Logics of Propositional Control

Jelle Gerbrandy
Dipartimento di Informatica
Torino, Italia
jelle@gerbrandy.com

ABSTRACT

The 'Cooperation Logic of Propositional Control', CL-PC, of van der Hoek and Wooldridge is a logic for reasoning about the ability of agents and groups of agents to obtain a certain state of affairs in a situation in which each of the agents controls a number of propositional variables.

We present a number of generalizations of this model, to represent situations in which agents only partially control the value of a variable, or cases in which agents share the control of a variable. We discuss and axiomatize some of these logics of 'partial control.' In addition, we show how this family of logics are closely connected to a body of work in mathematical logic: Cylindric Modal Logic.

Categories and Subject Descriptors

F.4.1 [**Mathematical Logic and Formal Languages**]: Mathematical Logic—*modal logic, model theory*; I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*coherence and coordination, multiagent systems*

General Terms

Theory

Keywords

agent and multi-agent architectures, cooperation and coordination among agents, logics for agent systems

1. INTRODUCTION

Cooperation logics are logics for reasoning about what agents and groups of agents in a multi-agent system can obtain. Formal models for cooperation logics have recently gained some attention; examples are Coalition Logic [10, 9], and ATL [1]. These logics provide ways of talking about what agents, alone or in groups, can obtain with or without the cooperation of the other agents.

A cooperation logic for a specific kind of multi-agent system is the *Coalition Logic of Propositional Control* (CL-PC) from [13]. CL-PC is a logic for reasoning about a situation in which a number of binary control variables is allocated to a group of agents. Given a certain *state* – that is, given that the values of the variables are set in a certain way – agents can then change the variables allocated to them to change the state into a new one.

One of the observations we make in this paper is that the CL-PC model is naturally seen as a *game* (in the game-theoretical sense): the model concerns a set of agents that each can choose between a number of *strategies* (a strategy is for an agent is a valuation for the propositional variables allocated to it), and a choice of strategy for each of the agents determines the final result (a valuation over all variables). Seen as a game, it is quite natural to generalize the kind of model used in CL-PC, which is exactly what we will do.

The model for CL-PC proposed by [13] is about a particular case of propositional control: each variable is controlled by exactly one agent, and all variables can be controlled independently of each other. It is natural to weaken one or both of these assumptions, and consider models in which several agents share the control over the same variables, or cases in which one cannot change a certain variable without changing another as well.

The paper consists of two parts. In section 2, we introduce a general model for reasoning about situations in which propositions get allocated to a number of agents, whose choices with regard to the values of 'their' variables determine their final value. We will call such situations 'games of propositional control.' We show how the logical formalism used in CL-PC can be extended to apply to this larger class of games in a straightforward way. The result of this exercise gives us a modal logic that has been studied in mathematical logic as 'cylindric modal logic' This is a useful observation as it gives us a number of properties that are shared by all logics of propositional control.

In section 3, we discuss a number of particular types of propositional games. One of these is, of course, the model for CL-PC from [13], but we will also consider other types of games, in which agents may *share* the control over certain variables, or games in which the choices of the agents are restricted in certain ways. We will give a number of sound and complete axiomatizations for reasoning for this type of games.

2. GAMES OF PROPOSITIONAL CONTROL

There are different ways of defining a model for representing the control of agents over a given set of propositional variables. One general schema is to define such a situation

AAMAS'06 May 8–12 2006, Hakodate, Hokkaido, Japan.

as a *game*. A specification of a game of propositional control, then, starts with an *allocation* of propositional variables to each of the agents. The agents each choose a valuation for the variables that are allocated to them; the *strategies* of an agent are valuations over the variables that he controls.

Because the valuations chosen by each of the agents may not determine the final values in an unambiguous way (because the choices may conflict), there is an *outcome function* that gives us, for each combination of choices of strategies, an *outcome*, which is a valuation for all propositional variables. Formally, we define a game of propositional control as:

DEFINITION 1. *A* game of propositional control *is a tuple*

$$\mathcal{M} = \langle \mathsf{Ag}, \mathsf{At}, \mathsf{alloc}, \mathsf{strategy}, \mathsf{outcome} \rangle$$

Where:

- $\mathsf{Ag} = \{1, 2, \ldots n\}$ *is a finite, non-empty set of agents*
- At *is a set of propositional variables (or 'atoms').*
- alloc *is an* allocation function *that assigns to each agent from* Ag *a subset of variables from* At.
- *A* strategy *for a player i is a valuation over the propositional variables in* $\mathsf{alloc}(i)$ *– that is, a function that assigns to each propositional variable in* $\mathsf{alloc}(i)$ *a truth value.*
- *The function* $\mathsf{strategy}$ *assigns a set of available strategies to each agent i.*
- *We say that a* state *of the game is a tuple $s = (s_1, \ldots s_n)$ consisting of a strategy s_i from* $\mathsf{strategy}(i)$ *for each agent i.*
- $\mathsf{outcome}$ *is an* outcome function *that assigns to each state s a valuation over all the propositional variables in* At.

As an example, consider a case with two agents, a website and a telephone operator, that together control a database with customer data of some company. Say that a variable credit_ok represents the information that the credit card rating of mr. X is fine, and another variable, webcustomer, represents the fact that mr. X has placed an order over the internet.

The variable credit_ok is controlled by both agents, in the sense that if either the telephone operator or the website get the information that there is something wrong with mr X's credit rating, they can change the value of credit_ok to 0. The credit rating can only be set to 1 if both agents agree. The variable webcustomer is completely under control of the website.

This situation can be represented as a game of propositional control. A game with two players is represented naturally as a matrix, in the following way:

website\ operator	credit_ok : 1	credit_ok : 0
credit_ok : 1 webcustomer : 1	credit_ok : 1 webcustomer : 1	credit_ok : 0 webcustomer : 1
credit_ok : 1 webcustomer : 0	credit_ok : 1 webcustomer : 0	credit_ok : 0 webcustomer : 0
credit_ok : 0 webcustomer : 1	credit_ok : 0 webcustomer : 1	credit_ok : 0 webcustomer : 1
credit_ok : 0 webcustomer : 0	credit_ok : 0 webcustomer : 0	credit_ok : 0 webcustomer : 0

Here, the top row describes the strategy choices of the telephone operator, while the left column describes the options for the website. The outcome of their choices is represented by the values in the cells of the matrix.

