Recurrent certainty in games of partial information

Anup Basil Mathew(IMSc) [1]
joint work with
Dietmar Berwanger(LSV, ENS-Cachan)

GImInAL - 7,10 Dec 2013

10 Dec 2013

[1] The author thanks LIA-INFORMEL for funding the visit to LSV during June-July 2013.

An 'Imperfect' Scenario

The Story

Suppose a teacher gives as homework the following problems

1. $\alpha \vee \gamma$
2. $\beta \vee \gamma$

The next day a student is asked one of these questions. For some reason the student only hears "…. $\vee\gamma$". How should the student answer knowing that the teacher is asking one of the homework problems?

How do we model this?

Suppose a teacher gives as homework the following problems

1. $\alpha \vee \gamma$
2. $\beta \vee \gamma$

The next day a student is asked one of these questions. For some reason the student only hears ".... $\vee\gamma$". How should the student answer knowing that the teacher is asking one of the homework problems?

How do we model this?

An 'Imperfect' Scenario

Model of the story

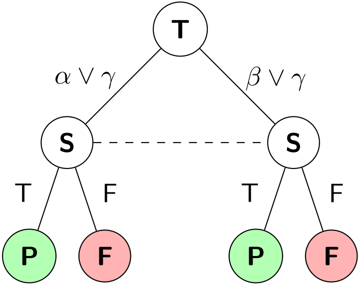

A finite **Extensive Form Game** among '$[n]$' players is described by

$$\mathcal{G} := (\mathcal{T}, turn, (\mathcal{I}_i)_{i \in [n]}, (\succeq_i)_{i \in [n]})$$

where

- $\mathcal{T}$ is a rooted action-labelled finite tree given by (v_0,V,E,l), $l : E \to A$ labels edges with actions from A
- $turn : V \to [n]$ gives the ownership of the nodes of the tree
- $\mathcal{I}_i$, the information partition of player i is a partiton of $\{v \in V | turn(v) = i\}$
- $\succeq_i$ gives preferences of player i over maximal paths(or *plays*) of $\mathcal{T}$.

- A **strategy** $\sigma_i : \mathcal{I}_i \to A$ for player i is a function that assigns an action to every information set $I_i \in \mathcal{I}_i$.
  A **strategy profile** $(\sigma_i)_{i \in [n]}$ is a tuple of strategies, one for each player.
- A **play** $v_0 a_0 v_1 a_1 \ldots a_{n-1} v_n$ **is consistent with strategy** σ_i, if for every v_i with $turn(v_i) = i$ $\sigma_i(I_i) = a_i$ where I_i in the unique partition containing v_i.

- A strategy(σ_i^{dom}) for player i is a **dominant strategy** if $\forall \sigma_i \in \Sigma_i, \forall \sigma_{-i} \in \Sigma_{-i}$, $(\sigma_i^{dom}, \sigma_{-i}) \succeq_i (\sigma_i, \sigma_{-i})$.
- If the preference relation is binary or **win-loss**, then dominant strategy is called winning strategy.
- A team or **coalition of players($S \subset [n]$) is said to have a dominant strategy** if there exists a strategy σ_i^{dom} for each player $i \in S$ such that
$\forall (\sigma_i)_{i \in S} \in (\Sigma_i)_{i \in S}$,
$\forall (\sigma_i)_{i \in [n] \setminus S} \in (\Sigma_{-i})_{i \in [n] \setminus S}$,
$\left((\sigma_i^{dom})_{i \in S}, (\sigma_i)_{i \in [n] \setminus S}\right) \succeq_S \left((\sigma_i)_{i \in S}, (\sigma_i)_{i \in [n] \setminus S}\right)$.
Note that here we assume all players in the coalition have the same preference relation over plays .

Extensive Form Games

Questions of interest

- Does there exist a **dominant strategy for player** i?
- Does there exist a **coordinated winning strategy for a coalition**?

- Does there exist a **dominant strategy for player** i?
- Does there exist a **coordinated winning strategy for a coalition**?

