Multi-Agent Systems

Albert-Ludwigs-Universität Freiburg

Bernhard Nebel, Felix Lindner, and Thorsten Engesser

Winter Term 2018/19

Recap

- Last session: Axioms of epistemic/doxastic logic, group knowledge (common knowledge, distributed knowledge).
- Today: Modeling changes of knowledge due to public communication and observations (muddy children puzzle).

Muddy Children Puzzle: Formulation

Consider n children playing outdoors together. Suppose k of them get mud on their foreheads. Each of the n children can see which of the other $n-1$ children are muddy or not, but, of course, can't be sure whether s/he is muddy.

1. The father shows up and announces: "At least one of you has mud on his/her forehead."
2. The father then asks: "Does any of you know whether s/he has mud on her/his forehead?"
3. After the k-th such question, all the k muddy children will answer "Yes!".

Interesting Questions

- Did the father tell the children anything new in the first announcement?
- Why is it that all the muddy children simultaneously know the answer to question (2) after exactly k rounds?

Base Case I

Case $k = 1$

- The muddy child only sees clean children. And all clean children see one muddy child.
  - Muddy child considers possible: 0 or 1 children are muddy.
  - Clean children consider possible: 1 or 2 children are muddy.
- After the father announces that at least one of them is muddy:
  - Muddy child considers possible: 1 muddy.
  - Clean children consider possible: 1 or 2 muddy.
- The father asks who knows to be muddy:
  - Muddy child knows!

Base Case II

Case $k = 2$

- The muddy children see exactly one muddy child. And all clean children see two muddy children.
  - Muddy children consider possible: 1 or 2 children are muddy.
  - Clean children consider possible: 2 or 3 children are muddy.
- After the father announces that at least one of them is muddy:
  - Muddy children consider possible: 1 or 2 muddy.
  - Clean children consider possible: 2 or 3 muddy.
- The father asks who knows to be muddy:
  - Nobody!
- Hence, there must be more than one muddy children.
  - Muddy children consider possible: 2 muddy.
  - Clean children consider possible: 2 or 3 muddy.
- The father asks who knows to be muddy:
  - The muddy children know!

Muddy Children: Initial

(reflexive edges omitted)

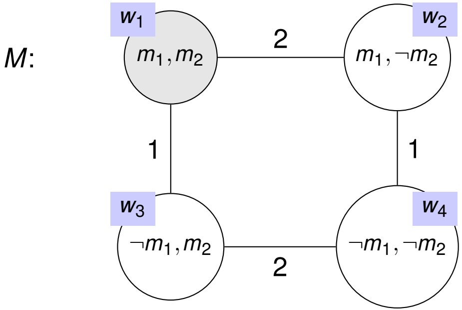

- $M, w_1 \models C_{\{1,2\}}(K_1 m_2 \vee K_1 \neg m_2)$
- $M, w_1 \models C_{\{1,2\}}(K_2 m_1 \vee K_2 \neg m_1)$
- $M, w_1 \models E_{\{1,2\}}(m_1 \vee m_2)$
- $M, w_1 \models \neg E^2_{\{1,2\}}(m_1 \vee m_2)$
- $M, w_1 \models \neg C_{\{1,2\}}(m_1 \vee m_2)$
- $M, w_1 \models D_{\{1,2\}}(m_1 \wedge m_2)$

Muddy Children: After First Announcement

(reflexive edges omitted)

Father:"At least one of you has mud on his/her forehead!"

M':

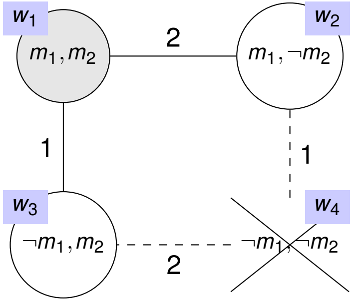

- $M', w_1 \models \neg K_1 m_1 \wedge \neg K_1 \neg m_1$
- $M', w_1 \models \neg K_2 m_2 \wedge \neg K_2 \neg m_2$
- $M', w_1 \models C_{\{1,2\}}(m_1 \vee m_2)$ ($\Rightarrow$announcement is informative)
- $\Rightarrow M', w_1 \models K_2(K_1 \neg m_2 \rightarrow K_1 m_1) \wedge K_1(K_2 \neg m_1 \rightarrow K_2 m_2)$

Muddy Children: After Question

(reflexive edges omitted)

Nobody answers “Yes” to father’s question “Does any of you know whether s/he has mud on her/his forehead?”

