \triangleleft

Then a modal language is defined with modalities of the form $[\alpha]$, 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 logic) leads to a state that satisfies φ . Formally,

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

An environment logic Λ^E is developed, which is sound and complete with respect to the class of environment models [14]. It contains seriality axioms and the **K** axiom for each modality $[\alpha]$, for α being consistent. For the details, the reader is referred to Sauro et al. [14]. The environment logic, can then be used for reasoning about the effects of concurrent actions.

Before agents are introduced into the environment, a separate agents module is developed for reasoning about the ability of (groups of) agents to perform actions. Each agent is assigned a set of atomic actions that he can perform and a group is assigned the set of actions its members can perform.

Definition 2.2 [Agents Model [14]] 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 $\text{act} : Ag \rightarrow \mathcal{P}(Ac)$ such that $\bigcup_{i \in Ag} \text{act}(i) = Ac$. For $G \subseteq Ag$, $\text{act}(G) := \bigcup_{i \in G} \text{act}(i)$. \triangleleft

We are not only interested in what atomic actions agents can perform but also in their abilities to enforce more complex actions. An agent language is developed with expressions $\langle\langle G \rangle\rangle\alpha$, 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 .

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

Then a cooperation logic for actions is developed, which is very much in the style of Coalition Logic [12] – the main difference being that it is concerned with the cooperative ability to force *actions*.

Definition 2.3 [Coalition Logic for Actions [14]] The coalition logic for actions Λ^A is defined to be the logic derived from the following set of axioms.

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

The rule of inference is modus ponens. \triangleleft

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 [14].

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

Definition 2.4 [Multi-agent System [14]] A multi-agent system (MaS) is a tuple

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

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

Now, we can reason about what states of affairs groups of agents can achieve by performing certain actions. The corresponding language contains all expressions of the environment logic and the cooperation logic for actions and additionally expressions for saying that a group has the power to achieve φ .

Definition 2.5 [Language for MaS [14]] The language for multi-agent systems \mathcal{L}_{cla} is generated by the following grammar:

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

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

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

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

The environment logic and the coalition logic for agents are combined by adding two interaction axioms.

Definition 2.6 [Cooperation Logic with Actions [14]] The cooperation logic with actions Λ^{CLA} combines the environment logic Λ^E and the coalition logic for actions Λ^A by adding

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

\triangleleft

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. [14]. Now, we proceed by adding an explicit representation of the agents' preferences to CLA.

3. Cooperation Logic with Actions and Preferences (CLA+P)

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 of what is covered in this section and detailed proofs, see [8].

3.1. Preference Logic

There are various ways how preferences can be added to a logic for cooperation and actions. We could for instance let the preferences of each agent 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 single agents 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 to one over formulas [15, 16].

Definition 3.1 [Preference Model [16]] 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 a reflexive and transitive relation and V is a propositional valuation. \triangleleft

As a language, we use a fragment of the basic preference language developed by van Benthem et al. [15]. It has strict and non-strict unary preference modalities for each agent.

Definition 3.2 [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:

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

\triangleleft

$\diamond^{\preceq_i}\varphi$ says that there is a state satisfying φ , and agent i prefers this state over the current one, i.e.

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

$\diamond^{\prec_i}\varphi$ is interpreted analogously. The preference relation \preceq should be reflexive and transitive and \prec should be its largest irreflexive subrelation. Thus, the following axiomatization is chosen.

Definition 3.3 [Preference Logic Λ^P] For a given set of agents Ag , let Λ^P be the logic generated by the following axioms for each agent $i \in Ag$:

For \diamond^{\preceq_i} and \diamond^{\prec_i} , we have duality axioms and **K**. For \diamond^{\preceq_i} , we also have reflexivity and transitivity axioms. Additionally, there are four axioms for the interaction between the strict and non-strict modalities:

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

The inference rules are modus ponens, necessitation and substitution of logical equivalents. \triangleleft

Note that transitivity for \diamond^{\prec_i} follows from the above axioms. We can show soundness and completeness of the preference logic. The fact that \prec is supposed to be the greatest irreflexive subrelation of \preceq can be dealt with by using the bulldozing technique. For the details, the reader is referred to van Benthem et al. [15].

Theorem 3.4 Λ^P is sound and complete with respect to the class of preference models.

