Public Announcements Logic

Nikolaos Nikolopoulos

ALMA
INTER-INSTITUTIONAL GRADUATE PROGRAM
«ALGORITHMS, LOGIC AND DISCRETE MATHEMATICS»

Overview

Overview

Motivation

Purpose: Formalization of agents' knowledge and beliefs
So far: Systems of modal logic where various conditions hold
How do we represent changes?

Motivation

Purpose: Formalization of agents' knowledge and beliefs
So far: Systems of modal logic where various conditions hold
How do we represent changes?

By modifying some initial conditions.. but is it enough?

Motivation

Purpose: Formalization of agents' knowledge and beliefs
So far: Systems of modal logic where various conditions hold
How do we represent changes?

By modifying some initial conditions.. but is it enough? NO

Motivation

Purpose: Formalization of agents' knowledge and beliefs
So far: Systems of modal logic where various conditions hold
How do we represent changes?

By modifying some initial conditions.. but is it enough? NO

Agents change their perspective of the model *dynamically*.
Goal: Model announcements (private or public) i.e. completely trustworthy, truthful statements, so that their consequences could be analyzed without "hard-wiring".

Motivation

Purpose: Formalization of agents' knowledge and beliefs
So far: Systems of modal logic where various conditions hold
How do we represent changes?

By modifying some initial conditions.. but is it enough? NO

Agents change their perspective of the model *dynamically*.
Goal: Model announcements (private or public) i.e. completely trustworthy, truthful statements, so that their consequences could be analyzed without "hard-wiring".
So what's this chapter all about?

Example (Naughty student)

Suppose we have *Tim* waiting to be punished outside of the principal's office. A teacher comes out of the office and says: "Tim you don't know it yet, but you are not going to be punished!"

We denote t agent corresponding to student *Tim* and p the fact that Tim is not going to be punished.

Example (Naughty student)

Suppose we have *Tim* waiting to be punished outside of the principal's office. A teacher comes out of the office and says: "Tim you don't know it yet, but you are not going to be punished!"

We denote t agent corresponding to student *Tim* and p the fact that Tim is not going to be punished.

We have the following situation:

$$\neg K_t p \xrightarrow[announcement]{p \wedge \neg K_t p} K_t p \wedge p$$

Note that (in the end) the negation of the announcement is true!

Example (Cheryl's birthday)

Suppose we have *Albert, Bernard and Cheryl* chatting online. Albert asks Cheryl for her birthdate. She replies with the following possible dates:

- $\{15, 16, 19\}$ of May
- $\{17, 18\}$ of June
- $\{14, 16\}$ of July
- $\{14, 15, 17\}$ of August

She also shares (privately) the month with *Albert* and the day with *Bernard*.

Is it possible for both of them *individually* to deduce Cheryl's birthday?

Example (Cheryl's birthday)

Suppose we have *Albert, Bernard and Cheryl* chatting online. Albert asks Cheryl for her birthdate. She replies with the following possible dates:

- $\{15, 16, 19\}$ of May
- $\{17, 18\}$ of June
- $\{14, 16\}$ of July
- $\{14, 15, 17\}$ of August

She also shares (privately) the month with *Albert* and the day with *Bernard*.

Is it possible for both of them *individually* to deduce Cheryl's birthday? NO

Examples: Cheryl's birthday cont'd

We could represent the problem's epistemic state (M_{cb}) as follows:

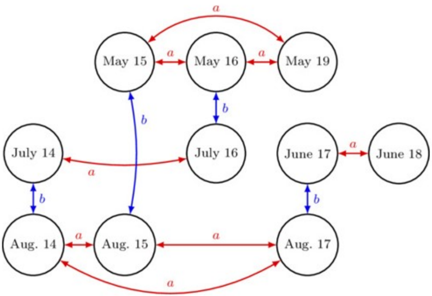

Observations:

- One world (node) for each possible date
- Each edge denotes possibility for a birthdate
- For each world reflexive property for both of these agents hold (we will deliberately omit these relationships for simplicity)
- Dates of the same day form an equivalence class for Bernard
- Dates of the same month form an equivalence class for Albert

Examples: Cheryl's birthday cont'd

We could represent the problem's epistemic state (M_{cb}) as follows:

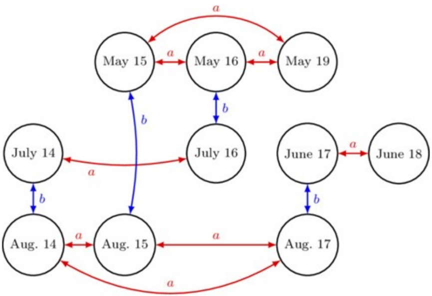

Observations:

