No Finite Model Property for Logics of Quantified Announcements

Hans VAN DITMARSCH[†] Tim FRENCH[††]
Rustam GALIMULLIN[†††]

[†] Open University of the Netherlands [††] University of Western Australia
[†††] University of Bergen

TARK 2021 (online), 25–27 June

Structure of the paper

The talk follows the classic story-telling device featured prominently in popular films of the 90s and early 00s.

We are shown the outcome. How did we get here?

Structure of the paper

No Finite Model Property for Logics of Quantified Annoucements

- What are Logics of Quantified Annoucements?
- What is the Finite Model Property? Why is it interesting?

PAL: The Existence of Santa

The Existence of Santa

Three children — Alice (a), Bobby (b), and Claire (c) — are arguing about the existence of Santa Claus. Alice knows that Santa exists (p), while Bobby and Claire consider it possible that Santa does not exist ($\neg p$). Also, children do not know about each other's beliefs.

The Existence of Santa

Three children — Alice (a), Bobby (b), and Claire (c) — are arguing about the existence of Santa Claus. Alice knows that Santa exists (p), while Bobby and Claire consider it possible that Santa does not exist ($\neg p$). Also, children do not know about each other's beliefs.

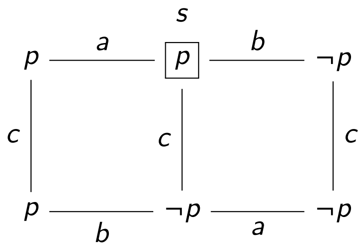

$M_s \models p$

$M_s \models \Box_a p \land \neg \Box_b p$

$M_s \not\models \Box_c p$

$M_s \models \Diamond_b \neg p$

The Existence of Santa

Three children — Alice (a), Bobby (b), and Claire (c) — are arguing about the existence of Santa Claus. Alice knows that Santa exists (p), while Bobby and Claire consider it possible that Santa does not exist ($\neg p$). Also, children do not know about each other's beliefs.

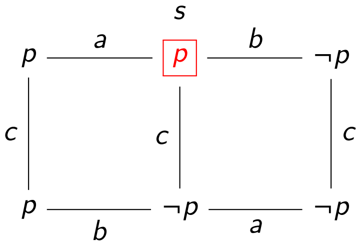

$M_s \models p$

$M_s \models \Box_a p \wedge \neg \Box_b p$

$M_s \not\models \Box_c p$

$M_s \models \Diamond_b \neg p$

PAL: The Existence of Santa

The Existence of Santa

Three children — Alice (a), Bobby (b), and Claire (c) — are arguing about the existence of Santa Claus. Alice knows that Santa exists (p), while Bobby and Claire consider it possible that Santa does not exist ($\neg p$). Also, children do not know about each other's beliefs.

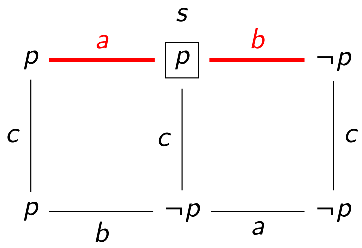

$M_s \models p$

$M_s \models \Box_a p \wedge \neg \Box_b p$

$M_s \not\models \Box_c p$

$M_s \models \Diamond_b \neg p$

PAL: The Existence of Santa

The Existence of Santa

Three children — Alice (a), Bobby (b), and Claire (c) — are arguing about the existence of Santa Claus. Alice knows that Santa exists (p), while Bobby and Claire consider it possible that Santa does not exist ($\neg p$). Also, children do not know about each other's beliefs.

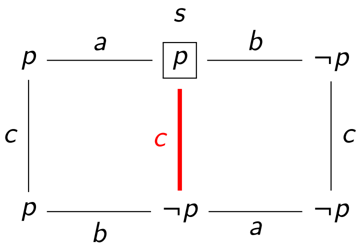

$M_s \models p$

$M_s \models \Box_a p \wedge \neg \Box_b p$

$M_s \not\models \Box_c p$

$M_s \models \Diamond_b \neg p$

The Existence of Santa

Three children — Alice (a), Bobby (b), and Claire (c) — are arguing about the existence of Santa Claus. Alice knows that Santa exists (p), while Bobby and Claire consider it possible that Santa does not exist ($\neg p$). Also, children do not know about each other's beliefs.