Proof. Follows from Theorem 3.9 in [15]. QED

3.2. Environment Logic with Preferences

As an intermediate step towards a logic for reasoning about cooperation, actions and preferences, we first combine the preference logic and the environment logic. The two models are combined by identifying their sets of states. Then the preferences of the agents range over the states of the environment. In such a system, the agents cannot act in the environment, but they can rather be seen as observers that observe the environment from the outside and have preferences over the states.

Definition 3.5 [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. \triangleleft

We combine the languages for the environment and the preferences and add expressions for saying that “every state accessible by an α transition is (strictly) preferred by agent i over the current state”.

Convention 3.6 In what follows, we write the symbol \triangleleft in statements that hold for both \preceq and \prec , each uniformly substituted for \triangleleft .

Definition 3.7 The language \mathcal{L}_{ep} contains all expressions of the environment language and the preference language and additionally formulas of the forms $[\alpha]^{\preceq_i} \top$ and $[\alpha]^{\prec_i} \top$, for α being an action expression.

Boolean combinations and expressions of the previously defined languages are interpreted in the standard way. For the newly introduced expressions, we have:

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

\triangleleft

Expressions of the form $[\alpha]^{\triangleleft_i} \top$ cannot be defined using just the preference language and the environment language. To see this, note that $[\alpha]^{\preceq_i} \top$ says that for every state accessible by an α -transition it holds that this same state is accessible by \preceq . Thus, we would have to be able to refer to particular states. Therefore, we add two inference rules for deriving the newly introduced expressions.

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

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

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

Theorem 3.8 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. $[\alpha]^{\preceq_i} \top \rightarrow (\Box^{\preceq_i} \varphi \rightarrow [\alpha] \varphi)$,
2. $[\alpha]^{\prec_i} \top \rightarrow (\Box^{\prec_i} \varphi \rightarrow [\alpha] \varphi)$.

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

Proof. Completeness follows from completeness of the sublogics and the closure under the new rules. QED

In the environment logic with preferences, the performance of concurrent actions changes the current state of the system also with respect to the ‘happiness’ of the agents: A transition from one state to another can also be a transition up or down in the preference orderings of the agents.

3.3. Cooperation Logic with Actions and Preferences

Now, agents are introduced 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 3.9 [Multi-agent System with Preferences] A multi-agent system with preferences (MaSP) M^{\preceq} is a tuple

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

where $\langle S, Ac, (\rightarrow)_{A \subseteq Ac}, V, Ag, \text{act} \rangle$ 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. \triangleleft

In order to get some intuitions about how MaSP's are related to other models of interaction, 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

$$\mathcal{G}_s = \langle Ag, (\mathcal{P}(\text{act}(i)))_{i \in Ag}, (\preceq_i)_{i \in Ag} \rangle$$

attached to it, where $\times_{i=1}^n A_i \preceq_i \times_{i=1}^n A'_i$ iff $t \preceq_i t'$ for $s \rightarrow_{\cup_{i \in Ag} A_i} t$ and $s \rightarrow_{\cup_{i \in Ag} A'_i} t'$.

For talking about the cooperative ability of agents with respect to preferences, we introduce two expressions saying that a group can force the system to move to a φ -state that some agent (strictly) prefers over the current one.

Definition 3.10 [Language \mathcal{L}_{cla+p}] The language \mathcal{L}_{cla+p} extends \mathcal{L}_{cla} by formulas of the form

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

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

$$M^{\preceq}, s \models \langle\langle G^{\preceq_i} \rangle\rangle \varphi \quad \text{iff} \quad \exists A \subseteq \text{act}(G) \text{ such that} \\ \forall B \subseteq \text{act}(Ag \setminus G), t \in S : \\ \text{if } s \rightarrow_{A \cup B} t, \text{ then } M^{\preceq}, t \models \\ \varphi \text{ and } s \prec_i t. \quad \triangleleft$$

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 $\langle\langle G^{\preceq_i} \rangle\rangle \varphi$ is equivalent to the existence of an action expression α that G can force and that has the property that all transitions of type α are guaranteed to lead to a φ -state preferred by agent i .