Suppose, for example, that the present state of the environment is given by the top-right box in the example above — the telephone operator did not approve the credit card rating of a web customer, but the website does — then the operator is able to change the value of credit_ok to 1 by changing her 'strategy.' However, she does not have this possibility if the present state is the one given by the outcome on the third row on the right: even if the values in the database are the same here, the telephone operator cannot, by herself, change the rating of the customer, because the website disapproves. They can, however, change the credit card rating *together* – by both changing the value of credit_ok to 1.

To reason about these kind of scenarios in a formal way, we use formulas of a simple modal language. Formulas of coalition logic are constructed from a given set of agents Ag and a given set of atomic variables At.

$$\phi ::= \top \mid p \mid \neg\phi \mid \phi \wedge \psi \mid \langle C \rangle \phi$$

Where $p \in \mathsf{At}$ and $C \subseteq \mathsf{Ag}$. The expression $\langle C \rangle \phi$ is a cooperation modality. The idea is that $\langle C \rangle \phi$ is true at some state just in case the agents in C can, by manipulating the values of the variables under their control, reach a state in which ϕ is true. We write $[C]$ for the dual of $\neg\langle C \rangle\neg$ of $\langle C \rangle$. The sentence $[C]\phi$ is true in case the agents in C cannot manipulate their variables in such a way that $\neg\phi$ becomes true: in other words, when the agents in C have no way of changing the truth value of ϕ from true to false.

DEFINITION 2 (INTERPRETATION).

$\mathcal{M}, s \models p$	*iff*	$\mathsf{outcome}(s)(p) = \mathsf{t}$
$\mathcal{M}, s \models \neg\phi$	*iff*	$\mathcal{M}, s \not\models \neg\phi$
$\mathcal{M}, s \models \phi \wedge \psi$	*iff*	$\mathcal{M}, s \models \phi$ *and* $\mathcal{M}, s \models \psi$
$\mathcal{M}, s \models \langle C \rangle \phi$	*iff*	*there is a state t such that* $s(i) = t(i)$*for each $i \notin C$, and* $\mathcal{M}, t \models \phi$

So, a sentence of the form $\langle C \rangle \phi$ is true at a state s when it holds that the agents in C can unilaterally change their choice of strategy in such a way that in the resulting state, ϕ is true. So, in a precise sense, a sentence $\langle C \rangle \phi$ expresses al *local ability* of a group of agents in a state s: it is true in a state s just in case the agents in C are able to change the situation so that ϕ becomes true.

One of the interesting insights in [13] is that we also express 'global' properties of ability of agents that have been studied in Pauly's Coalition Logic ([9], [10]) and in ATL ([1]) using these 'local' operators.

We say that a group C is α-effective for an outcome ϕ, and write $\langle\langle C \rangle\rangle \phi$, when the players in the group can orchestrate their strategies in such a way that no matter what the other players do, ϕ will be true in the outcome.

$$G, s \models \langle\langle C \rangle\rangle \phi \text{ iff } G, s \models \langle C \rangle [\mathsf{Ag} \setminus C] \phi$$

This definition says that a group C is α-effective for ϕ just in case there is a combination of strategies for the agents in C such that, no matter what the agents outside of C (that is, those in the set $\mathsf{Ag} \setminus C$) choose to do, the resulting state will always be one where ϕ is true.

The notion of α-ability is a *global* notion, in the sense that the truth of a sentence $\langle\langle C\rangle\rangle\phi$ is independent of the state in which it is evaluated — it either is true in all states in a game, or in none of them. So, the fact whether $\langle\langle G\rangle\rangle\phi$ is true expresses a property of the whole game rather than of a particular state in it.

To turn back to our example: we saw that the website had the power to change the value of the variables credit_ok at will. Indeed, the sentence $\langle\langle\mathsf{website}\rangle\rangle\mathsf{webcustomer}$ and the sentence $\langle\langle\mathsf{website}\rangle\rangle\neg\mathsf{webcustomer}$ are both true. Neither the website nor the operator is α-effective to approve the credit card rating of mr. X: $\neg\langle\langle\mathsf{website}\rangle\rangle\mathsf{credit\_ok}$ holds, and so does $\neg\langle\langle\mathsf{operator}\rangle\rangle\mathsf{credit\_ok}$, even if, for example, $\langle\mathsf{website}\rangle\mathsf{credit\_ok}$ is true in some states. However, together they are effective for setting credit_ok to 1, which is captured by the truth of $\langle\langle\{\mathsf{website},\mathsf{operator}\}\rangle\rangle\mathsf{credit\_ok}$.

2.1 The logic of games of propositional control

It turns out that the logic of games of propositional control is located at an intersection point between different areas of mathematical logic: it is the logic of finite variable fragments of first order logic without identity, which is closely related to diagonal-free cylindric algebra [15, 14, 8], while its models are products of $S5$-models [3]. For this reason, it is relatively well-studied.

From this work, we know that the logic of propositional games is not finitely axiomatizable. Moreover, the logic is not decidable. There is, however, a sound and complete axiomatization, but it uses a non-standard rule. These are negative results, but the logic of propositional games is still of interest, because it can serve as a kind of 'minimal logic of propositional control': we know that the logic of any particular subclass of models (such as CL-PC) will include this this logic, so we will discuss its main features.

The following set of axioms and rules is shown to be sound and complete for the class of all propositional games from [8] for the language of cylindric algebra, which [15] translates into modal logic. (In the following, if C is a singleton set $\{i\}$, we will write $\langle i\rangle$ instead of $\langle\{i\}\rangle$.)