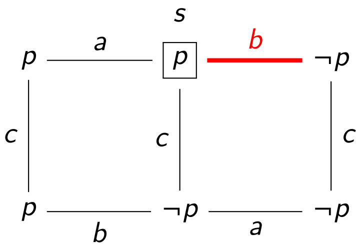

$M_s \models p$
$M_s \models \Box_a p \wedge \neg \Box_b p$
$M_s \not\models \Box_c p$
$M_s \models \Diamond_b \neg p$

What agents know may change once new fact becomes known.
$\psi := \neg p \rightarrow \Diamond_c p$

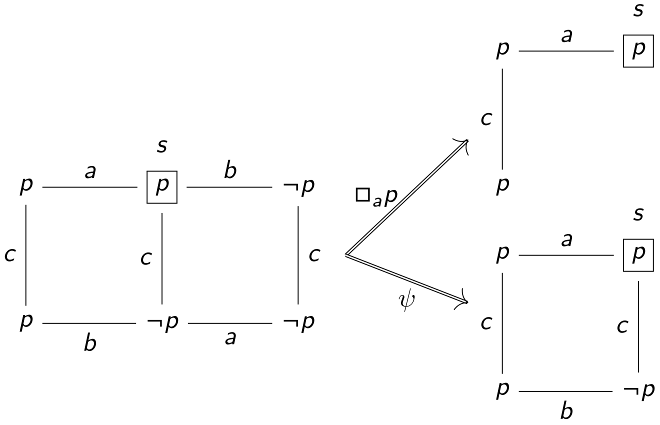

PAL: Learning about Santa

What agents know may change once new fact becomes known.
$\psi := \neg p \rightarrow \Diamond_c p$

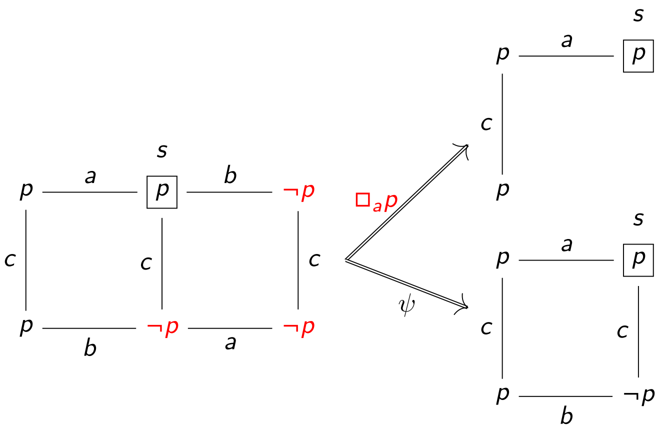

What agents know may change once new fact becomes known.
$\psi := \neg p \to \Diamond_c p$

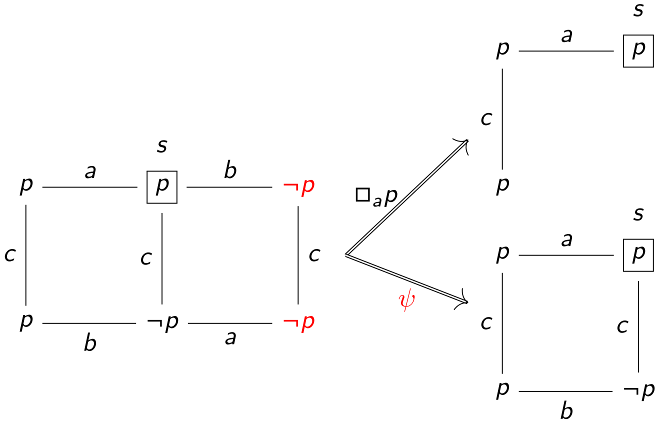

PAL: Syntax and Semantics

$$\mathcal{PAL} \ni \varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Box_a\varphi \mid [\varphi]\varphi$$
$$\mathcal{EL} \ni \varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Box_a\varphi$$

Definition (Semantics)

An announcement of φ in a pointed model $M_s = (S, \sim, V)$ results in an updated pointed model M_s^φ containing only φ-states:

- $S^\varphi = [\![\varphi]\!]_M$,
- $\sim_a^\varphi = \sim_a \cap (S^\varphi \times S^\varphi)$,
- $V^\varphi(p) = V(p) \cap S^\varphi$.

$$\begin{array}{lll} M_s \models \Box_a\varphi & \text{iff} & \forall t \in S : s \sim_a t \text{ implies } M_t \models \varphi \\ M_s \models [\varphi]\psi & \text{iff} & M_s \models \varphi \text{ implies } M_s^\varphi \models \psi \\ M_s \models \langle\varphi\rangle\psi & \text{iff} & M_s \models \varphi \text{ and } M_s^\varphi \models \psi \end{array}$$

PAL: Properties

For each $\varphi \in \mathcal{PAL}$ there is an equivalent $t(\psi) \in \mathcal{EL}$, where t is a translation function.

