Dynamic Coalition Logic: Granting and Revoking Dictatorial Powers

Rustam Galimullin[1] and Thomas Ågotnes[1,2]

[1] University of Bergen, Bergen, Norway
{rustam.galimullin, thomas.agotnes}@uib.no
[2] Southwest University, Chongqing, China

Abstract. One of the classic formalisms for reasoning about multi-agent coalitional ability is coalition logic (CL). In CL it is possible to express what a coalition can achieve in the next step no matter what agents outside of the coalition do at the same time. We propose an extension of CL with dynamic operators that allow us to grant dictatorial powers to agents or to revoke them. In such a way we are able to reason about the *dynamics of coalitional ability*. We also discuss some logical properties of the proposed formalisms and compare their relative expressive power.

Keywords: Dynamic Coalition Logic · Coalition Logic · Arrow Updates · Modal Logic.

1 Introduction

Reasoning about actions and abilities of single agents and groups of agents is ubiquitous in AI. The notable examples are classical and epistemic [6] multi-agent planning, game theory, software verification, and so on. One of the most recent challenges is verification of the safety of blockchains, and, in particular, smart contracts [11].

To simplify an example problem from [11], assume that there is a newly-founded company, and the initial block of the smart contract specifies that the board of directors consists of Alice, Bob, and Carol, and all financial decisions are made according to the majority rule. Apart from the board of directors, there are also employees, Dave and Ellen, with fewer privileges. Now assume that Bob was caught making suspicious transactions and, according to some financial regulation, they can no longer be on the board. Moreover, to substitute Bob, Ellen was promoted. The resulting new situation is recorded in the next block of the blockchain, where it is specified that Bob loses the right to make financial decisions, while Ellen obtains such a right.

Clearly, in the described scenario, it is vital to specify what an agent, or a group thereof, is able or unable to do. One of the most popular languages for reasoning about abilities of groups of agents is called *coalition logic* (CL) [15] (which can be considered as a Next-time fragment of alternating-time temporal logic [3]). CL extends propositional logic with constructs $\langle\langle C \rangle\rangle \varphi$ meaning that 'there is a joint action by agents from coalition C such that no matter what agents outside of the coalition do, φ holds after the execution of the joint action'.

While CL captures the abilities of agents to force certain outcomes, it provides only a static snapshot and thus is inadequate for the situations where new policies or regulations override agents' abilities. We thus propose the development and study of *dynamic coalition logic*, with dynamic operators in the spirit of dynamic epistemic logics[3] [7] that can modify or update the abilities of agents and coalitions. In the current paper we take the first step and focus on granting and revoking *dictatorial powers*, i.e. the ability of single agents to force an outcome.

To model updates of dictatorial powers, we borrow syntax and basic intuition from arrow update logic (AUL) [12], where constructs $U = \{(\chi_1, a_1, \psi_1), \ldots, (\chi_n, a_n, \psi_n)\}$ specify which belief relations should be preserved in a current model. AUL, being a dynamic epistemic logic [7], models such epistemic events as public and private announcements, lying, etc. Arrow updates were also used to reason about norms [13].

First, we consider *positive dictatorial dynamic coalition logic* (DDCL$^+$) that extends CL with updates $+U = \{(\chi_1, a_1, \psi_1)^+, \ldots, (\chi_n, a_n, \psi_n)^+\}$. In this case, $+U$ specifies between which states an agent should be granted the dictatorial power. In particular, $(\chi, a, \psi)^+$ means that agent a will be able to force any state where ψ is true, from any state where χ holds. In terms of models, this means that in the updated model there will be a set of new arrows satisfying the requirement. In this regard, DDCL$^+$ is slightly reminiscent of bridge logics [4]. However, in our case, the interpretation of arrows and the mechanism of adding relations are completely different.

Apart from the logic of granting dictatorial powers, we also study the logic of revoking such powers, which we call *negative dictatorial dynamic coalition logic* (DDCL$^-$). The logic extends CL with updates $-U = \{(\chi_1, a_1, \psi_1)^-, \ldots, (\chi_n, a_n, \psi_n)^-\}$ that, similarly to the updates of AUL, specify which dictatorial powers should be preserved, while all other such powers, not satisfying the specification, are removed.

In the paper, after we recall some background information about CL in Section 2, we present syntax and semantics of DDCL$^+$ and DDCL$^-$ (Section 3). In particular, we argue that updates $+U$ and $-U$ are not always executable. In Section 4, we study the expressivity of the logics. Specifically, in contrast to AUL, which has the same expressive power as the logic without arrow updates, we show that DDCL$^+$ and DDCL$^-$ are strictly more expressive than CL. Hence, again in contrast to AUL, there cannot be reduction axioms for the logics. Finally, we also show that DDCL$^+$ and DDCL$^-$ are incomparable. We discuss further research in Section 5.

[3] Dynamic epistemic logics with coalitional operators have been studied only in the setting of public announcements (see [8, 9]). These logics, however, are not strictly coalitional in the sense of [15] since they are defined on epistemic models, not concurrent game models.

2 Some Definitions and Notions of Coalition Logic

As the logics introduced in this paper are dynamic extensions of coalition logic, we first provide all the necessary background information on it (see [15, 14, 2]). Let P be a countable set of propositional variables, and A be a finite set of agents.

Definition 1. *The* language of coalition logic $\mathcal{CL}$ *is given recursively by the following grammar:*

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \langle\!\langle C \rangle\!\rangle \varphi$$

where $p \in P$ and $C \subseteq A$. Constructs $\langle\!\langle C \rangle\!\rangle \varphi$ are read 'coalition C can force φ'. We denote $A \setminus C$ as $\overline{C}$. The dual of $\langle\!\langle C \rangle\!\rangle \varphi$ is $[\![C]\!]\varphi := \neg\langle\!\langle C \rangle\!\rangle \varphi \neg\varphi$.