In this talk we are interested in determining the existense of coordinated winning strategy for a restricted class of games, namely **games of perfect recall**.

Let $X_i(v)$ denote the sequence of information sets of player i that are encountered on the path from v_0 to v.

A game is said to have **perfect recall** if for each player i, $X_i(v)$=$X_i(v')$ whenever $\{v, v'\} \subseteq I_i$ for some $I_i \in \mathcal{I}_i$.

Example of imperfect recall

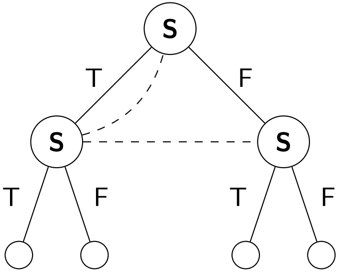

A game is said to have **perfect information** if for every player i, $\forall I_i \in \mathcal{I}_i$ it holds that $|I_i| = 1$.

Perfect information games are 'easy'. Why?

Because winning strategies of subgames can be composed to give winning strategy of the constituent game.

A game is said to have **perfect information** if for every player i, $\forall I_i \in \mathcal{I}_i$ it holds that $|I_i| = 1$.
Perfect information games are 'easy'. Why?
Because winning strategies of subgames can be composed to give winning strategy of the constituent game.

A game is said to have **perfect information** if for every player i, $\forall I_i \in \mathcal{I}_i$ it holds that $|I_i| = 1$.
Perfect information games are 'easy'. Why?
Because winning strategies of subgames can be composed to give winning strategy of the constituent game.

A game is said to have **perfect information** if for every player i, $\forall I_i \in \mathcal{I}_i$ it holds that $|I_i| = 1$.
Perfect information games are 'easy'. Why?
Because winning strategies of subgames can be composed to give winning strategy of the constituent game.

Imperfect information games are 'hard' because of the lack of compositionality.
Example

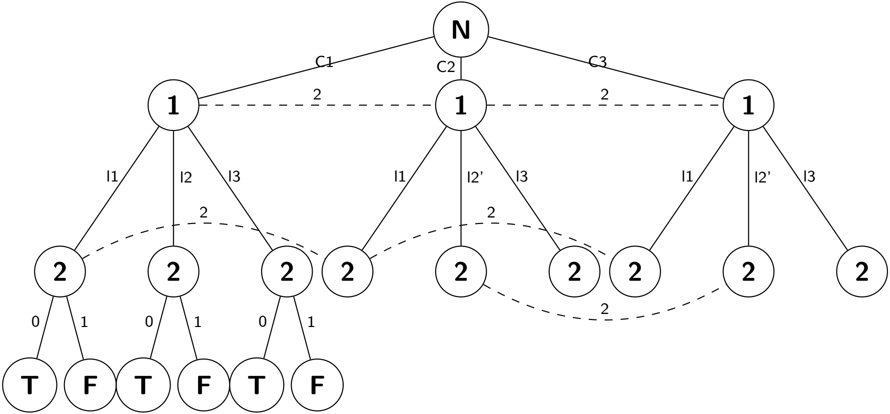

Infinite games via a finite representation.

Game graph

$$\mathcal{A} := (V, A, \delta, (\beta_i)_{i \in [n]}, turn, v_0)$$

where

- V is a finite set of graph nodes and A denotes the actions available to players.
- $\delta : V \times A \to V$ is the transition function on V.
- $\beta_i : V \to \quad B_i$ gives the observables for player i at each state in V where B_i is the set of observables for player $i \in [n]$.

Additionally we assume that the structure of the arena is common knowledge to all players and that the turn functions are 'layered'.

Infinite games via a finite representation.
Game graph

$$\mathcal{A} := (V, A, \delta, (\beta_i)_{i \in [n]}, turn, v_0)$$

where

- V is a finite set of graph nodes and A denotes the actions available to players.
- $\delta : V \times A \rightarrow V$ is the transition function on V.
- $\beta_i : V \rightarrow \quad B_i$ gives the observables for player i at each state in V where B_i is the set of observables for player $i \in [n]$.

