The Proof Theory of Common Knowledge

Thomas Studer

based on joint work with
Kai Brünnler, Samuel Bucheli, Gerhard Jäger, Mathis Kretz, Roman Kuznets, Grigori Mints

Institute of Computer Science
University of Bern
Switzerland

Roadmap

- Hilbert system for common knowledge
- Infinitary system based on an ω-rule
- Syntactic cut-elimination
- Infinite branches

Common knowledge

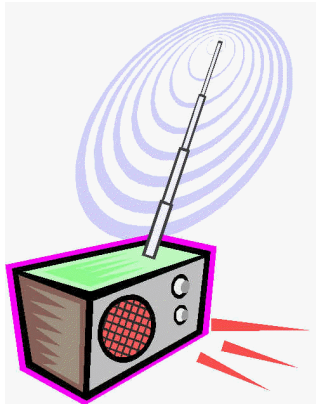

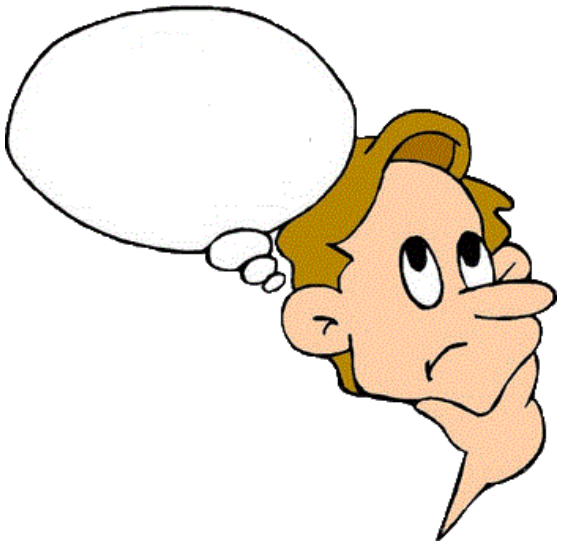

