

Model Checking for Coalition Announcement Logic

Rustam Galimullin¹, Natasha Alechina¹, and Hans van Ditmarsch²

¹ University of Nottingham, Nottingham, UK

`{rustam.galimullin,natasha.alechina}@nottingham.ac.uk`

² CNRS, LORIA, Univ. of Lorraine, France & ReLaX, Chennai, India

`hans.van-ditmarsch@loria.fr`

Abstract. Coalition Announcement Logic (CAL) studies how a group of agents can enforce a certain outcome by making a joint announcement, regardless of any announcements made simultaneously by the opponents. The logic is useful to model imperfect information games with simultaneous moves. We propose a model checking algorithm for CAL and show that the model checking problem for CAL is PSPACE-complete. We also consider a special positive case for which the model checking problem is in P. We compare these results to those for other logics with quantification over information change.

Keywords: Model Checking · Coalition Announcement Logic · Dynamic Epistemic Logic.

1 Introduction

In the multi-agent logic of knowledge we investigate what agents know about their factual environment and what they know about knowledge of each other [13]. (Truthful) Public announcement logic (PAL) is an extension of the multi-agent logic of knowledge with modalities for public announcements. Such modalities model the event of incorporating trusted information that is similarly observed by all agents [16]. The ‘truthful’ part relates to the trusted aspect of the information: we assume that the novel information is true.

In [2] the authors propose two generalisations of public announcement logic, GAL (group announcement logic) and CAL (coalition announcement logic). These logics allow for quantification over public announcements made by agents modelled in the system. In particular, the GAL quantifier $\langle G \rangle \varphi$ (parametrised by a subset G of the set of all agents A) says ‘there is a truthful announcement made by the agents in G , after which φ (holds)’. Here, the truthful aspect means that the agents in G only announce what they *know*: if a in G announces φ_a , this is interpreted as a public announcement $K_a \varphi_a$ such that a truthful announcement by agents in G is a conjunction of such known announcements. The CAL quantifier $\langle\!\langle G \rangle\!\rangle \varphi$ is motivated by game logic [15, 14] and van Benthem’s playability operator [7]. Here, the modality means ‘there is a truthful announcement

made by the agents in G such that no matter what the agents not in G simultaneously announce, φ holds afterwards'. In [2] it is, for example, shown that this subsumes game logic.

CAL has been far less investigated than other logics of quantified announcements – APAL [5] and GAL – although some combined results have been achieved [4]. In particular, model checking for CAL has not been studied. Model checking for CAL has potential practical implications. In CAL, it is possible to express that a group of agents (for example, a subset of bidders in an auction) can make an announcement such that no matter what other agents announce simultaneously, after this announcement certain knowledge is increased (all agents know that G have won the bid) but certain ignorance also remains (for example, the maximal amount of money G could have offered). Our model-checking algorithm may be easily modified to return not just ‘true’ but the actual announcement that G can make to achieve their objective. The algorithm and the proof of PSPACE-completeness build on those for GAL [1], but the CAL algorithm requires some non-trivial modifications. We show that for the general case, model checking CAL is in PSPACE, and also describe an efficient (PTIME) special case.

2 Background

2.1 Introductory Example

Two agents, a and b , want to buy the same item, and whoever offers the greatest sum, gets it. Agents may have 5, 10, or 15 pounds, and they do not know which sum the opponent has. Let agent a have 15 pounds, and agent b have 5 pounds. This situation is presented in Figure 1.

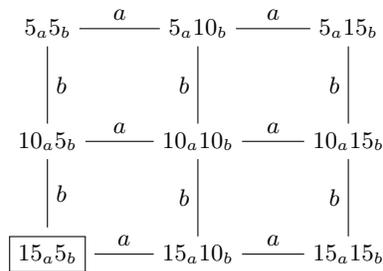


Fig. 1. Initial model $(M, 15_a 5_b)$

In the model (let us call it M), state names denote money distribution. Thus, $10_a 5_b$ means that agent a has 10 pounds, and agent b has 5 pounds. Labelled edges connect the states that a corresponding agent cannot distinguish. For example, in the actual state (boxed), agent a knows that she has 15 pounds, but

she does not know how much money agent b has. Formally, $(M, 15_a 5_b) \models K_a 15_a \wedge \neg(K_a 5_b \vee K_a 10_b \vee K_a 15_b)$. Note that edges represent equivalence relations, and in the figure we omit transitive and reflexive transitions.

Next, suppose that agents bid in order to buy the item. Once one of the agents, let us say a , announces her bid, she also wants the other agent to remain ignorant of the total sum at her disposal. Formally, we can express this goal as formula $\varphi ::= K_b(10_a \vee 15_a) \wedge \neg(K_b 10_a \vee K_b 15_a)$ (for bid 10 by agent a). Informally, if a commits to pay 10 pounds, agent b knows that a has 10 or more pounds, but she does not know the exact amount. If agent b does not participate in announcing (bidding), a can achieve the target formula φ by announcing $K_a 10_a \vee K_a 15_a$. In other words, agent a commits to pay 10 pounds, which denotes that she has at least that sum at her disposal. In general, this means that there is an announcement by a such that after this announcement φ holds. Formally, $(M, 15_a 5_b) \models \langle a \rangle \varphi$. The updated model $(M, 15_a 5_b)^{K_a 10_a \vee K_a 15_a}$, which is, essentially, a restriction of the original model to the states where $K_a 10_a \vee K_a 15_a$ holds, is presented in Figure 2.