Observation 3.11 Given a MaSP M^{\preceq} and a state s of its environment,

$$M^{\preceq}, s \models \langle\langle G^{\preceq_i} \rangle\rangle \varphi \quad \text{iff} \quad \text{there exists an action expression } \alpha \text{ such that } M^{\preceq}, s \models \\ \langle\langle G \rangle\rangle \alpha \wedge [\alpha] \varphi \wedge [\alpha]^{\preceq_i} \top.$$

Proof. Analogous to that of Observation 14 in [14]. For the left-to-right direction, use the action expression $\bigwedge \Phi(A, G) := \bigwedge (A \cup \{-a \mid a \in (\text{act}(G) \setminus A), a \notin \text{act}(Ag \setminus G)\})$. QED

Thus, formulas of the form $\langle\langle G^{\preceq_i} \rangle\rangle \varphi$ can be reduced to expressions of the sublogics. We also need new axioms establishing a relationship between the newly added formulas and the expressions of the sublogics.

Definition 3.12 [Cooperation Logic with Actions and Preferences Λ^{CLA+P}] Define Λ^{CLA+P} to be the smallest logic generated by the axioms of the cooperation logic with actions, the environment logic with preferences and

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

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

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

Theorem 3.13 The logic Λ^{CLA+P} is sound and complete with respect to the class of MaSP's.

3.4. Expressivity of CLA+P

The framework of CLA+P allows us to reason about coalitional power in an explicit way since we can express how groups of agents can achieve the truth of some formula, and moreover we can also express how coalitional power and actions relate to the agents' preferences.

In game theory, coalitional power has mostly been studied within coalitional games [11]. One of the most important solution concepts of coalitional games is the core, which is the set of outcomes the grand coalition can achieve that have the property that no coalition can achieve some other outcome that is strictly better for all its members.

In the framework of CLA+P, we can characterize the states in a model that have a very similar property: The formula $\hat{\psi}$ characterizes the set of states in which no group has the power of making the system move into a state that is strictly better for all its members:

$$\hat{\psi} := \bigwedge_{G \subseteq Ag} \bigwedge_{A \subseteq \text{act}(G)} \left(\bigvee_{i \in G} \neg [\bigwedge \Phi(A, G)]^{\prec_i} \top \right).$$

For the definition of $\bigwedge \Phi(A, G)$, see the proof of Observation 3.11.

In interactive situations, there can be different ways of how agents can achieve some result. These ways can consist of different plans that the agents can execute, and whereas all the actions or plans might lead to the same result φ , executing one plan might be better for the agents than executing another one. Being 'better' could be in the sense that one plan leads to an improvement of the situation for more agents than executing another plan does. In CLA+P, we can express that a group G can achieve φ in such a way that the situation improves for all its members:

$$\bigvee_{A \subseteq \text{act}(G)} \left(([\bigwedge \Phi(A, G)] \varphi) \wedge \bigwedge_{i \in G} [\bigwedge \Phi(A, G)]^{\preceq_i} \top \right).$$

Thus, the explicit representation of actions and preferences allows us to reason about how exactly a group would choose to achieve some result, assuming that the members make their decisions according to a certain solution concept.

Alternatively, executing one plan might be better than another one in the sense that it is cheaper. By having both actions and preferences in our framework, we can also express how actions and preferences interact and thereby our framework can also give rise to a formal model for cost-benefit analysis. In cost-benefit-analysis, decisions are made by comparing the expected cost of executing actions and the expected benefit.

4. Complexity of CLA+P

In this section, we investigate the complexity of the satisfiability problem of CLA+P. Let us start by trying to determine an upper bound.

4.1. Upper Bound for CLA+P

In order to establish an upper bound, it has to be shown that computing whether some formula is satisfiable can be done using a certain amount of time or space. The first step is to show that only a restricted class of models of CLA+P needs to be checked.

For a formula φ , let $Ag(\varphi)$ denote the set of agents occurring in φ . Now, we ask: Is any satisfiable formula φ also satisfiable in a MaSP whose set of agents is $Ag(\varphi)$? In Coalition Logic, the answer is negative due to formulas such as

$$\varphi' = \neg \langle \langle \{1\} \rangle \rangle p \wedge \neg \langle \langle \{1\} \rangle \rangle q \wedge \langle \langle \{1\} \rangle \rangle (p \vee q),$$

which can only be satisfied in coalition models with at least two agents [13]. However, as in CLA+P the underlying 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 newly introduced agent. k takes the role of all opponents of $Ag(\varphi)$ collapsed into one. This means that k gets the ability to perform exactly the actions that agents not occurring in φ can perform as a group. When transforming a model satisfying φ into one with set of agents $Ag(\varphi) \cup \{k\}$, we do not need to change the effects of actions or the abilities of agents in $Ag(\varphi)$. This is the main fact that makes the proof of Theorem 4.1 go through. Moreover, note that the preferences of agent k do not have any influence on the truth of φ since k does not occur in φ .