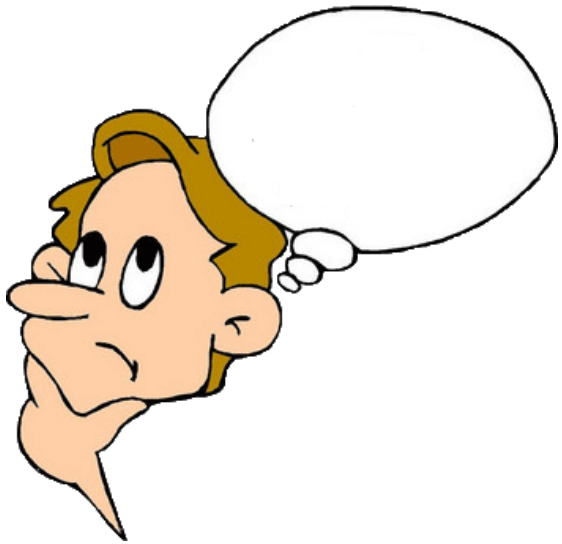

Common knowledge

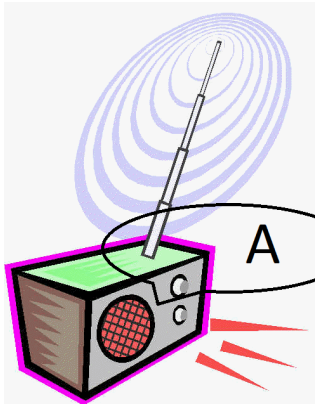

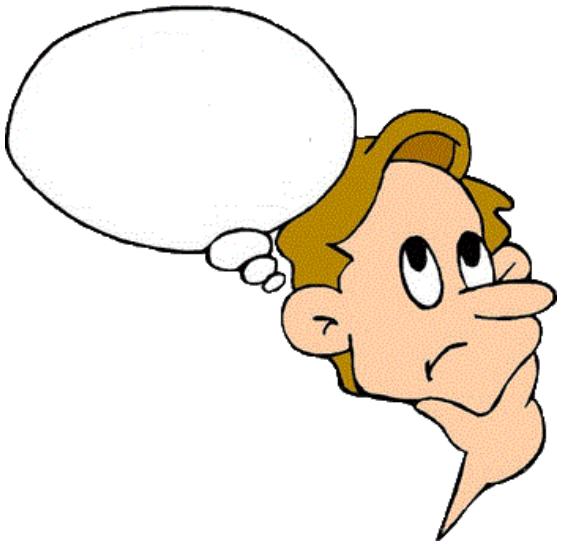

Common knowledge

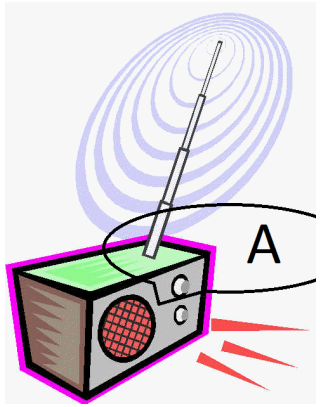

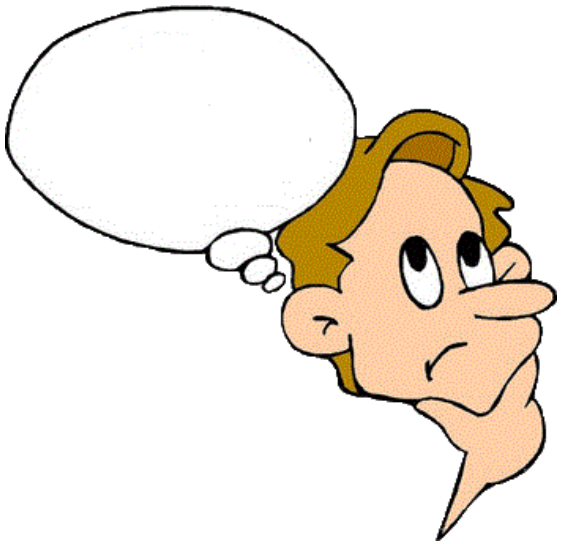

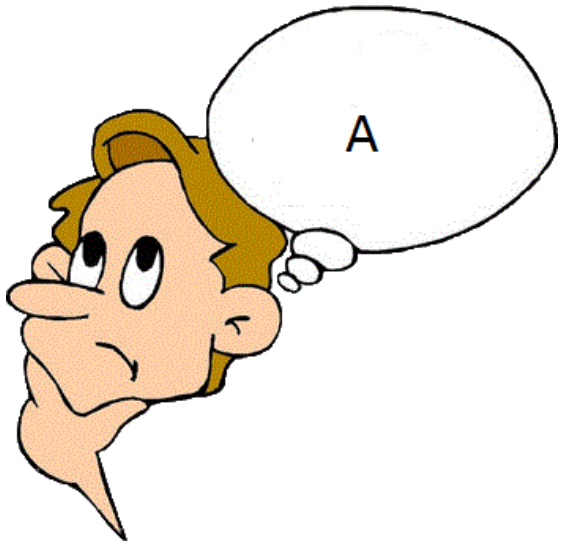

Common knowledge

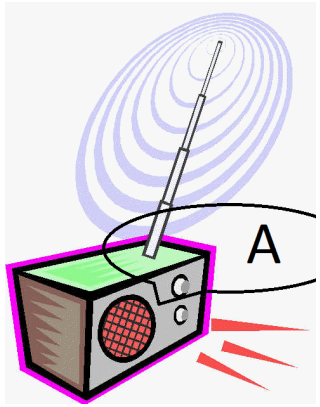

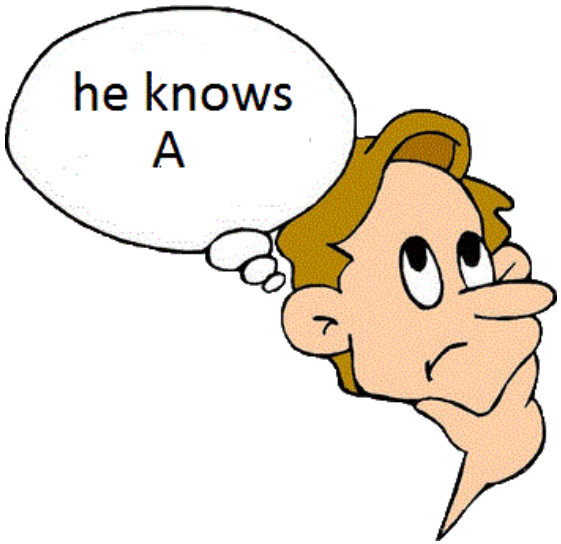

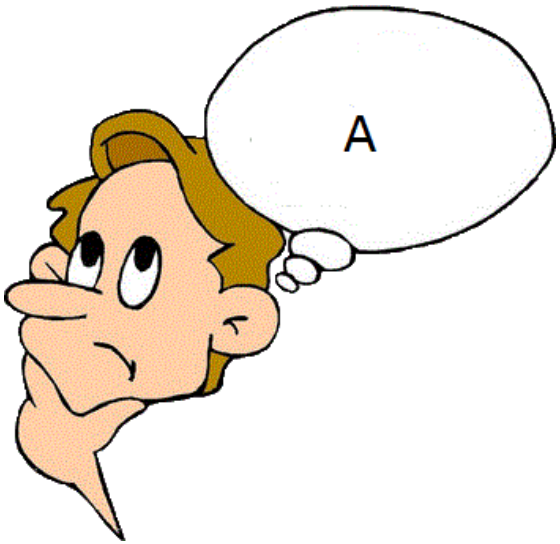