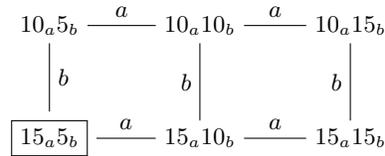


Fig. 2. Updated model $(M, 15_a 5_b)^{K_a 10_a \vee K_a 15_a}$

Indeed, in the updated model agent b knows that a has at least 10 pounds, but not the exact sum. The same holds if agent b announces her bid simultaneously with a in the initial situation. Moreover, a can achieve φ no matter what agent b announces, since b can only truthfully announce $K_b 5_b$, i.e. that she has only 5 pounds at her disposal. Formally, $(M, 15_a 5_b) \models \langle a \rangle \varphi$.

2.2 Syntax and Semantics of CAL

Let A denote a finite set of agents, and P denote a countable set of propositional variables.

Definition 1. *The language of coalition announcement logic \mathcal{L}_{CAL} is defined by the following BNF:*

$$\varphi, \psi ::= p \mid \neg\varphi \mid (\varphi \wedge \psi) \mid K_a \varphi \mid [\psi]\varphi \mid \llbracket G \rrbracket \varphi,$$

where $p \in P$, $a \in A$, $G \subseteq A$, and all the usual abbreviations of propositional logic and conventions for deleting parentheses hold. The dual operators are defined as follows: $\tilde{K}_a \varphi ::= \neg K_a \neg \varphi$, $\langle \psi \rangle \varphi ::= \neg [\psi] \neg \varphi$, and $\langle \llbracket G \rrbracket \rangle \varphi ::= \neg \llbracket G \rrbracket \neg \varphi$. Language \mathcal{L}_{PAL} is the language without the operator $\llbracket G \rrbracket \varphi$, and \mathcal{L}_{EL} is the pure epistemic language without the operators $[\psi]\varphi$ and $\llbracket G \rrbracket \varphi$.

Formulas of CAL are interpreted in epistemic models.

Definition 2. An epistemic model is a triple $M = (W, \sim, V)$, where W is a non-empty set of states, $\sim: A \rightarrow \mathcal{P}(W \times W)$ assigns an equivalence relation to each agent, and $V: P \rightarrow \mathcal{P}(W)$ assigns a set of states to each propositional variable. M is called finite if W is finite. A pair (M, w) with $w \in W$ is called a pointed model. Also, we write $M_1 \subseteq M_2$ if $W_1 \subseteq W_2$, \sim_1 and V_1 are results of restricting \sim_2 and V_2 to W_1 , and call M_1 a submodel of M_2 .

Definition 3. For a pointed model (M, w) and $\varphi \in \mathcal{L}_{EL}$, an updated model $(M, w)^\varphi$ is a restriction of the original model to the states where φ holds and to corresponding relations. Let $\llbracket \varphi \rrbracket_M = \{w : (M, w) \models \varphi\}$ where \models is defined below. Then $W^\varphi = \llbracket \varphi \rrbracket_M$, $\sim_a^\varphi = \sim_a \cap (\llbracket \varphi \rrbracket_M \times \llbracket \varphi \rrbracket_M)$ for all $a \in A$, and $V^\varphi(p) = V(p) \cap \llbracket \varphi \rrbracket_M$. A model which results in subsequent updates of (M, w) with formulas $\varphi_1, \dots, \varphi_n$ is denoted $(M, w)^{\varphi_1, \dots, \varphi_n}$.

Let \mathcal{L}_{EL}^G denote the set of formulas of the form $\bigwedge_{a \in G} K_a \varphi_a$, where for every $a \in G$ it holds that $\varphi_a \in \mathcal{L}_{EL}$. In other words, formulas of \mathcal{L}_{EL}^G are of the type ‘for all agents a from group\coalition G , a knows a corresponding φ_a .’

Definition 4. Let a pointed model (M, w) with $M = (W, \sim, V)$, $a \in A$, and formulas φ and ψ be given.³

$$\begin{aligned} (M, w) \models p & \quad \text{iff } w \in V(p) \\ (M, w) \models \neg \varphi & \quad \text{iff } (M, w) \not\models \varphi \\ (M, w) \models \varphi \wedge \psi & \quad \text{iff } (M, w) \models \varphi \text{ and } (M, w) \models \psi \\ (M, w) \models K_a \varphi & \quad \text{iff } \forall v \in W : w \sim_a v \text{ implies } (M, v) \models \varphi \\ (M, w) \models [\varphi] \psi & \quad \text{iff } (M, w) \models \varphi \text{ implies } (M, w)^\varphi \models \psi \\ (M, w) \models \llbracket G \rrbracket \varphi & \quad \text{iff } \forall \psi \in \mathcal{L}_{EL}^G \exists \chi \in \mathcal{L}_{EL}^{A \setminus G} : (M, w) \models \psi \rightarrow \langle \psi \wedge \chi \rangle \varphi \end{aligned}$$

The operator for coalition announcements $\llbracket G \rrbracket \varphi$ is read as ‘whatever agents from G announce, there is a simultaneous announcement by agents from $A \setminus G$ such that φ holds.’