Additionally we assume that the structure of the arena is common knowledge to all players and that the turn functions are 'layered'.

Extensive Form Games
Example

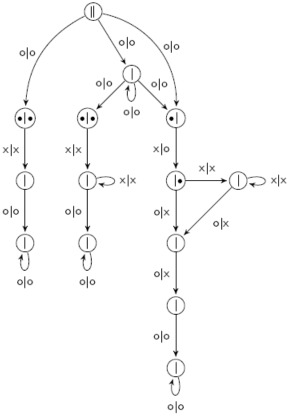

Figure : ⋆ Borrowed from [1]

$$\mathcal{A} \rightarrow \mathcal{G}$$

- Game tree is given by the finite plays on the game graph.
- For every play $\pi := v_0 a_0 v_1 a_1 \ldots$, we define an *i-projection of a play* π as follows $\beta_i(\pi) := \beta_i(v_0, a_0)\beta_i(v_1, a_1)\ldots$ where
$$\beta_i(v_i, a_i) := \begin{array}{ll} \beta_i(v_i)a_i & \text{if } turn(v_i) = i \\ \beta_i(v_i)\square & \text{otherwise.} \end{array}$$
This also gives us an **equivalence relation on finite/infinite plays** given by $\pi_1 \sim_i \pi_2$ iff $\beta_i(\pi_1) = \beta_i(\pi_2)$. We'll call this the *uncertainty relation*.
- Information Partition of player i i.e $\mathcal{I}_i$ is given by equivalence classes of $\sim_i$ over finite plays.

Extensive Form Games

Example

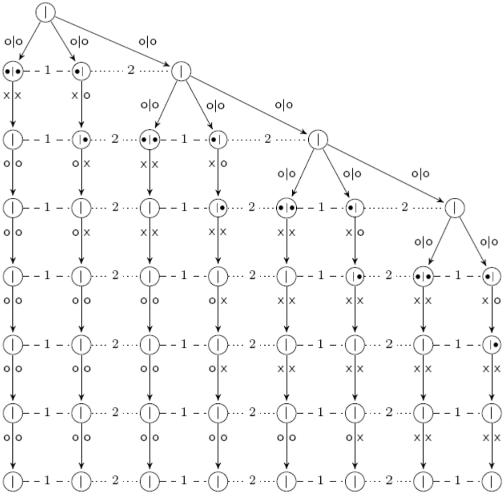

Figure : ⋆ Borrowed from [1]

Winning conditions i.e 'win-loss' preference relation on plays for the team/coalition is given by $W \subseteq Plays(\mathcal{A})$.
We choose the following finite representatiion of winning conditions via $\gamma : V \to \mathbb{N}$.

$$W := \{\pi \in Plays(\mathcal{A}) | \liminf_{i \to \infty} \gamma(v_i) isodd\}$$

We additionally impose the restriction that the **winning condition respects observational equivalence** i.e
$\forall \pi_1 \in W$, if for some $\pi \in Plays(\mathcal{A}), i \in [n]$ s.t $\beta_i(\pi_1) = \beta_i(\pi)$, then $\pi \in W$.

- Compositionality of strategies.
- Analysis upto a certain finite level is enough.
  Memoryless determinacy of parity and mean payoff games: a simple proof:*Henrik Bjorklund, Sven Sandberg, Sergei G. Vorobyov*.

- Compositionality of strategies.
- Analysis upto a certain finite level is enough.
  Memoryless determinacy of parity and mean payoff games: a simple proof:*Henrik Bjorklund, Sven Sandberg, Sergei G. Vorobyov*.

Existence of a winning strategy for a coalition is 'hard'.

- Peterson-Reif
- Pnueli-Rosner
- Berwanger-Kaiser

- Ever growing information sets.
  Not Really
  Since the number of states of the game graph are finite and states are all that determine future possible plays, any equivalence class can be 'shrunk' upto uniqueness.
  But then they are easy.