Formulas of coalition logic are interpreted on concurrent game models.

Definition 2. *A* concurrent game model *(CGM), or a* model, *is a tuple $M = (A, S, Act, act, out, L)$. A is a non-empty finite set of agents, and the subsets of A are called* coalitions. *S is a non-empty set of states, and Act is a non-empty set of actions.*

Function $act : A \times S \to 2^{Act} \setminus \emptyset$ assigns to each agent and each state a non-empty set of actions. A C-action at a state $s \in S$ is a tuple α_C such that $\alpha_C(i) \in act(i, s)$ for all $i \in C$. The set of all C-actions in s is denoted by $act(C, s)$. We will also write $\alpha_{C_1} \cup \alpha_{C_2}$ to denote a $C_1 \cup C_2$-action with $C_1 \cap C_2 = \emptyset$.

A tuple of actions $\alpha = \langle \alpha_1, \ldots, \alpha_k \rangle$ with $k = |A|$ is called an action profile. *An action profile is* executable *in state s if for all $i \in A$, $\alpha_i \in act(i, s)$. The set of all action profiles executable in s is denoted by $act(s)$. An action profile α* extends *a C-action α_C, written $\alpha_C \sqsubseteq \alpha$, if for all $i \in C$, $\alpha(i) = \alpha_C(i)$.*

Function out assigns to each state s and each $\alpha \in act(s)$ a unique output state. We write $Out(s, \alpha_C)$ for $\{out(s, \alpha) \mid \alpha \in act(s)$ and $\alpha_C \sqsubseteq \alpha\}$. Intuitively, $Out(s, \alpha_C)$ is the set of all states reachable by action profiles that extend some given C-action α_C. Finally, $L : S \to P$ is the valuation function.

We will also denote a CGM M with a designated, or current, state s as M_s.

Note that although in [15] the semantics of CL are given relative to effectivity models, we still can use CGMs as these types of models are semantically equivalent (see more on this topic in [10]).

Definition 3. *Let M_s be a pointed CGM. The semantics of CL are defined as follows:*

$M_s \models p$ *iff* $s \in L(p)$
$M_s \models \neg\varphi$ *iff* $M_s \not\models \varphi$
$M_s \models \varphi \wedge \psi$ *iff* $M_s \models \varphi$ *and* $M_s \models \psi$
$M_s \models \langle\!\langle C \rangle\!\rangle \varphi$ *iff* $\exists \alpha_C, \forall \alpha_{\overline{C}} : M_t \models \varphi$, *where* $t = out(s, \alpha_C \cup \alpha_{\overline{C}})$

Informally, the semantics of the coalition modality $\langle\langle C\rangle\rangle\varphi$ mean that in the current state of a given CGM there is a choice of actions by the members of coalition C such that no matter what the opponents from the anti-coalition $\overline{C}$ choose to do at the same time, φ holds after the execution of the corresponding action profile. Given φ and M, we define $[\![\varphi]\!]_M := \{s \in S \mid M_s \models \varphi\}$.

Definition 4. *We call a formula φ* valid *if for all M_s it holds that $M_s \models \varphi$.*

Definition 5. *Let $M = (A, S^M, Act^M, act^M, out^M, L^M)$ and $N = (A, S^N, Act^N, act^N, out^N, L^N)$ be two CGMs. A relation $Z \subseteq S^M \times S^N$ is called* bisimulation *if and only if for all $C \subseteq A$, $s_1 \in S^M$ and $s_2 \in S^N$, $(s_1, s_2) \in Z$ implies*

- *for all $p \in P$, $s_1 \in L^M(p)$ iff $s_2 \in L^N(p)$;*
- *for all $\alpha_C \in act^M(C, s_1)$, there exists $\beta_C \in act^N(C, s_2)$ such that for every $s'_2 \in Out^N(s_2, \beta_C)$, there exists $s'_1 \in Out^M(s_1, \alpha_C)$ such that $(s'_1, s'_2) \in Z$.*
- *The same as above with 1 and 2 swapped.*

If there is a bisimulation between M and M linking states s_1 and s_2, we call the pointed models bisimilar *($M_{s_1} \leftrightarrows N_{s_2}$).*

Theorem 1 ([1]). *Let M and N be CGMs such that $M \leftrightarrows N$ and there is a bisimilation between $s \in S^M$ and $t \in S^N$. Then for all $\varphi \in \mathcal{CL}$, $M_s \models \varphi$ iff $N_t \models \varphi$.*

Before we continue, we define an auxiliary set of forcing actions for each state and agent. Intuitively, an action is a forcing action if all action profiles it appears in lead to the same state.

Definition 6. *Let M be a CGM. The set of* forcing actions *for agent i and state s, denoted as $\mathfrak{f}(i, s)$, is defined as follows:*

$$\{\alpha_i \in act(i, s) \mid \forall \alpha, \beta \in act(s) : (\alpha_i \sqsubseteq \alpha \text{ and } \alpha_i \sqsubseteq \beta) \text{ implies } out(\alpha, s) = out(\beta, s)\}$$

Without loss of generality and to make the following technical presentation clearer, we assume that each action in the set of forcing actions is labelled with a pair of states it connects. Thus, elements of $\mathfrak{f}(i, s)$ are $a_i^{(s,t_1)}$, $b_i^{(s,t_2)}$, ...

3 Dictatorial Dynamic Coalition Logic

In this section we introduce the ways of granting and revoking dictatorial powers of agents. We borrow the syntax from arrow update logic [12].