An alternative equivalent semantics for $\llbracket G \rrbracket \varphi$ is conceivable in terms of model restrictions:

$$(M, w) \models \llbracket G \rrbracket \varphi \quad \text{iff } \forall \psi \in \mathcal{L}_{EL}^G \exists \chi \in \mathcal{L}_{EL}^{A \setminus G} : \\ (M, w) \models \psi \text{ implies } ((M, w) \models \chi \text{ and } (M, w)^{\psi \wedge \chi} \models \varphi)$$

The semantics for the ‘diamond’ version of coalition announcement operators are as follows:

$$(M, w) \models \langle \llbracket G \rrbracket \rangle \varphi \quad \text{iff } \exists \psi \in \mathcal{L}_{EL}^G \forall \chi \in \mathcal{L}_{EL}^{A \setminus G} : (M, w) \models \psi \wedge [\psi \wedge \chi] \varphi$$

³ For comparison, semantics for group announcement operator of the logic GAL mentioned in the introduction is $(M, w) \models [G] \varphi$ iff $\forall \psi \in \mathcal{L}_{EL}^G : (M, w) \models [\psi] \varphi$ and $(M, w) \models \langle G \rangle \varphi$ iff $\exists \psi \in \mathcal{L}_{EL}^G : (M, w) \models \langle \psi \rangle \varphi$.

Definition 5. We call formula φ valid if and only if for any pointed model (M, w) it holds that $(M, w) \models \varphi$. And φ is called satisfiable if and only if there is some (M, w) such that $(M, w) \models \varphi$.

Note that following [5, 1] we restrict formulas that agents in a group or coalition can announce to formulas of \mathcal{L}_{EL} . This allows us to avoid circularity in Definition 4.

2.3 Bisimulation

The basic notion of similarity in modal logic is bisimulation [8, Section 3].

Definition 6. Let two models $M = (W, \sim, V)$ and $M' = (W', \sim', V')$ be given. A non-empty binary relation $Z \subseteq W \times W'$ is called a bisimulation if and only if for all $w \in W$ and $w' \in W'$ with $(w, w') \in Z$:

- w and w' satisfy the same propositional variables;
- for all $a \in A$ and all $v \in W$: if $w \sim_a v$, then there is a v' such that $w' \sim'_a v'$ and $(v, v') \in Z$;
- for all $a \in A$ and all $v' \in W'$: if $w' \sim'_a v'$, then there is a v such that $w \sim_a v$ and $(v, v') \in Z$.

If there is a bisimulation between models M and M' linking states w and w' , we say that (M, w) and (M', w') are bisimilar.

Note that any union of bisimulations between two models is a bisimulation, and the union of all bisimulations is a maximal bisimulation.

Definition 7. Let model M be given. The quotient model of M with respect to some relation R is $M^R = (W^R, \sim^R, V^R)$, where $W^R = \{[w] \mid w \in W\}$ and $[w] = \{v \mid wRv\}$, $[w] \sim_a^R [v]$ iff $\exists w' \in [w], \exists v' \in [v]$ such that $w' \sim_a v'$ in M , and $V^R(p) = V(p)$.

Definition 8. Let model M be given. Bisimulation contraction of M (written $\|M\|$) is the quotient model of M with respect to the maximal bisimulation of M with itself. Such a maximal bisimulation is an equivalence relation.

Informally, bisimulation contraction is the minimal representation of M .

Definition 9. A model M is bisimulation contracted if M is isomorphic to $\|M\|$.

Proposition 1. $(\|M\|, w) \models \varphi$ iff $(M, w) \models \varphi$ for all $\varphi \in \mathcal{L}_{CAL}$.

Proof. By a straightforward induction on φ using the following facts: bisimulation contraction of a model is bisimilar to the model, bisimilar models satisfy the same formulas of \mathcal{L}_{EL} , and public announcements preserve bisimulation [11]. \square

3 Strategies of Groups of Agents on Finite Models

3.1 Distinguishing Formulas

In this section we introduce distinguishing formulas that are satisfied in only one (up to bisimulation) state in a finite model (see [9] for details). Although agents know and can possibly announce an infinite number of formulas, using distinguishing formulas allows us to consider only a finite one. This is done by associating strategies of agents with corresponding distinguishing formulas. Here and subsequently, all epistemic models are *finite* and *bisimulation contracted*. Also, without loss of generality, we assume that the set of propositional variables P is finite.

