

INTERACTIVE ENVIRONMENTS AND DISTRIBUTED CONTROL

Heinrich H. Nax

Bary S. R. Pradelski

&

hnax@ethz.ch

bpradelski@ethz.ch

April 3, 2017: Lecture 7



Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Lecture logic

Topic

- Use of social science (game theory) in control theory
- Literature and model overview
- Illustrations/ applications

Audience

- Background in optimization (not necessarily distributed control)
- The game theory you have learned so far


Appeal

- Different from centralized and top-down optimization techniques
- Important new area (interdisciplinary)
- Insights from and into social sciences and behavioral studies

Thanks and acknowledgements

- Special thanks to Jeff Shamma!

Game theory describes interactions



Stock market:

- Individuals (traders)
- Strategies (buy/sell)
- Outcome (profit/loss)

wired.co.uk

Fun and games:

- Players (hands)
- Strategies (rock-paper-scissors)
- Outcome (winner/loser)

Shamma

Game theory and (distributed) control...



Biology:

- Individuals (honeybees)
- Strategies (foraging nectar)
- Outcome (survival)

beecare.bayer.com

CONTROL THEORY:

- Distributed agents (turbines)
- Actions (orientation)
- System performance (energy)

studyindenmark.dk

Distributed/networked control systems

Broad based:

- *Non-human*: UVs, sensor networks, communication, ...
- *Human*: Communication, contagion, financial, ...
- *Hybrid*: Transportation, energy, ...

Game theory is on the rise in control theory...

did some accounting at recent IEEE Conference on Decision and Control:

- Networked Control I, II, III, IV, V, VI
- Agents and Autonomous Systems I, II, III, IV, V
- Distributed Control I, II, III
- Decentralized Dynamics and Optimization in Networks I, II, III
- Decentralized Control I, II, III
- Modeling, Coordination and Consensus in Multi-Component Systems I, II
- Distributed Coordination, Networked Interaction, and Games
- *Game Theory and Networked Systems (tutorial)*

Distributed control applications


- *Characteristics:*
 - Multiple decision making elements
 - Interdependency
 - No central authority
 - Distributed information
 - Collective performance



From the analyst's point of view, this constitutes a "game"!


Centralized versus distributed control

Optimization



VS

Decentralization



- *Distributed information*
 - Costly (time, energy, etc) communication
 - Not just “multi component”
 - Not just “graph structure”
- *Efficiency loss*
 - Tragedy of the commons
 - Price of Anarchy

Aims of today's lecture

- *Understand the common principles of distributed control applications:*
routing, flocking, formation, coverage, assignment, cooperation, ...
- *Walk through details of one application:* wind farm
- *Understand the game-theoretic parallels:* players, actions, outcomes



Required reading (please ask for more!):

Marden & Shamma, “Game theory and distributed control”, *Handbook of Game Theory IV*, Young & Zamir (Eds.), 2015.

Game theory and distributed systems

“... the study of mathematical models of conflict and cooperation between intelligent rational decision-makers.”

Myerson, *Game Theory*, 1991.


“...systems are characterized by decentralization in available information, multiplicity of decision makers, and individuality of objective functions for each decision maker.”

Saksena, O'Reilly, & Kokotovic, *Automatica*, 1984.

Game theory: distributed efficiency loss

Local objectives \neq collective objective

Braess Paradox



- 60 people from S to D
- No middle road:
NE – 90 mins
- With middle road:
NE – 119/120 mins

New road worsens congestion

Game theory: *Descriptive to prescriptive*

Descriptive vs *prescriptive* agenda

- Stylized models of societal situations
- Emphasis on new insights
- Not necessarily design tool
- Design only once behavior understood (auctions, matching, ...)

“The word model sounds more scientific than the word fable or tale, but I think we are talking about the same thing.”

Rubinstein, *Economic Fables*, 2012.

Game theory: Basic concepts

Elements:

- Players/Agents/Actors/Individuals
- Actions/Strategies/Choices/Decisions
- *Individual* preferences over *joint* choices (payoffs, utility functions)


Solution concept in a distributed environment

What to expect?

Nash Equilibrium.

Everyone's choice is a best response from an individual perspective *given* the choices of others.