- Ever growing information sets.
  Not Really
  Since the number of states of the game graph are finite and states are all that determine future possible plays, any equivalence class can be 'shrunk' upto uniqueness.
  But then they are easy.

- Ever growing information sets.
  Not Really
  Since the number of states of the game graph are finite and states are all that determine future possible plays, any equivalence class can be 'shrunk' upto uniqueness.
  But then they are easy.

- Ever growing information sets.
  Not Really
  Since the number of states of the game graph are finite and states are all that determine future possible plays, any equivalence class can be 'shrunk' upto uniqueness.
  But then they are easy.

Extensive Form Games

Why are imperfect information games hard(The real reason)

- Ever growing **Knowledge Hierarchies**.

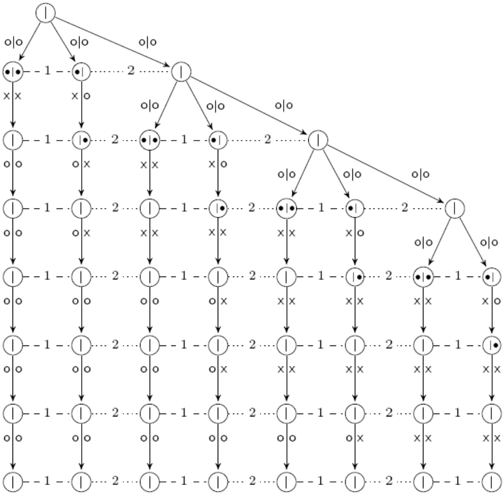

Figure : Borrowed from [1]

Tractable Classes of Imperfect Games

Investigation

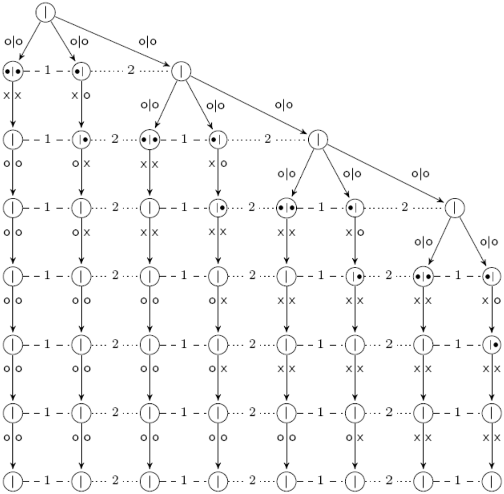

Figure : ⋆ Borrowed from [1]

Tractable Classes of Imperfect Games

Possible tweaks

- Limiting Information Sets for players - (A local property)
  'Recurring certainty' for players.
- Limiting Knowledge Hierarchies - (A 'locally global' property)
  'Hierarchic Knowledge' for the system. Studied since Peterson-Reif.

Tractable Classes of Imperfect Games

Possible tweaks

- Limiting Information Sets for players - (A local property)
  'Recurring certainty' for players.
- Limiting Knowledge Hierarchies - (A 'locally global' property)
  'Hierarchic Knowledge' for the system. Studied since Peterson-Reif.

Tractable Classes of Imperfect Games

Possible tweaks

- Limiting Information Sets for players - (A local property) 'Recurring certainty' for players.
- Limiting Knowledge Hierarchies - (A 'locally global' property) 'Hierarchic Knowledge' for the system. Studied since Peterson-Reif.

Tractable Classes of Imperfect Games

Possible tweaks

- Limiting Information Sets for players - (A local property) 'Recurring certainty' for players.
- Limiting Knowledge Hierarchies - (A 'locally global' property) 'Hierarchic Knowledge' for the system. Studied since Peterson-Reif.

Two questions about the classes:

- Given a game graph $\mathcal{A}$, can we determine whether $\mathcal{G}$ satisfies the above properties.
- Given that $\mathcal{A}$ satisfies the condition, can the coordinated winning strategy be determined.