Definition 10. *Let a finite epistemic model M be given. Formula $\delta_{S,S'}$ is called distinguishing for $S, S' \subseteq W$ if $S \subseteq \llbracket \delta_{S,S'} \rrbracket_M$ and $S' \cap \llbracket \delta_{S,S'} \rrbracket_M = \emptyset$. If a formula distinguishes state w from all other non-bisimilar states in M , we write δ_w .*

Proposition 2 ([9]). *Let a finite epistemic model M be given. Every pointed model (M, w) is distinguished from all other non-bisimilar pointed models (M, v) by some distinguishing formula $\delta_w \in \mathcal{L}_{EL}$.*

Given a finite model (M, w) , distinguishing formula δ_w is constructed recursively as follows:

$$\delta_w^{k+1} ::= \delta_w^0 \wedge \bigwedge_{a \in A} \left(\bigwedge_{w \sim_a v} \widehat{K}_a \delta_v^k \wedge K_a \bigvee_{w \sim_a v} \delta_v^k \right),$$

where $0 \leq k < |W|$, and δ_w^0 is the conjunction of all literals that are true in w , i.e. $\delta_w^0 ::= \bigwedge_{w \in V(p)} p \wedge \bigwedge_{w \notin V(p)} \neg p$.

Having defined distinguishing formulas for states, we can define distinguishing formulas for sets of states:

Definition 11. *Let some finite and bisimulation contracted model (M, w) , and a set S of states in M be given. A distinguishing formula for S is*

$$\delta_S ::= \bigvee_{w \in S} \delta_w.$$

3.2 Strategies

In this section we introduce strategies, and connect them to possible announcements using distinguishing formulas.

Definition 12. *Let $M/a = \{[w_1]_a, \dots, [w_n]_a\}$ be the set of a -equivalence classes in M . A strategy X_a for an agent a in a finite model (M, w) is a union of equivalence classes of a including $[w]_a$. The set of all available strategies of a is $S(a, w) = \{[w]_a \cup X_a : X_a \subseteq \bigcup M/a\}$. Group strategy X_G is defined as $\bigcap_{a \in G} X_a$ for all $a \in G$. The set of available strategies for a group of agents G is $S(G, w) = \{\bigcap_{a \in G} X_a : X_a \in S(a, w)\}$.*

Note, that for any (M, w) and $G \subseteq A$, $S(G, w)$ is not empty, since the trivial strategy that includes all the states of the current model is available to all agents.

Proposition 3. *In a finite model (M, w) , for any $G \subseteq A$, $S(G, w)$ is finite.*

Proof. Due to the fact that in a finite model there is a finite number of equivalence classes for each agent. \square

Thus, in Figure 1 of Section 2.1 there are three a -equivalence classes: $\{15_a5_b, 15_a10_b, 15_a15_b\}$, $\{10_a5_b, 10_a10_b, 10_a15_b\}$, and $\{5_a5_b, 5_a10_b, 5_a15_b\}$. Let us designate them by the first element of a corresponding set, i.e. 15_a5_b , 10_a5_b , and 5_a5_b . The set of all available strategies of agent a in $(M, 15_a5_b)$ is $\{15_a5_b, 15_a5_b \cup 10_a5_b, 15_a5_b \cup 5_a5_b, 15_a5_b \cup 10_a5_b \cup 5_a5_b\}$. Similarly, the set of all available strategies of agent b in $(M, 15_a5_b)$: $\{15_a5_b, 15_a5_b \cup 15_a10_b, 15_a5_b \cup 15_a15_b, 15_a5_b \cup 15_a10_b \cup 15_a15_b\}$. Finally, there is a group strategy for agents a and b that contains only two states – 15_a5_b and 10_a5_b . This strategy is an intersection of a 's $15_a5_b \cup 10_a5_b$ and b 's 15_a5_b , that is $\{15_a5_b, 15_a10_b, 15_a15_b, 10_a5_b, 10_a10_b, 10_a15_b\} \cap \{15_a5_b, 10_a5_b, 5_a5_b\}$.

Now we tie together announcements and strategies. Each of infinitely many possible announcements in a finite model corresponds to a set of states where it is true (a strategy). In a finite bisimulation contracted model, each strategy is definable by a distinguishing formula, hence it corresponds to an announcement. This allows us to consider finitely many strategies instead of considering infinitely many possible announcements: there are only finitely many non-equivalent announcements for each finite model, and each of them is equivalent to a distinguishing formula of some strategy.

Given a finite and bisimulation contracted model (M, w) and strategy X_G , a distinguishing formula δ_{X_G} for X_G can be obtained from Definition 11 as $\bigvee_{w \in X_G} \delta_w$.

Next, we show that agents know their strategies and thus can make corresponding announcements.