Definition 3 (logic of propositional games).

propositional logic *all validities of propositional logic*

reduction $\langle\{i_0,\ldots,i_m\}\rangle\phi \leftrightarrow \langle i_0\rangle\ldots\langle i_m\rangle\phi$

asynchronicity $\langle i\rangle\langle j\rangle\phi \rightarrow \langle j\rangle\langle i\rangle\phi$

normality $[i](\phi\rightarrow\psi)\rightarrow([i]\phi\rightarrow[i]\psi)$

reflexivity $\phi\rightarrow\langle i\rangle\phi$

symmetry $\phi\rightarrow[i]\langle i\rangle\phi$

transitivity $\langle i\rangle\langle i\rangle\phi\rightarrow\langle i\rangle\phi$

Plus the following rules:

$$\frac{\vdash\phi}{\vdash[i]\phi}\qquad\frac{\vdash\phi,\vdash\phi\rightarrow\psi}{\vdash\psi}$$

Define $\tau(\chi)$ *to be the sentence* $\neg\langle A\rangle[(\bigwedge_{i\in A}\langle i\rangle\chi)\wedge\neg\chi]$. *The logic contains the following non-standard rule:*

$$\frac{(p\wedge\tau(p\wedge\neg\phi))\rightarrow\phi}{\phi}$$

provided that p *does not occur in* ϕ

When there are only two agents, this rule can be replaced with the following axiom

confluence $\langle i\rangle[j]p\rightarrow[j]\langle i\rangle p$.

The reduction axiom says that the modality $\langle C\rangle$ can be defined in terms of $\langle i\rangle$-operators for $i\in C$ – the ability of a set of agents can be decomposed in the abilities of the agents separately. The asynchronicity axiom says that the order in which the $\langle i\rangle$-operators are applied does not matter.

The reflexivity axiom states that if ϕ is true, then any agent can 'change' the current state (by doing nothing) to make ϕ true. Note that symmetry, reflexivity and transitivity are known as the $S5$ axioms that are familiar from classical modal logic as the axioms that characterize an equivalence relation.

More complicated is the non-standard rule. A full explanation would lead too far for the purposes of this paper. We refer the reader to the works quoted – and in any case, we will not need the rule for axiomatizing the logics we introduce below.

Proposition 4 (completeness [8, 14]). *The logic of definition 3 is sound and complete with respect to the class of games of propositional control defined in definition 1.*

proof: This follows from the completeness results for cylindric modal logic. For some details of how to connect those results to games of propositional control, see the appendix. □

We repeat the following results

Proposition 5 ((un)decidability results [7, 5]).

- *If the number of agents is less than or equal to two, basic game logic is decidable.*
- *However, if there are more than 2 agents, basic game logic is not decidable.*

2.2 The Pareto Principle

The class of models defined in definition 1 is quite large. Many of the models that fall under the definition are hardly worthy of the name 'propositional control', as there need not be any systematic relation at all between the valuations that the agents choose and the actual outcome that results.

To do justice to the idea that the allocation function in a game of propositional control specifies some kind of control, we need some further constraints. There are many possible properties that are candidates for minimal conditions for a game to be called a game of propositional control, but one in particular stands out: the condition that if all agents that are allocated a given variable p choose the same value for p, then the outcome should reflect that choice.

Definition 6 (Pareto Principle).
Let $\mathcal{M}=\langle\mathsf{Ag},\mathsf{At},\mathsf{alloc},\mathsf{strategy},\mathsf{outcome}\rangle$ *be a game of propositional control.*

The Pareto principle holds for $\mathcal{M}$ *iff for each state* $s=(s_0,\ldots,s_n)$ *and for any* i *and* j *such that* $p\in\mathsf{alloc}(i)$ *and* $p\in\mathsf{alloc}(j)$ *it holds that if* $s_i(p)=s_j(p)$, *then* $\mathsf{outcome}(s)=s_i(p)$.

In the rest of this paper, we will only consider games for which the Pareto principle holds.

The Pareto principle is a property that is well known in social choice theory (which is where the name comes from). Social choice theory is concerned with finding a 'fair' way of defining a preference order among a given set of alternatives, given the preferences of a group of agents (typical examples are political elections and auctions). It is perhaps useful to remark here that we are dealing with a situation that in certain respects is more simple than the topic usually dealt with in social choice theory. Here, we are dealing with variables that can each have only one out of two possible truth values. In particular, Arrow's Theorem, a central result that establishes that a 'fair' social choice function does not exist, does not arise as such within our framework. There are, however, interesting results from social choice theory that apply to games of propositional control such as they are defined here (see, for example, [6] or [11]) that show that given certain assumptions about 'fairness', fair outcome functions do not exist. It seems, however, that these kind of concerns are tangential to the kind of problem that we are concerned with in this paper.

3. LOGICS OF PROPOSITIONAL CONTROL

We have talked about games of propositional control in a very general way. It is now time to be more specific. In the rest of the paper, we will study a number of subclasses of games of propositional control, and provide them with sound and complete axiomatizations. Even if the proofs are fairly involved (they can be found in the appendix), the axiomatizations themselves are relatively straightforward, and have axioms that correspond directly with certain properties of the models.

3.1 Logic of Exclusive control

One particular subclass of games of propositional control is one in which no single variable is allocated to two different players. We will call such games *games of exclusive control.* The kind of game that is studied in CL-PC, which was the main inspiration for this paper, is an example of such a game.

DEFINITION 7 (GAMES OF EXCLUSIVE CONTROL). *Let* $\mathcal{M} = (\mathsf{Ag}, \mathsf{At}, \mathsf{alloc}, \mathsf{strategy}, \mathsf{outcome})$ *be a game.*

- *An allocation function* alloc *is* exclusive *when no two agents have control over the same variable, i.e. when* $\mathsf{alloc}(i) \cap \mathsf{alloc}(j) = \emptyset$ *whenever* $i \neq j$.
- *A* game of exclusive control *is a game for in which the Pareto principle holds, and that based on an exclusive allocation function* alloc.
- *A* game of actual control *is a game in which the allocation of a variable p to an agent i implies that i has a choice available where p is false and one where p is true: for each i and* $p \in \mathsf{alloc}(i)$*, a strategy* s_i *such that* $s_i(p) = 1$*, and a strategy* s'_i *such that* $s'_i(p) = 0$.
- *A* game of full control *is a game where* $\mathsf{strategy}(i)$ *contains a strategy for* each *valuation of* $\mathsf{alloc}(i)$.