Theorem

PAL is sound and complete.

Theorem

PAL and EL are equally expressive.

Quantifying Over Announcements

We can shift our perspective from the effects of particular announcements to the question whether some information is attainable by any announcement.

The Big Three of Logics of Quantified Announcements:

- Arbitrary Public Announcement Logic (APAL)
- Group Announcement Logic (GAL)
- Coalition Announcement Logic (CAL)

APAL: Reachable Santa

(Almost) Strategic Santa

In the Santa example, is there an announcement such that Bobby learns that Santa does exist? Is it possible to make an announcement such that Bobby is certain about (non)existence of Santa in all possible worlds?

$\psi := \neg p \rightarrow \Diamond_a \Diamond_c p$

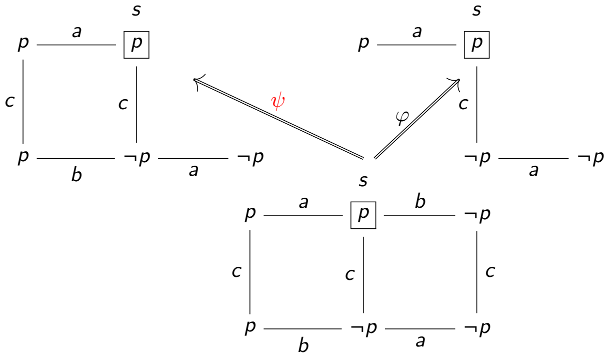

(Almost) Strategic Santa

In the Santa example, is there an announcement such that Bobby learns that Santa does exist? Is it possible to make an announcement such that Bobby is certain about (non)existence of Santa in all possible worlds?

$\varphi := \Box_a((p \rightarrow \Box_b p \land \Box_c p) \land (\neg p \rightarrow \Box_b \neg p \land \Box_c \neg p))$

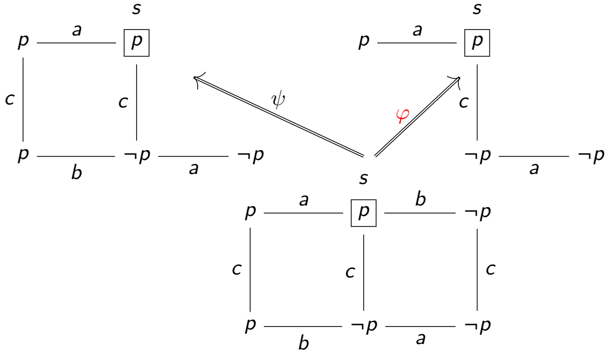

APAL: Syntax and Semantics

Arbitrary Public Announcement Logic (APAL) = PAL + $\{[!]\varphi, \langle ! \rangle \varphi\}$

$\langle ! \rangle \varphi$: 'there is a truthful public announcement such that φ holds in the resulting model'

$[!]\varphi$: 'after any public announcement, φ holds'

Definition (Semantics)

$$M_s \models [!]\varphi \quad \text{iff} \quad \forall \psi \in \mathcal{EL} : M_s \models [\psi]\varphi$$
$$M_s \models \langle ! \rangle\varphi \quad \text{iff} \quad \exists \psi \in \mathcal{EL} : M_s \models \langle \psi \rangle\varphi$$

GAL: (Almost) Strategic Santa

(Almost) Strategic Santa

Reconsider the Santa example from a more strategic point of view. Can Alice inform everyone that Santa does exist? Can she inform only Bobby and leave Claire in ignorance?

GAL: (Almost) Strategic Santa

(Almost) Strategic Santa

Reconsider the Santa example from a more strategic point of view. Can Alice inform everyone that Santa does exist? Can she inform only Bobby and leave Claire in ignorance?

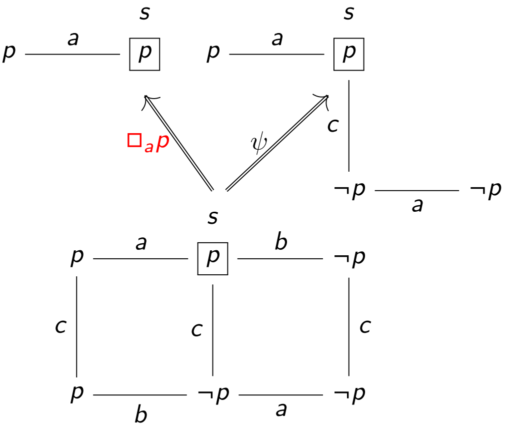

(Almost) Strategic Santa

Reconsider the Santa example from a more strategic point of view.