Nash equilibrium & descriptive agenda



“Keynes beauty contest”

- Choose number between 0 and 100
- Winner = Closest to 1/2 of average
- *NE: All pick 0*


Nash equilibrium & descriptive agenda

Game Elements	
Solution Concept	
Rationality	Perception


“Keynes beauty contest”

- Choose number between 0 and 100
- Winner = Closest to 1/2 of average
- *Individual best reply: pick 1/2 of what YOU THINK others' will play*

Repeated beauty contest




First round choices




Fourth vs third round

Nagel, “Unraveling in guessing games: An experimental study”, *AER*, 1995.

“Our” beauty contest from lecture 1



Nash equilibrium & descriptive agenda



“Keynes beauty contest”

- Choose number between 0 and 100
- Winner = Closest to $1/2$ of average
- *Long-run outcome: All pick 0, i.e. NE*

Learning/evolutionary games

Shift of focus:

- Away from solution concept—Nash equilibrium
- Towards how players might arrive to solution—i.e., *dynamics*

“The attainment of equilibrium requires a disequilibrium process.”

Arrow, 1987.

“The explanatory significance of the equilibrium concept depends on the underlying dynamics.”

Skyrms, 1992.

Distributed control: first, identify the target state; second, encourage dynamics that lead to it.

Literature

Monographs:

- Weibull, *Evolutionary Game Theory*, 1997.
- Young, *Individual Strategy and Social Structure*, 1998.
- Fudenberg & Levine, *The Theory of Learning in Games*, 1998.
- Samuelson, *Evolutionary Games and Equilibrium Selection*, 1998.
- Young, *Strategic Learning and Its Limits*, 2004.
- Sandholm, *Population Dynamics and Evolutionary Games*, 2010.

Surveys:

- Hart, “Adaptive heuristics”, *Econometrica*, 2005.
- Fudenberg & Levine, “Learning and equilibrium”, *Annual Review of Economics*, 2009.

Illustration: Fictitious play (1951)

Stages: $t = 0, 1, 2, \dots$

Each player:

- Maintain empirical frequencies (histograms) of opposing actions
- Forecasts (incorrectly) that others play according to observed empirical frequencies
- Selects an action that maximizes expected payoff

Bookkeeping:

$x^i(\cdot)$ = evolving empirical frequency of player i

Discrete-time:

$$x^i(t+1) = x^i(t) + \frac{1}{t+1} (x^i(t) - \mathbf{rand}[\beta^i(x^{-i}(t))])$$

Continuous-time:

$$\frac{dx^i}{dt} = -x^i + \beta^i(x^{-i})$$

Descriptive agenda analysis

Meta-theorem

For *[special structure games]* under *[specific dynamics]*, players exhibit *[asymptotic behavior]*.

Theorem

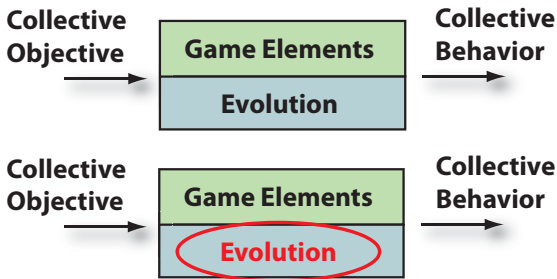
For *zero-sum games* under *fictitious play*, empirical frequencies *converge to NE*.

Many more...

Prescriptive agenda

Design degrees of freedom:

- Game elements: Players, Actions, Preferences
- Evolutionary dynamics: Online adaptation



Potential appeal:

- Distributed self-organization
- Adaptation to environment
- Resilience to disruptions


Prescriptive agenda in action

Theorem

For *potential games* under *restricted movement log linear learning*, joint actions “linger” at *potential maximizer*.

Distributed graph coverage



- Local movements
- Local information exchange
- Linger at maximal coverage



Marden and JSS, “Cooperative control and potential games”, 2009.

Yazicioglu, Egerstedt, and JSS, “A game theoretic approach to distributed coverage of graphs by heterogenous mobile agents”, 2013.


There and back again...



“The *explanatory significance* of the equilibrium concept depends on the underlying dynamics.”

Skyrms, 1992.

How to identify appropriate dynamics?