3.1 Granting Dictatorial Powers

Definition 7. *The* language of positive dictatorial dynamic coalition logic $\mathcal{DDCL}^+$ *is given by the following BNF:*

$$\begin{aligned} \varphi &::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \langle\langle C\rangle\rangle\varphi \mid [+U]\varphi \\ +U &::= (\varphi, a, \varphi)^+ \mid (\varphi, a, \varphi)^+, +U \end{aligned}$$

where $p \in P$, $a \in A$, and $C \subseteq A$. We will abuse the notation and treat the ist $+U := (\psi_1, a_1, \varphi_1)^+, \ldots, (\psi_n, a_n, \varphi_n)^+$ as the set $+U := \{(\psi_1, a_1, \varphi_1)^+, \ldots, (\psi_n, a_n, \varphi_n)^+\}$. The dual of $[+U]\varphi$ is $\langle +U \rangle \varphi := \neg[+U]\neg\varphi$.

The supposed meaning of $(\varphi, a, \psi)^+$ is as follows: in each φ-state, in the updated model, there will be a *new* action for agent a such that no matter which actions other agents choose, the target state is a ψ-state. In case of multiple φ- and ψ-states, we have a new action for each transition.

Example 1. Before giving the formal definition of the semantics, let us consider an example. In Figure 1, in model M there are three states, s, t, and u, and two agents, a and b. In s, agent a has three actions, a_0, a_1, and a_2, and she has the ability to decide which state will be next. Agent b does not have the ability to force anything in the model.

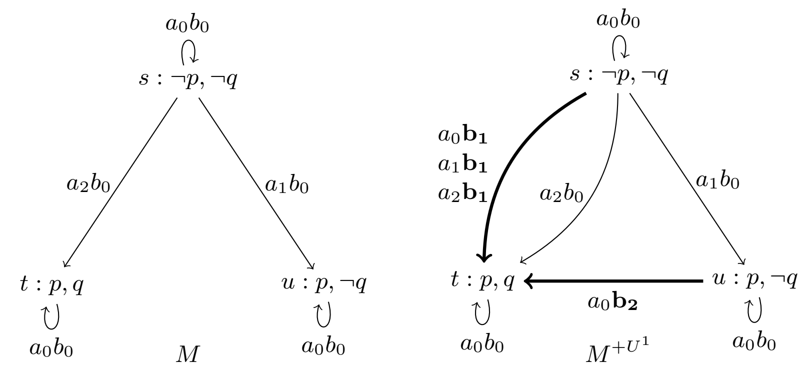

Fig. 1. Models M (left) and M^{+U^1} (right), where thick arrows depict new transitions, and new actions are in bold font.

To make the example more relatable, assume that p means that agent a has a cup of coffee, and q stands for b having coffee. Action a_1 signifies agent a pouring coffee just for herself; if she chooses a_2, then she pours coffee for b as well; and actions a_0 and b_0 are 'do nothing', or 'enjoy oneself' actions. It is clear from the figure that in s, where neither a nor b have coffee, a can choose to either get a cup for herself (transition to state u), or for both of them (transition to t), or just do nothing (self-loop in s). Agent b, being a polite guest of a, cannot do anything.

The set of forcing actions for a in s is $\mathfrak{f}(a, s) = \{a_0^{(s,s)}, a_1^{(s,u)}, a_2^{(s,t)}\}$ and the set of forcing actions for b in s is empty. In states t and u both of a and b have one forcing action each: a has a_0 and b has b_0.

Now, let us consider $+U^1 = \{(\neg q, b, q)^+\}$. Informally, we want to give agent b the power to get a cup of coffee whenever she does not have one. The result of updating our model with $+U^1$ is presented in Figure 1 on the right. In the figure, agent b gains two new actions (in bold font in the figure): b_1 in s, and b_2 in u. Then, for each action profile with b_1 or b_2 there is a transition (depicted by thick arrows) to state t, where q is true. This means that after the update, agent b has the dictatorial power to force q from any of the states where q does not hold. Formally, we have, for example, that $M_s \not\models \langle\!\langle b \rangle\!\rangle q$ and $M_s \models [+U^1]\langle\!\langle b \rangle\!\rangle q$. With the update, sets of forcing actions are changed as well. In state s, $\mathfrak{f}(a, s) = \{a_2^{(s,t)}\}$ and $\mathfrak{f}(b, s) = \{b_1^{(s,t)}\}$; in state t both a and b have $a_0^{(t,t)}$ and $b_0^{(t,t)}$ respectively; in state u, $\mathfrak{f}(b, u) = \{b_0^{(u,u)}, b_2^{(u,t)}\}$ and a does not have forcing actions.

As another example, consider $+U^2 = \{(\top, b, \neg p)^+\}$. We can imagine an informal reading as agent b gets rid of a's coffee no matter what, and precludes her from getting one if she hasn't got one yet (state s). The update of the initial model is depicted in Figure 2 on the left.

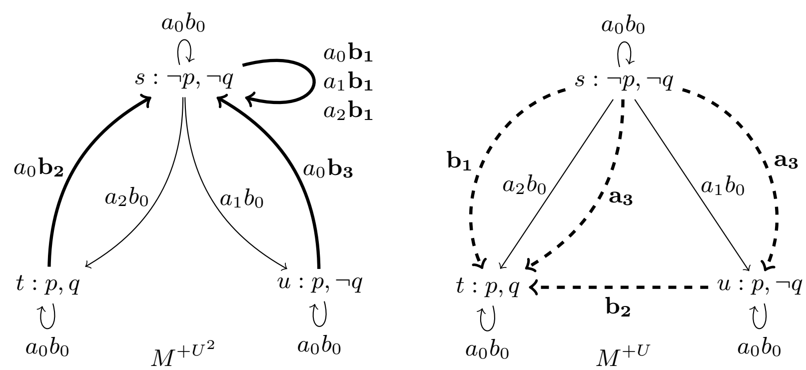

Fig. 2. Model M^{+U^2} (left) and a tentative updated model M^{+U} (right), where thick arrows depict new transitions, dashed arrows depict new tentative transitions required by update $+U$, and new actions are in bold font. Observe that in M^{+U} both a and b require dictatorial powers in s.

