Expressivity and Complexity of Reasoning about Coalitional Interaction: Extended Abstract

Cédric Dégremont Lena Kurzen

Abstract

This paper studies expressivity and complexity of normal modal logics for reasoning about coalitional power in multi-agent systems with agents that have preferences. A class of local and global notions relevant for reasoning about cooperation is identified. Many of these notions naturally correspond to concepts from game theory or social choice theory. The expressive power required to express the notions is determined via invariance results under various operations on different classes of Kripke models and frames. Our results apply to three different classes of models, each of them representing coalitional power from a slightly different perspective. One of these classes contains the models of a normal simulation of Pauly's Coalition Logic. A large class of known extended modal languages is specified and we show how the chosen notions can be expressed in fragments of this class. To determine how demanding reasoning about cooperation is in terms of computational complexity, we use complexity results for extended modal logics and obtain for each local notion an upper bound on the complexity of modal logics expressing it.

1 Introduction

Cooperation of agents is a major issue in fields such as computer science, economics and philosophy. The conditions under which coalitions are formed occur in various situations involving multiple agents. A single airline company e.g. cannot afford the cost of an airport runway whereas a group of companies can. Generally, agents can form groups in order to share complementary resources or because as a group they can achieve better results than individually. Various modal logic (ML) frameworks have been developed for reasoning about cooperation and coalitional power. Coalition Logic (**CL**) [Pau02] uses modalities of the form $\langle\!\langle C \rangle\!\rangle \phi$ saying that "coalition C has a joint strategy to ensure that ϕ". **CL** has neighborhood semantics but can be simulated on Kripke models [BHT07]. Another class of cooperation logics explicitly represents not only what results coalitions can achieve but also represents the strategies and actions by which groups can achieve something [WvdHW07, Bor07, GS07].

Another crucial concept for reasoning about interactive situations is that of *preferences*. It also received attention from modal logicians ([Gir08] surveys). Recent works (e.g. [ÅDvW07, ÅDvW09, Kur09]) propose different mixtures of cooperation and preference logics for. In such logics, many concepts from *game theory* (GT) and *social choice theory* (SCT) are commonly encountered. Depending on the situations to be modelled, different bundles of notions are important. Ability to express these notions – together with good computational behavior – make a logic appropriate for reasoning about cooperation.

1.1 Aim of our work

Rather than proposing a new logical framework, in the current work we analyze how demanding SCT and GT notions are for MLs in terms of expressivity and complexity. First, we choose three classes of models, each modelling cooperation an preferences in a different way. Then, we identify notions relevant for describing interactive situations and for each of the classes of models/frames we give satisfiability and validity invariance results as well as definability results for them, identifying the natural (extended) modal languages needed depending on the class of frames actually considered and the particular bundle of notions of interest. We draw some consequences about the complexity of reasoning about cooperation using ML. This way, our work makes it easier to identify safe and dangerous decision in terms of expressive power when designing MLs for cooperation and preferences. Our work allows us to compare classes of existing ML frameworks with respect to how suitable they are for reasoning about cooperation.

1.2 Methodology

We introduce three classes of ML models, each of which has been used to model cooperation. Then we identify a list of notions inspired by GT and SCT concepts for reasoning about coalitional power and preferences. Next, we look at how each notion can be interpreted in the models and determine the required expressive power for expressing it in modal languages. This is done both by investigating under which operations on models and frames the notions are invariant and by explicitly defining each notion in some (extended) modal language. This way, we obtain upper bounds on the complexity for satisfiability (SAT) and model checking (MC) problems of MLs able express the considered notions.

2 The Models: Three ways of modelling cooperation

This section gives an overview of the three different classes of models we consider in the present work. They correspond to models discussed in the literature which are used for modelling cooperation, each focusing on different aspects of coalitional power. We deliberately consider simplifying models or generalizations in order to avoid additional complexity due to assumptions about the models. This allows us to distinguish more clearly how the notions themselves are demanding and we can also evaluate from a high level perspective wich models are most appropriate for reasoning about certain aspects of cooperation.

The first class of models [DK08] focuses on preferences and their interaction with cooperation; it greatly simplifies the computation of coalitional powers because they are taken as primitives and are directly represented in the accessibility relation. The second class, *action-based coalitional models*, gives a natural account of coalitional power by representing it in terms of actions that agents can perform. The third class [BHT07] we consider are *power-based coalitional models*. Its focus lies on reasoning about and computing coalitional power itself, encoding the possible choices of a coalition as partitions of the state space. In each of the classes, preferences are represented as total preorders (TPO) over the states.

All our models are based on a finite set of agents $\mathtt{N}$. j usually ranges over $\mathtt{N}$. We denote the set of propositional letters by PROP and refer to the set of nominals by NOM; it is disjoint from PROP. A nominal is a propositional letter which is true in exactly one state. We usually let $p \in \text{PROP}$ and $i \in \text{NOM}$.

2.1 Coalition-labelled transition systems

A straightforward way to use Kripke models for reasoning about coalitional power is to focus on *sequential* systems, taking an accessibility relation for each coalition. The intuition is then that a group has the power to make the system move into exactly the states accessible by the relation.

Definition 1 ($\wp(\mathtt{N})$-LTS). *A $\wp(\mathtt{N})$-LTS (Labeled Transition Systems indexed by a finite set of coalitions $\wp(\mathtt{N})$) is of the form $\langle W, \mathtt{N}, \{ \xrightarrow{C} \mid C \subseteq \mathtt{N}\}, \{ \leq_j \mid j \in \mathtt{N}\}, V\rangle$, where $W \neq \emptyset$, $\mathtt{N} = \{1, \ldots, n\}$ for some $n \in \mathbb{N}$, $\xrightarrow{C} \subseteq W \times W$ for each $C \subseteq \mathtt{N}$, $\leq_j \subseteq W \times W$ for each $j \in \mathtt{N}$, and $V : \text{PROP} \cup \text{NOM} \to \wp W$, $|V(i)| = 1$ for each $i \in \text{NOM}$.*

W is the set of states, $\mathtt{N}$ a set of agents and $w \xrightarrow{C} v$ says that coalition C can change the state from w into v. Other interpretations are also possible, e.g. $w \xrightarrow{C} v$ could be interpreted as group preferences. $w \leq_j v$ means that j finds the v at least as good as w. $w \in V(p)$ means that p is true at w. Preferences are usually assumed to be total pre-orders (TPO). Let $\mathtt{TPO} - \wp(\mathtt{N}) - \mathtt{LTS}$ denote the class of $\wp(\mathtt{N}) - \mathtt{LTS}$ in which for each $j \in \mathtt{N}$, $\leq_j$ is a TPO. [DK08] has details.

2.2 Action-based coalitional models

In action-based coalitional models, coalitional power is represented in terms of actions. Agents can perform certain actions which then has the effect of changing the current state. The general idea is similar to that underlying some existing logics for cooperation, e.g. [Bor07, WvdHW07].