A game of actual control is a game in which each agent, for each propositional variable p allocated to him, has at least one choice available where p is true and one where p is false. In a game of full control each player has all combinations of truth values of propositional variables available to him — this reflects an assumption that the variables can all be set independently of each other.

Note that the outcome function outcome of an exclusive game is determined completely by the choices of the players and the assumption that the Pareto principle holds – since no variable is allocated to two agents, obviously, their choices cannot conflict. Exclusive games of full propositional control are studied in [4] as 'distributed evaluation games', and are the models of CL-PC that are studied in [13] and [12].

EXAMPLE 8. *Suppose a player i can control two bits of a register, represented by the values of the bits p and q. She can set either bit to 0 or to 1, as she likes, but she cannot set both bits to 0. A second player j controls r. We can represent the game in a matrix:*

$i \setminus j$	r	$\neg r$
p, q	p, q, r	$p, q, \neg r$
$p, \neg q$	$p, \neg q, r$	$p, \neg q, \neg r$
$\neg p, q$	$\neg p, q, r$	$\neg p, q, \neg r$

This is an exclusive game of actual control. The players do not have full control, since i has no strategy for $\neg p, \neg q$.

It turns out that logic of games of exclusive control, and their extensions to real and full exclusive control, are fairly straightforward extensions of the minimal game logic of the previous section.

Before we introduce these axioms an abbreviation will be useful. We follow [13] and say that an agent *locally controls* a sentence ϕ iff it holds (in a state) that $\langle i\rangle\phi \wedge \langle i\rangle\neg\phi$. This sentence is true in a state just in case i can unilaterally change her strategy in such a way that ϕ becomes true, but also has a strategy available where ϕ becomes false. We write $\mathsf{controls}(i, \phi)$ for the sentence $\langle i\rangle\phi \wedge \langle i\rangle\neg\phi$

The property of exclusive control is characterized by the following axiom, that states that if an agent i can change the truth value of a propositional variable p, then no other player can do so.

Exclusive Control (EC) $\mathsf{controls}(i, p) \rightarrow \neg\mathsf{controls}(j, p,)$ for $i \neq j$ and $p \in P$.

In games of actual control, there is, for each propositional variable p, at least one player that controls it:

Actual Control (AC) $\bigvee_{i \in A} \mathsf{controls}(i, p)$ for each $p \in P$.

Similarly, in games of full control, if an agent controls all variables from a given set X, then she has the ability to change the outcome in such a way that any consistent boolean combination of the variables in X is satisfied:

Full Control (FC) $\bigwedge_{p \in X} \mathsf{controls}(i, p) \rightarrow \langle i\rangle\phi_v$, where ϕ_v is the conjunction of literals true in a valuation v for X.

We will write, for example, that $EC, FC \vdash \phi$ just in case ϕ is derivable in the logic of games of propositional control of definition 3 from any instance of the axiom schemas EC and FC.

PROPOSITION 9 (COMPLETENESS RESULTS).

$EC \vdash \phi$ *iff* $G, s \models \phi$ *for all games G of exclusive control.*

$EC, AC \vdash \phi$ *iff* $G, s \models \phi$ *for all games G of exclusive, actual control.*

$EC, FC \vdash \phi$ *iff* $G, s \models \phi$ *for all games G of exclusive, full control.*

proof: See the appendix. It is perhaps important to note at this point that neither for this proposition, nor in any of the following ones, we need the full power of the logic of definition 3: we may drop the non-standard rule, and use the confluence axiom instead. □

3.2 Logic of Positive Control

If certain variables are allocated to more than one agent, we need a way of deciding what to do if one agent wants to set a variable to a different value than another. There are lots of options open. If there is an odd number of agents to which the variable is allocated, we can let the majority decide. In other scenarios we may have a constraint that a variable can only be set to a certain value with the consensus of all other players, such as in example in 2 in which the validity of the credit of a customer needs to be judged. In other situations, a constraint may be that if one of the players chooses a strategy in which p is true, then p is true in the outcome of the game.

Let us first consider games of the latter type:

DEFINITION 10 (GAMES OF POSITIVE CONTROL). *A game $\mathcal{M}$ is a* game of positive control *iff the Pareto principle holds, and it holds that if $s_i(p) = 1$ for some i, then* $\mathsf{outcome}(s)(p) = 1$.

Note that also in games of positive control, the Pareto principle implies that the outcome is completely determined by the strategies of the players.

Games of positive control have a particularly natural interpretation if we interpret the elements of the control variables $\mathcal{P}$ as *actions* that may be performed or not. Agents each have a repertoire of actions available that may overlap (or even be equal), and may choose to execute certain subsets of these actions (if we consider games of full control, they can execute any subset). The outcome such a game, then, consists of exactly the actions that have been executed. This behavior corresponds precisely to the definition of a game of positive control.

For example, if i chooses to execute p but not q, and j executes q, then the result will be that both p and q are executed. Such complex or concurrent actions can be described by a propositional formula; for example, $p \wedge \neg q$ is true of an outcome in which p is executed, but not q. The formula $\langle\langle i, j\rangle\rangle(p \wedge \neg q)$ then describes the game in which the group consisting of i and j has the ability to ensure that p is being executed, and to ensure that q is *not* done. In a game of positive control, it holds that $\langle\langle i, j\rangle\rangle(p \wedge \neg q)$ exactly when p is allocated to either i or j, while q is not allocated to any agent other than i or j.