Can she inform only Bobby and leave Claire in ignorance?

$\psi := \Box_a((p \to \Box_b p \land \Box_c p) \land (\neg p \to \Box_b \neg p \land \Box_c \neg p))$

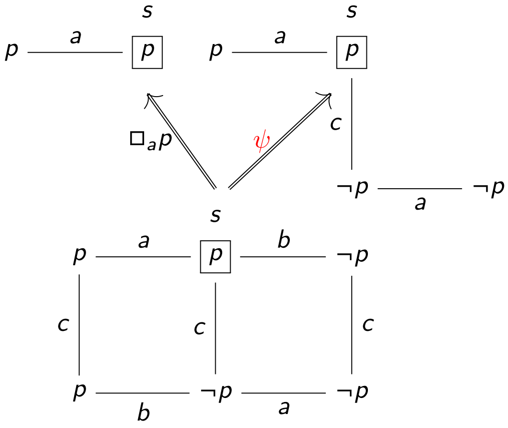

GAL: Syntax and Semantics

Group Announcement Logic (GAL) = PAL + $\{[G]\varphi, \langle G\rangle\varphi\}$

$\langle G\rangle\varphi$: 'agents from G have a joint announcement such that φ holds in the resulting model'

$[G]\varphi$: 'whatever agents from G announce, they cannot avoid φ'

Let $\psi_G := \bigwedge_{a\in G} \Box_a\psi_a$, where ψ_a is an epistemic formula (truthfulness).

Definition (Semantics)

$$M_s \vDash [G]\varphi \quad \text{iff} \quad \forall\psi_G : M_s \vDash [\psi_G]\varphi$$
$$M_s \vDash \langle G\rangle\varphi \quad \text{iff} \quad \exists\psi_G : M_s \vDash \langle\psi_G\rangle\varphi$$

CAL: Strategic Santa

Strategic Santa

Can Alice inform only Bobby that Santa exists if other children are allowed to spoil her announcements?

CAL: Strategic Santa

Strategic Santa

Can Alice inform only Bobby that Santa exists if other children are allowed to spoil her announcements? This is up to Bobby and Claire.

CAL: Strategic Santa

Strategic Santa

Can Alice inform only Bobby that Santa exists if other children are allowed to spoil her announcements? This is up to Bobby and Claire. They may keep silent (announce $\top$).

$\psi_1 := \psi_a \wedge \Box_b \top \wedge \Box_c \top$

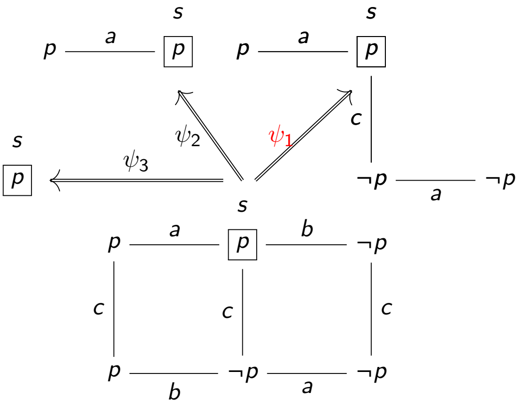

CAL: Strategic Santa

Strategic Santa

Can Alice inform only Bobby that Santa exists if other children are allowed to spoil her announcements? This is up to Bobby and Claire. Bobby may announce ψ_b. $\psi_2 := \psi_a \wedge \psi_b$

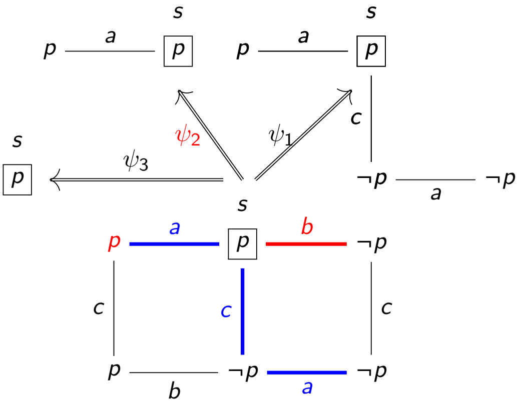

Strategic Santa

Can Alice inform only Bobby that Santa exists if other children are allowed to spoil her announcements? This is up to Bobby and Claire. Bobby may announce ψ_b. $\psi_3 := \psi_a \land \psi_b$