- One world (node) for each possible date
- Each edge denotes possibility for a birthdate
- For each world reflexive property for both of these agents hold (we will deliberately omit these relationships for simplicity)
- Dates of the same day form an equivalence class for Bernard
- Dates of the same month form an equivalence class for Albert

What changes if *Albert* makes the announcement

$$A_1 : \neg K_a p \wedge K_a \neg K_b p$$

Examples: Cheryl's birthday cont'd

Albert's public announcement (A_1) could be written as:

"I don't know Cheryl's birthday, but I also know that Bernard doesn't know either!"

This changes the previous epistemic state, since we are able to eliminate the worlds where :

$$K_a p \vee K_a K_b p$$

Examples: Cheryl's birthday cont'd

Albert's public announcement (A_1) could be written as:

"I don't know Cheryl's birthday, but I also know that Bernard doesn't know either!"

This changes the previous epistemic state, since we are able to eliminate the worlds where :

$$K_a p \vee K_a K_b p$$

Albert says two things that now become common knowledge. However, the content of $K_a p$ is *empty* because it doesn't rule out any possible world.

Examples: Cheryl's birthday cont'd

Albert's public announcement (A_1) could be written as:

"I don't know Cheryl's birthday, but I also know that Bernard doesn't know either!"

This changes the previous epistemic state, since we are able to eliminate the worlds where :

$$K_a p \vee K_a K_b p$$

Albert says two things that now become common knowledge.
However, the content of $K_a p$ is *empty* because it doesn't rule out any possible world.
On the contrary, that is not the case for Bertrand. Why?

Examples: Cheryl's birthday cont'd

Let's represent the new problem state (after A_1) as $M_{cb}[A_1]$. We only need to remove worlds where *Albert* would know that *Bernard* would know Cheryl's birthday. How will the situation look like now?

Examples: Cheryl's birthday cont'd

Let's represent the new problem state (after A_1) as $M_{cb}[A_1]$. We only need to remove worlds where *Albert* would know that *Bernard* would know Cheryl's birthday. How will the situation look like now?

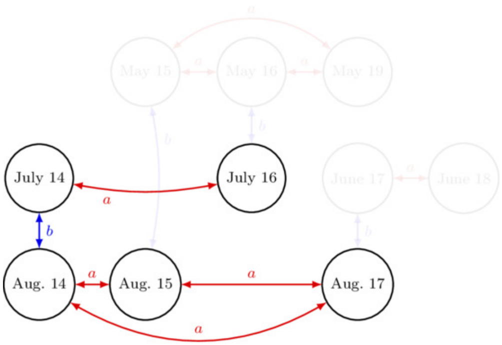

After hearing A_1, Bernard says: "I do know the date now!" A_2.
What restrictions apply now?

After hearing A_1, Bernard says: "I do know the date now!" A_2.
What restrictions apply now?
Apparently we could write A_2 as $K_b p$, so the new problem state $M_{cb}[A_1][A_2]$ looks like this:

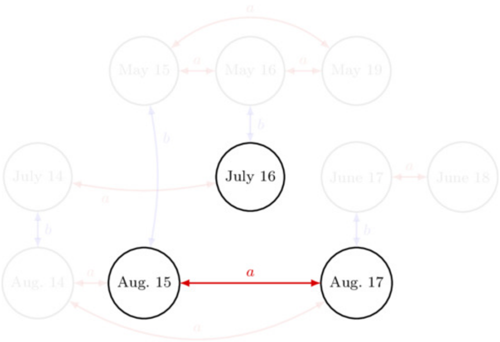

Finally, Albert announces that he also knows the date. So A_3 could be written as $K_a p$.
The only equivalence class of size 1 is "July 16", ergo it's the only date satisfying the model $M_{cb}[A_1][A_2][A_3]$!

