

Admissibility in Infinite Games

Dietmar Berwanger

EPFL, Lausanne

STACS 2007, Aachen

Strategic Interaction

Models of reactive systems:

- **zero-sum** two-player games: strict **competition**
 - optimal behaviour – guarantee a win;
 - determinacy neutralises rationality.
- **non-zero** sum, or n -player games: potential for **cooperation**
 - various solution concepts;
 - how to reason about how other players reason.

Strategic Interaction

Models of reactive systems:

- **zero-sum** two-player games: strict **competition**
 - optimal behaviour – guarantee a win;
 - determinacy neutralises rationality.
- **non-zero** sum, or n -player games: potential for **cooperation**
 - various solution concepts;
 - how to reason about how other players reason.

Understanding of **rationality** needed – borrow from Game Theory?

Yes, but with care.

Computational agents are special players

- non-terminating behaviour
 - open or infinite horizon
- extensive structure
 - state transition graphs
 - inherently sequential
- preplay commitment
- no steady state, no extraneous correlation
- particular payoff structures

Program

Identify criteria for **interactively optimal** behaviour for computational agents.

- adapt deductive solution concepts from noncooperative game theory,
- preserve automata-theoretic and logical foundations.

Program

Identify criteria for **interactively optimal** behaviour for computational agents.

- adapt deductive solution concepts from noncooperative game theory,
- preserve automata-theoretic and logical foundations.

Goals:

- **Describe** what to expect from the interaction of rational agents.
- **Prescribe** individually rational behaviour.
- **Design** mechanisms that promote favourable evolution.

The Minimax principle

Idea: play strategies that **guarantee** your security payoff:

$$\max_{s^0 \in S^0} \min_{s^{-0} \in S^{-0}} .$$

Problems: ignores that players are rational and aware of each other.

- ▶ inefficient and unstable solutions.

Equilibrium

Nash Equilibrium: a **self-enforcing** profile of strategies:

Each player, given the strategies of the others, should not have an alternate strategy that he strictly prefers.

Problems:

- Stable state assumption:
 - the strategies of the others are not given
- Multiplicity
 - coordination failure
- Dynamic inconsistency (in extensive games):
 - Non-credible threats

Rationality Obliges

- Strategy s **dominates** r , if
 - s is not worse than r , against *any* counter-strategy, and
 - s is strictly better, against *some* counter-strategy.
- **Admissibility** criterion: avoid dominated strategies
- Assume **common reasoning** about admissibility.
 - ▶ iterated elimination \leadsto **fixed point**

	L	R
T	+ - +	+ + -
B	+ + +	- - +
		X

	L	R
T	+ - +	- + -
B	+ + +	+ - -
		Y

▶ Reasoning...

Rationality Obliges

- Strategy s **dominates** r , if – given the player's belief
 - s is not worse than r , against *any* counter-strategy, and
 - s is strictly better, against *some* counter-strategy.
- **Admissibility** criterion: avoid dominated strategies
- Assume **common reasoning** about admissibility.
 - ▶ iterated elimination \leadsto **fixed point**

	L	R
T	+ - +	+ + -
B	+ + +	- - +

X

	L	R
T	+ - +	- + -
B	+ + +	+ - -

Y

- ▶ Player Row cannot decide between T and B.

Rationality Obliges

- Strategy s **dominates** r , if – given the player's belief
 - s is not worse than r , against *any* counter-strategy, and
 - s is strictly better, against *some* counter-strategy.
- **Admissibility** criterion: avoid dominated strategies
- Assume **common reasoning** about admissibility.
 - ▶ iterated elimination \leadsto **fixed point**

	L	R	
T	+ - +	+ + -	X
B	+ + +	- - +	

	L	R	
T	+ - +	- + -	Y
B	+ + +	+ - -	

- ▶ Player Column cannot decide between L and R.

Rationality Obliges

- Strategy s **dominates** r , if – given the player's belief
 - s is not worse than r , against *any* counter-strategy, and
 - s is strictly better, against *some* counter-strategy.
- **Admissibility** criterion: avoid dominated strategies
- Assume **common reasoning** about admissibility.
 - ▶ iterated elimination \leadsto **fixed point**

	L	R
T	+ - +	+ + -
B	+ + +	- - +

X

	L	R
T	+ - +	- + -
B	+ + +	+ - -

Y

- ▶ Player Matrix finds x to be better than y .