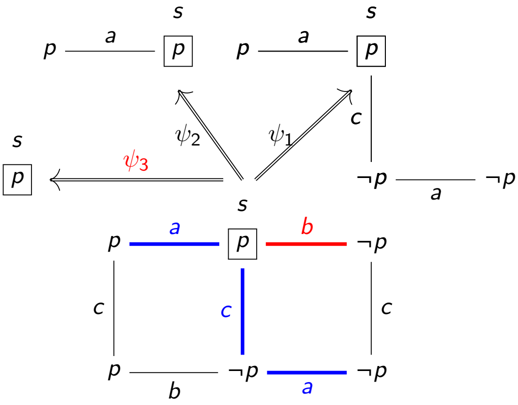

CAL: Syntax and Semantics

Coalition Announcement Logic (CAL) = PAL + $\{[\langle G\rangle]\varphi, \langle[G]\rangle\varphi\}$

$\langle[G]\rangle\varphi$: 'there exists a joint announcement by the agents from G such that no matter what other agents announce at the same time, φ holds '

$[\langle G\rangle]\varphi$: 'whatever agents from G announce, there is a simultaneous announcement by the anti-coalition such that φ holds'

Let $\psi_G := \bigwedge_{a \in G} \Box_a \psi_a$, where ψ_a is an epistemic formula.

Definition (Semantics)

$$
\begin{array}{lll}
M_s \models [\langle G\rangle]\varphi & \text{iff} & \forall \psi_G \exists \chi_{\overline{G}} : M_s \models \psi_G \to \langle \psi_G \land \chi_{\overline{G}} \rangle \varphi \\
M_s \models \langle[G]\rangle\varphi & \text{iff} & \exists \psi_G \forall \chi_{\overline{G}} : M_s \models \psi_G \land [\psi_G \land \chi_{\overline{G}}]\varphi
\end{array}
$$

Logics of Quantified Announcements: Properties

Recall...

- APAL: quantifying over all public announcements
- GAL: quantifying over public announcements that can be made by groups of agents
- CAL: quantifying over announcements by a coalition and the anti-coalition

Axiomatisation

APAL and GAL have (infinitary) complete axiomatisations. It is unknown whether finitary axiomatisations of the logics exist. A complete axiomatisation, finitary or infinitary, of CAL is also unknown.

Logics of Quantified Announcements: Properties

Recall...

- APAL: quantifying over all public announcements
- GAL: quantifying over public announcements that can be made by groups of agents
- CAL: quantifying over announcements by a coalition and the anti-coalition

Complexity

All of the logics are undecidable. Model checking is PSPACE-complete.

Logics of Quantified Announcements: Properties

Recall...

- APAL: quantifying over all public announcements
- GAL: quantifying over public announcements that can be made by groups of agents
- CAL: quantifying over announcements by a coalition and the anti-coalition

Expressivity

APAL and GAL are incomparable. CAL is not at least as expressive as either APAL and GAL.

Logics of Quantified Announcements: Properties

Recall...

- APAL: quantifying over all public announcements
- GAL: quantifying over public announcements that can be made by groups of agents
- CAL: quantifying over announcements by a coalition and the anti-coalition

What about the finite model property?

FMP

A logic has the FMP iff every formula of the logic that is true in some model is also true in a finite model.

We know that...

Finitary axiomatisation $\land$ FMP $\rightarrow$ Decidability

The implication can be satisfied in various ways.

The Finite Model Property

FMP

A logic has the FMP iff every formula of the logic that is true in some model is also true in a finite model.

We know that...

Finitary axiomatisation $\wedge$ FMP $\rightarrow$ Decidability

We also know that all APAL, GAL, and CAL are undecidable.

$\neg$Decidabitity $\rightarrow$ $\neg$Finitary axiomatisation $\vee$ $\neg$FMP

However, whether finitary axiomatisations of the logics exist is an open question.

The Finite Model Property

FMP

A logic has the FMP iff every formula of the logic that is true in some model is also true in a finite model.

Proof Idea

Present a formula of a logic. Show that there is an infinite model, where it is satisfied, and argue that the formula is not satisfied in any finite model.

Overview

- APAL, GAL, and CAL allow us to reason about whether there is a public announcement that achieves some epistemic goal.
- Quantification in APAL is over all formulas of EL.
- Quantification in GAL is over formulas of EL that are known to agents.
- There is double quantification in CAL: over announcements by a coalition, and over all responses by the anti-coalition to a coalition's announcement.
- None of APAL, GAL, and CAL have the finite model propety.