Overview

PAL framework: Definition

Definition (*PAL*: Public Announcement Logic)

Given a set of agents A and a set of atoms P, we define $\mathcal{L}_{KC[]}(A, P)$ by the BNF:

$$\phi \ ::= \ p \mid \neg\phi \mid (\phi \wedge \phi) \mid K_a\phi \mid C_B\phi \mid [\phi]\phi$$

and $\mathcal{L}_{K[]}(A, P)$ (without common knowledge) by the BNF:

$$\phi \ ::= \ p \mid \neg\phi \mid (\phi \wedge \phi) \mid K_a\phi \mid [\phi]\phi$$

where $a \in A$, $p \in P$ and $B \subseteq A$

$[\phi]\psi$ stands for "after announcement ϕ it holds that ψ"

Operators' duality

An announcement is always *truthful* & *public*.
However, often we might interpret $[\phi]\psi$ as:

"after update ϕ, it holds that ψ"

Operators' duality

An announcement is always *truthful* & *public*.
However, often we might interpret $[\phi]\psi$ as:

"after update ϕ, it holds that ψ"

Observe the duality :

- $\Box\phi \Rrightarrow [\phi]\psi$ (after *every* announcement ϕ, it holds that ψ)
- $\diamond\phi \Rrightarrow \langle\phi\rangle\psi$ (after *some* announcement ϕ, it holds that ψ)

Operators' duality

An announcement is always *truthful* & *public*.
However, often we might interpret $[\phi]\psi$ as:

"after update ϕ, it holds that ψ"

Observe the duality :

- $\Box\phi \Rrightarrow [\phi]\psi$ (after *every* announcement ϕ, it holds that ψ)
- $\Diamond\phi \Rrightarrow \langle\phi\rangle\psi$ (after *some* announcement ϕ, it holds that ψ)

That enables us to describe *unsuccessful updates* i.e. formulas that become false after their announcement. why didn't we add $\langle\cdot\rangle$ in the BNF before?

Quick Example

In the beginning, we saw the NAUGHTY STUDENT example, where the announcement $A: \ p \wedge \neg K_t p$ occurred.
That announcement became *false* after being expressed, since Tim now knows he is not going to be punished.
Using the above syntax one could describe this scenario as $\langle A \rangle \neg A$

Quick Example

In the beginning, we saw the NAUGHTY STUDENT example, where the announcement $A: \ p \land \neg K_t p$ occurred.
That announcement became *false* after being expressed, since Tim now knows he is not going to be punished.
Using the above syntax one could describe this scenario as $\langle A \rangle \neg A$

Also in the end of CHERYL'S BIRTHDAY example we demonstrated the restriction of our possible worlds after an announcement.
That's a crucial technique in *PAL*; semantics formalization follows.

Definition (Semantics of *PAL*)

Suppose that there is a set of agents A, atoms P and an epistemic model $M = \langle S, \sim, V \rangle$.

$$
\begin{array}{lr}
M, s \models p & s \in V_p \\
M, s \models \neg\phi & M, s \not\models \phi \\
M, s \models \phi \wedge \psi & (M, s \models \phi) \wedge (M, s \models \psi) \\
M, s \models K_a\phi & \forall t \in S \, : s \sim_a t \Rightarrow \; M, t \models \phi \\
M, s \models C_B\phi & \forall t \in S \, : s \sim_B t \Rightarrow \; M, t \models \phi \\
M, s \models [\phi]\psi & M, s \models \phi \Rightarrow M|\phi, s \models \psi
\end{array}
$$

$M|\phi = \langle S', \sim', V' \rangle$ with $S' = [\![\phi]\!]_M$, $\sim'_a = \sim_a \cap ([\![\phi]\!]_M \times [\![\phi]\!]_M)$ and $V_p{}' = V_p \cap [\![\phi]\!]_M$

Some interesting announcement properties follow:

- $\langle\phi\rangle\psi \rightarrow [\phi]\psi$ *functionality*
- $[\phi]\neg\psi \leftrightarrow (\phi \rightarrow \neg[\phi]\psi)$ *negation*
- $(\phi \rightarrow [\phi]\psi) \equiv (\phi \rightarrow \langle\phi\rangle\psi) \equiv [\phi]\psi$
- $\langle\phi\rangle\psi \equiv (\phi \wedge \langle\phi\rangle\psi) \equiv (\phi \wedge [\phi]\psi)$
- $[\phi \wedge [\phi]\psi]\chi \equiv [\phi][\psi]\chi$ *composition*

In $PAL [\phi]K_a\psi \not\equiv K_a[\phi]\psi$ in general because an announcement may not take place in an epistemic state. But if we somehow regulate an announcement or an expression based on its validity, then the equivalence holds, since it is now a total function.

In $PAL [\phi] K_a \psi \not\equiv K_a [\phi] \psi$ in general because an announcement may not take place in an epistemic state. But if we somehow regulate an announcement or an expression based on its validity, then the equivalence holds, since it is now a total function.

It is also true that $[\phi] K_a \psi \equiv \phi \rightarrow K_a [\phi] \psi$.