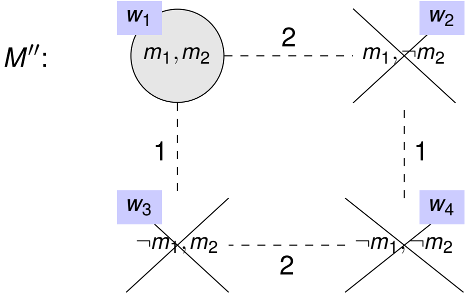

- $M'', w_1 \models C_{\{1,2\}}(m_1 \wedge m_2)$

Public Announcement Operator

$[!\varphi]\psi$: "After φ has been truthfully announced, ψ is the case."

- Semantics

$M,w \models [!\varphi]\psi$ iff. $M,w \not\models \varphi$, or else $M_\varphi, w \models \psi$

- M_φ is the relativation of M to the worlds where φ holds. The model $M_\varphi = (S', R', V')$ is given as follows:

$$S' = \{w \in S : M,w \models \varphi\} \quad (1)$$

$$R' = R|_{S' \times S'} \quad (2)$$

$$V'(p) = V(p) \cap S' \quad (3)$$

Public Announcement Operator

$[!\varphi]\psi$: "After φ has been truthfully announced, ψ is the case."

- Semantics

$M,w \models [!\varphi]\psi$ iff. $M,w \not\models \varphi$, or else $M_\varphi,w \models \psi$

- M_φ is the relativation of M to the worlds where φ holds. The model $M_\varphi = (S',R',V')$ is given as follows:

$$S' = \{w \in S : M,w \models \varphi\} \quad (1)$$

$$R' = R|_{S' \times S'} \quad (2)$$

$$V'(p) = V(p) \cap S' \quad (3)$$

Public Announcement Operator

$[!\varphi]\psi$: "After φ has been truthfully announced, ψ is the case."

- Semantics

$M, w \models [!\varphi]\psi$ iff. $M, w \not\models \varphi$, or else $M_\varphi, w \models \psi$

- M_φ is the relativation of M to the worlds where φ holds. The model $M_\varphi = (S', R', V')$ is given as follows:

$$S' = \{w \in S : M, w \models \varphi\} \quad (1)$$

$$R' = R|_{S' \times S'} \quad (2)$$

$$V'(p) = V(p) \cap S' \quad (3)$$

Muddy Children Puzzle: PAL: Initial

(reflexive edges omitted)

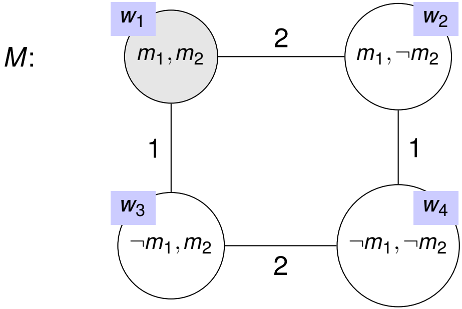

- To Show: $M, w_1 \models [!\varphi_1][!\varphi_2 \wedge \varphi_3]K_1 m_1 \wedge K_2 m_2$
- $\varphi_1 = m_1 \vee m_2$
- $\varphi_2 = (\neg K_1 m_1 \wedge \neg K_1 \neg m_1)$
- $\varphi_3 = (\neg K_2 m_2 \wedge \neg K_2 \neg m_2)$

Muddy Children Puzzle: PAL: After Announcement

(reflexive edges omitted)

$M, w_1 \models [!\varphi_1][!\varphi_2 \wedge \varphi_3]K_1 m_1 \wedge K_2 m_2$

iff. $M, w_1 \not\models \varphi_1$ or else $M_{\varphi_1}, w_1 \models [!\varphi_2 \wedge \varphi_3]K_1 m_1 \wedge K_2 m_2$