Common knowledge

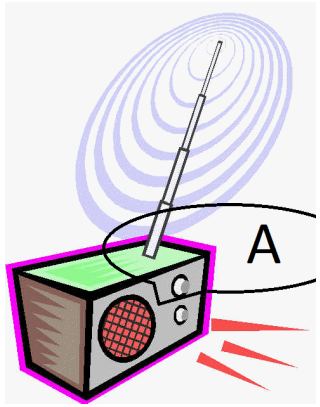

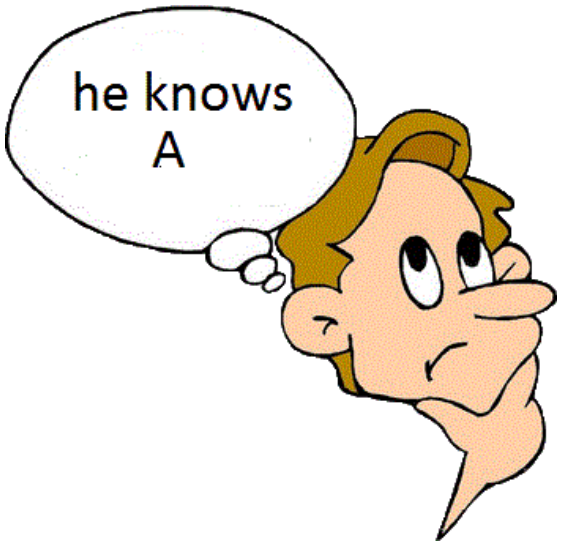

Common knowledge

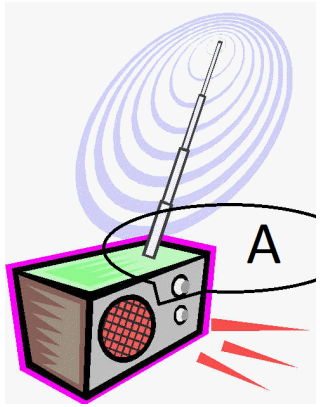

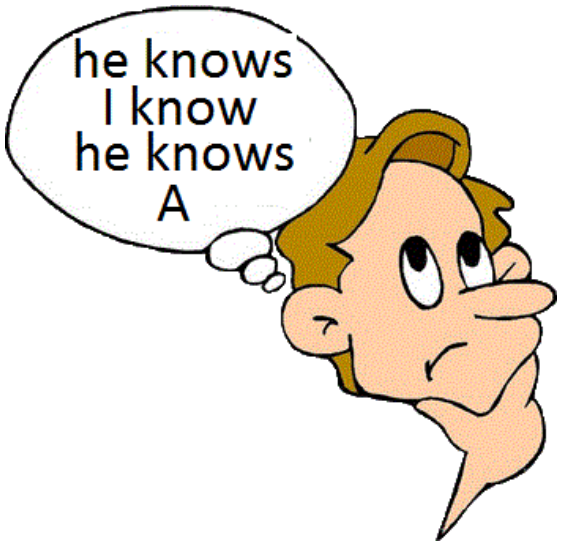

Common knowledge

Informally, common knowledge of a proposition A is defined as the infinitary conjunction

everybody knows A and

everybody knows that everybody knows A and

everybody knows that everybody knows that everybody knows A and

...