Theorem 4.1 *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 $M^{\preceq} = \langle S, Ac, (\rightarrow)_{A \subseteq Ac}, Ag, \text{act}, \{\preceq_i\}_{i \in Ag}, V \rangle$ satisfies φ . If $Ag \supset Ag(\varphi)$, we construct $M'^{\preceq'} = \langle S, Ac, (\rightarrow)_{A \subseteq Ac}, Ag(\varphi) \cup \{k\}, \text{act}', \{\preceq'_i\}_{i \in Ag(\varphi) \cup \{k\}}, V \rangle$, with $\text{act}'(k) = \bigcup_{j \in Ag \setminus Ag(\varphi)} \text{act}(j)$ and $\text{act}'(i) = \text{act}(i)$ 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, it can be shown that $M^{\preceq}, s \models \varphi$ iff $M'^{\preceq'}, s \models \varphi$. The interesting case is the one where φ is of the form $\langle \langle G \rangle \rangle \alpha$. Here, the claim follows from the definition of act' . The other cases involving coalition modalities follow. QED

Next, we would like to know how many actions a model needs in order to satisfy some formula. As an example, consider the formula

$$\varphi' = \langle \langle G \rangle \rangle (p \wedge q) \wedge \langle \langle G \rangle \rangle (\neg p \wedge q) \wedge \langle \langle G \rangle \rangle (\neg p \wedge \neg q).$$

It can only be satisfied in models with $|Ac| \geq 2$. The main task is to find “witnesses” for formulas of the form $\langle \langle G \rangle \rangle \psi$ in terms of concurrent actions that tell us how exactly G can achieve ψ . We can show that every satisfiable formula φ can be satisfied in a MaSP whose set of actions consists of the actions occurring in φ , one additional atomic action, and for every subformula of the forms $\langle \langle G \rangle \rangle \psi$ or $\langle \langle G^{\preceq_i} \rangle \rangle \psi$, one atomic action for each of G 's members. The one additional action is a dummy that serves for making sure that every agent can perform some action.

The key step in transforming a model satisfying a formula φ into one whose set of actions satisfies the above condition is to define the action distribution and the accessibility relations in an appropriate way. For every action expression α occurring in φ , we have to ensure that two states are accessible by an α -transition in the new model iff they were in the original one. Additionally, for any formula of the forms $\langle \langle G \rangle \rangle \psi$ or $\langle \langle G^{\preceq_i} \rangle \rangle \psi$, the set of actions that we introduced for that formula serves for making explicit how G can force φ . Note that we do not need to introduce any additional actions for making explicit how a group can force an action expression α . This results from the fact that in order to force α , agents only need to perform actions that already occur in α .

Theorem 4.2 *Every satisfiable formula $\varphi \in \mathcal{L}_{cla+p}$ is satisfiable in a MaSP with at most $|Ac(\varphi)| + (\sum_{\langle \langle G \rangle \rangle \psi \in \text{Sub}(\varphi)} |G|) + (\sum_{\langle \langle G^{\preceq_i} \rangle \rangle \psi \in \text{Sub}(\varphi)} |G|) + (\sum_{\langle \langle G^{\preceq_i} \rangle \rangle \psi \in \text{Sub}(\varphi)} |G|) + 1$ many actions.*

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

$$Ac' := Ac(\varphi) \cup \bigcup_{\langle\langle G \rangle\rangle\psi \in Sub(\varphi)} A_{G\psi} \\ \cup \bigcup_{\langle\langle G^{\preceq_i} \rangle\rangle\psi \in Sub(\varphi)} A_{G^{\preceq_i}\psi} \cup \\ \bigcup_{\langle\langle G^{\preceq_i} \rangle\rangle\psi \in Sub(\varphi)} A_{G^{\preceq_i}\psi} \cup \{\hat{a}\}.$$

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

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

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

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

For any $t \in T_{A'}^s$, we set $s \rightarrow_{A'} t$.

