

Coalitions and Announcements

Thomas Ågotnes
Department of Computer Engineering
Bergen University College
PB. 7030, N-5020 Bergen, Norway
tag@hib.no

Hans van Ditmarsch
Department of Computer Science
University of Otago
PO Box 56, Dunedin 9054, New Zealand
hans@cs.otago.ac.nz
and IRIT, Université Paul Sabatier, France

ABSTRACT

Two currently active strands of research on logics for multi-agent systems are dynamic epistemic logic, focusing on the epistemic consequences of actions, and logics of coalitional ability, focusing on what coalitions of agents can achieve by cooperating strategically. In this paper we make a first attempt to bridge these topics by considering the question: “what can a coalition achieve by public announcements?”. We propose, first, an extension of public announcement logic with constructs of the form $\langle G \rangle \varphi$, where G is a set of agents, with the intuitive meaning that G can jointly make an announcement such that φ will be true afterwards. Second, we consider a setting where all agents can make (truthful) announcements at the same time, and propose a logic with a construct $\langle\!\langle G \rangle\!\rangle \varphi$, meaning that G can jointly make an announcement such that no matter what the other agents announce, φ will be true. The latter logic is closely related to Marc Pauly’s Coalition Logic.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent Systems;
I.2.4 [Knowledge representation formalisms and methods]

General Terms

Theory

Keywords

Dynamic epistemic logic, coalition logic, public announcements

1. INTRODUCTION

Analysis of the dynamics of knowledge have received some attention recently (see [13] for an overview). Van Benthem [11] and Balbiani et al. [3] suggested an interpretation of the standard diamond where $\diamond \varphi$ means “there is an announcement after which φ .” This was in a setting going back to the Fitch-paradox [4]. Fitch observed that “there is an unknown truth” is inconsistent with “all truths are knowable”. Consider some such “ p is true and unknown”, formally $p \wedge \neg Kp$, and the requirement that all truths are knowable, “ $q \rightarrow \diamond Kq$ ”, for *that* truth. From $p \wedge \neg Kp$ and $(p \wedge \neg Kp) \rightarrow \diamond K(p \wedge \neg Kp)$ follows that at some future stage, $K(p \wedge \neg Kp)$ will be true: an inconsistency in a setting where knowl-

edge is at least introspective, because then both Kp and $\neg Kp$ will follow.

The new interpretation of the diamond \diamond in the Fitch setting firstly interprets $\diamond \varphi$ as ‘sometime later, φ ’, and secondly specifies this temporal specification as what may result of a specific event, namely a public announcement: ‘after some announcement, φ ’. In other words, the semantics is: $\diamond \varphi$ is true if and only if $\langle \psi \rangle \varphi$ is true for some ψ ; the expression $\langle \psi \rangle \varphi$ stands for ‘ ψ is true and after ψ is announced, φ is true.’ There are some restrictions on ψ . The resulting arbitrary announcement logic is axiomatisable and has various pleasing properties [3].

In the current contribution we investigate some variants of this arbitrary announcement logic. First, instead of $\diamond \varphi$ we use a more specific operator, namely $\langle G \rangle \varphi$. Here G is a subgroup of all agents that *simultaneously make a public announcement*. In other words, let $G = \{1, \dots, k\}$, then: $\langle G \rangle \varphi$ is true if and only if there exist formulae ψ_1, \dots, ψ_k such that $\langle K_1 \psi_1 \wedge \dots \wedge K_k \psi_k \rangle \varphi$ is true; now, the expression $\langle K_1 \psi_1 \wedge \dots \wedge K_k \psi_k \rangle \varphi$ stands for $K_1 \psi_1 \wedge \dots \wedge K_k \psi_k$ is true and after agents $1, \dots, k$, simultaneously (and truthfully) announce ψ_1, \dots, ψ_k , then φ is true’. Note that the remaining agents, not included in the set G of k agents, are not involved in making the announcement. The resulting logic is called *Group Announcement Logic (GAL)*.

We also introduce a different logic, *Coalition Announcement Logic (CAL)*, with an operator $\langle\!\langle G \rangle\!\rangle$ instead of $\langle G \rangle$ or \diamond . The formula $\langle\!\langle G \rangle\!\rangle \varphi$ is true if and only if the agents in G can simultaneously announce something which, regardless of what the remaining agents simultaneously announce, *still* guarantees that φ then becomes true.

Informally speaking, both $\langle G \rangle \varphi$ and $\langle\!\langle G \rangle\!\rangle \varphi$ expresses the fact that coalition G has the ability to make φ come about, or that G can achieve φ (under different assumptions about what the other agents do). Logics modelling the coalitional abilities of agents have been an active area of research in multi-agent systems in recent years, the most prominent frameworks being Pauly’s Coalition Logic [9] and Alur, Henzinger and Kupferman’s Alternating-time Temporal Logic [2]. The main constructs of these logics are indeed of the form $\langle G \rangle \varphi$ with the intuitive meaning that coalition G can achieve φ . In this paper we investigate these notions when the actions that can be performed are truthful public announcements. We present Group Announcement Logic and Coalition Announcement Logic in Sections 3 and 4, respectively. First, we briefly review the main concepts of public announcement logic and coalition logic.

2. BACKGROUND

The restricted space limits us to a somewhat terse review of key technical definitions; see the references for discussion and details.

2.1 Public Announcement Logic

Cite as: Coalitions and Announcements, Thomas Ågotnes and Hans van Ditmarsch, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. XXX-XXX.
Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

The language \mathcal{L}_{pal} of public announcement logic (PAL) [10] over a set of agents $N = \{1, \dots, n\}$ and a set of primitive propositions Θ is defined as follows, where i is an agent and $p \in \Theta$:

$$\varphi ::= p \mid K_i\varphi \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid [\varphi_1]\varphi_2$$

We write $\langle\varphi_1\rangle\varphi_2$ resp. $\hat{K}_i\varphi$ for the duals $\neg[\varphi_1]\neg\varphi_2$ and $\neg K_i\neg\varphi$.

A *Kripke structure* over N and Θ is a tuple $M = (S, \sim_1, \dots, \sim_n, V)$ where S is a set of states, $\sim_i \subseteq S \times S$ is an epistemic indistinguishability relation and is assumed to be an equivalence relation for each agent i , and $V : \Theta \rightarrow S$ assigns primitive propositions to the states in which they are true. A *pointed Kripke structure* is a pair (M, s) where s is a state in M . The interpretation of formulae in a pointed Kripke structure is defined as follows (the other clauses are defined in usual truth-functional way).

$$M, s \models K_i\varphi \text{ iff for every } t \text{ such that } s \sim_i t, M, t \models \varphi$$

$$M, s \models [\varphi]\psi \text{ iff } M, s \models \varphi \text{ implies that } M|\varphi, s \models \psi$$

where $M|\varphi = (S', \sim'_1, \dots, \sim'_n, V')$ such that $S' = \{s' \in S : M, s' \models \varphi\}$; $\sim'_i = \sim_i \cap (S' \times S')$; $V'(p) = V(p) \cap S'$.

The purely epistemic fragment of the language (i.e., formulae not containing public announcement operators $[\varphi]$) is denoted \mathcal{L}_{el} . It was already shown in Plaza's original publication on that logic [10] that the language of PAL is no more expressive than the purely epistemic fragment.