Definition 2 (N-ABC). *A $\mathtt{N}, (\mathtt{A_j})_{\mathtt{j \in N}}$-ABC (action-based coalitional model indexed by a finite set of agents $\mathtt{N}$ and a collection of finite sets of actions $(A_j)_{j \in N}$) is of the form $\langle W, \mathtt{N}, \{\xrightarrow{j,a} | j \in N, a \in A_j\}, \{\leq_j \mid j \in \mathtt{N}\}, V\rangle$, where $W \neq \emptyset$, $\mathtt{N} = \{1, \ldots, n\}$, for some $n \in \mathbb{N}$; for each $j \in N$ A_j is a finite set, $\xrightarrow{j,a} \subseteq W \times W$*

for each $j \in \mathtt{N}, a \in A_j$, $\leq_j \subseteq W \times W$ *is a* TPO *for each* $j \in \mathtt{N}$*, and* $V : \textsc{prop} \cup \textsc{nom} \to \wp(W)$, $|V(i)| = 1$ *for each* $i \in \textsc{nom}$*. Given a relation* $R \subseteq W \times W$*, we write* $R[w] := \{v \in W \mid wRv\}$*.*

W is the set of states, N a set of agents and $\xrightarrow{j,a}[w] \subseteq X$ means that at w agent j can guarantee by doing a that the next state will be in X. Thus, at w j can guarantee that the next state will be in X iff for some set Y,$X \supseteq Y \in \{\xrightarrow{j,a}[w] \mid a \in A_j\}$; ($Y$ will then be said to be in the *exact power* of j at w). Finally, we take powers to be additive, i.e. powers of coalitions to follow exactly from the powers of individuals. Therefore, at w coalition $C \subseteq \mathtt{N}$ can guarantee that the next state will be in X iff for some set Y we have $X \supseteq Y \in \{\bigcap_{j \in C} \xrightarrow{j,a_j}[w] \mid (a_1, \ldots, a_{|C|})\} \in \times_{j \in C} A_j\}$; ($Y$ is in the exact power of C at w). We let $\vec{C} := \times_{j \in C} A_j$. V and $\leq_j$ are as for $\mathtt{TPO} - \wp(\mathtt{N}) - \mathtt{LTS}$.

2.3 Power-based coalitional models

The following class of models that we present is (an extension of) a generalization of the models for the normal simulation of **CL** as presented in [BHT07] (with a preference ordering).

Definition 3 (PBC-Model)**.** *An* $\wp(\mathtt{N})$*-*PBC*-model (power based coalitional model indexed by a finite set of coalitions* $\wp(\mathtt{N})$*) is a tuple* $\langle W, \mathtt{N}, \{\sim_C \mid C \subseteq \mathtt{N}\}, F_{\mathbf{X}}, \{\leq_j \mid j \in \mathtt{N}\}, V\rangle$*, where* $W \neq \emptyset$*, each* $\sim_C \subseteq W \times W$ *is an equivalence relation,* $F_{\mathbf{X}} : W \to W$ *is a total function,* $\leq_j \subseteq W \times W$ *is a* TPO *for each* $j \in \mathtt{N}$*, and* $V : \textsc{prop} \cup \textsc{nom} \to \wp(W)$*,* $|V(i)| = 1$ *for each* $i \in \textsc{nom}$*.*

$W, \mathtt{N}$, $\leq_j$ and V are as in the previous model. $F_{\mathbf{X}}$ determines the actual course of actions in the system, i.e. if we are in w, then $F_{\mathbf{X}}(w)$ is the next state. The equivalence relation $\sim_C$ describes the power of coalition C. More precisely, $w \sim_C v$, this means that it is not in the power of C, to decide between w and v and thus neither whether we move to $F_{\mathbf{X}}(w)$ or $F_{\mathbf{X}}(v)$. Coalition C can on the other hand choose an equivalence class $[w]_{\sim_C}$ and thereby restrict the set of possible next states to $F_{\mathbf{X}}[[w]_{\sim_C}]$.

The models of the normal simulation **NCL** of **CL** are a special kind of PBC models.

Definition 4 (**NCL**-Independence)**.** *For every* $C \subseteq N$*,* $R_\emptyset \subseteq (R_C \circ R_{\overline{C}})$

Definition 5 (**NCL**-Model)**.** *An* ***NCL*** *model is a* N-PBC *model satisfying the following conditions:*
1. For all $C, D \subseteq \mathtt{N}$*, if* $D \subseteq C$*, then* $\sim_C \subseteq \sim_D$ *. 2.* **NCL***-Independence. 3.* $\sim_{\mathtt{N}} = id$*.*

3 Extended modal languages

For each type of models, we introduce a language from which we will actually consider different fragments.

3.1 Language interpreted on $\mathtt{N} - \mathtt{LTS}$.

$$\alpha ::= \ \leq_j \ \mid C \mid \alpha \cap \alpha \mid \overline{\alpha} \qquad \phi ::= \ p \mid i \mid x \mid \neg\phi \mid \phi \wedge \phi \mid \langle\alpha\rangle\phi \mid @_i\phi \mid @_x\phi \mid \ \downarrow x.\phi$$

where $j \in \mathtt{N}$, $C \in \wp(\mathtt{N}) - \{\emptyset\}$, $p \in \textsc{prop}$, $i \in \textsc{nom}$,$x \in \textsc{svar}$. SVAR is a countable set of variables.

Semantics. A program α is interpreted as a relation as indicated on the left. Formulas are interpreted together with an assignment $g : \textsc{svar} \to W$ as indicated (mostly) on the right. We skip booleans.

$\mathcal{M}, w, g \Vdash p$	iff	$w \in V(p)$	$\mathcal{M}, w, g \Vdash i$	iff	$w \in V(i)$
$R_{\leq_j}$	$=$	$\leq_j$	$\mathcal{M}, w, g \Vdash x$	iff	$w = g(x)$
R_C	$=$	$\xrightarrow{C}$	$\mathcal{M}, w, g \Vdash \langle\alpha\rangle\phi$	iff	$\exists v : wR_\alpha v$ and $\mathcal{M}, v, g \Vdash \phi$
$R_{\beta\cap\gamma}$	$=$	$R_\beta \cap R_\gamma$	$\mathcal{M}, w, g, \Vdash @_i\phi$	iff	$\mathcal{M}, v, g \Vdash \phi$ where $V(i) = \{v\}$
$R_{\overline{\beta}}$	$=$	$(W \times W) - R_\beta$	$\mathcal{M}, w, g, \Vdash @_x\phi$	iff	$\mathcal{M}, g(x), g \Vdash \phi$
			$\mathcal{M}, w, g, \Vdash \downarrow x.\phi$	iff	$\mathcal{M}, w, g[x := w] \Vdash \phi$

3.2 Language interpreted on ABC models.

The basic language for ABC is defined as follows (the extension with hybrid and boolean modal logic formulas is as for $\mathtt{N} - \mathtt{LTS}$).

$$\alpha ::= \ \leq_j \ \mid a_j \mid \alpha^{-1} \mid \alpha \cap \alpha \mid \overline{\alpha}$$

$$\phi := \ p \mid i \mid x \mid \neg\phi \mid \phi \wedge \phi \mid \langle\alpha\rangle\phi \mid \langle\alpha\rangle\phi \mid @_i\phi \mid @_x\phi \mid \ \downarrow x.\phi$$

where $j \in \mathtt{N}$, $a_j \in A_j$ (the set of actions available to j) and $p \in \textsc{prop}$.