Then we can show by induction on $\psi \in Sub(\varphi)$ that $M^{\preceq}, s \models \psi$ iff $M'^{\preceq}, s \models \psi$. QED

The next step is to show that every satisfiable formula φ is also 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 . Here, the irreflexivity of the strict preferences is causing problems and thus we restrict our investigations to formulas that do not involve strict preferences. We denote this fragment of \mathcal{L}_{cla+p} by $\mathcal{L}_{cla+p}^{\neq}$ and the corresponding fragment of CLA+P by CLA+P[≠].

Using the method of filtration [5], we show that any satisfiable formula $\varphi \in \mathcal{L}_{cla+p}^{\neq}$ is satisfiable in a model with exponentially many states. Note that formulas of the form $\langle\langle G \rangle\rangle\psi$ and $\langle\langle G^{\preceq_i} \rangle\rangle\psi$ correspond to formulas of the form $\bigvee_{A \subseteq act(G)} [\bigwedge \Phi(A, G)]\psi$ and $\bigvee_{A \subseteq act(G)} ([\bigwedge \Phi(A, G)]\psi \wedge$

$[\bigwedge \Phi(A, G)]^{\preceq_i} \top$), respectively – for $\bigwedge \Phi(A, G)$ as in the proof of Observation 3.11. During the filtration, the underlying agents model is not changed and therefore the truth of formulas of the form $\langle\langle G \rangle\rangle\alpha$ is preserved.

Theorem 4.3 *Every satisfiable $\varphi \in \mathcal{L}_{cla+p}^{\neq}$ is also satisfiable in a MaSP with $\leq 2^{|\varphi|}$ many states.*

Proof. Given that $M^{\preceq}, s \models \varphi$ for some $M^{\preceq} = \langle S, Ac, (\rightarrow)_{A \subseteq Ac}, Ag, act, \{\preceq_i\}_{i \in Ag}, V \rangle$, $s \in S$, we obtain $M^{f \preceq f} = \langle S, Ac, (\rightarrow^f)_{A \subseteq Ac}, Ag, act^f, \{\preceq_i^f\}_{i \in Ag}, V^f \rangle$ 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
1. $\forall [\alpha]\psi \in Sub(\varphi)$ such that $A \models \alpha$: if $M^{\preceq}, s \models [\alpha]\psi$, then $M^{\preceq}, t \models \psi$,
 2. $\forall [\alpha]^{\preceq_i} \top \in Sub(\varphi)$ such that $A \models \alpha$: if $M^{\preceq}, s \models [\alpha]^{\preceq_i} \top$, then $s \preceq_i t$,
 3. $\forall \langle\langle G \rangle\rangle\psi \in Sub(\varphi)$ such that $A \models \bigwedge \Phi(A', G)$ for some $A' \subseteq act(G)$: if $M^{\preceq}, s \models [\bigwedge \Phi(A', G)]\psi$, then $M^{\preceq}, t \models \psi$,
 4. $\forall \langle\langle G^{\preceq_i} \rangle\rangle\psi \in Sub(\varphi)$ such that $A \models \bigwedge \Phi(A', G)$ for some $A' \subseteq act(G)$: if $M^{\preceq}, s \models [\bigwedge \Phi(A', G)]\psi$ and $M^{\preceq}, s \models [\bigwedge \Phi(A', G)]^{\preceq_i} \top$, then $M^{\preceq}, t \models \psi$ and $s \preceq_i t$.
- $|s| \preceq_i^f |t|$ iff
1. $\forall \diamond^{\preceq_i} \psi \in Sub(\varphi)$: if $M^{\preceq}, t \models \psi \vee \diamond^{\preceq_i} \psi$ then $M^{\preceq}, s \models \diamond^{\preceq_i} \psi$,
 2. If there is some $[\alpha]^{\preceq_i} \top \in Sub(\varphi)$, then $s \preceq_i t$,
 3. If there is some $\langle\langle G^{\preceq_i} \rangle\rangle\psi \in Sub(\varphi)$, then $s \preceq_i t$.

$V^f(p) := \{|s| | M, s \models p\}$, for all propositional letters $p \in Sub(\varphi)$. We 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 the filtration does not change the underlying agents model.