Rationality Obliges

- Strategy s **dominates** r , if – given the player's belief
 - s is not worse than r , against *any* counter-strategy, and
 - s is strictly better, against *some* counter-strategy.
- **Admissibility** criterion: avoid dominated strategies
- Assume **common reasoning** about admissibility.
 - ▶ iterated elimination \leadsto **fixed point**

	L	R
T	+ - +	+ + -
B	+ + +	- - +
		X

	L	R
T	+ - +	- + -
B	+ + +	+ - -
		Y

- ▶ All players follow his reasoning, and discard the matrix Y .

Rationality Obliges

- Strategy s **dominates** r , if – given the player's belief
 - s is not worse than r , against *any* counter-strategy, and
 - s is strictly better, against *some* counter-strategy.
- **Admissibility** criterion: avoid dominated strategies
- Assume **common reasoning** about admissibility.
 - ▶ iterated elimination \leadsto **fixed point**

	L	R
T	+ - +	+ + -
B	+ + +	- - +

X

	L	R
T	+ - +	- + -
B	+ + +	+ - -

Y

- ▶ Player Row now finds T to be than B.

Rationality Obliges

- Strategy s **dominates** r , if – given the player's belief
 - s is not worse than r , against *any* counter-strategy, and
 - s is strictly better, against *some* counter-strategy.
- **Admissibility** criterion: avoid dominated strategies
- Assume **common reasoning** about admissibility.
 - ▶ iterated elimination \leadsto **fixed point**

	L	R
T	+ - +	+ + -
B	+ + +	- - +

X

	L	R
T	+ - +	- + -
B	+ + +	+ - -

Y

- ▶ Row discards the row B and Column follows his reasoning.

Rationality Obliges

- Strategy s **dominates** r , if – given the player's belief
 - s is not worse than r , against *any* counter-strategy, and
 - s is strictly better, against *some* counter-strategy.
- **Admissibility** criterion: avoid dominated strategies
- Assume **common reasoning** about admissibility.
 - ▶ iterated elimination \leadsto **fixed point**

	L	R
T	+ - +	+ + -
B	+ + +	- - +
		X

	L	R
T	+ - +	- + -
B	+ + +	+ - -
		Y

- ▶ Player Column now finds **R** to be better than **L**

Rationality Obliges

- Strategy s **dominates** r , if – given the player's belief
 - s is not worse than r , against *any* counter-strategy, and
 - s is strictly better, against *some* counter-strategy.
- **Admissibility** criterion: avoid dominated strategies
- Assume **common reasoning** about admissibility.
 - ▶ iterated elimination \leadsto **fixed point**

		L	R
T	+ - +	+ + -	
B	+ + +	- - +	
			X

		L	R
T	+ - +	- + -	
B	+ + +	+ - -	
			Y

- ▶ The predicted outcome is (T, R, X).

Infinite Sequential Games

- ▶ n - players interact to form an infinite path in a graph
each player has preferences over outgoing path – may be conciliated

Strategies: ω -trees s^i (colouring the graph unravelling) – **space** S^i uncountable

- ▶ *perfect information, win-or-lose objectives*

Particular case: **regular** games over **finite graphs**: $(\mathcal{G}, (W^i)_{i < n})$

\mathcal{G} game graph $(V = \dot{\cup}_{i < n} V_i, E)$, $W^i \in V^\omega$ regular set of paths through \mathcal{G}

Canonical case: **parity** games.

Weak Dominance – Admissibility

Fix strategy subspace $Q \subseteq S$ – frame of reference.

- A strategy $s \in S^0$ **dominates** $r \in S^0$ **on** Q^1 , if

$$u^0(s, t) \geq u^0(r, t) \quad \text{for all } t \in Q^1, \text{ and}$$

$$u^0(s, t) > u^0(r, t) \quad \text{for some } t \in Q^1.$$

- $s \in S^0$ is **admissible w.r.t.** Q , if no $s' \in Q^0$ dominates s on Q^1 .

Iterated Admissibility

Incorporate admissibility as principle of rationality.

Assume **common knowledge** of rationality – **stages**:

- $Q_0^i := S^i$
- $Q_{\alpha+1}^i := \{s \in Q_\alpha^i : s \text{ admissible w.r.t. } Q_\alpha\}$
- $Q_\lambda^i := \bigcap_{\alpha < \lambda} Q_\alpha^i$

► Reach deflationary fixed point Q_∞ when $Q_\alpha = Q_{\alpha+1}$.

Solution concept: Q_∞^i – iteratively admissible strategies

(Strategies $s \in Q_\alpha^i$ are α -admissible.)