In the updated model, agent a does not have any strategy to escape $\neg p$, or, formally, $M_s \models [+U^2][\![a]\!]\neg p$. The non-empty sets of forcing actions in the updated model are the following: $\mathfrak{f}(a, s) = \{a_0^{(s,s)}\}$, $\mathfrak{f}(b, s) = \{b_1^{(s,s)}\}$, $\mathfrak{f}(b, t) = \{b_0^{(t,t)}, b_2^{(t,s)}\}$, and $\mathfrak{f}(b, u) = \{b_0^{(u,u)}, b_3^{(u,s)}\}$.

Before we continue with the definition of the semantics, it should be noted that not every update can be implemented. For example, in our initial model in Figure 1, we may require the following update: $+U = \{(\neg q, b, q)^+, (\neg p, a, p)^+\}$. In this case, we have a clash of control while assigning a transition from s and t: both a and b should have the ability to force t. See the model on the right in Figure 2 for a representation of the problem.

Definition 8. *Let M be a CGM, and $+U$ be an update. We call $+U$ executable in M iff for all $(\varphi, i, \psi)^+, (\chi, j, \tau)^+ \in +U$: $[\![\varphi]\!]_M \cap [\![\chi]\!]_M = \emptyset$ whenever $i \neq j$.*

Informally, the definition says that an update is executable if it is not granting dictatorial powers to different agents in the same state.

Definition 9. *Let M_s be a CGM. The semantics of DDCL$^+$ extends Definition 3 with the following clause for updates:*

$$M_s \models [+U]\varphi \text{ iff } +U \text{ is executable in } M \text{ implies } M_s^{+U} \models \varphi$$

where $M^{+U} = (A, S, Act^{+U}, act^{+U}, out^{+U}, L)$ is the updated model.

To define $M^{+U} = (A, S, Act^{+U}, act^{+U}, out^{+U}, L)$, we first define the set of new *forcing actions for each agent i in each state s. Let $Pairs^{+U}_{(\varphi,i,\psi)^+} = \{s \mid M_s \models \varphi\} \times \{s \mid M_s \models \psi\}$ be all pairs of states between which we need to add transitions according to some $(\varphi, i, \psi)^+ \in +U$. The new set of forcing actions $\mathfrak{f}^{+U}(i, s)$ consists of actions $i_k^{(s,t)}$ for each $(s, t) \in Pairs^{+U}_{(\varphi,i,\psi)^+}$ and all $(\varphi, i, \psi)^+ \in +U$, where k starts with $|act(i, s)| + 1$ and is increased by 1 for each such (s, t). Intuitively, the set of new forcing actions is constructed according to $+U$ such that each new action has a unique ordinal number k.*

Then Act^{+U} is $Act \cup \bigcup_{i\in A,s\in S} \mathfrak{f}^{+U}(i, s)$. Function $act^{+U}(i, s) = act(i, s) \cup \mathfrak{f}^{+U}(i, s)$. Finally,

$$out^{+U}(\alpha, s) = \begin{cases} t, & \exists \alpha_i^{(s_1,s_2)} \sqsubseteq \alpha : \alpha_i^{(s_1,s_2)} \in \mathfrak{f}^{+U}(i, s) \text{ and } s_2 = t \\ out(\alpha, s), & \text{otherwise} \end{cases}$$

Intuitively, $out^{+U}(\alpha, s)$ takes the system into state t if there is a forcing action in α labelled with (s, t), and works as the original $out(\alpha, s)$ if there are no new forcing actions in α.

As the first step towards the systematic study of DDCL$^+$, we consider some valid and not valid formulas of the logic. Proofs are omitted due to the lack of space.

Proposition 1. *The following holds for formulas of DDCL$^+$.*

1. $\langle +U\rangle\varphi \wedge \langle +U\rangle\psi \leftrightarrow \langle +U\rangle(\varphi \wedge \psi)$ *is valid.*
2. $\langle +U^1\rangle\varphi \wedge \langle +U^2\rangle\varphi \rightarrow \langle +U^1 \cup +U^2\rangle\varphi$ *is not valid.*
3. $\langle +U^1\rangle\langle +U^2\rangle\varphi \rightarrow \langle +U^2\rangle\langle +U^1\rangle\varphi$ *is not valid.*
4. $\langle(\psi_1, a, \chi_1)^+, (\psi_2, b, \chi_2)^+\rangle\varphi \rightarrow \langle(\psi_1, a, \chi_1)^+\rangle\langle(\psi_2, b, \chi_2)^+\rangle\varphi$ *is not valid.*
5. $\langle(\psi_1, a, \chi_1)^+\rangle\langle(\psi_2, b, \chi_2)^+\rangle\varphi \rightarrow \langle(\psi_1, a, \chi_1)^+, (\psi_2, b, \chi_2)^+\rangle\varphi$ *is not valid.*
6. $\langle(\psi_1, a, \chi_1)^+, (\psi_2, b, \chi_2)^+\rangle\varphi \rightarrow \langle(\psi_1, a, \chi_1)^+\rangle\langle(\psi_2, b, \chi_2)^+\rangle\varphi$, *where ψ_1, χ_1, ψ_2, and χ_2 are propositional, is valid.*

Intuitively, property 1 states that positive updates are monotonic operators. That we cannot, in general, unite or commute updates is captured by items 2 and 3. The intuition behind possible counterexamples is that such actions may result in updates that are not executable. Properties 4 and 5 claim that we cannot decompose a single update into a series of consecutive ones, and vice versa. However, such a decomposition is possible if the starting and target states are specified by formulas of propositional logic, as claimed by item 6.

Remark 1. Items 4, 5, and 6 of Proposition 1 hint at possible interaction between DDCL^+ and a fragment thereof with only single-agent updates, where we grant dictatorial powers to one agent at a time. In particular, such updates are always executable: there is no clash of power if we consider only a single agent per update. Due to the lack of space, we leave the discussion of the fragment and how it relates to DDCL^+ for the future.