Recurrent certainty

Definitions

- **An information set** $I_i \in \mathcal{I}_i$ **for player** i **is called certain for player i** if every $\pi_1, \pi_2 \in I_i$, $\beta_i(\pi_1) = \beta_i(\pi_2) \Rightarrow \mathit{last}(\pi_1) = \mathit{last}(\pi_2)$ where $\mathit{last}(\pi)$ gives the position where the finite play π ends.
- We say that a play is *recurrently certain for player i* if there are infinitely many information sets of player i along the play which are certain.

- An *epistemic model* over an arena $\mathcal{A}$ is a kripke structure $\mathcal{K} = (K, (\sim_i)_{i\in[n]})$ where $K \subseteq \mathit{Plays}(\mathcal{A})$ and $(\sim_i)_{i\in[n]}$ is the uncertainty relation restricted to plays in K.

Theorem

Given a game graph $\mathcal{A}$, there is a decidable procedure to test whether recurrent certainty for player i is guaranteed along every play.

Why it works:

- Effective representation of information sets via i-projection of plays.
- Finite witness in plays for uncertain sets.

Corollary

Recurrent certainty implies periodic certainty.

Theorem

Given a game graph $\mathcal{A}$, there is a decidable procedure to test whether recurrent certainty for player i is guaranteed along every play.

Why it works:

- Effective representation of information sets via i-projection of plays.
- Finite witness in plays for uncertain sets.

Corollary

Recurrent certainty implies periodic certainty.

Theorem

Given a game graph $\mathcal{A}$, there is a decidable procedure to test whether recurrent certainty for player i is guaranteed along every play.

Why it works:

- Effective representation of information sets via i-projection of plays.
- Finite witness in plays for uncertain sets.

Corollary

Recurrent certainty implies periodic certainty.

Recurrent certainty

Determining winning strategy

Theorem

If given a game graph $\mathcal{A}$ with the guarantee of recurrent certainty for every play, then "existense of a winning strategy in$\mathcal{A}$ for a coaltion" is a decidable.

Why it works:

- From the corollary in the last section, there is a finite period in the order of the size of the graph within which every player is guaranteed to be certain atleast once. This is enough to prove that the knowledge hierarchies are bounded.
- Reduction to perfect information game with exponential blow-up.

Recurrent certainty

Determining winning strategy

Theorem

If given a game graph $\mathcal{A}$ with the guarantee of recurrent certainty for every play, then "existense of a winning strategy in$\mathcal{A}$ for a coaltion" is a decidable.

Why it works:

- From the corollary in the last section, there is a finite period in the order of the size of the graph within which every player is guaranteed to be certain atleast once. This is enough to prove that the knowledge hierarchies are bounded.
- Reduction to perfect information game with exponential blow-up.

- An epistemic model $\mathcal{K}$ is said to be $(i_1, .., i_n)$-*hierarchic* if it holds in $\mathcal{K}$ that $\sim_{i_1} \subseteq \sim_{i_2} ... \subseteq \sim_{i_n}$ where $(i_1, .., i_n)$ is some permutation of $[n]$. Therefore a sufficient condition for an epistemic model *not* to be hierarchic for $(i_1, .., i_n)$ is the existence of $\pi_1, \pi_2 \in \mathcal{K}$ such that for some $i_m, i_n \in [n]$ with , $\beta_i(\pi_1) = \beta_i(\pi_2) \Rightarrow last(\pi) = last(\pi_1)$
- We say that a play is *recurrently hierarchic* if there are infinitely many epistemic models along the play which are hierarchic.

Recurrent Hierarchicity

Testing Recurrent Hierarchicity

Theorem

Given a game graph $\mathcal{A}$, there is a decidable procedure to test whether recurrent hierarchicity is guaranteed along every play.

Why it works:

- Effective representation of epistemic models? We do this by adding a new player who is as uncertain as any player.
- Finite witness in plays for un-hierarchic epistemic models.

Corollary

Recurrent hierarchicity implies periodic hierarchicity.

Theorem

Given a game graph $\mathcal{A}$, there is a decidable procedure to test whether recurrent hierarchicity is guaranteed along every play.