This is equivalent to:
Common knowledge of A is the greatest fixed point of

$$\lambda X.\text{everybody knows } A \text{ and everybody knows } X.$$

The Language

$$A ::= p \mid \bar{p} \mid (A \vee A) \mid (A \wedge A) \mid \Diamond_i A \mid \Box_i A \mid \circledast A \mid \boxtimes A$$

Abbreviations:

$$\begin{aligned} \Box A &= \Box_1 A \wedge \ldots \wedge \Box_h A \\ \Diamond A &= \Diamond_1 A \vee \ldots \vee \Diamond_h A \\ \Box^n A &= \underbrace{\Box \ldots \Box}_{n-\text{times}} A \end{aligned}$$

Negation and implication are defined as usual

Coordinated Attack

Suppose two divisions of an army, located in different places, are about to attack their enemy. They have some means of communication but they may be unreliable (i.e. messages may get lost), and the only way to secure a victory is to attack simultaneously. The problem now is: how should generals G and H, who command the two divisions, coordinate their attack? Of course, general G could send a message with the time of attack to general H. We use the proposition del to denote the fact that the message with the time of attack has been delivered. Thus general H upon receiving the message, knows the time of attack, i.e., $\Box_H del$. However, since communication is unreliable, general G considers it possible that his message has not been delivered, i.e., $\neg\Box_G\Box_H del$. But if general H sends an acknowledgment, he in turn cannot be sure whether the acknowledgment has been delivered, i.e., $\neg\Box_H\Box_G\Box_H del$. Hence yet another acknowledgment is needed, and so on.

Coordinated Attack, formally

They attack only if both know that they attack, i.e.,

$$\circledast(att \to \Box att).$$

Thus by $A \wedge \circledast(A \to \Box A) \to \circledast A$ we obtain

$$att \to \circledast att. \tag{1}$$

A further reasonable assumption is that it is common knowledge that neither general attacks unless the message with the time of attack has been delivered, i.e.,

$$\circledast(att \to del).$$

Thus by $\circledast A \wedge \circledast(A \to B) \to \circledast B$ we obtain

$$\circledast att \to \circledast del. \tag{2}$$

Taking (1) and (2) together we obtain

$$att \to \circledast del.$$

They attack only if message delivery is common knowledge.

The Hilbert System H_R

(TAUT) all instances of propositional tautologies

$$\text{(MP)}\ \frac{A \quad A \to B}{B}$$

$$\text{(K)}\quad \Box_i A \wedge \Box_i (A \to B) \to \Box_i B \qquad \text{(NEC)}\ \frac{A}{\Box_i A}$$

$$\text{(CCL)}\quad \boxtimes A \to (\Box A \wedge \Box \boxtimes A)$$

$$\text{(I-R)}\ \frac{B \to (\Box A \wedge \Box B)}{B \to \boxtimes A}$$

Theorem

H_R *is a sound and complete deductive system for common knowledge.*

Induction axiom

We can replace (I-R) with

$$\circledast A \wedge \circledast(A \to B) \to \circledast B$$

$$\frac{A}{\circledast A}$$

$$A \wedge \circledast(A \to \Box A) \to \circledast A$$

The ω-rule: System G_C

$$\Gamma, p, \bar{p} \qquad \wedge\ \frac{\Gamma, A \quad \Gamma, B}{\Gamma, A \wedge B} \qquad \vee\ \frac{\Gamma, A, B}{\Gamma, A \vee B}$$

$$\Box_i\ \frac{\Gamma, \circledast\Delta, A}{\Diamond_i\Gamma, \circledast\Delta, \Box_i A, \Sigma}$$

$$\boxtimes\ \frac{\Gamma, \Box^k A \quad \text{for all } k \geq 1}{\Gamma, \boxtimes A} \qquad \circledast\ \frac{\Gamma, \circledast A, \Diamond A}{\Gamma, \circledast A}$$

Theorem

G_C is a sound and complete deductive system for common knowledge.