2.2 Coalition Logic

The language of coalition logic [9] over N and Θ is defined as follows, where $G \subseteq N$ and $p \in \Theta$:¹

$$\varphi ::= p \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \langle G \rangle\varphi$$

An *effectivity function* over a set of states S is a function $E : 2^N \rightarrow 2^{2^S}$, associating to every coalition the sets of states for which the coalition is effective. An effectivity function E is *N-maximal* iff for all X , if $\bar{X} \notin E(\emptyset)$ then $X \in E(N)$; *outcome monotonic* iff for all $X \subseteq X' \subseteq S$, if $X \in E(G)$ then $X' \in E(G)$; *superadditive* iff for all X_1, X_2, G_1, G_2 such that $G_1 \cap G_2 = \emptyset$, $X_1 \in E(G_1)$ and $X_2 \in E(G_2)$ imply that $X_1 \cap X_2 \in E(G_1 \cup G_2)$; *regular* iff for all G and all X , if $X \in E(G)$ then $\bar{X} \notin E(\bar{G})$; *playable* iff (i) for all $G, \emptyset \notin E(G)$; (ii) for all $G, S \in E(G)$; (iii) E is N -maximal; (iv) E is outcome-monotonic; (v) E is superadditive. Playability is a key notion in coalition logic.

A *coalition model* is a tuple $\mathcal{M} = (S, E, V)$, where E associates a playable effectivity function $E(s)$ over S to each state s , and V is a propositional valuation over S . Alternatively, we can view a coalition model as assigning a neighbourhood function $E(G) : S \rightarrow 2^{2^S}$ to each coalition G . The interpretation of a formula in the combination of a coalition model and a state is defined as follows (the other clauses as usual), where $\varphi^{\mathcal{M}} = \{t : \mathcal{M}, t \models \varphi\}$:

$$\mathcal{M}, s \models \langle G \rangle\varphi \text{ iff } \varphi^{\mathcal{M}} \in E(s)(G)$$

3. GROUP ANNOUNCEMENT LOGIC

The main construct of the language of *Group Announcement Logic* (GAL) is $\langle G \rangle\varphi$, intuitively meaning that there is some announcement the group G can truthfully make after which φ will be true.

¹We here use $\langle G \rangle$ where Pauly [9] uses $[G]$; to interpret the diamond rather than the box as inclusion in the neighbourhood, contrary to the common practice in neighbourhood semantics. The reason for this is that it will make the relationship between coalition logic and other logics we discuss here clearer.

3.1 Formal Language and Semantics

Assume a set N of n agents and a set Θ of infinitely many primitive propositions. The language \mathcal{L}_{gal} of GAL is defined by extending the language of PAL with a new operator $\langle G \rangle$ for each coalition G :

$$\varphi ::= p \mid K_i\varphi \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \langle G \rangle\varphi \mid [\varphi_1]\varphi_2$$

where i is an agent, G is a set of agents and $p \in \Theta$. We write $[G]\varphi$ for the dual $\neg\langle G \rangle\neg\varphi$ and $\langle i \rangle\varphi$ for $\langle \{i\} \rangle\varphi$. We will henceforth call the fragment without public announcement operators \mathcal{L}_g .

The interpretation of formulae in a pointed Kripke structure is defined by extending the definition for PAL with the following clause for the new operator:

$$M, s \models \langle G \rangle\varphi \text{ iff for each agent } i \in G \text{ there is a formula } \psi_i \in \mathcal{L}_{el} \text{ such that } M, s \models \langle \bigwedge_{i \in G} K_i\psi_i \rangle\varphi$$

We get the following meaning for the dual:

$$M, s \models [G]\varphi \text{ iff for every formula } \psi_i \in \mathcal{L}_{el} \text{ for every agent } i \in G, M, s \models [\bigwedge_{i \in G} K_i\psi_i]\varphi$$

The definitions can alternatively be written out in more detail:

$$M, s \models \langle G \rangle\varphi \text{ iff for each agent } i \in G \text{ there is a formula } \psi_i \in \mathcal{L}_{el} \text{ such that } M, s \models \bigwedge_{i \in G} K_i\psi_i \text{ and } M|\bigwedge_{i \in G} K_i\psi_i, s \models \varphi$$

$$M, s \models [G]\varphi \text{ iff for every formula } \psi_i \in \mathcal{L}_{el} \text{ for every agent } i \in G, M, s \models \bigwedge_{i \in G} K_i\psi_i \text{ implies that } M|\bigwedge_{i \in G} K_i\psi_i, s \models \varphi$$

Observe that $\langle G \rangle$ quantifies only over purely epistemic formulae. The reason for this is as follows. First, in the semantics of $\langle G \rangle\varphi$ the formulae ψ_i in $\bigwedge_{i \in G} K_i\psi_i$ cannot be unrestricted \mathcal{L}_{gal} formulas, as that would make the definition circular: such a ψ_i could then be the formula $\langle G \rangle\varphi$ itself that we are trying to interpret. We therefore avoid quantifying over formulae containing $\langle G \rangle$ operators. However, as the public announcement logic is equally expressive as the purely epistemic language, the semantics obtained by quantifying over the fragment of the language without $\langle G \rangle$ operators is the same as the semantics obtained by quantifying only over epistemic formulae. In other words, we have that

$$M, s \models \langle G \rangle\varphi \text{ iff for each agent } i \in G \text{ there is a formula } \psi_i \in \mathcal{L}_{pal} \text{ such that } M, s \models \langle \bigwedge_{i \in G} K_i\psi_i \rangle\varphi$$

As usual, a formula φ is *valid*, $\models \varphi$, iff $M, s \models \varphi$ for all M and s .

3.2 Expressivity

GAL includes PAL, but is it more expressive? Or, can $\langle G \rangle\varphi$ already be expressed in PAL? The following theorem shows that the answer to the latter question is “no”.

THEOREM 1. *GAL is strictly more expressive than PAL.*


PROOF. The proof is very similar to the proof that arbitrary public announcement logic is strictly more expressive than PAL [3]. GAL is obviously at least as expressive as PAL. For the strictness part, consider the formula

$$\xi = \langle a, b, c \rangle (K_a p \wedge \neg K_b K_a p)$$

Assume that there is a PAL formula equivalent to (i.e., true in exactly the same states) ξ . Since PAL is no more expressive than standard epistemic logic, there is an \mathcal{L}_{el} formula ψ equivalent to ξ . ψ can only contain a finite number of atoms, so let q be an atom not occurring in ψ . Consider the following model M where agents b and c knows about p but a does not, and where p is currently true:

$$\bullet_1^p \dots \overset{a}{\dots} \dots \bullet_0^{\neg p}$$

Consider furthermore the model M' :



where access for c is the identity relation. It is easy to see that $(M, 1) \not\models \xi$. For the second model, we have that $(M', 10) \models \langle K_a \top \wedge K_b \top \wedge K_c (p \vee q) \rangle (K_a p \wedge \neg K_b K_a p)$, and thus that $(M', 10) \models \xi$. However, $(M, 1)$ and $(M', 10)$ are bisimilar with respect to the epistemic language not including atom q , thus ψ cannot discern between these two pointed models and cannot be equivalent to ξ . \square

3.3 Logical Properties

To sharpen the intuition about the logic we mention some relevant validities, with particular attention to interaction between group announcement and epistemic modal operators.

PROPOSITION 2 (ELEMENTARY VALIDITIES).