Why it works:

- Effective representation of epistemic models? We do this by adding a new player who is as uncertain as any player.
- Finite witness in plays for un-hierarchic epistemic models.

Corollary

Recurrent hierarchicity implies periodic hierarchicity.

Theorem

Given a game graph $\mathcal{A}$, there is a decidable procedure to test whether recurrent hierarchicity is guaranteed along every play.

Why it works:

- Effective representation of epistemic models? We do this by adding a new player who is as uncertain as any player.
- Finite witness in plays for un-hierarchic epistemic models.

Corollary

Recurrent hierarchicity implies periodic hierarchicity.

Theorem

Given a game graph $\mathcal{A}$, there is a decidable procedure to test whether recurrent hierarchicity is guaranteed along every play.

Why it works:

- Effective representation of epistemic models? We do this by adding a new player who is as uncertain as any player.
- Finite witness in plays for un-hierarchic epistemic models.

Corollary

Recurrent hierarchicity implies periodic hierarchicity.

Recurrent Hierarchicity

Determining winning strategy

Theorem

If given a game graph $\mathcal{A}$ with the guarantee of recurrent hierarchicity, then "existense of a winning strategy in $\mathcal{A}$ for a coalition" is a decidable.

Why it works:

- We have some notion of homomorphism between epistemic models that preserves "strategic properties". Additionally for hierarchic epistemic models there is a homomorphically equivalent epistemic model that is bounded in the size of the graph game.
- Reduction to perfect information game with non-elementary blow-up.

Recurrent Hierarchicity

Determining winning strategy

Theorem

If given a game graph $\mathcal{A}$ with the guarantee of recurrent hierarchicity, then "existense of a winning strategy in $\mathcal{A}$ for a coalition" is a decidable.

Why it works:

- We have some notion of homomorphism between epistemic models that preserves "strategic properties". Additionally for hierarchic epistemic models there is a homomorphically equivalent epistemic model that is bounded in the size of the graph game.
- Reduction to perfect information game with non-elementary blow-up.

- Why imperfect information games are hard.
- Limiting knowledge hierarchies is the way to go.
- Following this approach to the problem we have two tractable classes:
    - 'Recurrent certainity'.
    - 'Recurrent hierarchicity'

- Why imperfect information games are hard.
- Limiting knowledge hierarchies is the way to go.
- Following this approach to the problem we have two tractable classes:
  - 'Recurrent certainity'.
  - 'Recurrent hierarchicity'

Conclusion
Summary

- Why imperfect information games are hard.
- Limiting knowledge hierarchies is the way to go.
- Following this approach to the problem we have two tractable classes:
  - 'Recurrent certainity'.
  - 'Recurrent hierarchicity'

Conclusion

Critique

- Observable winning conditions.
- Perfect Recall in a distributed setting.
- In "every" known tractable class, the proof via reduction to perfect information games.

- Observable winning conditions.
- Perfect Recall in a distributed setting.
- In "every" known tractable class, the proof via reduction to perfect information games.

- Observable winning conditions.
- Perfect Recall in a distributed setting.
- In "every" known tractable class, the proof via reduction to perfect information games.

- Public announcement makes determining N.E easy. Why?
- What are the "strategic properties" preserved by this approach?
- More natural winning conditions.

- Public announcement makes determining N.E easy. Why?
- What are the "strategic properties" preserved by this approach?
- More natural winning conditions.

- Public announcement makes determining N.E easy. Why?
- What are the "strategic properties" preserved by this approach?
- More natural winning conditions.

References

- D. Berwanger and . Kaiser. Information Tracking in Games on Graphs. Journal of Logic, Language and Information, 2010
- D. Berwanger, . Kaiser, and B. Puchala. A Perfect-Information Construction for Coordination in Games. In 31st IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS '11
- Gary L. Peterson and John H. Reif, Multiple-Person Alternation. 20th Annual IEEE Symposium on Foundations of Computer Science, San Juan, Puerto Rico, October 1979, pp. 348-363.

Questions/Suggestions/Critique

Thank You.