M_{φ_1}:

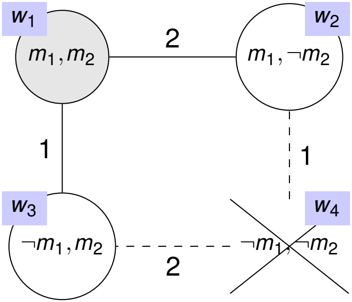

- $M_{\varphi_1}, w_1 \models C_{\{1,2\}}(m_1 \vee m2)$

(reflexive edges omitted)

$M_{\varphi_1}, w_1 \models [!\varphi_2 \wedge \varphi_3] K_1 m_1 \wedge K_2 m_2$

iff. $M_{\varphi_1}, w_1 \not\models \varphi_2 \wedge \varphi_3$ or else $M_{\varphi_{1_{\varphi_2 \wedge \varphi_3}}} K_1 m_1 \wedge K_2 m_2$

$M_{\varphi_{1_{\varphi_2 \wedge \varphi_3}}}$:

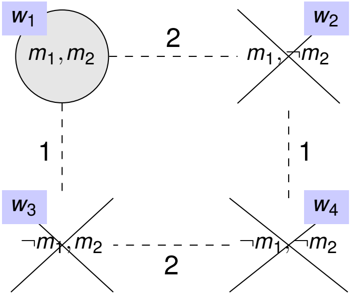

- $M_{\varphi_{1_{\varphi_2 \wedge \varphi_3}}}, w_1 \models C_{\{1,2\}}(m_1 \wedge m2)$

Moore's Paradox

- Interestingly, $[!\varphi]\varphi$ is not valid in general.
- Indeed, $[!(p \land \neg Kp)]\neg(p \land \neg Kp)$ is valid. This is related to Moore's paradox saying one cannot know sentences of the form "φ is true and I don't know φ."
  - Let M be a model and w a world in it.
  - Assume $M, w \models p \land \neg Kp$.
  - Let N be the relativation of M, $M_{p \land \neg Kp}$.
  - Because $N, w \models p \land \neg Kp$, there must be a successor of w, w', such that $N, w' \neg p$. But as w' is in N, it must also be the case that $N, w' \models p \land \neg Kp$.
  - Contradiction!

Translation of PAL to Epistemic Logic

Theorem (cf., [1])

For every formula φ with public announcement operator there is a equivalent formula $t(\varphi)$ without public announcement operator.

- $t(p) = p$
- $t(\neg\varphi) = \neg t(\varphi)$
- $t(\varphi \wedge \psi) = t(\varphi) \wedge t(\psi)$
- $t(K_i\varphi) = K_i t(\varphi)$
- $t([!\varphi]p) = t(\varphi) \rightarrow p$
- $t([!\varphi)\neg\psi) = t(\varphi \rightarrow \neg[!\varphi]\psi)$
- $t([!\varphi](\psi \wedge \chi)) = t([!\varphi]\psi \wedge [!\varphi]\chi)$
- $t([!\varphi]K_i\psi) = t(\varphi \rightarrow K_i[!\varphi]\psi)$
- $t([!\varphi][!\psi]\chi) = t([!(\varphi \wedge [!\varphi]\psi)]\chi)$

$\Rightarrow$PAL does not introduce something really new. But it makes modeling public announcements easier.

Summary and Outlook

- Public communication and observations change what is common knowledge among agents $\Rightarrow$This kind of dynamics can be modeled using the Public Announcement Operator.
- Public Announcement Logic can be translated to Epistemic Logic.
- The approach can be generalized to updating epistemic models due to arbitary actions (not only announcements). $\Rightarrow$Planning based on Dynamic Epistemic Logic is a research area in our group.

Literature

- L. S. Moss, Dynamic Epistemic Logic, Chapter 6, In H. van Dithmarschen, J. Y. Halpern, W. van der Hoek, B. Kooi (eds.) Handbook of Epistemic Logic, College Publications, 2015.
- Y. Shoham, K. Layton-Brown, Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations, Cambridge University Press, 2009.