By definition of $S_{Sub(\varphi)}$, $|S_{Sub(\varphi)}| \leq 2^{|\varphi|}$. QED

Applying the constructions in the proofs of Theorems 4.1, 4.2 and 4.3 successively, we obtain the following:

Corollary 4.4 *Every satisfiable formula $\varphi \in \mathcal{L}_{cla+p}^{\neq}$ is satisfiable in a MaSP of size exponential in $|\varphi|$ satisfying the*

conditions $|Ag| \leq |Ag(\varphi)| + 1$ and $|Ac| \leq |Ac(\varphi)| + \sum_{\langle\langle G \rangle\rangle \psi \in Sub(\varphi)} |G| + (\sum_{\langle\langle G^{\neq i} \rangle\rangle \psi \in Sub(\varphi)} |G|) + 1$.

Having non-deterministically guessed a model of size exponential in $|\varphi|$, we can check in time exponential in $|\varphi|$ whether this model satisfies φ . This then gives us a NExpTime upper bound.

Theorem 4.5 *The satisfiability problem of $CLA+P^{\neq}$ is in NExpTime.*

Proof. Given φ , we non-deterministically choose a model M^{\neq} of size exponential in $|\varphi|$ satisfying the conditions $|Ag| \leq |Ag(\varphi)| + 1$ and $|Ac| \leq |Ac(\varphi)| + \sum_{\langle\langle G \rangle\rangle \psi \in Sub(\varphi)} |G| + (\sum_{\langle\langle G^{\neq i} \rangle\rangle \psi \in Sub(\varphi)} |G|) + 1$. Then, given this model, we can check in time $\mathcal{O}(|\varphi| ||M^{\neq}||)$, for $||M^{\neq}||$ being the size of M^{\neq} , whether M^{\neq} 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 the satisfiability problem of $CLA+P^{\neq}$ is in NExpTime. QED

This section has shown that the satisfiability problem of $CLA+P^{\neq}$ is in NExpTime. As the reader might expect, it has however a rather high computational complexity caused by the environment logic. The next section shows that the satisfiability problem of the environment logic is already NExpTime-hard and therefore adding agents as actors and preferences does not increase the complexity significantly.

4.2. Lower Bound

Establishing a NExpTime lower bound for the satisfiability problem of $CLA+P$ can be done by reducing that of the Boolean modal logic \mathbf{K}_m^{\cup} [10] to it, which is known to be NExpTime-complete [9].

Models of \mathbf{K}_m^{\cup} have a set of accessibility relations R_1, \dots, R_m and the associated language \mathcal{L}_m^{\cup} that is used for talking about the models contains corresponding basic modal parameters $\mathcal{R}_1, \dots, \mathcal{R}_m$. Using the operations \neg and \cup , more complex modal parameters can be built. The modalities then run along the corresponding sets of accessibility relations in the models.

Then a model M of \mathbf{K}_m^{\cup} with set of states W can be translated into an environment model $\tau_1(M)$ with set of states $W \cup \{u\}$ for some newly introduced state u and set of actions $Ac = \{a_1, \dots, a_m\}$. The accessibility relation $(\rightarrow)_{Ac}$ is defined as

$$w \rightarrow_A w' \text{ iff } A = \{a_i \mid (w, w') \in \mathcal{R}_i\} \text{ or } w' = u.$$

Thus, u is accessible from everywhere by any transition \rightarrow_A . This ensures that each \rightarrow_A is serial. Formulas $\varphi \in \mathcal{L}_m^{\cup}$ can be translated into $\tau_2(\varphi) \in \mathcal{L}_e$ in a straightforward way: Inside the modalities, modal parameters \mathcal{R}_i are translated into atomic actions a_i and complex parameters are translated into action expressions (\neg and \cup correspond to \neg and \vee respectively).

Theorem 4.6 *For any formula $\varphi \in \mathcal{L}_m^{\cup}$ and any model M of \mathbf{K}_m^{\cup} , for any state w in M :*

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

The reduction is polynomial and hence the satisfiability problems of $CLA+P^{\neq}$ and $CLA+P$ are NExpTime-hard.

Corollary 4.7 *The satisfiability problem of $CLA+P^{\neq}$ is NExpTime-complete.*

This section has shown that the satisfiability problem of $CLA+P$ without strict preferences is NExpTime-complete. This rather high complexity is due to the environment logic which itself is already NExpTime-complete. Adding agents with nonstrict preferences as actors to the environment logic does not increase the complexity significantly. Due to the undefinability of irreflexivity extending the complexity results of $CLA+P^{\neq}$ to full $CLA+P$ cannot be done using standard techniques such as filtration as we did in Theorem 4.3.