The problem of cut-elimination

$$\dfrac{\dfrac{\Box_i \dfrac{\overset{\pi_1}{\vartriangle}}{A, \Gamma, \circledast \bar{B}}}{\Box_i A, \Diamond_i \Gamma, \Sigma, \circledast \bar{B}} \quad \boxtimes \dfrac{\vdots \quad \overset{\pi_{2k}}{\vartriangle} \quad \vdots}{\Box^k B, \Delta \qquad\qquad \boxtimes B, \Delta} \; 1 \leq k < \omega}{\Box_i A, \Diamond_i \Gamma, \Sigma, \Delta}\,\text{cut}$$

Typical cut-elimination procedure yields:

$$\dfrac{\dfrac{\overset{\pi_1}{\vartriangle}}{A, \Gamma, \circledast \bar{B}} \quad \boxtimes \dfrac{\vdots \quad \overset{\pi_{2k}}{\vartriangle} \quad \vdots}{\Box^k B, \Delta \qquad\qquad \boxtimes B, \Delta} \; 1 \leq k < \omega}{\Box_i \dfrac{A, \Gamma, \Delta}{\Box_i A, \Diamond_i \Gamma, \Sigma, \Diamond_i \Delta}}\,\text{cut}$$

System D_C

Nested sequents:

- make $\Box_i$ a structural rule
- allow deep application of rules

Ex: $A, B, [C, [D]_i]_j$ corresponds to $A \vee B \vee \Box_j(C \vee \Box_i D)$

$$\Gamma\{p, \bar{p}\} \qquad \wedge \frac{\Gamma\{A\} \quad \Gamma\{B\}}{\Gamma\{A \wedge B\}} \qquad \vee \frac{\Gamma\{A, B\}}{\Gamma\{A \vee B\}}$$

$$\Box_i \frac{\Gamma\{[A]_i\}}{\Gamma\{\Box_i A\}} \qquad \Diamond_i \frac{\Gamma\{\Diamond_i A, [\Delta, A]_i\}}{\Gamma\{\Diamond_i A, [\Delta]_i\}}$$

$$\circledast \frac{\Gamma\{\Box^k A\} \quad \text{for all } k \geq 1}{\Gamma\{\circledast A\}} \qquad \diamondsuit \frac{\Gamma\{\diamondsuit A, \Diamond^k A\}}{\Gamma\{\diamondsuit A\}}$$

Properties of D_C

Lemma (Structural rules and invertibility)

(i) The rules necessitation, weakening and contraction are admissible for system D_C.
(ii) All rules in D_C *are invertible for* D_C.

Theorem (Cut-elimination for the deep system)

If $\mathsf{D}_\mathsf{C} \vdash^{\alpha}_{\omega \cdot n} \Gamma$, *then* $\mathsf{D}_\mathsf{C} \vdash^{\varphi_1^n(\alpha)}_{0} \Gamma$.

Theorem (Cut-elimination for the shallow system)

If $\mathsf{G}_\mathsf{C} \vdash^{\alpha}_{\omega \cdot n} \Gamma$, *then* $\mathsf{G}_\mathsf{C} \vdash^{\omega \cdot (\varphi_1^n(\omega \cdot \alpha)+1)}_{0} \Gamma$

Theorem (Upper bounds)

If A *is a valid formula, then* $\mathsf{D}_\mathsf{C} \vdash^{<\varphi_2 0}_{0} A$ *and* $\mathsf{G}_\mathsf{C} \vdash^{<\varphi_2 0}_{0} A$.