1. $\langle G \rangle p \rightarrow p$ and $\langle G \rangle \neg p \rightarrow \neg p$. *Announcements cannot change the truth value of atomic propositions.*
2. $\langle \emptyset \rangle \varphi \leftrightarrow [\emptyset] \varphi \leftrightarrow \varphi$. *The empty group is powerless*
3. $\langle K_{j_1} \psi_{j_1} \wedge \dots \wedge K_{j_k} \psi_{j_k} \rangle \varphi \rightarrow \langle \{j_1, \dots, j_k\} \rangle \varphi$
4. $\varphi \rightarrow \langle G \rangle \varphi$ *(Dual of) truth axiom*

3.3.1 What can be achieved by sequences of announcements?

Intuitively, $\langle G \rangle \varphi$ means that G can achieve a situation where φ is true in “one step”, by making a joint announcement. One can easily imagine situations where it could be interesting to reason about what a group can achieve by making *repeated* announcements, i.e., by a sequence of announcements, one after the other. A general example is a conversation over an open channel. A particularly interesting property is: “there is some sequence, of arbitrary length, of announcements by G which will ensure that φ becomes true”.

For arbitrary announcement logic, the validity of the principle $\Box \varphi \rightarrow \Box \Box \varphi$ follows from the simple observation that a sequence of two announcements ψ and χ is equivalent to the single announcement of $\psi \wedge [\psi] \chi$. Less obvious is that $[G] \varphi \rightarrow [G][G] \varphi$ is also valid.

PROPOSITION 3. $\models [G] \varphi \rightarrow [G][G] \varphi$

PROOF. The diamond version $\langle G \rangle \langle G \rangle \varphi \rightarrow \langle G \rangle \varphi$ of this validity makes clear that the requirement is that two successive announcements by the agents in G simultaneously can also be seen as a single announcement by the agents in G simultaneously. The proof is by a purely *structural* observation. This is surprising: we will reflect after the proof on the difficulties of a more obvious attempt to use the language and its semantics.

Consider two successive announcements $\bigwedge_{i \in G} K_i \varphi_i$ and $\bigwedge_{i \in G} K_i \psi_i$. Let a Kripke structure M and a state s in M be given such that $M, s \models \bigwedge_{i \in G} K_i \varphi_i$, and similarly $\bigwedge_{i \in G} K_i \psi_i$ is true in state s in the restriction of M to the $\bigwedge_{i \in G} K_i \varphi_i$ -states: $M|_{\bigwedge_{i \in G} K_i \varphi_i}, s \models \bigwedge_{i \in G} K_i \psi_i$. Given some agent j , the denotation of $K_j \varphi_j$ is a union N of j -equivalence classes in M that includes the j -class containing actual state s . The partition for agent j in the resulting model $M|_{\bigwedge_{i \in G} K_i \varphi_i}$ is clearly the intersection of the domain of $M|_{\bigwedge_{i \in G} K_i \varphi_i}$ with that union N . Let that new partition be N' . The subsequent announcement $K_j \psi_j$ by agent j must *again* be a union of j -equivalence classes that includes actual j -class containing actual state s , but now in the model $M|_{\bigwedge_{i \in G} K_i \varphi_i}$. This must therefore be

a subset N'' of N' . But that means that agent j could initially already have announced a proposition of which the denotation is a union of j -equivalence classes of which the $M|_{\bigwedge_{i \in G} K_i \varphi_i}$ restriction is that N'' ! In other words, agent j initially already knew something stronger than his first announcement, but he only revealed that in his second announcement. As this holds for all agents, all agents could initially have announced that stronger proposition. Therefore, for each agent i there must be formulas χ_i such that the sequence of announcements $\bigwedge_{i \in G} K_i \varphi_i$ and $\bigwedge_{i \in G} K_i \psi_i$ is the same as a single announcement $\bigwedge_{i \in G} K_i \chi_i$ in the sense that for arbitrary φ : $\langle \bigwedge_{i \in G} K_i \varphi_i \rangle \langle \bigwedge_{i \in G} K_i \psi_i \rangle \varphi$ is equivalent to $\langle \bigwedge_{i \in G} K_i \chi_i \rangle \varphi$. This shows that $\langle G \rangle \langle G \rangle \varphi \leftrightarrow \langle G \rangle \varphi$ as required. \square

The curious part of this proof is that it does not *construct* the required χ_i . It merely shows that they *exist*. It is currently even unclear to us what they are. It may be instructive to elaborate. First, consider two consecutive announcements by a single agent only. Using the validity for public announcements $[\varphi][\psi] \chi \leftrightarrow [\varphi \wedge [\varphi] \psi] \chi$, we then have that for two successive $K_i \varphi$ and $K_i \psi$: $K_i \varphi \wedge [K_i \varphi] K_i \psi$ is equivalent to $K_i \varphi \wedge (K_i \varphi \rightarrow K_i [K_i \varphi] \psi)$ which is equivalent using some further propositional steps to $K_i (\varphi \wedge [K_i \varphi] \psi)$. This has the proper form of an announcement known by i . For the multi-agent case the succession of group announcements $\bigwedge_{i \in G} K_i \varphi_i$ and $\bigwedge_{i \in G} K_i \psi_i$ can be rewritten to (we omit details) $\bigwedge_{i \in G} (K_i \varphi_i \wedge (\bigwedge_{j \in G} K_j \varphi_j \rightarrow K_i [\bigwedge_{j \in G} K_j \varphi_j] \psi_i))$. (The problematic part of this formula, for further reduction, is the $\bigwedge_{j \in G-i} K_j \varphi_j$ part of $\bigwedge_{j \in G} K_j \varphi_j$.) This does not imply an *obvious* formula of form $\bigwedge_{i \in G} K_i \chi_i$.

We thus get yet another interpretation of $\langle G \rangle \varphi$; namely exactly the property alluded to above: “there is some sequence, of arbitrary length, of announcements by G which will ensure that φ becomes true”.

COROLLARY 4. $M, s \models \langle G \rangle \varphi$ iff there is a sequence $\alpha_1, \dots, \alpha_k$ for some $k \geq 1$ where for every $1 \leq j \leq k$, $\alpha_j = \{\psi_i^j : i \in G\}$ is a set containing one formula for each agent in G , such that $M, s \models \langle \bigwedge_{i \in G} K_i \psi_i^1 \rangle \dots \langle \bigwedge_{i \in G} K_i \psi_i^k \rangle \varphi$

In Section 3.5 we discuss a security protocol example involving sequences of announcements.

3.3.2 Interaction axioms

For arbitrary announcement logic we have that $K_i \Box \varphi \rightarrow \Box K_i \varphi$, but not the other way round. Now, we can do more.

PROPOSITION 5. *Validities, for arbitrary i and G :*

1. $K_i [i] \varphi \leftrightarrow [i] K_i \varphi$
2. $K_i [G] \varphi \rightarrow [G] K_i \varphi$ *(but not the other way round)*

PROOF.

1. $\Box K_i \varphi \rightarrow K_i \Box \varphi$ is false because after going to an i -accessible state the subsequent model restriction may exclude the actual state. But for singleton announcements we have an equivalence. This is because any i -accessible state will be contained in the actual i -class.
2. As for arbitrary announcement logic. Note that agent i may but need not be in group G . \square

We now proceed to a more systematic treatment of validities.