3.2 Revoking Dictatorial Powers

Apart from granting dictatorial powers, another way of updating coalitional abilities is by revoking such powers. Similarly to DDCL^+, we approach this problem from the perspective of arrow updates.

Definition 10. *The* language of negative dictatorial dynamic coalition logic $\mathcal{DDCL}^-$ *is given by the following grammar:*

$$\begin{aligned} \varphi \;\; &::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \langle\!\langle C \rangle\!\rangle \varphi \mid [-U]\varphi \\ -U &::= (\varphi, a, \varphi)^- \mid (\varphi, a, \varphi)^-, -U \end{aligned}$$

where $p \in P$, $a \in A$, *and* $C \subseteq A$. *We treat the list* $-U$ *as the set* $-U := \{(\psi_1, a_1, \varphi_1)^-, \dots, (\psi_n, a_n, \varphi_n)^-\}$. *The dual of* $[-U]\varphi$ *is* $\langle -U \rangle\varphi := \neg[-U]\neg\varphi$.

The idea behind $-U$ updates is similar to how arrow updates work in AUL [12]: the list $-U$ specifies which dictatorial powers should be *preserved*. In other words, for each χ-state where there is a local dictatorial agent a forcing a ψ-state, we check whether there is a corresponding $(\chi, a, \psi)^-$ in $-U$. If yes, then we leave the corresponding arrows as they are; if not, we delete the corresponding arrows.

Example 2. Recall model M^{+U^2} in Figure 2 with the non-empty sets of forcing actions $\mathfrak{f}(a,s) = \{a_0^{(s,s)}\}$, $\mathfrak{f}(b,s) = \{b_1^{(s,s)}\}$, $\mathfrak{f}(b,t) = \{b_0^{(t,t)}, b_2^{(t,s)}\}$, and $\mathfrak{f}(b,u) = \{b_0^{(u,u)}, b_3^{(u,s)}\}$.

Now, assume that the abuse of power by b after update $+U^2$ was not tolerated in the office, and the new policy was issued specifying that once a has got a coffee, she can enjoy it in peace. The corresponding update is $-U^3 = \{(p, b, p)^-, (\neg p, b, \neg p)^-\}$, where the first clause preserves self-loops in states t and u, and the second clause preserves some of the self-loops in state s. The result of updating M^{+U^2} with $-U^3$ is shown in Figure 3. The new sets of forcing actions are $\mathfrak{f}(a,s) = \{a_0^{(s,s)}\}$, $\mathfrak{f}(b,s) = \{b_1^{(s,s)}\}$, $\mathfrak{f}(a,t) = \{a_0^{(t,t)}\}$, $\mathfrak{f}(b,t) = \{b_0^{(t,t)}\}$, $\mathfrak{f}(a,u) = \{a_0^{(u,u)}\}$, and $\mathfrak{f}(b,u) = \{b_0^{(u,u)}\}$. Now it holds that $M_t^{+U^2} \models \langle\!\langle b \rangle\!\rangle \neg p \wedge [-U^3][\![b]\!]p$.

Similarly to the case of $+U$, we need to be careful with $-U$ since we do not want to end up in a situation, where an agent does not have any actions in some state. Indeed, consider the further update of the model from Figure 3 with $-U^4 = \{(p \wedge \neg p, a, p \wedge \neg p)^-\}$ meaning that we are required to revoke all dictatorial powers from all the agents (since none of the states satisfy $p \wedge \neg p$). The resulting model would look like the one presented in Figure 3 on the right.

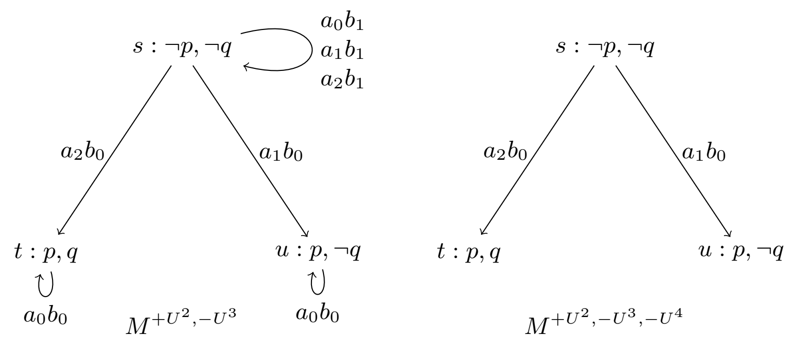

Fig. 3. Model $M^{+U^2,-U^3}$ (left) and a tentative updated model $M^{+U^2,-U^3,-U^4}$ (right).

In the resulting tentative model there are no actions in states t and u, hence the structure in Figure 3 is not a CGM at all. To tackle this issue, we, once again, require a corresponding condition of executability.

We first specify which forcing actions should be *preserved* in $\mathfrak{f}(i, s)$ according to $-U$ in a given CGM M. Set $\mathfrak{f}^{-U}(i, s)$ of forcing actions to be preserved is defined as

$$\mathfrak{f}^{-U}(i, s) = \{\alpha_i^{(s,t)} \in \mathfrak{f}(i, s) \mid \exists(\varphi, i, \psi)^- \in -U : M_s \models \varphi \text{ and } M_t \models \psi\}$$

Intuitively, a forcing action $\alpha^{(s,t)}$ should be preserved if there is a $(\varphi, i, \psi)^- \in -U$ such that s satisfies φ, and t satisfies ψ.

Definition 11. *Let M be a CGM, and $-U$ be an update. We call $-U$ executable in M iff for all $i \in A$ and $s \in S$ at least one of the following conditions is true:*

- $|\mathfrak{f}^{-U}(i, s)| \neq 0$
- $\exists \alpha_i \in act(i, s) : \alpha_i \notin \mathfrak{f}(i, s)$

In other words, executability of $-U$ means that for all forcing actions, either we have a clause in $-U$ that allows us to preserve at least one forcing action from a state, or there are other, non-forcing, actions by the agent in the state. This ensures that agents do not run out of actions as a result of an update.