Cut-elimination on one slide

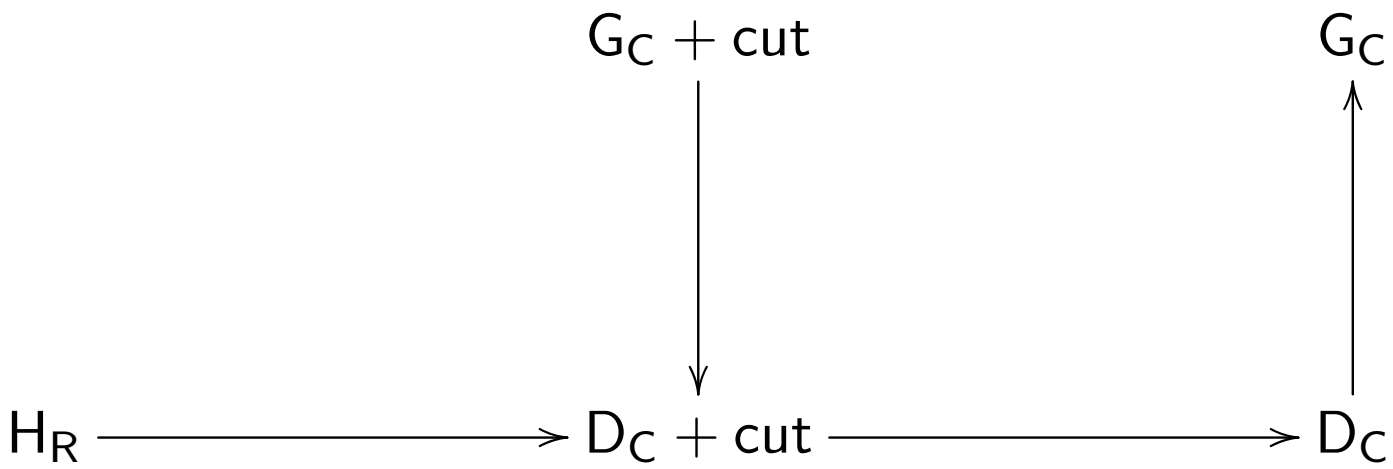

Infinite branches

The infinitary system S:

$$\Gamma, p, \bar{p} \qquad \wedge\ \frac{\Gamma, A \quad \Gamma, B}{\Gamma, A \wedge B} \qquad \vee\ \frac{\Gamma, A, B}{\Gamma, A \vee B}$$

$$\Box_i\ \frac{\Gamma, A}{\Diamond_i \Gamma, \Box_i A, \Sigma}$$

$$\circledast\ \frac{\Gamma, \Box A \wedge \Box \circledast A}{\Gamma, \circledast A} \qquad \diamondsuit\!\!\!\!\circledast\ \frac{\Gamma, \Diamond A \vee \Diamond \diamondsuit\!\!\!\!\circledast A}{\Gamma, \diamondsuit\!\!\!\!\circledast A}$$

Global condition: every infinite branch contains a $\circledast$-thread, i.e. there is a $\circledast A$ unfolded infinitely often.

An S-proof for the induction axiom

$$
\dfrac{
  \dfrac{\text{(ax')}}{\neg A, A}
  \;\text{(}\Box\text{)}
  \quad\text{leads to}\quad
  \Diamond\neg A, \diamondast(A \wedge \Diamond\neg A), \Box A
  \qquad
  \dfrac{\dfrac{\dfrac{\dfrac{
      \dfrac{\text{(ax')}}{\neg A, A, \diamondast(A \wedge \Diamond\neg A), \boxast A}
      \quad
      \dfrac{\vdots}{\neg A, \Diamond\neg A, \diamondast(A \wedge \Diamond\neg A), \boxast A}\,(\boxast)
    }{\neg A, A \wedge \Diamond\neg A, \diamondast(A \wedge \Diamond\neg A), \boxast A}\,(\wedge)
  }{\Diamond\neg A, \Diamond(A \wedge \Diamond\neg A), \Diamond\diamondast(A \wedge \Diamond\neg A), \Box\boxast A}\,(\Box)
  }{\Diamond\neg A, \Diamond(A \wedge \Diamond\neg A) \vee \Diamond\diamondast(A \wedge \Diamond\neg A), \Box\boxast A}\,(\vee)
  }{\Diamond\neg A, \diamondast(A \wedge \Diamond\neg A), \Box\boxast A}\,(\diamondast)
}{
  \dfrac{\Diamond\neg A, \diamondast(A \wedge \Diamond\neg A), \Box A \wedge \Box\boxast A}{\Diamond\neg A, \diamondast(A \wedge \Diamond\neg A), \boxast A}\,(\boxast)
}\,(\wedge)
$$

Completeness for S

Let $\mathcal{T}$ be a proof search tree for Γ. Define an infinite game on it where player I tries to show that Γ is provable.