Given a coalition $C \subseteq \mathtt{N}$, we let $\vec{C} := \times_{j\in C} A_j$. And given some profile of actions $\vec{c} \in \vec{C}$ we often write $\bigcap \vec{c}$ to stand for $(\bigcap_{a_j \in |\vec{c}|} a_j)$.

$$\begin{array}{lcl} R_{a_j} & = & \xrightarrow{j,a} \\ R_{\alpha^{-1}} & = & \{(v,w) | wR_\beta v\} \end{array}$$

Since there are no new clauses, we only give a few to give the intuition.

$$\begin{array}{lll} \mathcal{M},w,g \Vdash \langle a_j\rangle\phi & \text{iff} & \text{there is } v \text{ with } w\xrightarrow{j,a}v \text{ and } \mathcal{M},v,g \Vdash \phi \\ \mathcal{M},w,g \Vdash \langle \leq_j\rangle\phi & \text{iff} & \text{there is } v \text{ with } w \leq_j v \text{ and } \mathcal{M},v,g \Vdash \phi \\ \mathcal{M},w,g \Vdash \langle \alpha\rangle\phi & \text{iff} & \exists v : wR_\alpha v \text{ and } \mathcal{M},v,g \Vdash \phi \end{array}$$

3.3 Language interpreted on PBC/NCL models.

Here, the standard language $\mathcal{L}_{\mathbf{NCL}}$ for PBC and **NCL** is given as it is defined in [BHT07]. The extension with hybrid and boolean modal logic formulas are as for $\mathtt{N}-\mathtt{LTS}$.

$$\alpha ::= \ \leq_j \ | \ \alpha^{-1} \ | \ \alpha \cap \alpha$$

$$\phi := \ p \ | \ i \ | \ x \ | \neg\phi \ | \ \phi \wedge \phi \ | \ \langle C\rangle\phi \ | \ \mathbf{X}\phi \ | \ \langle\alpha\rangle\phi \ | \ @_i\phi \ | \ @_x\phi \ | \ \downarrow x.\phi$$

where $j \in \mathtt{N}$ (the set of agents), $C \in \wp(\mathtt{N})$ and $p \in \textsc{prop}$.

$$\begin{array}{lll} \mathcal{M},w,g \Vdash \langle C\rangle\phi & \text{iff} & \text{there is } v \text{ with } w \sim_C v \text{ and } \mathcal{M},v,g \Vdash \phi \\ \mathcal{M},w,g \Vdash \mathbf{X}\phi & \text{iff} & \mathcal{M},F_{\mathbf{X}}(w),g \Vdash \phi \\ \mathcal{M},w,g \Vdash \langle \leq_j\rangle\phi & \text{iff} & \text{there is } v \text{ with } w \leq_j v \text{ and } \mathcal{M},v,g \Vdash \phi \\ \mathcal{M},w,g \Vdash \langle \alpha\rangle\phi & \text{iff} & \exists v : wR_\alpha v \text{ and } \mathcal{M},v,g \Vdash \phi \end{array}$$

4 Thinking about coalitional power

In this section, we analyze the coalitional power as modelled in PBC, **NCL** and **CL**. We investigate the relations between standard assumptions about coalitional power. Then we consider the translation from **CL** to **NCL** [BHT07] and show how **NCL** models can be translated into corresponding **CL** models.

4.1 On the relation between PBC models and NCL models

Both PBC and **NCL** can be used for reasoning about coalitional power. We say that coalition C can force a set X at state w iff at w it can guarantee that the next state is in X. Put it differently, C can force X if some subset of X is in the exact power of C at w. In Section 5, we discuss global properties that can be reasonable assumptions about the coalitional powers. One kind of assumptions reflecting the independence of agents is generally assumed in the literature (cf. [Pau02, BHT07, BPX01]). To be precise, we distinguish between two assumptions about powers and show how they relate. Let $P_w(X)$ denote the collection of exact powers of coalition C at w; informally it contains the possible sets of states coalition C can choose from. Let $\overline{C} = \mathtt{N} \setminus C$ and $\overline{X} = W \setminus X$. Independence of coalitions says that whatever choices two disjoint coalitions make, there will be a next state that resulting from these choices.

Definition 6 (Independence of coalitions (IC))**.** *$\forall w$, if $C \cap D = \emptyset$ then $\forall X \in P_C(w)$ $\forall Y \in P_D(w)$ we have $X \cap Y \neq \emptyset$.*

The next condition says that if C can ensure that the next state is in X, $\overline{C}$ cannot ensure that is not.

Definition 7 (Condition about complementary coalitions (CCC))**.** *$\forall w$, $\forall X$, if $\exists X'$ such that $X \supseteq X' \in P_C(w)$, then there is no Y such that $\overline{X} \supseteq Y \in P_{\overline{C}}(w)$.*

Coalition monotonicity says that if a coalition can achieve something then so can all supersets of it.

Definition 8 (Coalition monotonicity (CM))**.** *$\forall w$ $\forall X$, if $C \subseteq D$ and $\exists Y$ such that $X \supseteq Y \in P_C(w)$, then $\exists Z$ such that $X \supseteq Z \in P_D(w)$.*

Fact 1. *IC implies CCC.*

Proof. Take some arbitrary w and write $P(C)$ for $P_C(w)$. Assume that $X \supseteq X' \in P(C)$ (a). Now assume for contradiction that here is some Y such that $\overline{X} \supseteq Y \in P(\overline{C})$ (b). Since $C \cap \overline{C} = \emptyset$, by Independence $Y \cap X' \neq \emptyset$ (c). Then by (c) and (b), $\overline{X} \cap X' \neq \emptyset$ (d). But (d) and (a) implies $\overline{X} \cap X \neq \emptyset$, a contradiction. Thus by reductio from (b), there is no Y such that $\overline{X} \supseteq Y \in P_{\overline{C}}(w)$. □

Fact 2. *CCC + CM implies IC.*

Proof. Assume that $C \cap D = \emptyset$, i.e. $D \subseteq \overline{C}$ (e). Take $X \in P(C)$ (f), $Y \in P(D)$ (g). Assume for contradiction that $X \cap Y = \emptyset$ (h), i.e. $Y \subseteq \overline{X}$ (i). From (f) and CCC it follows that there is no Z such that $\overline{X} \supseteq Y \in P(\overline{C})$ (j). But then by (j), (e) and CM there is no Z such that $\overline{X} \supseteq Y \in P(D)$ (k). But (k) contradicts (g). Thus by reductio from (h), $X \cap Y \neq \emptyset$ (m). Since X and$Y$ were arbitrary, by generalization from (m) that if $C \cap D \emptyset$ then $\forall X \in P_C(w)\ \forall Y \in P_D(w)$ we have $X \cap Y \neq \emptyset$. □

Note that on PBC models, CC is actually the following:

$$\forall w [\forall X \text{ if } \exists v (v \in R_\emptyset[w] \wedge R_C[v] \subseteq X), \text{ then } \neg \exists t (t \in R_\emptyset[w] \wedge R_{\overline{C}}[t] \subseteq \overline{X})] \quad (1)$$

For the specific case of PBC models which are **NCL** models, [BHT07] take the following condition.

Definition 9 (**NCL**-Independence)**.** *For every* $C \subseteq \mathtt{N}$, $R_\emptyset \subseteq (R_C \circ R_{\overline{C}})$

This assumptions as a natural modal axiomatization and on power based models the following holds:

Fact 3. *CCC is equivalent to* **NCL**-*Independence.*

Proof. From right to left. Assume that (1) does not hold. Then we have some w, X and v such that $v \in R_\emptyset[w]$ (2) and $R_C[v] \subseteq X$ (3), and moreover we have some t such that $t \in R_\emptyset[w]$ (4) and $R_{\overline{C}}[v] \subseteq \overline{X}$ (5). But by (3) and (5), $R_C[v] \cap R_{\overline{C}}[v] = \emptyset$ (6). By (2), (4) we have $v, t \in R_\emptyset[w]$ but since $R_\emptyset$ is an equivalence relation, it follows that $(v, t) \in R_\emptyset$ (7). But (7) and (6) together implies the negation of **NCL**-Independence.