Also positive control is characterized by a straightforward *axiom of positive control*:

Positive Control (POS) $p \rightarrow \bigvee_{i \in \mathsf{Ag}} [\mathsf{Ag} \setminus \{i\}]p$

That this axiom is sound can be seen as follows. If p is true, then by the Pareto principle, there must be an agent, say i, that has chosen a strategy in which p is true. Then, by the definition of a game of positive control, there is nothing that the other players can do about it, so $[\mathsf{Ag} \setminus \{i\}]p$ must be true.

PROPOSITION 11 (COMPLETENESS). *$POS \vdash \phi$ iff $\mathcal{M}, s \models \phi$ in all games of positive control.*

proof: See the appendix. □

3.3 Logic of Consensus Games

The mirror image of the games of positive control are games in which the truth-value of a propositional variable is determined by *consensus*: a propositional variable is true only if all players choose that proposition to be true.

DEFINITION 12 (CONSENSUS GAMES). *A game is a* consensus game *iff* $\mathsf{outcome}(s)(p) = 1$ *iff* $s_i(p) = 1$ *for each i such that* $p \in \mathsf{alloc}(i)$

A consensus game might also be called a game of 'negative' control, in which a variable is false in the outcome when at least one of the agents wants it to be false.

Consensus Axiom (CON) $\neg p \rightarrow \bigvee_{i \in N} [\mathsf{Ag} \setminus \{i\}]\neg p$.

The idea is similar to the axiom of positive control: if the outcome is $\neg p$, then, because of the Pareto principle, at least one of the players must have chosen a strategy in which p is false, which implies, by definition of the consensus game, that no matter what the other players will do, p will be false in the outcome.

EXAMPLE 13. *A simple safety mechanism is that two players must both give their consent for some action to be carried out (a customer and a seller confirm some cash transaction; two people lift a heavy table togethere). One very simple way is to describe this by the following game, where each player can either confirm p or not, and the result will be p only if both players give their consent:*

$a \setminus b$	p	$\neg p$
p	p	$\neg p$
$\neg p$	$\neg p$	$\neg p$

The minimal game logic together with the axiom of consensus is sound and complete with respect to all consensus games.

PROPOSITION 14. *$CON \vdash \phi$ iff $\mathcal{M}, s \models \phi$ in all consensus games $\mathcal{M}$*

proof: Analogous to the proof of proposition 11, which can be found in the appendix. □

4. CONCLUSIONS

We have defined a general model of 'propositional control' inspired by the Coalition Logic of Propositional Control of [13], and connected this work with work on mathematical logic (in particular with Cylindric Modal Logic) and with work on 'judgment aggregation' in social choice theory.

We then studied in more detail some particular examples of families of games of propositional control – in particular, we studied different forms of games of exclusive control – in which agents either have some or all strategies available, and provided sound and complete axiomatization for these logics. We also looked at games of 'positive control', in which a vote of one player suffices to make a proposition true, and at 'consensus games' in which a vote of a single player suffices to make a propositional variable false, and provides sound and complete axiomatization for these logics.

A number of open problems and questions are raised by this paper. We mention a few. First of all, there is a general question. Cylindrical Modal Logic, which we proposed as a kind of 'minimal logic of propositional control' is not decidable, and has a 'strange' axiomatization with a non-standard rule. In contrast, the particular examples of logics we studied in section 3 are much better behaved – they are clearly decidable, and do not need special kinds of axiomatization. It is not really clear why this is the case, or where the 'borderline' is between decidability and undecidability. The answer to this question might lie with the class of games for which the Pareto principle holds true.

A more specific question regards the complexity of the logics discussed in this paper; we only know that the satisfiability and the model checking problems of the logics for games of full exclusive control are PSPACE-complete ([13]). (Satisfiabily and model-checking have the same complexity because of the strong restrictions on the kind of models that are allowed, which means that there are relatively few models to consider in the satisfaction problem. We do not expect the same results for our logics.) At last, there may be other types of logics of propositional control over and above the ones discussed in this paper that might be interesting to study.

To conclude, we list the main completeness results of this paper in a table:

class of games	characterizing axioms
exclusive	EC
exclusive, actual control	EC, AC
exclusive, full control	EC, AC, FC
games of positive control	POS
consensus games	CON

5. REFERENCES

[1] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Lecture Notes in Computer Science*, 1536:23–60, 1998.

[2] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Theoretical Tracts in Computer Science. Cambridge University Press, 2001.

[3] D. Gabbay and V. B. Shehtman. Products of modal logics, part i. *Logic Journal of the IGPL*, 6(1):73–146, 1998.

[4] P. Harrenstein. *Logic in Conflict. Logical Explorations in Strategic Equilibrium*. PhD thesis, Utrecht University, 2004.

[5] R. Hirsch, I. Hodkinson, and A. Kurucz. On modal logics between KxKxK and S5xS5xS5. *Journal of Symbolic Logic*, 67:221–234, 2002.

[6] C. List and P. Pettit. Aggregating sets of judgements: An impossibility result. *Economics and Philosophy*, 18:89–110, 2002.

[7] R. Maddux. The equational theory of ca_3 is undecidable. *Journal of Symbolic Logic*, 45(2):311–316, 1980.

[8] S. Mikulas. Gabbay-style calculi. In H. Wansing, editor, *Proof Theory of Modal Logic*, volume 2 of *Applied Logic Series*, pages 243–252. Kluwer Academic Publishers, 1996.

[9] M. Pauly. *Logic for Social Software*. PhD thesis, University of Amsterdam, 2001. ILLC Dissertation Series 2001-10.

[10] M. Pauly. A modal logic for coalitional power in games. *Journal of Logic and Computation*, 12(1):149–166, 2002.

[11] M. Pauly and M. van Hees. Logical constraints on judgment aggregation. *Journal of Philosophical Logic*, 2003. forthcoming.

[12] W. van der Hoek and M. Wooldridge. On the dynamics of delegation, cooperation, and control: a logical account. In *AAMAS '05: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 701–708, New York, NY, USA, 2005. ACM Press.

[13] W. van der Hoek and M. Wooldridge. On the logic of co-operation and propositional control. *Artificial Intelligence*, 64(1–2):81–119, 2005.