Proposition 4. *Let agent a have strategy X_a in some finite bisimulation contracted (M, w) . Then $(M, w) \models K_a \delta_{X_a}$. Also, let $X_G ::= X_a \cap \dots \cap X_b$ be a strategy, then $(M, w) \models K_a \delta_{X_a} \wedge \dots \wedge K_b \delta_{X_b}$, where $a, \dots, b \in G$.*

Proof. We show just the first part of the proposition, since the second part follows easily. By the definition of a strategy, $X_a = [w_1]_a \cup \dots \cup [w_n]_a$ for some $[w_1]_a, \dots, [w_n]_a \in M/a$. For every equivalence class $[w_i]_a$ there is a corresponding distinguishing formula $\delta_{[w_i]_a}$. Since for all $v \in [w_i]_a$, $(M, v) \models \delta_{[w_i]_a}$ (by Proposition 2), we have that $(M, v) \models K_a \delta_{[w_i]_a}$. The same holds for other equivalence classes of a including the one with w , and we have $(M, w) \models K_a \delta_{X_a}$. \square

The following proposition (which follows from Propositions 2 and 4) states that given a strategy, corresponding public announcement yields exactly the model with states specified by the strategy.

Proposition 5. *Given a finite bisimulation contracted model $M = (W, \sim, V)$ and a strategy X_a , $W^{K_a \delta_{X_a}} = X_a$. More generally, $W^{K_a \delta_{X_a} \wedge \dots \wedge K_b \delta_{X_b}} = X_G$, where $a, \dots, b \in G$.*

So, we have tied together announcements and strategies via distinguishing formulas. From now on, we may abuse notation and write M^{X_G} , meaning that M^{X_G} is an update of model M by a joint announcement of agents G that corresponds to strategy X_G .

Now, let us reformulate semantics for group and coalition announcement operators in terms of strategies.

Proposition 6. *For a finite bisimulation contracted model (M, w) we have that*

$$(M, w) \models \langle\!\langle G \rangle\!\rangle \varphi \text{ iff } \exists X_G \in S(G, w) \forall X_{A \setminus G} \in S(A \setminus G, w) : (M, w)^{X_G \cap X_{A \setminus G}} \models \varphi.$$

Proof. By Propositions 4 and 5, each strategy corresponds to an announcement. Each true announcement is a formula of the form $K_a \psi_a \wedge \dots \wedge K_b \psi_b$ where ψ_a is a formula which is true in every state of some union of a -equivalence classes and corresponds to a strategy. Similarly for announcements by groups. Hence we can substitute quantification over formulas with quantification over strategies in the truth definitions. \square

Definition 13. *Let some finite bisimulation contracted model (M, w) and G be given. A maximally informative announcement is a formula $\psi \in \mathcal{L}_{EL}^G$ such that $w \in W^\psi$ and for all $\psi' \in \mathcal{L}_{EL}^G$ such that $w \in W^{\psi'}$ it holds that $W^\psi \subseteq W^{\psi'}$. For finite models such an announcement always exists [3]. We will call the corresponding strategy X_G the strongest strategy on a given model.*

Intuitively, the strongest strategy is the smallest available strategy. Note that on a bisimulation contracted model, the strongest strategy of agents G is $X_G = [w]_a \cap \dots \cap [w]_b$ for $a, \dots, b \in G$, that is agents' strategies consist of the single equivalence classes that include the current state.

4 Model Checking for CAL

Employing strategies allows for a rather simple model checking algorithm for CAL. We switch from quantification over infinite number of epistemic formulas, to quantification over a finite set of strategies (Section 4.1). Moreover, we show that if the target formula is a positive PAL formula, then model checking is even more effective (Section 4.2).

4.1 General Case

First, let us define the model checking problem.

Definition 14. Let some model (M, w) and some formula φ be given. The model checking problem is the problem to determine whether φ is satisfied in (M, w) .

Algorithm 1 takes a finite model M , a state w of the model, and some $\varphi_0 \in \mathcal{L}_{CAL}$ as an input, and returns *true* if φ_0 is satisfiable in the model, and *false* otherwise.

Algorithm 1: $mc(M, w, \varphi_0)$

```

1: case  $\varphi_0$ :
2:    $p$  : if  $w \in V(p)$  then return true else return false;
3:    $\neg\varphi$  : if  $mc(M, w, \varphi)$  then return false else return true;
4:    $\varphi \wedge \psi$  : if  $mc(M, w, \varphi)$  and  $mc(M, w, \psi)$  then return true else return false;
5:    $K_a\varphi$  :
       check = true
       for all  $v$ 's s.t.  $w \sim_a v$ 
           if  $\neg mc(M, v, \varphi)$  then check = false
       return check
6:    $[\psi]\varphi$  : compute the  $\psi$ -submodel  $M^\psi$  of  $M$ 
       if  $w \in W^\psi$  then return  $mc(M^\psi, w, \varphi)$  else return true;
7:    $\langle\langle G \rangle\rangle\varphi$  : compute  $(\|M\|, w)$  and sets of strategies  $S(G, w)$  and  $S(A \setminus G, w)$ 
       for all  $X_G \in S(G, w)$ 
           check = true
           for all  $X_{A \setminus G} \in S(A \setminus G, w)$ 
               if  $\neg mc(\|M\|^{X_G \cap X_{A \setminus G}}, w, \varphi)$  then check = false
           if check then return true
       return false.