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 [14] which is combined with a preference logic [15, 16]. The resulting logic $CLA+P$, which is shown to be sound and complete, allows us to make explicit how groups of agents can achieve certain results. Additionally, we can 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 its members in the sense that after the achievement all of them are better 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.

The satisfiability problem of CLA+P without strict preferences is shown to be NExpTime-complete. It remains to be investigated whether the same holds for CLA+P. From a computational viewpoint, it seems to be appealing to change the environment logic in order to decrease computational complexity.

There are two immediate ways to extend the logic developed in this paper. First of all, we can follow the ideas of Ågotnes et al. [3] and add a restricted form of quantification that allows statements of the form $\langle\langle P^{\preceq_i} \rangle\rangle\psi$ saying that there is some group G that has property P and $\langle\langle G^{\preceq_i} \rangle\rangle\psi$.

Moreover, it might be promising to develop a cooperation logic with actions and preferences based on a logic for reasoning about complex plans such as the one developed by Gerbrandy and Sauro [7].

Acknowledgements

I am very grateful to Ulle Endriss, Eric Pacuit, Luigi Sauro, Johan van Benthem and Olivier Roy for their help and their useful comments. This work is funded by the Dutch Organization for Scientific Research (NWO) for the project TACTICS, grant number 612.000.525.

References

- [1] T. Ågotnes, P. E. Dunne, W. van der Hoek, and M. Wooldridge. Logics for coalitional games. In *LORI '07: Proceedings of the Workshop on Logic, Rationality and Interaction*, Beijing, China, 2007. to appear.
- [2] T. Ågotnes, W. van der Hoek, and M. Wooldridge. On the logic of coalitional games. In *AAMAS '06: Proceedings of the fifth International Joint Conference on Autonomous Agents and Multi-agent Systems*, pages 153–160, Hakodate, Japan, 2006.
- [3] T. Ågotnes, W. van der Hoek, and M. Wooldridge. Quantified coalition logic. In *IJCAI '07: Proceedings of the twentieth international joint conference on Artificial Intelligence*, pages 1181–1186, Hyderabad, India, 2007.
- [4] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Lecture Notes in Computer Science*, 1536:23–60, 1998.
- [5] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Number 53 in Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, UK, 2001.
- [6] S. Borgo. Coalitions in action logic. In *IJCAI '07: Proceedings of the twentieth international joint conference on Artificial Intelligence*, pages 1822–1827, Hyderabad, India, 2007.
- [7] J. Gerbrandy and L. Sauro. Plans in cooperation logic: a modular approach. In *Proceedings of the IJCAI Workshop on Nonmonotonic Reasoning, Action and Change (NRAC 2007)*, Hyderabad (India), 2007.
- [8] L. Kurzen. Logics for Cooperation, Actions and Preferences. Master's thesis, Universiteit van Amsterdam, the Netherlands, 2007.
- [9] C. Lutz and U. Sattler. The complexity of reasoning with boolean modal logics. In F. Wolter, H. Wansing, M. de Rijke, and M. Zakharyashev, editors, *Advances in Modal Logics Volume 3*. CSLI Publications, Stanford, 2001.
- [10] C. Lutz, U. Sattler, and F. Wolter. Modal logics and the two-variable fragment. In *Annual Conference of the European Association for Computer Science Logic CSL'01*, LNCS, Paris, France, 2001. Springer Verlag.
- [11] M. J. Osborne and A. Rubinstein. *A course in game theory*. MIT Press, Cambridge, MA, 1994.
- [12] M. Pauly. A modal logic for coalitional power in games. *Journal of Logic and Computation*, 12(1):149–166, 2002.
- [13] M. Pauly. On the complexity of coalitional reasoning. *International Game Theory Review*, 4:237–254, 2002.
- [14] L. Sauro, J. Gerbrandy, W. van der Hoek, and M. Wooldridge. Reasoning about action and cooperation. In *AAMAS '06: Proceedings of the fifth International Joint Conference on Autonomous Agents and Multi-agent Systems*, pages 185–192, Hakodate, Japan, 2006.
- [15] J. van Benthem, O. Roy, and P. Girard. Everything else being equal: A modal logic approach to ceteris paribus preferences, 2007.
- [16] J. van Benthem, S. van Otterloo, and O. Roy. Preference logic, conditionals and solution concepts in games. In *Festschrift for Krister Segerberg*. University of Uppsala, 2005.