From left to right. Assume that **NCL**-Independence does not hold. It follows that we have some w and z such that $z \in R_\emptyset[w]$ (8) and $(w, z) \notin R_C \circ R_{\overline{C}}$ (9). By (9) and the fact that for every C, R_C is an equivalence relation, $R_C[w] \cap R_{\overline{C}}[z] = \emptyset$ (10). Then $R_{\overline{C}}[z] \subseteq (W \setminus R_C[w])$ (11). Assume for contradiction that (1) holds. Instantiating w by w and X by $R_C[w]$ in (1) we get: if $\exists v (v \in R_\emptyset[w] \wedge R_C[v] \subseteq R_C[w])$, then $\neg \exists t (t \in R_\emptyset[w] \wedge R_{\overline{C}}[t] \subseteq \overline{R_C[v]})]$ (12). By reflexivity of $R_\emptyset$, $(w \in R_\emptyset[w] \wedge R_C[w] \subseteq R_C[w])$ (13). By (12) and existential generalization of (13), $\neg \exists t (t \in R_\emptyset[w] \wedge R_{\overline{C}}[t] \subseteq \overline{R_C[v]})]$ (14). But from (8) and (11), $(z \in R_\emptyset[w] \wedge R_{\overline{C}}[z] \subseteq \overline{R_C[v]})]$ (15). But (15) contradicts (14), thus (1) does not hold. □

4.2 The relation between NCL and CL

In this section, we take a closer look at the relation between **CL** and its normal simulation **NCL**. At this point, we will only briefly recall the semantics of **CL**. For the details of, we refer the reader to [Pau02].

Definition 10 (**CL**-Model)**.** *A* **CL**-*model is a pair* $((\mathtt{N}, W, E), V)$ *where* $\mathtt{N}$ *is a set of agents,* $S \neq \emptyset$ *is a set of states,* $E : W \to (\wp(\mathtt{N}) \to \wp(\wp(W)))$ *is called an effectivity structure. It satisfies the conditions of* ***playability****:*

- *Liveness:* $\forall C \subseteq \mathtt{N} : \emptyset \notin E(C)$,
- *Termination:* $\forall C \subseteq \mathtt{N} : W \in E(C)$,
- $\mathtt{N}$-*maximality.* $\forall X \subseteq W : (W \setminus X \notin E(\emptyset) \Rightarrow X \in E(N))$
- *Outcome monotonicity.* $\forall X \subseteq X' \subseteq W, C \subseteq N : (X \in E(C) \Rightarrow X' \in E(C))$,
- *Superadditivity.* $\forall X_1, X_2 \subseteq W, C_1, C_2 \subseteq \mathtt{N} : ((C_1 \cap C_2 = \emptyset \ \&\ X_1 \in E(C_1) \ \&\ X_2 \in E(C_2)) \Rightarrow X_1 \cap X_2 \in E(C_1 \cup C_2))$.

$V : \text{PROP} \to \wp(W)$ *is a propositional valuation function.*

The language $\mathcal{L}_{\mathbf{CL}}$ of **CL** is a standard modal language with a modality $\langle\!\langle C \rangle\!\rangle$ for each $C \subseteq \mathtt{N}$. The intended meaning of $\langle\!\langle C \rangle\!\rangle \phi$ is "coalition C has the power of achieving that ϕ". The semantics is as follows:

$$M, w \vDash \langle\!\langle C \rangle\!\rangle \phi \text{ iff } [\![\phi]\!]_M \in E(w)(C).$$

In what follows, we will write $E_w(C)$ for $E(w)(C)$.

Let us now give a brief overview of **NCL**, the normal simulation of **CL**. In [BHT07], a translation τ from the language $\mathcal{L}_{\mathbf{CL}}$ of **CL** to that of **NCL** is given such that for all $\phi \in \mathcal{L}_{\mathbf{CL}}$ it holds that ϕ is satisfiable in a model for **CL** iff $\tau(\phi)$ is satisfiable in a model of **NCL**. τ is defined as follows:

$$\tau(p) = p, \tau(\langle\!\langle C \rangle\!\rangle \phi) = \langle \emptyset \rangle [C] \mathbf{X} \tau(\phi)$$

The main result is then that ϕ is a theorem of **CL** iff $\tau(\phi)$ is one of **NCL**. Via completeness of **CL** and soundness of **NCL**, it follows then that whenever $\tau(\phi)$ is satisfied in a model of **NCL**, then there is a **CL** model that satisfies ϕ.

We want to make this result more explicit and show how we can translate each pointed model (M, w) of **NCL** into a model $f(M, w)$ of **CL** such that for all $\phi \in \mathcal{L}_{\mathbf{CL}}$, $(M, w) \Vdash \tau(\phi)$ iff $f(M, w) \Vdash \phi$.

Proposition 1. *For all $\phi \in \mathcal{L}_{\mathbf{CL}}$, if $\tau(\phi)$ is satisfiable in a pointed model M, w of* **NCL***, then ϕ is satisfiable in a model $f(M, w)$ of* **CL***.*

Proof. We define f as follows. For $M = \langle W, \mathtt{N}, \{\sim_C \mid C \subseteq \mathtt{N}\}, F_X, \{\leq_j \mid j \in \mathtt{N}\}, V \rangle$, $f(M) := \langle \mathtt{N}, (W, E), V \rangle$, where

$$E_w(C) := \{\{Y | Y \supseteq F_X[[w']_{\sim_C}]\} | w' \in [w]_{\sim_\emptyset}\}.$$

First, we show that $f(M)$ is a **CL** model, by showing that E is playable. Liveness follows from the totality of F_X. Termination. Follows from the closure of $E_w(C)$ under supersets. For $\mathtt{N}$-maximality, let $X \subseteq W$ such that $W \setminus X \notin E_w(\emptyset)$. Then there is some $w' \in [w]_{\sim_\emptyset}$ such that $F_X(w') \in X$. Since $X \supseteq F_X[\{w'\}] = \{F_X(w')\}$, $X \in E_w(\mathtt{N})$. Outcome-monotonicity follows from the closure of $E_w(C)$ under supersets. For superadditivity, let $X_1, X_2 \subseteq W, C_1, C_2 \subseteq \mathtt{N}$, such that $C_1 \cap C_2 = \emptyset$. Assume that $X_1 \in E_w(C_1)$ and that $X_2 \in E_w(C_2)$. This implies that for each $i \in \{1, 2\}$, there is some $w_i \in [w]_{\sim_\emptyset}$ such that $X_i \supseteq F_X[[w_i]_{\sim_{C_i}}]$. We have that $E_w(C_1 \cup C_2) = \{\{Y | Y \supseteq F_X[[w']_{\sim_{C_1 \cup C_2}}]\} | w' \in [w]_{\sim_\emptyset}\}$. Thus, we have to show there is some $w^+ \in [w]_{\sim_\emptyset}$ such that $X_1 \cap X_2 \supseteq F_X[[w^+]_{\sim_{C_1 \cup C_2}}]$. We have that $w_1 \sim_\emptyset w_2$. Thus, $w_1 \sim_{C_1} \circ \sim_{\overline{C_1}} w_2$ and since $C_1 \cap C_2 = \emptyset$ and thus $C_2 \subseteq \overline{C_1}$, we have that $\sim_{\overline{C_1}} \subseteq \sim_{C_2}$, which then implies that $w_1 \sim_{C_1} \circ \sim_{C_2} w_2$. Then there is some w^+ such that $w_1 \sim_{C_1} w^+$ and $w^+ \sim_{C_2} w_2$. Thus, $w^+ \in [w_1]_{\sim_{C_1}} \cap [w_2]_{\sim_{C_2}}$ and therefore $[w^+]_{\sim_{C_1}} = [w_1]_{\sim_{C_1}}$ and $[w^+]_{\sim_{C_2}} = [w_2]_{\sim_{C_2}}$. Since $\sim_{C_1 \cup C_2} \subseteq (\sim_{C_1} \cap \sim_{C_2})$, $[w^+]_{\sim_{C_1 \cup C_2}} \subseteq [w^+]_{\sim_{C_1}} \cap [w^+]_{\sim_{C_2}}$. Hence, $F_X[[w^+]_{\sim_{C_1 \cup C_2}}] \subseteq X_1 \cap X_2$, and thus $X_1 \cap X_2 \subseteq E_w(C_1 \cup C_2)$.