3.4 Towards an Axiomatisation

We present a sound axiomatisation of group announcement logic which we also conjecture is complete. The formulation refers to the necessity forms as introduced by Goldblatt [5]. A *necessity form* contains a unique occurrence of a special symbol \sharp . If ψ is

such a necessity form and $\varphi \in \mathcal{L}_{gal}$, then $\psi(\varphi)$ is obtained from ψ by substituting \sharp in ψ for φ . The *necessity forms* are inductively defined as follows. Let $\varphi \in \mathcal{L}_{gal}$. Then: \sharp is a nec. form; if ψ is a nec. form then $(\varphi \rightarrow \psi)$ is a nec. form; if ψ is a nec. form then $[\varphi]\psi$ is a nec. form; if ψ is a nec. form then $K_a\psi$ is a nec. form. P_φ denotes the set of primitive propositions occurring in φ . The axiomatisation **GAL** is given in Table 1.

instant. of prop. tautologies	
$K_a(\varphi \rightarrow \psi) \rightarrow (K_a\varphi \rightarrow K_a\psi)$	distr. of kn. over impl. truth
$K_a\varphi \rightarrow \varphi$	
$K_a\varphi \rightarrow K_aK_a\varphi$	positive introspection
$\neg K_a\varphi \rightarrow K_a\neg K_a\varphi$	negative introspection
$[\varphi]p \leftrightarrow (\varphi \rightarrow p)$	atomic permanence
$[\varphi]\neg\psi \leftrightarrow (\varphi \rightarrow \neg[\varphi]\psi)$	ann. and negation
$[\varphi](\psi \wedge \chi) \leftrightarrow ([\varphi]\psi \wedge [\varphi]\chi)$	ann. and conjunction
$[\varphi]K_a\psi \leftrightarrow (\varphi \rightarrow K_a[\varphi]\psi)$	ann. and knowledge
$[\varphi][\psi]\chi \leftrightarrow [\varphi \wedge [\varphi]\psi]\chi$	ann. composition
$[B]\varphi \rightarrow [\bigwedge_{a \in B} K_a\psi_a]\varphi$ where $\psi_a \in \mathcal{L}_{el}$	group and specific ann.
From φ and $\varphi \rightarrow \psi$, infer ψ	modus ponens
From φ , infer $K_a\varphi$	nec. of knowledge
From φ , infer $[\psi]\varphi$	nec. of announcement
From φ , infer $[B]\varphi$	nec. of arb. ann.
From $\varphi([p]\psi)$, infer $\varphi([B]\psi)$	deriving group ann.
where $p \notin P_\varphi \cup P_\psi$	(a.k.a. $R([B])$)

Table 1: The sound axiomatisation GAL

PROPOSITION 6. *The axiomatisation GAL is sound.*

PROOF. All axioms and rules are sound. For the soundness of the rule $R([B])$, we refer to soundness of the rule $R(\square)$ in [3]. \square

3.5 Example: Security Protocols

In security protocols a sender and a receiver attempt to communicate a secret to each other without an eavesdropper learning it. A very powerful eavesdropper is one that intercepts all communications. This creates the setting where sender, receiver, and eavesdropper are three agents that can be modelled in a multi-S5 system and where all communications are public announcements by sender and receiver. One specific example of such a setting is known as the Russian Cards Problem [12]. The setting is one where a pack of all different cards are distributed over the three ‘players’, where every player only knows his own cards, where sender and receiver have an informational advantage over the eavesdropper because they hold more cards, and where the ‘secrets’ that should not be divulged are card ownership. Posed as a riddle it looks as follows—Anne and Bill are sender and receiver, Cath the eavesdropper:

From a pack of seven known cards $\{0, 1, 2, 3, 4, 5, 6\}$ Anne and Bill each draw three cards and Cath gets the remaining card. How can Anne and Bill openly (publicly) inform each other about their cards, without Cath learning from any of their cards who holds it?

To simplify matters, assume that Anne has drawn $\{0, 1, 2\}$, that Bill has drawn $\{3, 4, 5\}$ and that Cath therefore has card 6. The initial Kripke model describing this setting consists of all possible card deals (valuations). In that model an epistemic class for an agent can be identified with the hand of cards of that agent. For example, given that Anne holds $\{0, 1, 2\}$, she cannot distinguish the four deals—allow us to use some suggestive notation—012.345.6, 012.346.5, 012.356.4, and 012.456.3 from one another.

Given that all announcements that can be made by a player are known by that player, they consist of unions of equivalence classes for that player and can therefore be identified with sets of alternative hands for that player. One solution is where Anne says “My hand of cards is one of 012, 034, 056, 135, 246” and where Bill after that says “My hand of cards is one of 345, 125, 024” (the last is equivalent in that information state to Bill saying “Cath has card 6”). Anne and Bill in fact execute a protocol here, not in the sense of sets of sequences of announcements but in the sense of functions from local states of agents to nondeterministic choice between announcements. For example, Anne is executing “given cards i, j, k , the first of my five hands is that actual hand ijk ; the second of my five hands to announce is ikl where k, l are chosen from the five remaining cards; the third is imn where m, n are the remaining two cards; etc...; shuffle the hands before announcing them.”

We can describe this solution in logic. Agent a stands for Anne, b for Bill, and c for Cath. Let q_i stand for ‘agent i holds card q ’ and let klm_i stand for $k_i \wedge l_i \wedge m_i$. The information requirements are that Anne learns Bill’s cards (the conjunction suggests all hands): $\bigwedge_{ijk}(ijk_b \rightarrow K_aijk_b)$ (*one*); and that Bill learns Anne’s cards: $\bigwedge_{ijk}(ijk_a \rightarrow K_bijk_a)$ (*two*). And the safety / security requirements are that Cath does not learn the ownership of any card (except her own card): $\bigwedge_{q=0}^6((q_a \rightarrow \neg K_cq_a) \wedge (q_b \rightarrow \neg K_cq_b))$ (*three*). All protocols are finite, because the model is finite and all informative announcements result in actual model restriction. But it is unclear how long such protocols need to be. *That* uncertain but finite length cannot be described in public announcement logic. But (overlooking the intermediate requirements for safety for the moment) the *existence* of a protocol can be described in group announcement logic, because the diamond in $\langle ab \rangle \varphi$ may refer to arbitrarily finite length protocols taking place between sender a and receiver b in the presence of other agents, such as the eavesdropper, as was discussed in Section 3.3.1.

Let’s see how this works for the length-two protocol above that solves the Russian Cards Problem. First, we model the solution in public announcement logic. In the solution, first Anne announces $012_a \vee 034_a \vee 056_a \vee 135_a \vee 246_a$ (*anne*). Then Bill announces $345_b \vee 125_b \vee 024_b$ (*bill*). After these two announcements the solution requirements are satisfied.

$$012.345.6 \models \langle K_aanne \rangle \langle K_bbill \rangle (one \wedge two \wedge three)$$

Now instead of Anne and Bill taking turns in speaking, we can also see them as operating simultaneously, where Bill ‘thinks’ true during Anne’s announcement and vice versa. We then get

$$012.345.6 \models \langle K_aanne \wedge K_b \top \rangle \langle K_a \top \wedge K_bbill \rangle (one \wedge two \wedge three)$$

In GAL we thus have

$$012.345.6 \models \langle ab \rangle \langle ab \rangle (one \wedge two \wedge three)$$

The significance of this is that it says that there *exists* an announcement a can make such that after that there *exists* an announcement b can make, after which the goal will be achieved. This is exactly the type of property one would want to model check against the description of a system – one would typically want to check whether there exist *some* successful announcement rather than checking whether a particular announcement is successful.

By Proposition 3 we also get that

$$012.345.6 \models \langle ab \rangle (one \wedge two \wedge three)$$

By successively replacing two announcements by a single announcement we can of course capture any finite-length execution sequence between a sender and receiver in this way as ‘some terminal condi-

tion φ sender and receiver can achieve by collaborating’ expressed in $\langle ab \rangle \varphi$ form.