Definition 12. *Let M_s be a CGM. The semantics of DDCL$^-$ extends Definition 3 with the following clause for updates:*

$$M_s \models [-U]\varphi \textit{ iff } -U \textit{ is executable in } M \textit{ implies } M_s^{-U} \models \varphi$$

where $M^{-U} = (A, S, Act^{-U}, act^{-U}, out^{-U}, L)$ is the updated model.

We denote by $\mathfrak{f}^-(i, s)$ the set $\mathfrak{f}(i, s) \setminus \mathfrak{f}^{-U}(i, s)$ of forcing actions of agent i in state s to be removed *from the model. Then function $act^{-U}(i, s) = act(i, s) \setminus \mathfrak{f}^-(i, s)$, and the updated set of executable action profiles is $act^{-U}(s)$. Set Act^{-U} is $\bigcup_{i \in A, s \in S} act^{-U}(i, s)$. Finally, $out^{-U}(\alpha, s)$ is restricted to those action profiles α that are in $act^{-U}(s)$.*

One of the side-effects of considering forcing actions in the setting of ability updates is that if for a given model there are no forcing actions, then all $-U$'s are executable and none of $-U$ has any effect on the model. This result follows directly from the definition of the semantics.

Proposition 2. *Let M be a CGM such that for all $i \in A$ and $s \in S$, $\mathfrak{f}(i,s) = \emptyset$. Then for any $-U$ and $\varphi \in \mathcal{DDCL}^-$ it holds that $M_s \models [-U]\varphi$ iff $M_s \models \varphi$.*

Similarly to DDCL$^+$, we mention some properties of DDCL$^-$.

Proposition 3. *The following holds for formulas of DDCL$^-$.*

1. $\langle -U \rangle \varphi \wedge \langle -U \rangle \psi \leftrightarrow \langle -U \rangle (\varphi \wedge \psi)$ *is valid.*
2. $\langle -U^1 \rangle \varphi \wedge \langle -U^2 \rangle \varphi \to \langle -U^1 \cup -U^2 \rangle \varphi$ *is not valid.*
3. $\langle -U^1 \rangle \langle -U^2 \rangle \varphi \to \langle -U^2 \rangle \langle -U^1 \rangle \varphi$ *is not valid.*

Property 1 states that negative updates are monotonic, while items 2 and 3 say that in general we cannot take a union of negative updates or change the order of their application. The counterexamples can be provided by exploiting the fact that negative updates may become not executable once united or applied in a different order.

4 Expressivity

Definition 13. *Let φ and ψ be formulas. We say that they are* equivalent *if for all M_s it holds that $M_s \models \varphi$ iff $M_s \models \psi$.*

Definition 14. *Let $\mathcal{L}_1$ and $\mathcal{L}_2$ be two languages. We say that $\mathcal{L}_1$* is at least as expressive as *$\mathcal{L}_2$ ($\mathcal{L}_2 \leqslant \mathcal{L}_1$) if and only if for all $\varphi \in \mathcal{L}_2$ there is an equivalent $\psi \in \mathcal{L}_1$. If $\mathcal{L}_1$ is* not *at least as expressive as $\mathcal{L}_2$, we write $\mathcal{L}_2 \nleqslant \mathcal{L}_1$. If $\mathcal{L}_2 \leqslant \mathcal{L}_1$ and $\mathcal{L}_1 \nleqslant \mathcal{L}_2$, we write $\mathcal{L}_2 < \mathcal{L}_1$ and say that $\mathcal{L}_1$* is strictly more expressive than *$\mathcal{L}_2$. Finally, if $\mathcal{L}_1 \nleqslant \mathcal{L}_2$ and $\mathcal{L}_2 \nleqslant \mathcal{L}_1$, we say that $\mathcal{L}_1$ and $\mathcal{L}_2$ are incomparable.*

We first show that both logics, DDCL$^+$ and DDCL$^-$, are strictly more expressive than $\mathcal{CL}$, and thus, in contrast to the situation with AUL [12], positive and negative updates cannot be eliminated.

Proposition 4. *$\mathcal{CL} < \mathcal{DDCL}^+$ and $\mathcal{CL} < \mathcal{DDCL}^-$.*

Proof. The fact that both $\mathcal{DDCL}^+$ and $\mathcal{DDCL}^-$ are at least as expressive as $\mathcal{CL}$ follows from the fact that $\mathcal{CL} \subseteq \mathcal{DDCL}^+$ and $\mathcal{CL} \subseteq \mathcal{DDCL}^-$.

For $\mathcal{DDCL}^+ \nleqslant \mathcal{CL}$, consider models M_s and N_s in Figure 4. There is only one agent a with the only available action a_0. It is immediate that M_s and N_s are bisimilar and thus cannot be distinguished by any formula of $\mathcal{CL}$. At the same time, $\langle (p, a, \neg p)^+ \rangle \langle\!\langle a \rangle\!\rangle \neg p$ is true in N_s (with the resulting updated model N_s^{+U}) and false in M_s (as there are no states where $\neg p$ would hold).

In order to show that $\mathcal{DDCL}^- \nleqslant \mathcal{CL}$, we will use a technical trick that some $-U$'s are not executable in some models. Consider $\langle (p, a, p)^- \rangle \langle\!\langle a \rangle\!\rangle p \in \mathcal{DDCL}^-$,

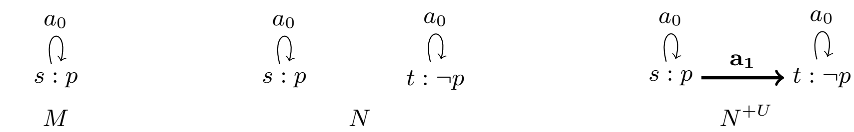

Fig. 4. Models, from left to right, M_s, N_s, and N_s^{+U}. New actions are in bold font.