[14] Y. Venema. Cylindric modal logic. *Journal of Symbolic Logic*, 60:591–623, 1995.

[15] Y. Venema. Rectangular games. *Journal of Symbolic Logic*, 63(4):1549–1564, 1998.

Appendix

In this appendix, we provide some formal background about the relation between games of propositional control and standard models of modal logic, and proofs of the main theorems.

Games of propositional control can be seen as a kind of Kripke models. The following definition makes the connection precise.

DEFINITION 15 (KRIPKE MODEL INDUCED BY A GAME). *Given a game* $\mathcal{M} = \langle \mathsf{Ag}, \mathsf{At}, \mathsf{alloc}, \mathsf{strategy}, \mathsf{outcome}\rangle$ *The* Kripke model induced by M *is a tuple* $(W, (\longrightarrow_i)_{i\in\mathsf{Ag}}, V)$, *where:*

- W *is the set of states of* $\mathcal{M}$ *(that is, the Cartesian product of the sets* $\mathsf{strategy}(i)$*)*
- $s \longrightarrow_i s'$ *iff* $s_j = s'_j$ *for each* $j \neq i$, *and*
- $V(s) = \mathsf{outcome}(s)$.

In the following proofs, we will often confuse a game and the Kripke model induced by it, and write, for example, that $s \longrightarrow_i s'$ in a game $\mathcal{M}$ when s and s' are two states that differ at most in the strategy of i.

This connection between propositional games and Kripke models preserves the truth of sentences.

OBSERVATION 16. *If* s *is a state of a propositional game* $\mathcal{M}$, *and* $\mathcal{K}$ *is the Kripke model induced by* $\mathcal{M}$, *then* $\mathcal{M}, s \models \phi$ *iff* $\mathcal{K}, s \models \phi$ *(using the standard definition of modal logic).*

With these observations, we have a way of transforming games into Kripke models that preserves the truth value of sentences. This gives us a way of applying known results about Kripke models to games of propositional control. In particular, we can tie the logic of games to the following definition and proposition come from [15], [8].

DEFINITION 17 (CUBES). *A* cube *is a Kripke model* $(U^{|\mathsf{Ag}|}, (\longrightarrow_i)_{i\in\mathsf{Ag}}, V)$, *with* $s \longrightarrow_i s'$ *iff* $s_j = s'_j$ *for each* $j \neq i$.

PROPOSITION 18 ([15], [8]). *The set of cubes is axiomatized by the minimal logic of propositional control of definition 3.*

To show that the logic of cubes is the same as the logic of games of propositional control, we need to show that each cube corresponds to a game, and each game corresponds to a cube. The Kripke models induced by a game of propositional control are not precisely cubes, however: they are more like 'rectangles' (if the players number two) or like '(hyper)-bricks' (if there are more than 2 players). 'Cubes' are a subclass of these kind of models; those induced by games in which all players have the same number of strategies available. We will not provide the formal details here, but it is fairly straightforward to transform a Kripke model that is shaped like a 'brick' into an equivalent model that is shaped like proper 'cube' (at least, if the number of agents is finite). Perhaps abusing the geometrical metaphor, one possible way to make a cube from a brick is to add, where needed, a new layer to the 'long' sides of the brick until the model has become a cube – to not affect the validities in the model, this new layer can simply be a copy of the side to which it is added.

In the completeness proofs that follow, we use the notion of a *bisimulation*.

Definition 19 (bisimulation). *A pointed Kripke model is a pair (K, s), where K is a Kripke model, and s is a state in K.*

A relation $\sim$ between pointed Kripke models is a bisimulation *iff whenever $(K, s) \sim (K', s')$, it holds that:*

I. $V(s) = V'(s')$

II. If $s \longrightarrow_i t$, then there is a t' with $s' \longrightarrow'_i t'$ and $(K, t) \sim (K', t')$

III. If $s' \longrightarrow'_i t'$, then there is a t with $s \longrightarrow_i t$ and $(K, t) \sim (K', t')$

We say that two pointed Kripke models are bisimilar iff there exists a bisimulation connecting the two models.

A basic result says that two bisimilar Kripke models satisfy the same sentences. If it does not lead to confusion, we will write $s \sim s'$ for $(K, s) \sim (K', s')$. Also, we will be slightly sloppy, and talk about bisimilarity of a Kripke model and a game, meaning that the first Kripke model is bisimilar to the model induced by the game.

Proposition 9. *$EC \vdash \phi$ iff $\mathcal{M}, s \models \phi$ in all games $\mathcal{M}$ of exclusive control.*

proof: Soundness is just a matter of checking that the axioms are valid in a game of exclusive control.

A sketch of the completeness proof before starting: first we construct a canonical Kripke model in the standard way ([2]), then we show that each world in the canonical model is bisimilar to a state in (the Kripke model induced by) a game. This means that all sentences in the maximal consistent set in the canonical model are true in the related game, and thus that any consistent set can be satisfied in some game.

The canonical model $K^{can} = (W, \longrightarrow_i, V)$ has as its domain all maximal consistent sets (mcs) Σ. We set $V(\Sigma)(p) = 1$ iff $p \in \Sigma$, and the accessibility relations are given by $\Sigma \longrightarrow_i \Gamma$ iff for all ϕ: if $[i]\phi \in \Sigma$ then $\phi \in \Gamma$. It is straightforward to prove that the *truth lemma* holds: $K^{\mathsf{can}}, \Sigma \models \phi$ iff $\phi \in \Sigma$.

The interesting part of the proof is the construction of a game G_Σ corresponding to each mcs Σ. First, we define an allocation function alloc_Σ by setting $\mathsf{alloc}_\Sigma(i) = \{p \mid \mathsf{controls}(i, p) \in \Sigma\}$. Because Σ includes the axiom of exclusive control, we know that alloc_Σ is exclusive.