The above is merely *one* execution sequence. The ability to exchange a secret depends on the ability to produce such a sequence by (typically) random choice between many other protocol executions. Also, the initial state may not be the hand of cards $\{0, 1, 2\}$ but any other hand. For example, a different execution to “My hand of cards is one of 012, 034, 056, 135, 246” is “My hand of cards is one of 012, 035, 046, 134, 256”, to which corresponds a different subsequent *b*-announcement. But that would simply constitute a different way to achieve postcondition $one \wedge two \wedge three$ similarly captured by the truth of $\langle ab \rangle (one \wedge two \wedge three)$ in state 012.345.6 of the initial information state. And conditionalising for Anne’s means shifting to the model perspective: the requirement that “ $\langle ab \rangle (one \wedge two \wedge three)$ is valid in the initial card deal model” corresponds to “there is a protocol for Anne and Bill to safely exchange their secrets”. Clearly, this also holds for protocols consisting of more than two announcements, and for secrets and safety requirements other than $one \wedge two \wedge three$.

One must be careful when interpreting the meaning the existence of sequences of announcements. If we can replace the two successive announcements: Anne says “My hand of cards is one of 012, 034, 056, 135, 246” after which Bill says “My hand of cards is one of 345, 125, 024”, by a single one, does that not mean that all protocols can be reduced to length 1? And what would in this case that single simultaneous announcement be? Well: as both agents are announcing facts and not knowledge, their single announcement is simply the conjunction of their successive announcements. As the second one for Anne and the first one for Bill was ‘true’ (vacuous), this means that they could *simultaneously* have made their successive announcements: Anne says “My hand of cards is one of 012, 034, 056, 135, 246” and simultaneously Bill says “My hand of cards is one of 345, 125, 024”. Unfortunately, even though this indeed solves the problem, the agents do not know the public consequences of their joint action merely from the public consequences of their individual part in it. More concretely: a different execution of the protocol for Anne, when she holds cards $\{0, 1, 2\}$, is the announcement “My hand of cards is one of 012, 035, 046, 134, 256”. From that with Bill’s above announcement Cath can deduce straightaway that the card deal is 012.345.6. And, obviously, Bill does not know whether Anne is going to announce the original or the alternative set of five hands, and any of many others. In epistemic terms: although the knowledge postconditions are met by the above simultaneous announcements, the agents do not know that before the announcement, and consider it possible that an infelicitous combination instead may occur. The subtleties of ability under incomplete information are further discussed in Section 3.6 below.

We conclude, that $\langle ab \rangle (one \wedge two \wedge three)$ is a necessary but not sufficient property to ensure the existence of a safe protocol to realize security postconditions. But this is already an achievement: we do not know of a logical language that can express this so succinctly. In public announcement logic, for a *finite* model we would need a very long, deeply nested modal formula; where ‘deep’ is really deep: in the order of the number of states of the model. For infinite models, it cannot be expressed at all in that logic.


3.6 Announcements and Ability

Our initial intuitive interpretation of a formula of the form $\langle C \rangle \varphi$ was that coalition C has the *ability* to make φ come about by making some public announcement. We now have a better understanding of group announcement logic; let us discuss to what extent that intuition is precise.

Recent work on strategy logics have illuminated the fact that there are many subtly different notions of ability in the context of incomplete information [6, 8, 1, 7] (see [7] for a recent summary). [7, p. 433] discuss three levels of ability in general strategy logics, which we now discuss counterparts to in the special context of truthful public announcements². For simplicity we consider a singleton coalition $\{a\}$. What does it mean that agent a has the ability to make a goal φ come about by making a public announcement? Let us begin with the weakest form of ability.

Being able to, but not necessarily knowing it.


The formula $\langle a \rangle \varphi$ means that there is something which a knows, and if the fact that a knows it is announced, φ is a consequence. However, it might be the case that a *doesn’t know this*, i.e., that $K_a \langle a \rangle \varphi$ is not true. As an example, first observe that $\langle K_a \psi_a \rangle \varphi \rightarrow K_a \langle K_a \psi_a \rangle \varphi$ is not a principle of public announcement logic. As a counter-example take state s of the following model



and take $\psi_a = \top$ and $\varphi = p$. However, this does not mean that a cannot achieve φ in all her accessible states by some *other* announcements (possibly different ones in different states). But in group announcement logic, we have in the model above that $s \models \langle a \rangle p$ (a can announce $K_a \top$), but $t \not\models \langle a \rangle p$ and thus, $s \models \neg K_a \langle a \rangle p$. So, $\langle a \rangle \varphi \rightarrow K_a \langle a \rangle \varphi$ is not a principle of group announcement logic. This is a first illustration of the fact that we must be careful when using the term “ability”: in some (but not necessarily all) circumstances it might be counter-intuitive to say that a has the ability make φ come about, when she is not aware that she is; when she cannot discern between the actual situation and a situation in which she does not have this ability.

Being able to, knowing that, but not knowing how.

Consider the following model and formula:



$$\varphi = K_2 q \wedge (\neg K_2 p \vee \hat{K}_1 (K_2 p \wedge \neg K_2 q))$$

If we take the current state to be s , we have a situation where 1 is able to make φ come about and where she in addition knows this; a stronger type of ability than in the example above. Formally: $s \models \langle 1 \rangle \varphi$, because $s \models \langle K_1 q \rangle \varphi$, and $t \models \langle 1 \rangle \varphi$ because $t \models \langle K_1 p \rangle \varphi$. Thus, $s \models K_1 \langle 1 \rangle \varphi$. However, we argue, it might still be counter-intuitive to say that 1 can make φ come about in this situation. The reason is that she has to use *different announcements in indiscernible states*. Observe that $s \models \langle K_1 p \rangle \neg \varphi$ and $t \models \langle K_1 q \rangle \neg \varphi$: while the same announcements can be made in both states, they don’t have the same consequences. In fact, there exists *no* single announcement agent 1 can make which will ensure that φ will be true in both s and t . To see this, we can enumerate the possible models resulting from 1 making an announcement in s or t . Because such a model must include 1’s equivalence class $\{s, t\}$, there are four possibilities. First, the starting model itself (e.g., 1 announces a tautology), in which φ does not hold in s . Second, the model where only state u is removed (e.g., 1 announces $K_1 p$), in

²This special case is considerably different from the case of arbitrary actions which is normally studied in the context of strategy logics. In particular, the state transition for a given action (announcement) is *deterministic*, which implies, e.g., that the question of whether strategies are *uniform* or not is moot.

which φ does not hold in s (as we saw above). Third, the model where only state v is removed (e.g., 1 announces K_1q), in which φ does not hold in t (as we saw above). Fourth, the model where both u and v are removed, in which φ holds in neither s nor t .

Since agent 1 cannot discern state s from state t , she has the ability to make φ come about only in the sense that she depends on guessing the correct announcement. In other words, she can make φ come about, knows that she can make φ come about, but does not know *how* to make φ come about.

Being able to, knowing that, knowing how.

Thus, we can formulate a strong notion of the ability of a to achieve φ by public announcements: there exists a formula ψ such that a knows ψ and in any state a considers possible, $\langle K_a\psi_a \rangle\varphi$ holds.

Compare this version of ability, “there is an announcement which a knows will achieve the goal”, with the previous version above, “ a knows that there is an announcement which will achieve the goal”. We can call these notions *knowing de re*, and *knowing de dicto*, respectively, that the goal can be achieved, following [8] who use the same terminology (after the corresponding notion used in quantified modal logic) for general strategy logics. Note, however, that it is not *prima facie* clear that there is a distinction between these notions in GAL, because of the intimate interaction between knowledge and possible actions (announcements), but the model and formula above show that there indeed is.