This shows that $f(M)$ is a **CL** model. Now, we show that for all $\phi \in \mathcal{L}_{\mathbf{CL}}$, for an **NCL** model M, $M, w \Vdash \tau(\phi)$ iff $f(M, w) \Vdash \phi$. This is done by induction on ϕ. The only interesting case is that where $\phi := \langle\!\langle C \rangle\!\rangle \psi$. Let $M, w \Vdash \langle \emptyset \rangle [C] \mathbf{X} \tau(\psi)$. Then there is some $w' \in [w]_{\sim_\emptyset}$ such that for all $w'' \in [w']_{\sim_C}$, $M, F_X(w'') \Vdash \tau(\psi)$. By induction hypothesis, $f(M, F_X(w'')) \Vdash \psi$. Now, we want to show that $[\![\psi]\!]_{f(M,w)} \in E_w(C)$. This follows from the fact that for all $w'' \in [w']_{\sim_C}$, $f(M, F_X(w'')) \Vdash \psi$. For the other direction, let $f(M, w) \Vdash \langle\!\langle C \rangle\!\rangle \psi$. Then, there is some $X \in E_w(C)$ such that $X \subseteq [\![\psi]\!]_{f(M,w)}$. By definition of $f(M, w)$, this means that there is some $w' \in [w]_{\sim_\emptyset}$ such that $X \supseteq F_X[[w']_{\sim_C}]$. Since by inductive hypothesis, $[\![\tau(\psi)]\!]_{M,w} = [\![\psi]\!]_{f(M,w)}$, $X \subseteq [\![\tau(\psi)]\!]_{M,w}$. Hence, $M, w \Vdash \langle \emptyset \rangle [C] \mathbf{X} \tau(\psi)$. □

5 GT and SCT notions: how demanding are they?

In this section, we briefly sketch the notions that we consider in our work. In the most basic case, reasoning about cooperative interaction considers what coalitions of agents can achieve and what individuals prefer. Using these elements, more elaborated notions can be built. We consider natural counterparts of SCT and GT notions and are interested in local notions i.e. properties of a particular state in a particular system, i.e. properties of pointed models $\mathcal{M}, w$. Moreover, we consider global notions, which are properties of classes of systems: we are interested in the class of frames that a property characterizes. With respect to content, apart from notions describing only coalitional powers or preferences, we consider stability and effectivity concepts. We give a few representative notions. For more see [DK08].

Local Notions

1. Coalition C can guarantee that at the next state p is true.
2. There is a state i finds at least as good where p holds.
3. C can guarantee that the next state is one j finds a.l.a.g. as the current one.
4. There is a state that all agents in C prefer but coalition C cannot achieve it.
5. The current state is (strongly) *Nash* stable, i.e. no agent has the power to guarantee that the next state will be one that she strictly prefers (finds at least as good) to the current one.
6. Strong Local Dictator
7. The current state is weakly *Pareto*-efficient.
8. The current state is *Pareto*-efficient.
9. The current state is strongly *Pareto*-efficient.

Notions 1 and 2 describe basic concepts about cooperation and preferences, respectively. 4 is a property about the limits of the power of the grand coalition. It says that even though there is a state that would be better for everybody, even if all agents work together they cannot achieve that this will be the next state. 5 is a stability notion saying that no agent has an incentive to change the current state by himself.

Global Notions

1. Only coalitions containing a majority of `N` have nontrivial power.
2. Coalition monotonicity: if D is a subset of C then for all X, if D can force X then so can C.
3. Conditions about consistency of complementary coalitations
4. One agent is a strong local dictator in every state (*strong dictator*). $\exists j \forall w(\forall C \forall X(X \in P_C(w) \rightarrow (X \subseteq (\leq_j [w]))))$
5. Coalitions can achieve only what all its members prefer.
6. If j can guarantee that the next state is one j strictly prefers to the current one then for every set X and for all coalition C, if C cannot force the next state to be in X but $C \cup \{i\}$ can, then j finds all states in X at least as good as the current one.

2 is an important global property which says that if a group can achieve something, then so can every superset of agents. This property is assumed in **CL** and **NCL** and holds in `ABC` models. 5 can be a reasonable restriction on coalitions' power. It says that only achieve results that can be achieved without making a member worse off.

5.1 Representative Results

We give the strongest invariance results we found. We consider the operations ($\cap$-)Bisimulation, disjoint unions (`DU`), generated submodel (-frame)(`GSM` (`GSF`)) and bounded morphic images (`BMI`). We then define the notions using extended modal languages (defined in the obvious way given the similarity type considered) for which the model-theory is well-understood. (**X** scans F_X.) Finally we draw conclusions about complexity in terms of upper bounds (UB) on SAT and on the Combined Complexity of MC.

5.2 Operations on models and frames and Invariance

Expressivity of MLs is usually characterized by invariance results. Definitions of relevant relations and operations between models follow. Let τ be a finite modal similarity type with only binary relations. Let $\mathcal{M} = \langle W, (R_k)_{k\in\tau}, V\rangle$ and $\mathcal{M}' = \langle W', (R'_k)_{k\in\tau}, V'\rangle$ be models of similarity type τ.

Definition 11 (Bisimulations)**.** *A bisimulation between $\mathcal{M}$ and $\mathcal{M}'$ is a non-empty binary relation $Z \subseteq W \times W'$ fulfilling the following conditions:*

`AtomicHarmony`	*For every $p \in$ PROP, wZw' implies $w \in V(p)$ iff $w' \in V'(p)$.*
`Forth`	$\forall k \in \tau$ *, if wZw' & $R_k wv$ then $\exists v' \in W'$ s.t. $R'_k w'v'$ & vZv'.*
`Back`	$\forall k \in \tau$ *, if wZw' & $R'_k w'v'$ then $\exists v \in W$ s.t. $R_k wv$ & vZv'.*

In a nutshell $\cap$-Bisimulations require that `Back` and `Forth` also hold for the intersection (resp. the converse) of the relations. We now define bounded morphisms, generated subframes and disjoint unions.