In $PAL [\phi]K_a\psi \not\equiv K_a[\phi]\psi$ in general because an announcement may not take place in an epistemic state. But if we somehow regulate an announcement or an expression based on its validity, then the equivalence holds, since it is now a total function.

It is also true that $[\phi]K_a\psi \equiv \phi \rightarrow K_a[\phi]\psi$.

It is also possible to remove announcements from logical expressions altogether, using the following rules:

$$
\begin{aligned}
[\phi]p &\leftrightarrow (\phi \rightarrow p) \\
[\phi](\psi \wedge \chi) &\leftrightarrow ([\phi]\psi \wedge [\phi]\chi) \\
[\phi](\psi \rightarrow \chi) &\leftrightarrow ([\phi]\psi \rightarrow [\phi]\chi) \\
[\phi]\neg\psi &\leftrightarrow (\phi \rightarrow \neg[\phi]\psi) \\
[\phi]K_a\psi &\leftrightarrow (\phi \rightarrow K_a[\phi]\psi) \\
[\phi][\psi]\chi &\leftrightarrow [\phi \wedge [\phi]\psi]\chi)
\end{aligned}
$$

It is also possible to remove announcements from logical expressions altogether, using the following rules:

$$
\begin{aligned}
[\phi]p &\leftrightarrow (\phi \rightarrow p) \\
[\phi](\psi \wedge \chi) &\leftrightarrow ([\phi]\psi \wedge [\phi]\chi) \\
[\phi](\psi \rightarrow \chi) &\leftrightarrow ([\phi]\psi \rightarrow [\phi]\chi) \\
[\phi]\neg\psi &\leftrightarrow (\phi \rightarrow \neg[\phi]\psi) \\
[\phi]K_a\psi &\leftrightarrow (\phi \rightarrow K_a[\phi]\psi) \\
[\phi][\psi]\chi &\leftrightarrow [\phi \wedge [\phi]\psi]\chi)
\end{aligned}
$$

Since announcements could always be replaced, one might wonder do we really need to use them?
The answer is yes, because they guide our intuition throughout changes in epistemic states.
We shall not forget that *abstraction* is a crucial technique towards the improvement of readability and understanding in general.

that's why we don't go around writing machine code, right?

Let's add *common knowledge* into our mix!

Let's add *common knowledge* into our mix!
At first, observe that announcing ϕ doesn't make ϕ common knowledge.

Let's add *common knowledge* into our mix!
At first, observe that announcing ϕ doesn't make ϕ common knowledge.

One might consider expanding $[\phi]K_a\psi \leftrightarrow (\phi \rightarrow K_a[\phi]\psi)$ into:

$$[\phi]C_A\psi \leftrightarrow (\phi \rightarrow C_A[\phi]\psi)$$

Let's add *common knowledge* into our mix!
At first, observe that announcing ϕ doesn't make ϕ common knowledge.

One might consider expanding $[\phi]K_a\psi \leftrightarrow (\phi \rightarrow K_a[\phi]\psi)$ into:

$$[\phi]C_A\psi \leftrightarrow (\phi \rightarrow C_A[\phi]\psi)$$

This formula is **invalid** but the following holds :

$$\left.\begin{array}{ll} \chi \rightarrow [\phi]\psi & \text{valid} \\ (\chi \wedge \phi) \rightarrow E_B\chi & \text{valid} \end{array}\right\} \chi \rightarrow [\phi]C_B\psi \textit{ valid}$$

Let's add *common knowledge* into our mix!
At first, observe that announcing ϕ doesn't make ϕ common knowledge.

One might consider expanding $[\phi]K_a\psi \leftrightarrow (\phi \rightarrow K_a[\phi]\psi)$ into:

$$[\phi]C_A\psi \leftrightarrow (\phi \rightarrow C_A[\phi]\psi)$$

This formula is **invalid** but the following holds :

$$\left.\begin{array}{ll} \chi \rightarrow [\phi]\psi & \text{valid} \\ (\chi \wedge \phi) \rightarrow E_B\chi & \text{valid} \end{array}\right\} \chi \rightarrow [\phi]C_B\psi \textit{ valid}$$

Finally, $[\phi]\psi$ *valid* $\leftrightarrow$ $[\phi]C_B\psi$ *valid*

Naturally, we might believe that since an announcement is *public* and *truthful* by design, the statement being announced remains true afterwards. Unfortunately, this is **false**.

We've already seen such examples e.g. NAUGHTY STUDENT. Obviously, the success depends on a) the formula and b) the epistemic state. Let's formalize this notion.