An application

- *A wind farm:*
 - Each windmill takes a directional orientation and a blade angle
 - Depending on wind direction, this leads to an energy production for each windmill
 - The central authority (for simplicity) aims to maximize the energy total
 - For larger wind farms the centralized control approach has proven unsuccessful



Marden et al. 2013. “A Model-Free Approach to Wind Farm Control Using Game Theoretic Methods”. *IEEE Transactions on Control Systems Technology* 21(4):
“Each turbine does not have access to the functional form of the power generated by the wind farm. This is because the aerodynamic interaction between the turbines is poorly understood. [...] Each turbine may not have access to the choices of other turbines. This is because of the lack of a suitable communication system.”

Bee intermezzo

- *Bees:*
 - Bees fly to different patches of flowers foraging for nectar
 - If nectar per flower is abundant (high payoffs), bees continue in the current patch with high probability
 - If a series of flowers yields low payoff, bees fly far away to a new patch




Rule governs the behavior of bees (Thuijsman et al. JTB 1995)

Shown to be a successful foraging strategy at the population level (implementing NE – Young 2009, even total payoff maximizing NE – Pradelski and Young 2012).

More formally:

- *Game:*

- Players $i = 1, 2, \dots, n$
- Finite strategy set $A_i = \{a_i, b_i, \dots, k_i\}$
- Joint strategy space $A = \prod_i A_i$
- Payoffs $u_i : A \rightarrow R$




How do you get windmills to play this game –giving them private utility functions– so as to maximize total energy production?

The single turbine:

- *Game (given a certain wind direction):*

- Players $i = 1, 2, \dots, n$
(windmills/turbines)
- Finite strategy set $A_i = \{a_i, b_i, \dots, k_i\}$
(orientations)
- Joint strategy space $A = \prod_i A_i$ (wind park configuration)
- Payoffs $u_i : A \rightarrow R$ (own energy production)



The learning rule (pseudo code)

1. Initialize. $t = 0, 1$: each turbine i select a random (benchmark) orientation \bar{a}_i^t resulting in power u_i^t
- Windmill 'moods'. $t + 1 > 1$:
 - if $a_i^t \neq a_i^{t-1}$ or $u_i^t \geq u_i^{t-1}$, windmill 'content'
 - if $a_i^t = a_i^{t-1}$ and $u_i^t < u_i^{t-1}$, windmill 'discontent'
- 2a. Benchmark update. $t + 1 > 1$:
 - if 'content',
 - keep or switch benchmark according to higher payoff
 - if 'discontent',
 - keep old benchmark
2. Action update. $t + 1 > 1$:
 - if 'content',
 - play a_i^t with (high probability) $1 - \epsilon$ and *RAND* with ϵ
 - if 'discontent',
 - windmill plays *RAND* with probability 1

Performance

Theorem. For any desired probability $p < 1$, there exists $\epsilon > 0$ such that, for sufficiently large iterations, total power generated is maximal with at least probability p .


Intuition:

- A series of experiments leads to states with ever higher welfare until someone's payoff goes down.
- That individual becomes discontent, and his searching may cause other agents to become discontent.
- Eventually the discontent agents settle into a new all-content state, where the settling probability increases with the overall welfare of the state.

Alternative approaches: cooperative control

- *Game:*

- Players $i = 1, 2, \dots, n$
- Finite strategy set $A_i = \{a_i, b_i, \dots, k_i\}$
- Joint strategy space $A = \prod_i A_i$
- Payoffs $u_i : A \rightarrow R$ (**total** energy production)



Making windmills play this game –giving them altruistic utility functions– will also maximize total energy production.

Harmony intermezzo

Prisoner's dilemma:

- Recall the difference between the prisoner's dilemma and the harmony game:
 - defection dominant strategy in **prisoner's dilemma**
 - cooperation dominant strategy in **harmony game**

		Confess A	Stay quiet A
Confess B		-6 -6	0 -10
Stay quiet B		0 -10	-2 -2

How to transform a prisoner's dilemma into a harmony game by adding altruism...