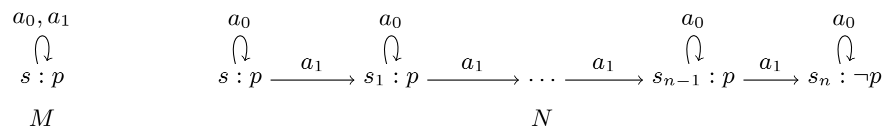

Fig. 5. Models M_s (left) and N_s (right).

and assume towards a contradiction that there is an equivalent $\psi \in \mathcal{CL}$ with $|\psi| = n$.

Consider models M_s and N_s in Figure 5. The former has only one state with a loop, and the latter is a chain of length $n+1$ with $\neg p$ being the case only at the farthest state from s. Although M and N are not bisimilar, it is clear from the construction of the models that both of them satisfy the same ψ up to modal depth n: there is simply not enough modal depth to witness a difference. On the other hand, $M_s \models \langle (p, a, p)^- \rangle \langle\!\langle a \rangle\!\rangle p$ (all the forcing actions remain intact), and $N_s \not\models \langle (p, a, p)^- \rangle \langle\!\langle a \rangle\!\rangle p$. Indeed, since the update requires us to remove all forcing actions that do not conform to $(p, a, p)^-$, we have to remove the loop at the last state s_n, which results in s_n being without any actions, and thus the whole update is not executable in N_s.

Another non-obvious question is whether DDCL$^+$ and DDCL$^-$ are different. We show that it is indeed the case, and, in particular, that the logics are incomparable.

Theorem 2. $\mathcal{DDCL}^- \not\leq \mathcal{DDCL}^+$.

Proof. Consider models M_s and N_s in Figure 6. Observe that they are bisimilar, and thus satisfy the same formulas of $\mathcal{CL}$. Moreover, it can be argued that the models also satisfy the same formulas of $\mathcal{DDCL}^+$. Indeed, we can reason by induction that adding a forcing transition in one model, adds an equivalent forcing transition in the other model. Intuitively, some $+U$ is executable in one model if and only if it is executable in the other model, and no new forcing arrow can take us to a non-bisimilar state.

Updates $-U$ depend, on the other hand, on the sets of forcing actions. For M, the sets of forcing actions are $\mathfrak{f}(a, s) = \{a_0^{(s,s)}\}$, $\mathfrak{f}(b, s) = \{b_0^{(s,s)}\}$, $\mathfrak{f}(a, t) = \{a_0^{(t,t)}\}$, and $\mathfrak{f}(b, t) = \{b_0^{(t,t)}\}$. The thing to notice here is that both agents have forcing actions in both states. For model N, the non-empty sets of forcing actions are $\mathfrak{f}(a, s) = \{a_0^{(s,s)}, a_1^{(s,t)}\}$ and $\mathfrak{f}(a, t) = \{a_0^{(t,t)}, a_1^{(t,s)}\}$. Since there are no forcing actions for agent b, we can exploit it with $-U$'s. In particular, consider $\langle (p, a, p)^- \rangle p$,

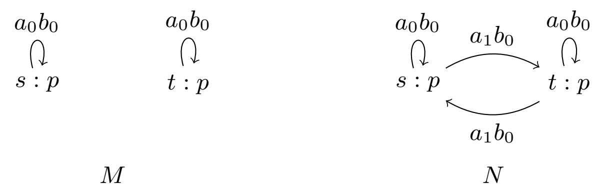

Fig. 6. Models M_s (left) and N_s (right).

which intuitively orders to preserve only a's forcing actions. It is easy to see that $\langle(p, a, p)^-\rangle$ is not executable in M_s, and hence $M_s \not\models \langle(p, a, p)^-\rangle p$. At the same time $N_s \models \langle(p, a, p)^-\rangle p$, since updating N_t with $\langle(p, a, p)^-\rangle$ leaves the model intact.

Theorem 3. *$\mathcal{DDCL}^+ \not\leq \mathcal{DDCL}^-$.*

Proof. Consider models M_s and N_s in Figure 7. Observe that model N is actually a disjoint union of two models. Moreover, M_s and N_s are bisimilar.

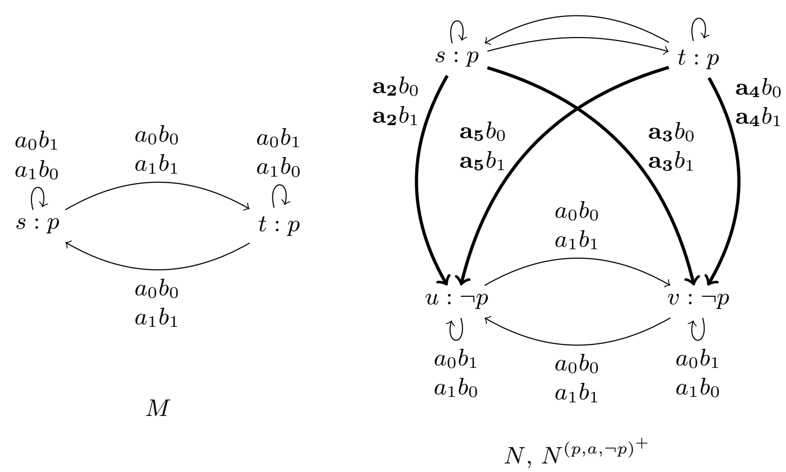

Fig. 7. Models M_s (on the left), N_s (on the right minus thick transitions), and $N_s^{(p,a,\neg p)^+}$ (including thick transitions). Transitions between states s and t in models N and $N^{(p,a,\neg p)^+}$ are exactly like in M, and thus some labels are omitted for readability.