We introduce some notation: if v is a valuation over a finite set X, let ϕ_v be the conjunction of all literals of X true in v. The strategies $\mathsf{strategy}_\Sigma(i)$ of a player i in G_Σ are exactly those valuation functions v on the set $\mathsf{alloc}_\Sigma(i)$ such that $\langle i\rangle\phi_v \in \Sigma$. In the game G_Σ, let s_Σ be the state such that $\phi_{s_\Sigma(i)} \in \Sigma$ for each i. Such a state is guaranteed to exist by the reflexivity axiom.

We claim that there is a bisimulation $\sim$ between maximal consistent sets in the canonical model K^{can} and states in the set of all (Kripke models induced by) exclusive games, given by:

$$\Sigma \sim (G, s) \text{ iff } (G, s) = (G_\Sigma, s_\Sigma).$$

To make our proof run a bit more smoothly, we will write that $(G, s) \longrightarrow_i (G', s')$ iff $G = G'$ and $s \longrightarrow_i s'$ in G. We check the three bisimulation clauses:

I. From the definition of G_Σ and the assumption that the Pareto principle holds for G_Σ, it follows that $\Sigma \models p$ iff $(G_\Sigma, s_\Sigma) \models p$.

II. Suppose that $\Sigma \longrightarrow_i \Gamma$. We claim that $(G_\Sigma, s_\Sigma) \longrightarrow_i (G_\Gamma, s_\Gamma)$. It follows from lemma 21 that $G_\Sigma = G_\Gamma$, and that for each $p \notin \mathsf{alloc}(i)$, $p \in \Sigma$ iff $p \in \Gamma$. That means that $s_\Sigma(j) = s_\Gamma(j)$ for each $j \neq i$, and that therefore $s_\Sigma \longrightarrow_i s_\Gamma$.

III. For the other direction, suppose that $(G_\Sigma, s_\Sigma) \longrightarrow_i (G_\Sigma, s')$. Let $s = s_\Sigma$. By definition of G_Σ, we have that $\langle i\rangle\phi_{s'_i} \in \Sigma$. That means that there is a Γ in the canonical modal such that $\Sigma \longrightarrow_i \Gamma$ and $\phi_{s'_i} \in \Gamma$. By exclusivity and lemma 20, it follows that $[i]\phi_{s_j} \vee [i]\neg\phi_{s_j} \in \Gamma$ for each $j \neq i$. Since by definition, $\phi_{s_j} \in \Sigma$, it follows that $\phi_{s_j} \in \Gamma$, and therefore that $s' = s_\Gamma$.

The proofs of completeness for the logics with the axioms AC and FC is straightforward given these results. □

Lemma 20 (derived theorems).

1. $\vdash \langle i\rangle[j]\phi \rightarrow [j]\langle i\rangle\phi$ *(confluence)*
2. $EC \vdash \mathsf{controls}(i, p) \rightarrow [j]\mathsf{controls}(i, p)$
   Moreover, if ϕ is propositional and X is the set of propositional variables occurring in ϕ, it holds that:
3. $\bigwedge_{p\in X} \neg\mathsf{controls}(i, p) \vdash [i]\phi \vee [i]\neg\phi$
4. $EC, \bigwedge_{p\in X} \mathsf{controls}(j, p), \langle j\rangle\phi \vdash [i]\langle j\rangle\phi$

proof:

2. In $S5$, it follows from $\mathsf{controls}(i, p)$ that $[i]\mathsf{controls}(i, p)$. Using the axiom EC, we have then that $[i]([j]p \vee [j]\neg p)$, which is equivalent (using reflexivity and propositional logic) to $[i]((p \rightarrow [j]p) \wedge (\neg p \rightarrow [j]\neg p))$. Since we have that $\langle i\rangle p \wedge \langle i\rangle\neg p$, it follows in $S5$ that $\langle i\rangle[j]p \wedge \langle i\rangle[j]\neg p$. Using confluence, we conclude that $[j]\langle i\rangle p \wedge [j]\langle i\rangle\neg p$, and hence that $[j]\mathsf{controls}(i, p)$.

3. Prove this by induction on the structure of ϕ.

4. Suppose that $\bigwedge_{p\in X} \mathsf{controls}(j, p)$. With (3) and EC, we have that $\vdash \bigwedge_{p\in X} \mathsf{controls}(j, p) \rightarrow ([i]\phi \vee [i]\neg\phi)$, and with necessitation and some further reasoning, we conclude that $[j]([i]\phi \vee [i]\neg\phi)$. Since $[i]\phi \vee [i]\neg\phi$ is equivalent in $S5$ to $(\phi \rightarrow [i]\phi) \wedge (\neg\phi \rightarrow [i]\neg\phi)$, we have in particular that $[j](\phi \rightarrow [i]\phi)$. Together with our assumption that $\langle j\rangle\phi$, it follows that $\langle j\rangle[i]\phi$. Using confluence, then, $[i]\langle j\rangle\phi$. □

Lemma 21. *Let $\longrightarrow_i$ be the accessibility relation in the canonical model, and let alloc_Σ be defined as above. It holds that:*

1. *If $\Sigma \longrightarrow_i \Gamma$ then $\mathsf{alloc}_\Sigma = \mathsf{alloc}_\Gamma$*

2. *If* $\Sigma \longrightarrow_i \Gamma$ *and* $p \notin \mathsf{alloc}_\Sigma(i)$*, then* $p \in \Sigma$ *iff* $p \in \Gamma$

3. *If* $\Sigma \longrightarrow_i \Gamma$ *and* s_j *is a valuation on* $\mathsf{alloc}_\Sigma(j)$*, then* $\langle j \rangle \phi_{s_j} \in \Sigma$ *implies that* $\langle j \rangle \phi_{s_j} \in \Gamma$ *for each* j*.*

proof:

1. Suppose $\Sigma \longrightarrow_i \Gamma$, and that $\mathsf{controls}(j,p) \in \Sigma$. Then, by lemma 20, $[i]\mathsf{controls}(j,p) \in \Sigma$, so $\mathsf{controls}(j,p) \in \Gamma$. Using symmetry, we can conclude that also if $\mathsf{controls}(j,p) \in \Gamma$, then $\mathsf{controls}(j,p)) \in \Sigma$. So $\mathsf{alloc}_\Sigma = \mathsf{alloc}_\Gamma$.