Definition (The secret of success)

We call a formula *successful*, if it becomes common knowledge after being announced and*unsuccessful* otherwise.
Moreover, if $\phi \in \mathcal{L}_{KC[]}$ and state $(M, s) : \; M \in \mathcal{S}5$, then:

- ϕ *successful formula* $\leftrightarrow [\phi]\phi$ valid
- ϕ *unsuccessful formula* otherwise
- ϕ *successful update* $(M, s) \leftarrow \; M, s \models \langle\phi\rangle\phi$
- ϕ *unsuccessful update* $(M, s) \leftrightarrow \; M, s \models \langle\phi\rangle\neg\phi$

Definition (The secret of success)

We call a formula *successful*, if it becomes common knowledge after being announced and*unsuccessful* otherwise.
Moreover, if $\phi \in \mathcal{L}_{KC[]}$ and state (M, s) : $M \in \mathcal{S}5$, then:

- ϕ *successful formula* $\leftrightarrow [\phi]\phi$ valid
- ϕ *unsuccessful formula* otherwise
- ϕ *successful update* $(M, s) \leftarrow \; M, s \models \langle\phi\rangle\phi$
- ϕ *unsuccessful update* $(M, s) \leftrightarrow \; M, s \models \langle\phi\rangle\neg\phi$

Also it holds that if :

$$\phi \textit{ successful} \leftrightarrow C_A\phi \textit{ successful} \leftrightarrow \; (\phi \rightarrow [\phi]C_A\phi) \textit{ valid}$$

Success Guarantee

We defined *success* in terms of announcement validity. However, the necessary conditions for formulas' success is not known.

Success Guarantee

We defined *success* in terms of announcement validity. However, the necessary conditions for formulas' success is not known.
Observe that :

$\phi,\ \psi$ *successful* then $\neg\phi, \phi \wedge \psi, \phi \rightarrow \psi, [\phi]\psi$ not always *successful*

Finally, $C_A\phi$ successful $\forall\ \phi \in \mathcal{L}_{KC[]}(A, P)$
Note that despite the validity of $[C_A\phi]C_A\phi$, this might not hold $\forall\ B \subseteq A$.

When announcing formulas that are *public knowledge*, no restriction between states occur. So when does the truth of a formula change?

When announcing formulas that are *public knowledge*, no restriction between states occur. So when does the truth of a formula change?

Given that an announcement may change/restrict the available epistemic states of the model, that depends if the formula is affected by the restriction taking place.
Can we somehow create *truth-preserving* formulas?

Definition

We call $\mathcal{L}^0_{KC[]}(A, P)$ *language of preserved formulas* all these formulas ϕ where:

$$\phi ::= p \mid \neg p \mid (\phi \wedge \phi) \mid (\phi \vee \phi) \mid K_a\phi \mid C_B\phi \mid [\neg\phi]\phi$$

Every submodel in $\mathcal{L}^0_{KC[]}(A, P)$ is *truth-preserving*

Definition

We call $\mathcal{L}^0_{KC[]}(A, P)$ *language of preserved formulas* all these formulas ϕ where:

$$\phi ::= p \mid \neg p \mid (\phi \wedge \phi) \mid (\phi \vee \phi) \mid K_a\phi \mid C_B\phi \mid [\neg\phi]\phi$$

Every submodel in $\mathcal{L}^0_{KC[]}(A, P)$ is *truth-preserving*

Also if $\phi \in \mathcal{L}^0_{KC[]}$, $\psi \in \mathcal{L}_{KC[]}$, it is valid that $\phi \rightarrow [\psi]\phi$

Definition

We call $\mathcal{L}^0_{KC[]}(A, P)$ *language of preserved formulas* all these formulas ϕ where:

$$\phi ::= p \mid \neg p \mid (\phi \wedge \phi) \mid (\phi \vee \phi) \mid K_a\phi \mid C_B\phi \mid [\neg\phi]\phi$$

Every submodel in $\mathcal{L}^0_{KC[]}(A, P)$ is *truth-preserving*

Also if $\phi \in \mathcal{L}^0_{KC[]}$, $\psi \in \mathcal{L}_{KC[]}$, it is valid that $\phi \rightarrow [\psi]\phi$

Finally, every $\phi \in \mathcal{L}^0_{KC[]}$ is successful (the contrary does not hold!)

- van Ditmarsch, H et al. "Dynamic Epistemic Logic" (2008) : 67-88
- Nebel, B et al. "Dynamic Epistemic Logic: Public Announcements" (2019)
- Baltag, A et al. "Dynamic Epistemic Logic", The Stanford Encyclopedia of Philosophy (Winter 2016 Edition), Edward N. Zalta (ed.)