It can be argued that all of these notions of ability are *useful*, but it is of importance to discern between them.

More than one agent.

In the case of more than one agent, there are even more subtleties. In particular, what does it mean that a group know how to achieve something, i.e., which joint announcement will be effective? That everybody knows it? That they have common knowledge of it?

In [8] it is argued that the answer depends on the situation. It might be the case that the agents have common knowledge (although they then need some resolution mechanism for cases when there are more than one effective announcement, in order to coordinate); that every agent knows the effective announcement; that the agents have distributed knowledge about the effective announcement and thus can pool their knowledge together to find out what they should do; that a particular agent (the “leader”) knows the effective announcement and can communicate it to the others.


3.7 Coalition Logic and Neighbourhood Sem.

We now discuss the relationship between GAL and coalition logic, and neighbourhood semantics for GAL in general.

PROPOSITION 7. *When $G_1 \cap G_2 = \emptyset$, the following is not valid:*

$$(S) \quad (\langle G_1 \rangle\varphi_1 \wedge \langle G_2 \rangle\varphi_2) \rightarrow \langle G_1 \cup G_2 \rangle(\varphi_1 \wedge \varphi_2)$$

PROOF. Let M be the following model.



We have that $M, s \models \langle 1 \rangle(K_2p \wedge \neg K_1q)$ (1 can announce K_1p), and that $M, s \models \langle 2 \rangle(K_1q \wedge \neg K_2p)$ (2 can announce K_2q), but not that $\langle 1, 2 \rangle(K_2p \wedge \neg K_1q \wedge K_1q \wedge \neg K_2p)$. \square

The following is immediate, as coalition logic contains (S).

COROLLARY 8. *Coalition logic is not contained in group announcement logic.*

Let us now discuss the meaning of the group announcement logic operator $\langle G \rangle$ in the context of neighbourhood semantics. For simplicity we will here only consider the language without the regular

public announcement operators: for the rest of this section the language \mathcal{L}_g is assumed.

In order to be able to interpret the epistemic modalities, we must extend coalition models with indistinguishability relations. An *epistemic coalition model (ECM)* is a tuple $\mathcal{M} = (S, E, V, \sim_1, \dots, \sim_n)$, where (S, E, V) is a coalition model and each \sim_i is an epistemic indistinguishability relation. We can use ECMs to interpret our coalition operators exactly as in coalition logic:

$$\mathcal{M}, s \models_e \langle G \rangle\varphi \Leftrightarrow \varphi^{\mathcal{M}} \in E(s)(G)$$

(and as usual for the other operators; we use the subscript on \models_e to discern between the two interpretations). We can thus say that a pointed Kripke structure (M, s) and a pointed ECM (\mathcal{M}, s') are *equivalent* iff they satisfy the same formulae, i.e., if for all φ ,

$$M, s \models \varphi \Leftrightarrow \mathcal{M}, s' \models_e \varphi \quad (1)$$

COROLLARY 9. *It is not the case that for every pointed Kripke structure there is an equivalent pointed epistemic coalition model.*

PROOF. It is easy to see that if the (S) axiom does not hold in a pointed epistemic effectivity model, then the underlying effectivity function is not superadditive. \square

While we cannot use playable effectivity functions to interpret GAL formulae, let us see if we can relax the playability requirements to obtain an equivalent neighbourhood semantics. Let a *general epistemic coalition model (GECM)* be defined as a ECM, except that the effectivity function does not need to be playable.

First, let us look at equivalence on the model level: we say that a Kripke structure M and a GECM \mathcal{M} are equivalent when $M \models \varphi$ iff $\mathcal{M} \models_e \varphi$ (i.e., when $M, s \models \varphi$ for all states s in M iff $\mathcal{M}, s' \models_e \varphi$ for all states s' in \mathcal{M}) for all formulae φ .

PROPOSITION 10. *There does not exist an equivalent general epistemic coalition model for every Kripke structure.*

PROOF. Take a Kripke model M with two states s, t ; p is true only in s ; agent 2 can discern between s and t while agent 1 cannot. Let $\gamma = \langle 2 \rangle(K_1p \vee K_1\neg p)$. Suppose that the epistemic effectivity model \mathcal{M} is equivalent to M . We have that $M, s \models \gamma$ (2 can announce K_2p) and that $M, t \models \gamma$ (2 can announce $K_2\neg p$) and thus that $M \models \gamma$ and, from equivalence, that $\mathcal{M} \models_e \gamma$. Let s' be an arbitrary state in \mathcal{M} . Since $\mathcal{M}, s' \models_e \gamma$, we have that $(K_1p \vee K_1\neg p)^{\mathcal{M}} \in E(s')(2)$. But we also have that $M \models \neg(K_1p \vee K_1\neg p)$, so $\mathcal{M} \models_e \neg(K_1p \vee K_1\neg p)$ and $(K_1p \vee K_1\neg p)^{\mathcal{M}} = \emptyset \in E(s')(2)$. But we also have that $\perp^{\mathcal{M}} = \emptyset$, so $\mathcal{M}, s' \models_e \langle 2 \rangle\perp$. Since s' was arbitrary, we have that $\mathcal{M} \models_e \langle 2 \rangle\perp$, which again implies that $M \models \langle 2 \rangle\perp$ which is not the case. Thus, M and \mathcal{M} cannot be equivalent. \square

A consequence of Proposition 10 is that given a Kripke model M , there does not always exist a GECM \mathcal{M} such that for every state s in M there exists a state s' in \mathcal{M} such that (M, s) and (\mathcal{M}, s') satisfy exactly the same formulae, and vice versa (if that was the case, then M and \mathcal{M} would also be equivalent on the model level). In particular, we cannot in general define an effectivity function (with *any* properties) over the state space of a given Kripke model and get an equivalent interpretation of formulae at all the states.

As the counter-example in the proof of Proposition 10 shows, the Kripke structure might not contain “enough” states. But in fact, any Kripke structure can be *extended* in a very simple way to a structure over which we can define an equivalent effectivity function, without changing satisfiability at any of the original states. Simply take the *power model* consisting of the union of all subsets of the

original Kripke model. Formally, the power model \hat{M} of a Kripke model M is defined as follows. A *submodel* of a Kripke model M is a model where the states are a subset of the states in M , and the valuation function and indistinguishability relations are restrictions of those in M to the state space of the submodel. The power model \hat{M} of M is obtained by taking the disjoint union of M and every proper submodel of M , after renaming the states of the proper submodels such that the state spaces are disjoint. A simple example of a power model is shown in Figure 1. We say that two pointed Kripke models are equivalent iff they satisfy the same formulae.




Figure 1: A Kripke model M and its power model \hat{M}

PROPOSITION 11. For any state s of any Kripke structure M , (M, s) and (\hat{M}, s) are equivalent.

PROOF. Immediate, since there is no access between the disjoint subsets in the power model. \square

Given a Kripke structure M , we define the induced GEKM

$$\mathcal{M}^M = (\hat{S}, E, \hat{V}, \sim_1, \dots, \sim_n)$$

such that $\hat{M} = (\hat{S}, \hat{V}, \sim_1, \dots, \sim_n)$ is the power model of M and

$$X \in E(s)(G) \Leftrightarrow \exists \varphi : \varphi^{\hat{M}} \subseteq X \text{ and } \hat{M}, s \models \langle G \rangle \varphi$$

for any G and $s \in \hat{S}$.