1. at any $(\Box')$ node, player I chooses one of the children,
2. at any $(\wedge)$ node, player II chooses one of the children,

Such a game results in a path in $\mathcal{T}$. Finite path: player I wins if the path ends in an axiom. Infinite path: player I wins if the path contains a $\boxtimes$-thread.

Theorem

1. *There is a winning strategy for player I if and only if there is an* S*-proof for* Γ *contained in* $\mathcal{T}$*.*
2. *There is a winning strategy for player II if and only if there is an* $\mathsf{S}_{\mathsf{Dis}}$*-disproof for* Γ *contained in* $\mathcal{T}$*.*
3. *The game is determined, i.e. one of the players has a winning strategy.*

Completeness for S

Theorem

S *is a complete deductive system for common knowledge.*

Proof. Let A be a formula that is not provable in S.
The proof search tree for A does not contain a proof for A.
There is no winning strategy for player I.
There must be a winning strategy for player II.
The proof search tree for A contains a $\mathsf{S}_{\mathsf{Dis}}$-disproof for A.
That disproof induces a counter model for A.

The situation for μ

H_μ is a Hilbert system for the modal μ-calculus

Theorem

H_μ *is a sound and complete deductive system for the* μ*-calculus.*

Proof: very involved

The situation for μ

H_μ is a Hilbert system for the modal μ-calculus

Theorem

H_μ *is a sound and complete deductive system for the* μ*-calculus.*

Proof: very involved

G_μ is a Gentzen system (with an ω-rule) for the modal μ-calculus

Theorem

G_μ *is a sound and complete deductive system for the* μ*-calculus.*

Proof of soundness: uses finite model property
Proof of completeness: canonical model construction

The situation for μ (2)

D_μ is a nested sequent system (with an ω-rule) for the modal μ-calculus

Theorem

1. D_μ *is a sound and complete deductive system for the* $\nu\Box$*-fragment (aka continuous fragement).*
2. D_μ *enjoys syntactic cut-elimination.*
3. D_μ *is* not *complete for the modal* μ*-calculus.*

Proofs:

1. Syntactic embedding of the $\nu\Box$-fragment of G_μ
2. Standard
3. Counter example: accessible part may be larger than ω, i.e. the valid formula $\Box(\mu X.\Box X) \to \mu X.\Box X$ is not derivable.

The situation for μ (3)

S_μ is a system with infinite proof branches for the modal μ-calculus

Theorem

S_μ is a sound and complete deductive system for the μ-calculus.

Proof: using determinacy

Finitary Systems

Lemma (Small model property)

There is a function f such that if a formula A is satisfiable, then there exists a model of size at most $f(A)$.

Definition (The system $\mathsf{G_C}^{<\omega}$)

The system $\mathsf{G_C}^{<\omega}$ is defined by replacing the ω-rule in the system $\mathsf{G_C}$ by the rule

$$\frac{\Gamma, \Box^k A \quad \text{for all } 1 \leq k \leq f(\bigvee \Gamma \vee \circledast A)}{\Gamma, \circledast A, \Sigma}$$

Finitary Systems

Lemma (Small model property)

There is a function f such that if a formula A is satisfiable, then there exists a model of size at most $f(A)$.

Definition (The system $\mathsf{G}_\mathsf{C}{}^{<\omega}$)

The system $\mathsf{G}_\mathsf{C}{}^{<\omega}$ is defined by replacing the ω-rule in the system G_C by the rule

$$\frac{\Gamma, \Box^k A \quad \text{for all } 1 \leq k \leq f(\bigvee \Gamma \vee \circledast A)}{\Gamma, \circledast A, \Sigma}$$

Other possibilities

- Use induction rule instead of ω-rule (AlberucciJäger05)
- Reformulate focus games as sequent calculi (BrünnlerLange08)
- Tableau systems (AbateGoréWidman07,GorankoShkatov08)

Why is it so difficult?

Theorem

The logic of common knowledge lacks Craig interpolation.

New ideas are needed to design a nice finitary cut-free system.

Thank you!