Definition 12 (BM). *$f : W \rightarrow W'$ is a bounded morphism from $\mathcal{M}$ to $\mathcal{M}'$ iff:*

AtomicHarmony	*For every $p \in$ PROP, $w \in V(p)$ iff $f(w) \in V'(p)$.*
R − homomorphism	*$\forall k \in \tau$, if $R_k wv$ then $R' f(w) f(v)$.*
Back	*$\forall k \in \tau$, if $R'_k f(w) v'$ then $\exists v \in W$ s.t. $f(v) = v'$ and $R_k wv$.*

Definition 13 (Generated Submodel). *We say that that $\mathcal{M}'$ is a generated submodel (GSM) of $\mathcal{M}$ iff $W' \subseteq W$, $\forall k \in \tau$, $R'_k = R_k \cap (W' \times W')$, $\forall p \in$ PROP , $V'(p) = V(p) \cap (W' \times W')$ and if $w \in W'$ and Rwv then $v \in W'$.*

Definition 14 (Disjoint Unions). *Let $(\mathcal{M}_j)_{j \in J}$ be a collection of models with disjoint domains. Define their disjoint union $\biguplus_j \mathcal{M}_j = \langle W, R, V \rangle$ as the union of their domains and relations, and define for each $p \in$ PROP, $V(p) := \bigcup_j V_j(p)$.*

Definition 15 (Invariance). *A property of pointed models $\Phi(X, y)$ is* invariant under λ-Bisimulations *iff whenever there exists a λ-bisimulation Z between $\mathcal{M}$ and $\mathcal{M}'$ such that $(w, w') \in Z$, then $\Phi(\mathcal{M}, w')$ holds iff $\Phi(\mathcal{M}', w')$ holds. Invariance for other operations is defined similarly.*

We now consider closure conditions. First, we consider bounded morphic images (BMI) of frames. BM on frames are obtained by dropping AtomicHarmony in Def. 12. A class of frames is closed under BMI iff it is closed under *surjective* BM. Next, we consider closure under generated subframes (GSF) – the frame-analogue to GSM (cf. Def. 13). We also check if properties *reflect* GSF. A property ϕ *reflects* GSF if whenever for every frame $\mathcal{F}$, it holds that every GSF of $\mathcal{F}$ has property ϕ, then so does $\mathcal{F}$. We also consider closure under taking disjoint unions (DU) of frames, which are defined in the obvious way.

5.2.1 Pointed Model Definability

Simple coalitional power (Notion 1).

	$\wp$(N) − LTS	ABC	PBC	**NCL**
Strongest Invariance	Bisimulation	$\cap$-Bisimulation	Bisimulation	Bisimulation
Definability	$\langle C \rangle p$	$\bigvee_{\vec{a_j} \in \vec{C}} [\bigcap \vec{a_j}] p$	$\langle \emptyset \rangle [C] \mathbf{X} p$	$\langle \emptyset \rangle [C] \mathbf{X} p$
UB on SAT	PSPACE	PSPACE	PSPACE	NEXPTIME [BGH+08]
UB on MC	PTIME	PTIME	PTIME	PTIME

Remark: PSPACE-membership on ABC models should be taken carefully when comparing it to other logics. Indeed when e.g. set of actions of same size are given to every agents, the length of above formula is exponential in the number of agents.

Simple preference (Notion 2).

	$\wp$(N) − LTS	TPO − $\wp(N)$ − LTS	ABC	PBC
Strongest Invariance	Bisimulation	Bisimulation	Bisimulation	Bisimulation
Definability	$\langle \leq_j \rangle p$	$\langle \leq_j \rangle p$	$\langle \leq_j \rangle p$	$\langle \leq_j \rangle p$
UB on SAT	PSPACE	PSPACE	PSPACE	PSPACE
UB on MC	PTIME	PTIME	PTIME	PTIME

Remark: NP-membership is out of reach even when assuming preferences to be TPOs since we consider the complexity of the full logic of cooperation and preferences and not only of the preference part.

***C* can guarantee that the next state is one *j* finds a.l.a.g. as the current one.** (Notion 3).

	$\wp$(N)LTS	ABC	PBC
Strongest Invariance	$\cap$-Bisimulation	GSM and DU	GSM and DU
Definability	$\langle C \cap \leq_j \rangle \top$	$\bigvee_{\vec{a_j} \in \vec{C}} (\downarrow x.[\bigcap \vec{a_j}](\downarrow y.@_x \langle \leq_j \rangle y))$	$\downarrow x.\langle \emptyset \rangle [C] \mathbf{X} \downarrow y.@_x \langle \leq_j \rangle y$
UB on SAT	PSPACE [DLNN91]	Π^0_1	Π^0_1
UB on MC	PTIME[Lan06]	PSPACE	PSPACE

Note that in $\wp$(N) − LTS, we only need the intersection, whereas for ABC and PBC, we have to express that the states accessible by one relation are a subset of the states accessible by another relation.

There is a state that all members of coalition *C* prefer but *C* cannot achieve it (Notion 4).

	$\wp(\mathtt{N}) - \mathtt{LTS}$	ABC	PBC
Strongest Invariance	GSM and DU	GSM and DU	GSM and DU
Definability	$\langle(\bigcap_{j\in C} \leq_j) \cap \overline{(\xrightarrow{C})}\rangle\top$	$\downarrow x.\langle\bigcap_{j\in C} \leq_j\rangle \downarrow y.@_x \bigwedge_{\vec{a_j}\in\vec{C}}\langle\bigcap \vec{a_j}\rangle\neg x$	$\downarrow x.\langle\bigcap_{j\in C} \leq_j\rangle \downarrow y.@_x\neg[C]\mathbf{X}y$
UB on SAT	EXPTIME	Π^0_1	Π^0_1
UB on MC	PTIME	PSPACE	PSPACE

Nash-stability (Notion 5).

	$\wp(\mathtt{N}) - \mathtt{LTS}$	ABC	PBC
Strongest Invariance	GSM and DU	GSM and DU	GSM and DU
Definability	$\bigwedge_{j\in\mathtt{N}} \downarrow x.[j\cap \leq_j]\langle\leq_i\rangle x$	$\bigwedge_{j\in N}\bigwedge_{a_j\in A_j} \downarrow x.\langle a_j\rangle\langle\leq\rangle x$	$\bigwedge_{j\in N} \downarrow x.[\emptyset]\langle\{j\}\rangle\mathbf{X}\langle\leq\rangle x$
UB on SAT	Π^0_1	EXPTIME	EXPTIME
UB on MC	PSPACE	PSPACE	PSPACE

Remarks: The notions of Nash-stability used in $\wp(\mathtt{N}) - \mathtt{LTS}$ and ABC/PBC/**NCL** models are strongly related but mathematically slightly different. Note also that our logical definition of Nash-stability with respect to ABC models crucially draws on the fact that the preference relation is a total pre-order.

Strong Nash-stability (Notion 5, strong version).

	$\wp(\mathtt{N}) - \mathtt{LTS}$	ABC	PBC
Strongest Invariance	$\cap$-Bisimulation	GSM and DU	GSM and DU
Definability	$\bigwedge_{j\in\mathtt{N}}[i\cap \leq_j]\bot$	$\neg\bigvee_{j\in N}\bigvee_{a_j\in A_j} \downarrow x.[a_j]\langle\leq^{-1}\rangle x$	$\neg\bigvee_{j\in N} \downarrow x.\langle\emptyset\rangle[\{i\}]\langle\leq^{-1}\rangle x$
UB on SAT	PSPACE	Π^0_1	Π^0_1
UB on MC	PTIME	PSPACE	PSPACE

Remark: Same remarks as for Nash-stability.

Existence of a strong local dictator (Notion 6).

	$\wp(\mathtt{N}) - \mathtt{LTS}$	ABC	PBC	**NCL**
Strongest Invariance	GSM and DU?	GSM and DU	GSM and DU	–
Definability	$[C \cap \overline{\leq_j}]\bot$	$[a \cap \overline{\leq_j}]\bot$	$\downarrow x.[\emptyset][C]X \downarrow y.@_x\langle\leq_j\rangle y$	$\top$
UB on SAT	EXPTIME	EXPTIME	Π^0_1	NEXPTIME
UB on MC	PTIME	PTIME	PTIME	PTIME

weak *Pareto*-efficiency (Notion 7).