THEOREM 12. For any M and state s in \hat{M} , (\hat{M}, s) and (\mathcal{M}^M, s) are equivalent.

Thus, while we cannot get equivalence on the model level even with general effectivity functions (Prop. 10), we can get equivalence on the level of pointed models: given a pointed Kripke model (M, s) , the pointed GEKM (\mathcal{M}^M, s) satisfies exactly the same formulae.

EXAMPLE 13. Let M be the model in Figure 1. We have that

$$E(s)(2) = \{ X \subseteq \{s, t, s_1, t_1\} : s \in X \text{ or } s_1 \in X \}$$

In other words: in s 2 can make an announcement such that any formula true in s or s_1 (and maybe in some other states as well) is still true after the announcement. We have for example in s that agent 2 can make an announcement such that 1 gets to know p . In the two different semantics: $M, s \models K_2 p$ and $M | K_2 p, s \models K_1 p \Rightarrow M, s \models \langle 2 \rangle K_1 p$; and $\{s_1\} \in E(s)(2) \Rightarrow (K_1 p)^{\mathcal{M}^M} \in E(s)(2) \Rightarrow \mathcal{M}^M, s \models_e \langle 2 \rangle K_1 p$. Here are some other examples of members of $E(s)(2)$ with corresponding announcements (by agent 2) and a formula φ true after the announcement such that $\varphi^{\hat{M}}$ is the set:

Set	Announcement	Formula
$\{s\}$	\top	$p \wedge \neg K_1 p$
$\{s, s_1\}$	\top	$K_2 p$
$\{s, t_1\}$	\top	$p \wedge \neg K_1 p \vee \neg p \wedge K_1 \neg p$
$\{s_1\}$	$K_2 p$	$K_1 p$
$\{s_1, t_1\}$	$K_2 p$	$K_1 p \vee K_1 \neg p$
...		

Let us consider the properties of the induced effectivity function. We call an effectivity function *recursive* if it is the case that a coalition G is effective for a set of outcomes X if and only if it is effective for the set of outcomes in which it is effective for X : $X \in E(s)(G)$ iff $\{s' \mid X \in E(s')(G)\} \in E(s)(G)$.

PROPOSITION 14. The effectivity function of the induced GEKM is N -maximal, outcome-monotonic and recursive.

4. COALITION ANNOUNCEMENT LOGIC

In group announcement logic $\langle G \rangle \varphi$ means that G can jointly make some announcement such that φ will be true. We now introduce and briefly discuss *Coalition Announcement Logic (CAL)*, in which $\llbracket G \rrbracket \varphi$ means that G can jointly make some announcement such that no matter what announcements the other agents make, φ will become true.

4.1 Language and Semantics

The language \mathcal{L}_{cal} of CAL over a set N of n agents and a set Θ of primitive propositions is defined by the following grammar:

$$\varphi ::= p \mid K_i \varphi \mid \neg \varphi \mid \varphi_1 \wedge \varphi_2 \mid \llbracket G \rrbracket \varphi \mid [\varphi_1] \varphi_2$$

where G is a set of agents and $p \in \Theta$. We use \mathcal{L}_c to refer to the fragment without public announcement operators. The semantics is defined as follows (the other clauses as before).

$$M, s \models \llbracket G \rrbracket \varphi \text{ iff for every agent } i \in G \text{ there exists a formula } \psi_i \in \mathcal{L}_{el} \text{ such that for every formula } \psi_j \in \mathcal{L}_{el} \text{ for each of the agents } j \notin G \text{ we have that } M, s \models \bigwedge_{i \in G} K_i \psi_i \wedge [K_1 \psi_1 \wedge \dots \wedge K_n \psi_n] \varphi$$

$\llbracket G \rrbracket$ denotes a $\exists \forall$ pattern of quantifiers. Note that in the definition the second quantifier is over all possible formulae for agents outside G , but the use of the ‘‘box’’ version of the public announcement operator ensures that only the formulae actually known by those agents plays a role. To understand the semantics better it might be instructive to look at the dual $\llbracket G \rrbracket \varphi \equiv \neg \llbracket G \rrbracket \neg \varphi$:

$$M, s \models \llbracket G \rrbracket \varphi \text{ iff for all formulae } \psi_i \in \mathcal{L}_{el} \text{ for every agent } i \in G \text{ there is a formula } \psi_j \in \mathcal{L}_{el} \text{ for each of the agents } j \notin G, \text{ such that } M, s \models \bigwedge_{i \in G} K_i \psi_i \text{ implies that } M, s \models \langle K_1 \psi_1 \wedge \dots \wedge K_n \psi_n \rangle \varphi$$

Thus, $\llbracket G \rrbracket \varphi$ means that coalition G cannot avoid φ – no matter what she (truthfully) announces, the rest of the agents can choose some announcement such that φ will be true.

The formula $\llbracket G \rrbracket \varphi$ cannot be expressed in PAL:

THEOREM 15. CAL is strictly more expressive than PAL.

PROOF. With $N = \{a, b, c\}$, observing that $M, s \models \llbracket N \rrbracket \varphi$ iff $M, s \models \langle N \rangle \varphi$, the proof of Theorem 1 also proves this case. \square

4.2 Properties

If the grand coalition can make an announcement after which φ follows, then it has the ability to achieve φ :

$$(PAN) \quad \langle K_1 \psi_1 \wedge \dots \wedge K_n \psi_n \rangle \varphi \rightarrow \llbracket N \rrbracket \varphi$$

is a valid schema. To see this, observe the meaning of $\llbracket N \rrbracket$:

$$M, s \models \llbracket N \rrbracket \varphi \text{ iff there exist formulae } \psi_1, \dots, \psi_n \in \mathcal{L}_{el} \text{ such that } M, s \models \langle K_1 \psi_1 \wedge \dots \wedge K_n \psi_n \rangle \varphi$$

Note that (PAN) does not hold in general if we substitute other coalitions for N .

If the empty coalition has the ability to achieve φ , then φ will be true after any public announcement by the grand coalition:

$$(PA\emptyset) \quad \llbracket \emptyset \rrbracket \varphi \rightarrow [K_1 \psi_1 \wedge \dots \wedge K_n \psi_n] \varphi$$

is a valid schema. To see this, observe the meaning of $\llbracket \emptyset \rrbracket$:

$$M, s \models \llbracket \emptyset \rrbracket \varphi \text{ iff for all formulae } \psi_1, \dots, \psi_n \in \mathcal{L}_{el}, M, s \models [K_1 \psi_1 \wedge \dots \wedge K_n \psi_n] \varphi$$

Here is one interaction axiom relating coalitional and epistemic operators:

PROPOSITION 16. $\langle\!\langle G \rangle\!\rangle \hat{K}_i \varphi \rightarrow \hat{K}_i \langle\!\langle G \rangle\!\rangle \varphi$ is valid, for any i and G

The properties mentioned so far relate ability to public announcements and to knowledge. Moving on to more general properties of ability, we have the following axioms and rules of coalition logic [9], all valid and validity preserving:

$$\begin{array}{l} (\perp) \neg \langle\!\langle G \rangle\!\rangle \perp \quad (\top) \langle\!\langle G \rangle\!\rangle \top \\ (N) (\neg \langle\!\langle \emptyset \rangle\!\rangle \neg \varphi \rightarrow \langle\!\langle N \rangle\!\rangle \varphi) \quad (M) \langle\!\langle G \rangle\!\rangle (\varphi \wedge \psi) \rightarrow \langle\!\langle G \rangle\!\rangle \psi \\ (S) (\langle\!\langle G_1 \rangle\!\rangle \varphi_1 \wedge \langle\!\langle G_2 \rangle\!\rangle \varphi_2) \rightarrow \langle\!\langle G_1 \cup G_2 \rangle\!\rangle (\varphi_1 \wedge \varphi_2) \quad G_1 \cap G_2 = \emptyset \\ (E) \frac{\varphi \leftrightarrow \psi}{\langle\!\langle G \rangle\!\rangle \varphi \leftrightarrow \langle\!\langle G \rangle\!\rangle \psi} \quad (MP) \frac{\varphi \cdot \varphi \rightarrow \psi}{\psi} \end{array}$$

So coalition logic is embedded in coalitional public announcement logic; every theorem of the former is a theorem of the latter. In addition, unlike in coalition logic we have the following (valid):

$$(P) \quad \langle\!\langle G \rangle\!\rangle p \leftrightarrow p$$

We immediately get a lower bound on the complexity of the satisfiability problem for CAL.