```

Now, we show correctness of the algorithm.

Proposition 7. Let (M, w) and $\varphi \in \mathcal{L}_{CAL}$ be given. Algorithm $mc(M, w, \varphi)$ returns *true* iff $(M, w) \models \varphi$.

Proof. By a straightforward induction on the complexity of φ . We use Proposition 6 to prove the case for $\langle\langle G \rangle\rangle$:

\Rightarrow : Suppose $mc(M, w, \langle\langle G \rangle\rangle\varphi)$ returns *true*. By line 7 this means that for some strategy X_G and all strategies $X_{A \setminus G}$, $mc(\|M\|^{X_G \cap X_{A \setminus G}}, w, \varphi)$ returns *true*. By the induction hypothesis, $(\|M\|, w)^{X_G \cap X_{A \setminus G}} \models \varphi$ for some X_G and all $X_{A \setminus G}$, and $(\|M\|, w) \models \langle\langle G \rangle\rangle\varphi$ by the semantics.

\Leftarrow : Let $(\|M\|, w) \models \langle\langle G \rangle\rangle\varphi$, which means that there is some strategy X_G such that for all $X_{A \setminus G}$, $(\|M\|, w)^{X_G \cap X_{A \setminus G}} \models \varphi$. By the induction hypothesis, the latter holds iff for some X_G and for all $X_{A \setminus G}$, $mc(\|M\|^{X_G \cap X_{A \setminus G}}, w, \varphi)$ returns *true*. By line 7, we have that $mc(\|M\|, w, \langle\langle G \rangle\rangle\varphi)$ returns *true*.

Proposition 8. Model checking for CAL is PSPACE-complete.

Proof. All the cases of the model checking algorithm apart from the case for $\langle\langle G \rangle\rangle$ require polynomial time (and polynomial space as a consequence). The case for $\langle\langle G \rangle\rangle$ iterates over exponentially many strategies. However each iteration can be computed using only polynomial amount of space to represent $(\|M\|, w)$ (which contains at most the same number of states as the input model M) and the result of the update (which is a submodel of $(\|M\|, w)$) and make a recursive call to check whether φ holds in the update. By reusing space for each iteration, we can compute the case for $\langle\langle G \rangle\rangle$ using only polynomial amount of space.

Hardness can be obtained by a slight modification of the proof of PSPACE-hardness of the model-checking problem for GAL in [1]. The proof encodes satisfiability of a quantified boolean formula as a problem whether a particular GAL formula is true in a model corresponding to the QBF formula. Since the encoding uses only two agents: an omniscient g and a universal i , we can replace $[g]$ and $\langle g \rangle$ with $\langle\langle g \rangle\rangle$ and $\langle\langle g \rangle\rangle$ (since i 's only strategy is equivalent to \top) and obtain a CAL encoding. \square

4.2 Positive Case

In this section, we demonstrate the following result: if in a given formula of \mathcal{L}_{CAL} subformulas within scopes of coalition announcement operators are positive PAL formulas, then complexity of model checking is polynomial. Allowing coalition announcement modalities to bind positive formulas only is a natural restriction because any positive formula that is distributed knowledge among the agents in G can be made common knowledge by a group or coalition announcement by G . Formally, $(M, w) \models_{D_G} \varphi$ implies $(M, w) \models \langle\langle G \rangle\rangle C_G \varphi$, where D_G stands for distributed knowledge which is interpreted by the intersection of all \sim_a relations, and C_G stands for common knowledge which is interpreted by the transitive and reflexive closure of the union of all \sim_a relations. See [10, 12], and also [?] where this is called *resolving* distributed knowledge. In other words, positive epistemic formulas can always be moved from the distributed knowledge to the common knowledge of the agents by cooperative communication. Negative formulas do not have this property. For example, it can be distributive knowledge of agents a and b that p and $\neg K_b p$: $D_{a,b}(p \wedge \neg K_b p)$. However it is impossible to achieve common knowledge of this formula: $C_G(p \wedge \neg K_b p)$ is inconsistent, since it implies both $K_b p$ and $\neg K_b p$. Going back to the example in section 2.1, it is distributed knowledge of a and b that $K_a 15a$ and $K_b 5b$. Both formulas are positive and can be made common knowledge if a and b honestly report the amount of money they have. However it is also distributed knowledge that $\neg K_a 5b$ and $\neg K_b 15a$. The conjunction

$$K_a 15a \wedge K_b 5b \wedge \neg K_a 5b \wedge \neg K_b 15a$$

is distributed knowledge, but cannot be made common knowledge for the same reason as above.