	$\wp(\mathtt{N}) - \mathtt{LTS}$	$\mathtt{TPO} - \wp(N) - \mathtt{LTS}$	ABC	PBC
Strongest Invariance	GSM and DU	GSM and DU	GSM and DU	GSM and DU
Definability	$\downarrow x.[\bigcap_{j\in\mathtt{N}} \leq_j]\bigvee_{j\in\mathtt{N}}\langle\leq_i\rangle x$	as $\wp(N) - \mathtt{LTS}$	as $\wp(N) - \mathtt{LTS}$	as $\wp(N) - \mathtt{LTS}$
UB on SAT	Π^0_1	Π^0_1	Π^0_1	Π^0_1
UB on MC	PSPACE	PSPACE	PSPACE	PSPACE

***Pareto*-efficient** (Notion 8).

	$\wp(\mathtt{N}) - \mathtt{LTS}$	$\mathtt{TPO} - \wp(N) - \mathtt{LTS}$	ABC	PBC
Strongest Invariance	GSM and DU	GSM and DU	GSM and DU	GSM and DU
Definability	$\neg \downarrow x.\langle\bigcap_{j\in N} \leq_j\rangle(\bigvee_{j\in\mathtt{N}}[\leq_j]\neg x)$	as $\wp(\mathtt{N}) - \mathtt{LTS}$	as $\wp(\mathtt{N}) - \mathtt{LTS}$	as $\wp(\mathtt{N}) - \mathtt{LTS}$
UB on SAT	Π^0_1	Π^0_1	Π^0_1	Π^0_1
UB on MC	PSPACE	PSPACE	PSPACE	PSPACE

strong *Pareto*-efficiency (Notion 9).

	$\wp(\mathtt{N}) - \mathtt{LTS}$	$\mathtt{TPO} - \wp(N) - \mathtt{LTS}$	ABC	PBC
Strongest Invariance	$\cap$-Bisimulation	$\cap$-Bisimulation	$\cap$-Bisimulation	$\cap$-Bisimulation
Definability	$[\bigcap_{j\in\mathtt{N}} \leq_j]\bot$	$[\bigcap_{j\in\mathtt{N}} \leq_j]\bot$	$[\bigcap_{j\in\mathtt{N}} \leq_j]\bot$	$[\bigcap_{j\in\mathtt{N}} \leq_j]\bot$
UB on SAT	PSPACE	PSPACE	PSPACE	PSPACE
UB on MC	PTIME	PTIME	PTIME	PTIME

The crucial difference between Notions 7 and 8 on the one hand and Notion 9 on the other hand is that (weak) Pareto efficiency talks about strict preference whereas strong Pareto efficiency does not. This results in a difference in required expressive power and in complexity.

5.2.2 Frame Definability

Majority (Notion 1).

	$\wp(\mathtt{N})-\mathtt{LTS}$	ABC	PBC	**NCL**
Strongest closure	GSF, DU, BMI	GSF (BMI?)	GSF, DU, BMI	GSF, DU, BMI
Definability	$\bigwedge_{C:\lvert C\rvert<\lvert\mathtt{N}\rvert/2}[C]\bot$	$\bigwedge_{C:\lvert C\rvert<\lvert\mathtt{N}\rvert/2}(\mathtt{E}p \rightarrow \bigwedge_{\vec{a_j}\in\vec{C}}\langle\bigcap\vec{a_j}\rangle p)$	$\bigwedge_{C:\lvert C\rvert<\lvert\mathtt{N}\rvert/2}([C]p \leftrightarrow [\emptyset]p)$	$\bigwedge_{C:\lvert C\rvert<\lvert\mathtt{N}\rvert/2}([C]p \rightarrow [\emptyset]p)$
UB on SAT	PSPACE	EXPTIME	PSPACE	PSPACE
UB on MC	PTIME	PTIME	PTIME	PTIME

Coalition monotonicity (Notion 2).

	$\wp(\mathtt{N})-\mathtt{LTS}$	ABC	PBC	**NCL**
Strongest closure	GSF, DU, BMI	–	GSF, DU, BMI	–
Definability	$\langle C\rangle p \rightarrow \langle C\cup D\rangle p$	$\top$	$[C]p \rightarrow [C\cup D]p$	$\top$
UB on SAT	PSPACE	PSPACE	PSPACE	NEXPTIME [BGH+08]
UB on MC	PTIME	PTIME	PTIME	PTIME

Consistency of complementary coalitions (Notion 3).

	$\wp(\mathtt{N})-\mathtt{LTS}$	ABC	PBC	**NCL**
Strongest closure	GSF, DU, BMI	GSF, DU, BMI	GSF, DU, BMI	–
Definability	$\langle C\rangle p \rightarrow [N\setminus C]p$	$(\bigvee_{\vec{c}\in\vec{C}}(\langle\bigcap_{a_c\in\lvert\vec{c}\rvert} a_c\rangle\top \wedge [\bigcap a_c]p)) \rightarrow \bigwedge_{\vec{d}\in(\vec{\mathtt{N}\setminus C})}\langle\bigcap_{a_d\in\lvert\vec{d}\rvert} a_d\rangle p$	$\langle\emptyset\rangle[C]Xp \rightarrow [\emptyset]\langle(\mathtt{N}\setminus C)\rangle\mathbf{X}p$	$\top$
UB on SAT	PSPACE	PSPACE	PSPACE	NEXPTIME
UB on MC	PTIME	PTIME	PTIME	PTIME

Existence of a strong global director (Notion 4).

	$\wp(\mathtt{N})-\mathtt{LTS}$	ABC	PBC	**NCL**
Strongest closure	GSF, DU, BMI	GSF, DU, BMI	GSF, DU, BMI	–
Definability	$\bigvee_{j\in\mathtt{N}}(\langle C\rangle p \rightarrow \langle\leq_j\rangle p)$	$\bigvee_{j\in N}\bigwedge_{k\in N}\bigwedge_{a_k\in A_k}(\langle a_k\rangle p \rightarrow \langle\leq_j\rangle p)$	$\langle\emptyset\rangle\langle C\rangle Xp \rightarrow \langle\leq_j\rangle p$	$\top$
UB on SAT	PSPACE	PSPACE	PSPACE	NEXPTIME
UB on MC	PTIME	PTIME	PTIME	PTIME

Coalitions can achieve only what all its members prefer. (Notion 5).

	$\wp(\mathtt{N})-\mathtt{LTS}$	ABC	PBC	**NCL**
Strongest Closure	GSF, DU, BMI	GSF, DU	GSF, DU, BMI	GSF, DU, BMI
Definability	$\langle C\rangle\phi \rightarrow \bigwedge_{j\in\mathtt{N}}\langle\leq_i\rangle\phi$	$\bigwedge_{\vec{a_j}\in\vec{C}}(\langle\bigcap\vec{a_j}\rangle\phi \rightarrow \langle\bigcap_{j\in C}\leq_j\rangle\phi)$	$\langle\emptyset\rangle[C]\mathbf{X}\phi \rightarrow \langle\leq_j\rangle\phi$	$\langle\emptyset\rangle[C]\mathbf{X}\phi \rightarrow \langle\leq_j\rangle\phi$
UB on SAT	PSPACE[HM92]	PSPACE	PSPACE	PSPACE
UB on MC	PTIME[FL79]	PTIME	PTIME	PTIME

Individual Rationality (Notion 6).