THEOREM 17. *The sat. problem for CAL is PSPACE-hard.*

PROOF. Follows immediately from the fact the satisfiability problem for coalition logic is PSPACE-hard [9] and can be reduced to satisfiability in CAL since coalition logic is included in CAL. \square

4.3 Example

For a brief example, consider once more the model M in Figure 1, where agent 1 cannot distinguish between p and not p , but agent 2 can, and where this is common knowledge. Suppose that p is actually true. In CAL, we now have that 1 cannot guarantee either knowledge of p or persisting ignorance of p . Clearly, agent 1 cannot make any informative announcement, but agent 2 can choose whether just to announce ‘true’ or to announce p . In the first case, 1’s ignorance of p persists, whereas in the second case, 1 will learn that p . In other words, agent 1 has hardly any ‘powers’ at all. Formally, we have that $M, s \not\models \langle\!\langle 1 \rangle\!\rangle K_1 p$ because $M, s \models \langle\!\langle K_1 \top \wedge K_2 \top \rangle\!\rangle \neg K_1 p$ but also that $M, s \not\models \langle\!\langle 1 \rangle\!\rangle (\hat{K}_1 p \wedge \hat{K}_1 \neg p)$ because $M, s \models \langle\!\langle K_1 \top \wedge K_2 p \rangle\!\rangle K_1 p$. On the other hand, agent 2’s powers are, of course, much larger. He can decide whether to let agent 1 learn p or not, a feat they obviously also can achieve together; therefore $M \models \langle\!\langle 2 \rangle\!\rangle (K_1 p \vee K_1 \neg p)$ as well as $M \models \langle\!\langle 1, 2 \rangle\!\rangle (K_1 p \vee K_1 \neg p)$.

The security protocol example, when applied to the Russian Cards Problem, is another application of coalition announcement logic – but for this specific application in a less interesting sense: obviously, whatever information eavesdropper Cath publicly provides will be to her disadvantage. Therefore, whatever sender and receiver (Anne and Bill) can achieve with Cath keeping quiet although still listening (as in GAL) they will also be able to achieve with Cath simultaneously announcing at will (as in CAL). In other words, we still have that $012.345.6 \models \langle\!\langle ab \rangle\!\rangle (one \wedge two \wedge three)$.

4.4 Neighbourhood Semantics

Neighbourhood semantics is particularly interesting for CAL, because of the close connection to coalition logic. Again, we will here only consider the language \mathcal{L}_c without regular public announcement operators.

The equivalent of Proposition 11 still holds CAL:

$$M, s \models \varphi \Leftrightarrow \hat{M}, s \models \varphi$$

for any state s in M and any φ , where \hat{M} is the power model of M . In this case we define the induced general epistemic coalition model of a given Kripke structure M to be $\mathcal{M}^M = (\hat{S}, E, \hat{V}, \sim_1, \dots, \sim_n)$

where $\hat{M} = (\hat{S}, \hat{V}, \sim_1, \dots, \sim_n)$ is the power model of M and

$$X \in E(s)(G) \Leftrightarrow \begin{cases} \exists \varphi : \varphi^{\hat{M}} \subseteq X \text{ and } \hat{M}, s \models \langle\!\langle G \rangle\!\rangle \varphi & G \neq N \\ \forall \varphi : (\varphi^{\hat{M}} \subseteq \hat{S} \setminus X \Rightarrow \hat{M}, s \not\models \langle\!\langle \emptyset \rangle\!\rangle \varphi) & G = N \end{cases}$$

for any G and $s \in \hat{S}$. We can now prove the equivalent to Theorem 12, and consider the properties of the induced effectivity function.

THEOREM 18. *For any M and state s in \hat{M} , (\hat{M}, s) and (\mathcal{M}^M, s) are equivalent.*

PROPOSITION 19. *The effectivity function in the induced GEKM is playable.*

5. CONCLUSIONS

Our goal has been to shed some light on the notion of coalitional ability by public announcements. Clearly, there are plenty of open questions. Of particular interests are complete axiomatisations of the two logics, and a formal comparison of their relative expressiveness, as well as a comparison with arbitrary announcement logic [3] – in particular in regards to expressiveness. We have seen that each of the two logics can be interpreted using general epistemic coalition models – complete descriptions of the properties of the two classes of such models are however missing. The neighbourhood semantics can be extended to the language including the public announcement operators, but that would require some kind of identification of the elements of the effectivity function with actual announcements. We note that a similar construction can be used to get a neighbourhood semantics for arbitrary announcement logic. Issues of computational complexity should also be studied. Another future direction is to add common knowledge.

6. REFERENCES

- [1] T. Ågotnes. Action and knowledge in alternating-time temporal logic. *Synthese (KRA)*, 149(2):377–409, 2006.
- [2] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *JACM*, 49:672–713, 2002.
- [3] P. Balbiani, A. Baltag, H. van Ditmarsch, A. Herzig, T. Hoshi, and T. D. Lima. What can we achieve by arbitrary announcements? A dynamic take on Fitch’s knowability. In *the Proceedings of TARK XI*, pages 42–51, 2007.
- [4] B. Brogaard and J. Salerno. Fitch’s paradox of knowability. In E. Zalta, editor, *The Stanford Encycl. of Philosophy*. 2004.
- [5] R. Goldblatt. *Axiomatising the Logic of Computer Programming*. Springer-Verlag, 1982.
- [6] W. Jamroga. Some remarks on alternating temporal epistemic logic. In *Proc. FAMAS’03*, pages 133–140, 2003.
- [7] W. Jamroga and T. Ågotnes. Constructive knowledge: what agents can achieve under imperfect information. *Journal of Applied Non-Classical Logics*, 17(4):423–475, 2007.
- [8] W. Jamroga and W. van der Hoek. Agents that know how to play. *Fundamenta Informaticae*, 63:185–219, 2004.
- [9] M. Pauly. A modal logic for coalitional power in games. *Journal of Logic and Computation*, 12(1):149–166, 2002.
- [10] J. Plaza. Logics of public communications. In *Proc. of the 4th Intern. Symp. on Methodologies for Intel. Systems*, pp. 201–216. Oak Ridge National Laboratory, 1989.
- [11] J. van Benthem. What one may come to know. *Analysis*, 64(2):95–105, 2004.
- [12] H. van Ditmarsch. The russian cards problem. *Studia Logica*, 75:31–62, 2003.
- [13] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*. Springer, 2007.