Definition 15. *The language \mathcal{L}_{PAL^+} of the positive fragment of public announcement logic PAL is defined by the following BNF:*

$$\varphi, \psi ::= p \mid \neg p \mid (\varphi \wedge \psi) \mid (\varphi \vee \psi) \mid K_a \varphi \mid [\neg\psi]\varphi,$$

where $p \in P$ and $a \in A$.

Definition 16. *Formula φ is preserved under submodels if for any models M_1 and M_2 , $M_2 \subseteq M_1$ and $(M_1, w) \models \varphi$ implies $(M_2, w) \models \varphi$.*

A known result that we use in this section states that formulas of \mathcal{L}_{PAL^+} are preserved under submodels [12]. We also need the following special fact:

Proposition 9. *$\langle\!\langle G \rangle\!\rangle\varphi \leftrightarrow \langle\!\langle A \setminus G \rangle\!\rangle\varphi$ is valid for positive φ on finite bisimulation contracted models.*

Proof. The left-to-right direction is generally valid and we omit the proof. Suppose that $(M, w) \models \langle\!\langle A \setminus G \rangle\!\rangle\varphi$. By Proposition 6, we have that for all $X_{A \setminus G}$, there is some X_G such that $(M, w)^{X_{A \setminus G} \cap X_G} \models \varphi$. This implies that $(M, w)^{\top_{A \setminus G} \cap X_G} \models \varphi$ for the trivial strategy $\top_{A \setminus G}$ and some X_G . The latter is equivalent to $(M, w)^{X_G} \models \varphi$. Since φ is positive (and hence preserved under submodels), $(M, w)^{X'_G} \models \varphi$, where X'_G is the strongest strategy of G . The latter implies (again, due to the fact that φ is positive) that for all updates of the form $X'_G \cap X_{A \setminus G}$ (since they generate a submodel of $(M, w)^{X'_G}$), we also have $(M, w)^{X'_G \cap X_{A \setminus G}} \models \varphi$. And this is $(M, w) \models \langle\!\langle G \rangle\!\rangle\varphi$ by Proposition 6. \square

Now we are ready to deal with model checking for the positive case.

Proposition 10. *Let $\varphi \in \mathcal{L}_{CAL}$ be a formula such that all its subformulas ψ that are within scopes of $\langle\!\langle G \rangle\!\rangle$ belong to fragment \mathcal{L}_{PAL^+} . Then the model checking problem for CAL is in P.*

Proof. For this particular case we modify Algorithm 1 by inserting the following instead of the case on line 7.

$\langle\!\langle G \rangle\!\rangle\varphi$: compute $(\|M\|, w)$ and $(\|M\|^{X_G}, w)$, where X_G corresponds to the strongest strategy of G ,
if $mc(\|M\|^{X_G}, w, \varphi)$ **then return true else return false.**

For all subformulas of φ_0 , the algorithm calls are in P. Consider the modified call for $\langle\!\langle G \rangle\!\rangle\varphi$. It requires constructing a single update model given a specified strategy, which is a simple case of restricting the input model to the set of states in the strategy. This can be done in polynomial time. Then we call the algorithm on the updated model for φ , which by assumption requires polynomial time. \square

Now, let us show that the algorithm is correct.

Proposition 11. *Let (M, w) and $\varphi \in \mathcal{L}_{PAL^+}$ be given. The modified algorithm $mc(M, w, \varphi)$ returns true iff $(M, w) \models \varphi$.*

Proof. By induction on φ . We show case $\langle\langle G \rangle\rangle\varphi$:

\Rightarrow : Suppose that $mc(M, w, \langle\langle G \rangle\rangle\varphi)$ returns *true*. This means that $mc(\|M\|^{X_G}, w, \varphi)$ returns *true*, where X_G is the strongest strategy of G . By the induction hypothesis, we have that $(\|M\|, w)^{X_G} \models \varphi$. Since φ is positive, for all stronger updates $X_G \cap X_{A \setminus G}$ it holds that $(\|M\|, w)^{X_G \cap X_{A \setminus G}} \models \varphi$, which is $(\|M\|, w) \models \langle\langle G \rangle\rangle\varphi$ by Proposition 6. Finally, the latter model is bisimilar to (M, w) and hence $(M, w) \models \langle\langle G \rangle\rangle\varphi$.

\Leftarrow : Let $(M, w) \models \langle\langle G \rangle\rangle\varphi$. By Proposition 6 this means that there is some X_G such that for all $X_{A \setminus G}$: $(M, w)^{X_G \cap X_{A \setminus G}} \models \varphi$. Set of all $X_{A \setminus G}$'s also includes the trivial strategy $\top_{A \setminus G}$, and we have $(M, w)^{X_G \cap \top_{A \setminus G}} \models \varphi$, which is equivalent to $(M, w)^{X_G} \models \varphi$. Since φ is positive and hence preserved under submodels, $(M, w)^{X'_G} \models \varphi$, where X'_G is the strongest strategy of G . By the induction hypothesis, we have that $mc(\|M\|^{X'_G}, w, \varphi)$ returns *true*. And by line 7 of the modified algorithm, we conclude that $mc(\|M\|, w, \langle\langle G \rangle\rangle\varphi)$ returns *true*. \square

The case of $\langle\langle G \rangle\rangle\varphi$ is resolved by translating the formula into $\langle\langle A \setminus G \rangle\rangle\varphi$, which is allowed by Proposition 9.