The models are constructed in such a manner that the sets of forcing actions in all states for all agents of both models are empty. Hence, given an arbitrary formula φ of $\mathcal{DDCL}^-$ we can use Proposition 2 to get a translation $t(\varphi)$ into an equivalent formula $\mathcal{CL}$. Finally, since M_s and N_s agree on formulas of $\mathcal{CL}$ we can conclude that they also agree on all formulas of $\mathcal{DDCL}^-$.

Now consider $\langle (p, a, \neg p)^+ \rangle \langle\!\langle a \rangle\!\rangle \neg p \in \mathcal{DDCL}^+$. Since there are no states that satisfy $\neg p$ in M, updating the model with $\langle (p, a, \neg p)^+ \rangle$, which is executable in M, yields exactly the same model. Because there are no $\neg p$-states, we have $M_s \not\models \langle (p, a, \neg p)^+ \rangle \langle\!\langle a \rangle\!\rangle \neg p$.

On the other hand, there are states satisfying $\neg p$ in N, and the update of N with $\langle (p, a, \neg p)^+ \rangle$ is shown in Figure 7 on the right including thick transitions. It is clear that $N_s^{(p,a,\neg p)^+} \models \langle\!\langle a \rangle\!\rangle \neg p$, and hence $N_s \models \langle (p, a, \neg p)^+ \rangle \langle\!\langle a \rangle\!\rangle \neg p$.

5 Discussion

We presented two dynamic extensions CL that allow us to reason about the dynamics of coalitional ability. The first extension, DDCL$^+$, deals with granting dictatorial powers to single agents. The second extension, DDCL$^-$, reasons about revoking dictatorial powers. We showed that both formalisms are strictly more expressive than CL, and that they are mutually incomparable.

Since this work is just the first step towards dynamic coalition logic, there is a plethora of open questions and further research directions. For example, it is not clear how to combine granting and revoking dictatorial powers together in the same update. Apart from that, the next natural step is reasoning about granting powers to coalitions, rather than to single agents, i.e. we will consider $(\chi, C, \psi)^+$ and $(\chi, C, \psi)^-$ in the future. The challenge here is that while we may want to grant a coalition some forcing power, we may also want that none of the members of the coalition has such a power on their own.

Another exciting avenue of further research is comparing the complexities of the model-checking problems of DDCL$^+$ and DDCL$^-$. Yet another immediate next step is providing sound and complete axiomatisations of the logics. However, it seems particularly difficult as axiomatisations of many well-known relation changing logics are still unknown [4, 5]. Finally, there is also a conceptual subtlety worth exploring. In our definition of forcing actions we called an action forcing if for all action profiles it appears in, the outcome state is the same. In other words, forcing actions in our interpretation force *single states*. This is a reasonable interpretation of forcing/dictatorship, but it is not the only one. Another possible interpretation of a forcing action is that the action forces a (not necessarily singleton) set of φ-states. We plan to investigate this notion of the forcing action in the future.

Acknowledgments

We thank the anonymous reviewers of the paper.

References

1. Ågotnes, T., Goranko, V., Jamroga, W.: Alternating-time temporal logics with irrevocable strategies. In: Samet, D. (ed.) Proceedings of the 11th TARK. pp. 15–24 (2007). https://doi.org/10.1145/1324249.1324256

2. Ågotnes, T., Goranko, V., Jamroga, W., Wooldridge, M.: Knowledge and ability. In: van Ditmarsch, H., Halpern, J.Y., van der Hoek, W., Kooi, B. (eds.) Handbook of Epistemic Logic, pp. 543–589. College Publications (2015)
3. Alur, R., Henzinger, T.A., Kupferman, O.: Alternating-time temporal logic. Journal of the ACM **49**, 672–713 (2002). https://doi.org/10.1145/585265.585270
4. Areces, C., Fervari, R., Hoffmann, G.: Relation-changing modal operators. Logic Journal of the IGPL **23**(4), 601–627 (2015). https://doi.org/10.1093/jigpal/jzv020
5. Aucher, G., van Benthem, J., Grossi, D.: Modal logics of sabotage revisited. Journal of Logic and Computation **28**(2), 269–303 (2018). https://doi.org/10.1093/logcom/exx034
6. Bolander, T., Andersen, M.B.: Epistemic planning for single and multi-agent systems. Journal of Applied Non-Classical Logics **21**(1), 9–34 (2011). https://doi.org/10.3166/jancl.21.9-34
7. van Ditmarsch, H., van der Hoek, W., Kooi, B.: Dynamic Epistemic Logic, Synthese Library, vol. 337. Springer (2008)
8. Galimullin, R.: Coalition announcements. Ph.D. thesis, University of Nottingham, UK (2019)
9. Galimullin, R.: Coalition and relativised group announcement logic. Journal of Logic, Language and Information **https://doi.org/10.1007/s10849-020-09327-2** (2021)
10. Goranko, V., Jamroga, W.: Comparing semantics of logics for multi-agent systems. Synthese **139**(2), 241–280 (2004). https://doi.org/10.1023/B:SYNT.0000024915.66183.d1
11. Herlihy, M., Moir, M.: Blockchains and the logic of accountability: Keynote address. In: Grohe, M., Koskinen, E., Shankar, N. (eds.) Proceedings of the 31st LICS. pp. 27–30. ACM (2016). https://doi.org/10.1145/2933575.2934579
12. Kooi, B., Renne, B.: Arrow update logic. Review of Symbolic Logic **4**(4), 536–559 (2011). https://doi.org/10.1017/S1755020311000189
13. Kuijer, L.B.: An arrow-based dynamic logic of norms. In: Gutierrez, J., Mogavero, F., Murano, A., Wooldridge, M. (eds.) Proceedings of the 3rd SR. pp. 1–11 (2015)
14. Pauly, M.: Logic for Social Software. Ph.D. thesis, ILLC, University of Amsterdam, The Netherlands (2001)
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