2. From exclusivity EC it follows that if $p \notin \mathsf{alloc}_\Sigma(i)$, then $[i]p \vee [i]\neg p \in \Sigma$. From there it is an easy step to conclude that $p \in \Sigma$ implies $p \in \Gamma$.

3. Clearly, it holds $\mathsf{controls}(j,p)$ for all propositional variables occurring in ϕ_{s_j}. Suppose $\langle j \rangle \phi_{s_j} \in \Sigma$. It follows from lemma 20 that $[i]\langle j \rangle \phi_{s_j} \in \Sigma$, and therefore that $\langle j \rangle \phi_{s_j} \in \Gamma$. □

Proposition 11. *$POS \vdash \phi$ iff $\mathcal{M}, s \models \phi$ in all games of positive control $\mathcal{M}$*

proof: The structure of the completeness proof is similar to that of proposition 9. We construct the canonical model K^{can}, and construct games G_Σ for each maximal consistent set Σ, and show that the model (G_Σ, s_Σ) is bisimilar to (K^{can}, Σ).

First, we need some notation. If v is a valuation over a finite set X, ϕ_v^+ is the conjunction of all propositional variables of X that are true in v, i.e. $\phi_v^+ = \bigwedge\{p \mid v(p) = 1\}$. We will write $[-i]\phi$ for $[\mathsf{Ag} \setminus \{i\}]\phi$.

To construct our game G_Σ, we set $\mathsf{alloc}(i)$ to be the set of all variables for any i. So the strategies $\mathsf{strategy}_\Sigma(i)$ for agent i consist of valuations over *all* variables, and is given by $\mathsf{strategy}_\Sigma(i) = \{v \mid \langle i \rangle [-i] \bigwedge\{p \mid v(p) = 1\} \in \Sigma\}$. So, $\mathsf{strategy}_\Sigma(i)$ contains a strategy for each set of variables that i can *guarantee* to be true.

We now claim that the following relation $\sim$ is a bisimulation:

$$(K^{can}, \Sigma) \sim (G, s) \text{ iff } G = G_\Sigma \text{ and for each } i\text{:} \quad s_i(p) = 1 \Leftrightarrow [-i]p \in \Sigma.$$

We need to show that the three bisimulation clauses hold.

I. Suppose $G_\Sigma, s \models p$. Then there must be an agent i with $s_i(p) = 1$. By definition of $\sim$, it holds that $[-i]p \in \Sigma$. By necessitation, $p \in \Sigma$, and by construction of the canonical model that $K^{can}, \Sigma \models p$.

For the other direction, if $p \in \Sigma$, then, with the axiom of positive control, $\bigvee_{i \in N}[-i]p \in \Sigma$. Because Σ is maximal, we have that $[-i]p \in \Sigma$ for some i. But then, by definition of $\sim$, $s_i(p) = 1$, and so, $G_\Sigma, s \models p$.

II. Suppose that $\Sigma \longrightarrow_i \Gamma$ and that $\Sigma \sim G_\Sigma, s$. By standard $S5$-reasoning, if $\langle i \rangle \phi \in \Sigma$, then so is $[i]\langle i \rangle \phi$, which means that $\langle i \rangle \phi \in \Gamma$. By symmetry, the other direction works as well. Since the definition of the game G_Σ depends only on sentences from Σ that are of the form $\langle i \rangle \phi$, so we can conclude that $G_\Sigma = G_\Gamma$.

Consider now the state s' for which it holds for each j that $s'_j(p) = 1$ iff $[-j]p \in \Gamma$. We show that $s \longrightarrow_i s'$. To see that $s_j = s'_j$ for $j \neq i$, note that if $[-j]p \in \Sigma$ then, by $S5$ and asynchronicity, $[i][-j]p \in \Sigma$, and therefore that $[-j]p \in \Gamma$; with symmetry, it follows that $s_j = s'_j$. We also need to show that s'_i is a strategy in G_Σ. This follows from the observation that $[-i] \bigwedge\{p \mid s'_i(p) = 1\} \in \Gamma$, and therefore that $\langle i \rangle [-i] \bigwedge\{p \mid s'_i(p) = 1\} \in \Sigma$.

III. Suppose that $\Sigma \sim (G_\Sigma, s)$, and that $s \longrightarrow_i s'$. We need to find a Γ such that $\Sigma \longrightarrow_i \Gamma$ and $\Gamma \sim (G_\Sigma, s')$. It is enough to show that the set $\{[-j]p \mid s'_j(p) = 1 \text{ and } j \in N\} \cup \{\phi \mid [i]\phi \in \Sigma\}$ is consistent; we can then extend it to a maximal consistent set to obtain our wanted Γ. Because $s_j = s'_j$ for $j \neq i$, it must hold that $\{[-j]p \mid s'_j(p) = 1 \text{ and } j \neq i\} \subseteq \Sigma$, and by $S5$ and asynchronicity, we have that $[i][-j]p \mid s'_j(p) = 1 \text{ and } j \neq i\} \subseteq \Sigma$. So, it is enough to show that $\{[-i]p \mid s'_i(p) = 1\} \cup \{\phi \mid [i]\phi \in \Sigma\}$ is consistent. Suppose it is not. By definition of G_Σ, we have that $\langle i \rangle [-i] \bigwedge\{p \mid s'_i(p) = 1\} \in \Sigma$. But then it would follow that $\langle i \rangle \bot \in \Sigma$, which is a contradiction in $S5$.

□

Proposition 14. *$CON \vdash \phi$ iff $\mathcal{M}, s \models \phi$ in all consensus games $\mathcal{M}$*

proof: Analogous to the completeness proof of the logic of positive control. Define the set of strategies for i in the game G_Σ as $\mathsf{strategy}_\Sigma(i) = \{v \mid \langle i \rangle [-i] \bigwedge\{p \mid v(p) = 0\} \in \Sigma\}$, and then adapt the proof accordingly. □