Issues

- **Procedural** soundness
 - non-emptiness $Q_\alpha \neq \emptyset$;
 - progress guarantee.
- **Metatheory**
 - representation independence;
 - stability under irrational deviations;
 - concept justifiable by player introspection.
- **Computational** aspects
 - automatic recognition and synthesis of IAS;
 - restriction to manageable strategies.

Answers

In the general case

- **One-step soundness:** Each strategy eliminated in stage α is dominated by an $\alpha + 1$ -admissible strategy.
- **Non-stagnation:** Every sequence of $3n$ stages solves a new subgame.

Answers

In the general case

- **One-step soundness:** Each strategy eliminated in stage α is dominated by an $\alpha + 1$ -admissible strategy.
- **Non-stagnation:** Every sequence of $3n$ stages solves a new subgame.

Regular games over finite graphs

- **Foundedness:** Games stabilise in a finite number of stages.
- **Finite-state compatibility:** The set of IA strategies is regular.

Key: Value Characterisation, Truncation

A strategy subspace $Q \subseteq S$ associates values to rooted subgames:

$$\text{value}^i(Q, p) = \begin{cases} +1 & \text{player } i \text{ has a winning strategy in } (\Gamma, Q) \upharpoonright_p \\ -1 & \text{player } i \text{ loses all plays in } (\Gamma, Q) \upharpoonright_p \\ 0 & \text{otherwise} \end{cases}$$

Characterisation

s is admissible w.r.t. $Q \iff \text{value}^i(Q, p) = \text{value}^i(\{s\}, p)$.

Truncation

Substitution of equivalent subgames preserves α -admissibility.

One-step Soundness

Theorem

Each strategy eliminated in stage α is dominated by an $(\alpha + 1)$ -admissible strategy.

Proof:

- Set out with $r \in Q_\alpha^i \setminus Q_{\alpha+1}^i$.
- Construct ascending domination chain $r = s_0, s_1, \dots$
if s_ℓ does not attain value of Q at a position of depth ℓ , shift to some strategy from Q^i that does
- The pointwise limit σ of this sequence is fine:
 - ▶ $\sigma \in Q_\alpha^0$,
 - ▶ σ is $(\alpha + 1)$ -admissible, and
 - ▶ σ dominates r on Q_α^1 .

Non-Stagnation

Theorem

Every sequence of $3n$ stages solves a new subgame.

When is a subgame $\Gamma \upharpoonright_p$ solved?

- **For player i :** at stage α , such that for every $\beta \geq \alpha$:

$$p \text{ reachable in } Q_\beta^i \implies Q_\beta^i \upharpoonright_p = Q_\alpha^i \upharpoonright_p .$$

- **Globally:** maximum over all these α .

Examples:

- p is found to be winning/losing/unreachable for player i
- strategies in $Q_\alpha \upharpoonright_p$ are self-supporting

Non-Stagnation – Proof idea

Two-player case: some position is solved for player 0:

Assume $\text{value}^0(Q_\beta, \cdot)$ remains unchanged for $\alpha \leq \beta \leq \alpha + 3$.

Then eliminated strategies are either

(A) non-winning from positions of $\text{value}^0(Q, p) = +1$

- ▶ can occur only at stage α .

(B) non-pending from positions of $\text{value}^0(Q, p) = 0$

- ▶ triggered at stage $\alpha + 1$ by eliminations of type (A) in stage α , and
- ▶ in stage $\alpha + 2$ due to eliminations of type (B) in stage $\alpha + 1$
- ▶ and never more.

n -players case: recurse into solved subgames (involve one player less).

Finite-state compatibility

Theorem

*In games on finite graphs with **regular** winning conditions, the sets of IA strategies are regular.*

Proof: Reduction to multi-**parity** games

- For regular Q^i , the set of admissible strategies is regular.
- Iteration terminates after after $3n|G|$ stages:
 - ▶ parity games are prefix independent;
 - ▶ for bisimilar positions p, p' (reachable!) $Q_\alpha \upharpoonright_p = Q_\alpha \upharpoonright_{p'}$;
 - ▶ if p closes, then all bisimilar ancestors close;
 - ▶ there are at most $|G|$ bisimulation classes.

Conclusion & Outlook

- Admissibility generalises to infinite games of perfect information.
- The procedure is sound and robust.

Work in progress

- foundation of **backwards induction** for perfect information games;
- **forward inductive** solution for infinite signalling games;

Future

- tractable automata construction;
- synthesisable strategies;