5 Concluding Remarks

We have shown that the model checking problem for CAL is PSPACE-complete, just like the one for GAL [1] and APAL [5]. However, in a special case when formulas within scopes of coalition modalities are positive PAL formulas, the model checking problem is in P. The same result would apply to GAL and APAL; in fact, in those cases the formulas in the scope of group and arbitrary announcement modalities can belong to a larger positive fragment (the positive fragment of GAL and of APAL, respectively, rather than of PAL). The latter is due to the fact that GAL and APAL operators are purely universal, while CAL operators combine universal and existential quantification, and CAL does not appear to have a non-trivial positive fragment extending that of PAL.

There are several interesting open questions. For example, the relative expressivity of GAL and CAL is still an open question. It is also not known what is the model checking complexity for coalition logics with more powerful actions like private announcements [6].

Acknowledgements

We thank anonymous IJCAI 2018 and KI 2018 referees for constructive comments, and IJCAI 2018 referees for finding an error in the earlier version of this paper.

References

1. Ågotnes, T., Balbiani, P., van Ditmarsch, H., Seban, P.: Group announcement logic. *Journal of Applied Logic* **8**(1), 62–81 (2010). <https://doi.org/10.1016/j.jal.2008.12.002>

2. Ågotnes, T., van Ditmarsch, H.: Coalitions and announcements. In: Padgham, L., Parkes, D.C., Müller, J.P., Parsons, S. (eds.) 7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008), Estoril, Portugal, May 12-16, 2008, Volume 2). pp. 673–680. IFAAMAS (2008). <https://doi.org/10.1145/1402298.1402318>
3. Ågotnes, T., van Ditmarsch, H.: What will they say? - public announcement games. *Synthese* **179**(Supplement-1), 57–85 (2011). <https://doi.org/10.1007/s11229-010-9838-8>
4. Ågotnes, T., van Ditmarsch, H., French, T.S.: The undecidability of quantified announcements. *Studia Logica* **104**(4), 597–640 (2016), <https://doi.org/10.1007/s11225-016-9657-0>
5. Balbiani, P., Baltag, A., van Ditmarsch, H., Herzig, A., Hoshi, T., de Lima, T.: ‘Knowable’ as ‘known after an announcement’. *Review of Symbolic Logic* **1**(3), 305–334 (2008). <https://doi.org/10.1017/S1755020308080210>
6. Baltag, A., Moss, L.S., Solecki, S.: The logic of public announcements and common knowledge and private suspicions. In: *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-98)*, Evanston, IL, USA, July 22-24, 1998. pp. 43–56 (1998)
7. van Benthem, J.: *Logic in Games*. MIT Press (2014)
8. Blackburn, P., van Benthem, J.: Modal logic: A semantic perspective. In: Blackburn, P., van Benthem, J., Wolter, F. (eds.) *Handbook of Modal Logic*, pp. 1–84. Elsevier (2006)
9. van Ditmarsch, H., Fernández-Duque, D., van der Hoek, W.: On the definability of simulation and bisimulation in epistemic logic. *Journal of Logic and Computation* **24**(6), 1209–1227 (2014). <https://doi.org/10.1093/logcom/exs058>
10. van Ditmarsch, H., French, T., Hales, J.: Positive announcements. *CoRR abs/1803.01696* (2018), <http://arxiv.org/abs/1803.01696>
11. van Ditmarsch, H., van der Hoek, W., Kooi, B.: *Dynamic Epistemic Logic*, *Synthese Library*, vol. 337. Springer (2008). <https://doi.org/10.1007/978-1-4020-5839-4>
12. van Ditmarsch, H., Kooi, B.: The secret of my success. *Synthese* **153**(2), 339–339 (2006). <https://doi.org/10.1007/s11229-006-8493-6>
13. Hintikka, J.: *Knowledge and belief. An introduction to the logic of the two notions*. Cornell University Press, Ithaca, NY (1962)
14. Parikh, R.: The logic of games and its applications. In: Karplinski, M., van Leeuwen, J. (eds.) *Topics in the Theory of Computation, Annals of Discrete Mathematics*, vol. 24, pp. 111 – 139. Elsevier Science, Amsterdam (1985). [https://doi.org/https://doi.org/10.1016/S0304-0208\(08\)73078-0](https://doi.org/https://doi.org/10.1016/S0304-0208(08)73078-0)
15. Pauly, M.: A modal logic for coalitional power in games. *Journal of Logic and Computation* **12**(1), 149–166 (2002). <https://doi.org/10.1093/logcom/12.1.149>
16. Plaza, J.: Logics of public communications (reprint of 1989’s paper). *Synthese* **158**(2), 165–179 (2007). <https://doi.org/10.1007/s11229-007-9168-7>