	$\wp(\mathtt{N}) - \mathtt{LTS}$	ABC	PBC
Strongest Closure	GSF and DU	GSF and DU	GSF and DU
Definability	$[p \wedge \langle i \rangle q \wedge \langle \leq_j \rangle (q \wedge \langle \leq_i \rangle \neg p)] \rightarrow \bigwedge_{j \in C \subseteq \mathtt{N}} [(\langle C \rangle r \wedge \bigwedge_{D \subseteq C \setminus i} \neg \langle D \rangle r) \rightarrow \langle \leq_i \rangle (r \wedge \neg \langle \leq_i \rangle p)]$	$(i \wedge \bigvee_{a_j \in A_j} [a_j][\leq_j] \neg i \wedge \bigvee_{\vec{d} \in (C \cup \vec{\{j\}})} ([\bigcap \vec{d}]p) \rightarrow (\bigvee_{\vec{c} \in \vec{C}} [\bigcap \vec{c}]p) \vee [\bigcap \vec{d}]((p \wedge k) \rightarrow @_i \langle \leq_j \rangle k)))$	$(i \wedge \langle \emptyset \rangle [j] \mathbf{X} [\leq_j] \neg i \wedge \langle \emptyset \rangle [C \cup j] X p) \rightarrow (\langle \emptyset \rangle [C] X p \vee \vee [\emptyset][C \cup j] X ((p \wedge k) \rightarrow @_i \langle \leq_j \rangle k))$
UB on SAT	PSPACE	PSPACE	PSPACE
UB on MC	PTIME	PTIME	PTIME

5.3 More general observations

For purely preference based notions such as efficiency notions, reasoning about *strong* efficiency concepts (which contain the non-existence of a state which is *at least as good*) is computationally easier than reasoning about *weak* versions of these concepts (containing the non-existence of a state which is *strictly better*). The same holds for reasoning about stability notions in $\wp(\mathtt{N}) - \mathtt{LTS}$. For reasoning about stability-like notions on ABC and PBC models the opposite is true. Concerning preferences, we observe that having strict preferences as a primitive computationally simplifies reasoning about preferences. In the case of $\wp(\mathtt{N}) - \mathtt{LTS}$, all the notions for which we couldn't find decidable logics, involve strict preferences.

6 Conclusion

Our work analyzes the complexity and expressive power needed for reasoning about cooperation using (extended) modal logics. We considered three different classes of models representing different ways how cooperation has been modelled in the literature. In $\wp(\mathtt{N}) - \mathtt{LTS}$, coalitional power is a primitive and directly represented in the accessibility relations for coalitions. In the second class of models, ABC, coalitional power of agents is represented in terms of the actions that agents can perform and whose performance then leads to the system changing its current state. The power of a coalition is computed from the powers of the individuals in it. In power-based coalitional models (PBC), every states belongs to a partition indicating the choice currently be made by the coalition. We extended the models with a representation of agents' preferences as total preorders over the set of states.

Then we identified a wide range of notions relevant for reasoning about cooperation and preferences of agents. The notions range from very simple expressions about the power to achieve some proposition and expressions saying that an agent prefers a state where some proposition holds, to more complex notions combining agents' preferences and (coalitional) powers such as *individual rationality* and various stability and efficiency notions. In general, we distinguis between local notions, describing properties of a state of a system, and global notions describing properties of classes systems.

For each of the classes of models we defined a family of extended modal languages.

For each of the notions and class of models, we gave invariance/closure results. Moreover, we gave explicit definability results, showing how each of the local notions can be defined for each of the classes of models. For the global notions, we showed how these properties can be axiomatized.

Our main results concerning our aim of determining how demanding GT and SCT notions are, can be summarized as follows:

- We can clearly see how the decision to take the preferences to be total preorders influences the required expressive power to express some of the notions involving preferences. In $\wp\mathtt{N} - \mathtt{LTS}$ not having strict preferences as a primitive has the effect of making so called "strong" versions of stability and efficiency notions easier to express than their "weak" notions. For ABC and PBC on the other hand, we observed the opposite effect: "weak" notions are easier to express than their "strong" versions. This is due to the fact that in these models, the strong versions require the converse modality, whereas the weak ones do not.

- The complexity results have to be taken with some caution, since they crucially depend on the parameters. In ABC e.g., many of the formulas defining the notions we considered are actually exponential in the number of actions.

- Comparing the results for the different classes models, we can see that local notions involving the power of a coalition to make the system move to a state preferred by some agent such as Notion 3 and strond Nash-stability, we can see that for $\wp(\mathtt{N}) - \mathtt{LTS}$ this can be done using logics with SAT in PSPACE, whereas for ABC and PBC, we could only find logics with undecidable SAT.

Our results help to make design choices concerning the models and languages to be used when developing modal logics for reasoning about (preference-based) cooperation in multi-agent systems. Based on the GT and SCT notions that one would like to express and the constraints on the computational tractability, appropriate choices can then be made.

References

[ÅDvW07] Thomas Ågotnes, Paul E Dunne, Wiebe van der Hoek, and Michael Wooldridge. Logics for coalitional games. In *LORI '07*, Beijing, China, 2007.

[ÅDvW09] Thomas Ågotnes, Paul E Dunne, Wiebe van der Hoek, and Michael Wooldridge. Reasoning about coalitional games. *Artificial Intelligence*, 173(1):45–79, 2009.

[BGH+08] Philippe Balbiani, Olivier Gasquet, Andreas Herzig, Franois Schwarzentruber, and Nicolas Troquard. Coalition games over Kripke semantics. In C. Dégremont, L. Keiff, and H. Rückert, editors, *Dialogues, Logics and Other Strange Things – Essays for Shahid Rahman*, pages 11–32. College Publications, 2008.

[BHT07] Jan Broersen, Andreas Herzig, and Nicolas Troquard. Normal Coalition Logic and its conformant extension. In Dov Samet, editor, *TARK'07*, pages 91–101. PUL, 2007.

[Bor07] Stefano Borgo. Coalitions in action logic. In *IJCAI '07: Proceedings of the twentieth international joint conference on Artificial Intelligence*, pages 1822–1827, 2007.

[BPX01] Nuel Belnap, Michael Perloff, and Ming Xu. *Facing the future: Agents and choices in our indeterminist world.* Oxford University Press, Oxford, 2001.

[DK08] Cédric Dégremont and Lena Kurzen. Modal logics for preferences and cooperation: Expressivity and complexity. ILLC PP-2008-39, University of Amsterdam, 2008.

[DLNN91] Francesco M. Donini, Maurizio Lenzerini, Daniele Nardi, and Werner Nutt. The complexity of concept languages. In *KR*, pages 151–162, 1991.

[FL79] M. J. Fischer and R. E. Ladner. Propositional dynamic logic of regular programs. *J. Comput. Syst. Sci.*, 1979.

[Gir08] Patrick Girard. *Modal Logic for Preference Change.* PhD thesis, Stanford, 2008.

[GS07] Jelle Gerbrandy and Luigi Sauro. Plans in cooperation logic: a modular approach. In *Proceedings of the IJCAI Workshop on Nonmontonic Reasoning, Action and Change (NRAC 2007)*, 2007.

[HM92] Joseph Y. Halpern and Yoram Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54(3):319–379, 1992.

[Kur09] Lena Kurzen. Reasoning about cooperation, actions and preferences. *Synthese: Knowledge, Rationality and Action*, 2009. to appear.

[Lan06] Martin Lange. Model checking pdl with all extras. *J. Applied Logic*, 4(1):39–49, 2006.

[Pau02] Marc Pauly. A modal logic for coalitional power in games. *JLC*, 12(1):149–166, 2002.

[WvdHW07] Dirk Walther, Wiebe van der Hoek, and Michael Wooldridge. Alternating-time temporal logic with explicit strategies. In *TARK '07: Proceedings of the 11th conference on Theoretical aspects of rationality and knowledge*, pages 269–278. ACM, 2007.