Harmony intermezzo

		Confess	Stay quiet
		A	A
Confess	B	-6	0
Stay quiet	B	-10	-2

		Defect	Cooperate
		A	A
Defect	B	10-6	10-10
Cooperate	B	10-10	10-2

Harmony intermezzo

		Confess		Stay quiet	
		A	A	A	A
Confess	B	-6	-6	0	-10
Stay quiet	B	-10	0	-2	-2

		Defect		Cooperate	
		A	A	A	A
Defect	B	4	4	10	0
Cooperate	B	0	10	8	8

Harmony intermezzo

- *Now each player cares for self and other the same way:*
 - Write ϕ_S for payoff for self
 - Write ϕ_O for payoff of other
 - Assume $u_i(\phi_S, \phi_O) = \phi_S + \phi_O$
 - i.e. altruism/other-regarding concern

		Defect		Cooperate	
		A		A	
Defect	B	4+4	4+4	10+0	0+10
	B	0+10	10+0	8+8	8+8
Cooperate	B	0+10	10+0	8+8	8+8
	B	0+10	10+0	8+8	8+8

Harmony intermezzo

- *Now each player cares for self and other the same way:*
 - Write ϕ_S for payoff for self
 - Write ϕ_O for payoff of other
 - Assume $u_i(\phi_S, \phi_O) = \phi_S + \phi_O$
 - i.e. altruism/other-regarding concern

		Defect	Cooperate
		A	A
Defect	B	8	10
Cooperate	B	10	16

Now any dynamic that implements Nash equilibrium in this modified harmony game would maximize total payoffs...

Stag hunt game

Stag hunt dilemma:

- Think of the following *coordination game*:
 - there are two actions: *safe* and *risky*
 - one equilibrium is when **both players play safe**
 - another is when **both players play risky**
 - **risky leads to higher total payoffs**

		Risky	Safe
		A	A
Risky	B	5, 5	0, 4.5
Safe	B	4.5, 0	2, 2

Adding altruism...

Stag hunt modified

- *Now each player cares for self and other the same way:*
 - Write ϕ_S for payoff for self
 - Write ϕ_O for payoff of other
 - Assume $u_i(\phi_S, \phi_O) = \phi_S + \phi_O$
 - i.e. altruism/other-regarding concern

		Risky A	Safe A
Risky B		5+5	0+4.5
Safe B		4.5+0	2+2

Stag hunt modified

- *Now each player cares for self and other the same way:*
 - Write ϕ_S for payoff for self
 - Write ϕ_O for payoff of other
 - Assume $u_i(\phi_S, \phi_O) = \phi_S + \phi_O$
 - i.e. altruism/other-regarding concern

		Risky	Safe
		A	A
Risky	B	10	4.5
	B	4.5	4

Now risky-risky is the unique Nash equilibrium.

Harmony intermezzo

- *Now each player cares for self and other the same way:*
 - Write ϕ_S for payoff for self
 - Write ϕ_O for payoff of other
 - Assume $u_i(\phi_S, \phi_O) = \phi_S + \phi_O$
 - i.e. altruism/other-regarding concern

		Risky	Safe
		A	A
Risky	B	10	4.5
	B	4.5	4

Now any dynamic that implements Nash equilibrium in this game would maximize total payoffs...

But this need not always work

		Risky A	Safe A
Risky B		5	0
Safe B		3	2

		Risky A	Safe A
Risky B		10	3
Safe B		3	4

*Now **not** any dynamic that implements Nash equilibrium in this game would maximize total payoffs – we are back to a selection problem!*

Differences in information

- ***Own energy only:***
 - $u_i(\phi_S) = \phi_S$
 - no information necessary about structure of the game
 - program dynamic offline
 - very specific dynamics will work
 - dynamic requires no feedback
- ***Total energy:***
 - e.g. $u_i(\phi_S) = \phi_S + \phi_O$
 - need to understand structure of the game in order to identify which specification will generate desired equilibria
 - more general class of dynamics will work
 - program dynamic offline
 - dynamic requires feedback about energy total as game continues

Which approach is better depends on the application.


Required reading

Marden & Shamma, “Game theory and distributed control”, *Handbook of Game Theory IV*, Young & Zamir (Eds.), 2015.

Abstract. Game theory has been employed traditionally as a modeling tool for describing and influencing behavior in societal systems. Recently, game theory has emerged as a valuable tool for controlling or prescribing behavior in distributed engineered systems. The rationale for this new perspective stems from the parallels between the underlying decision making architectures in both societal systems and distributed engineered systems. In particular, both settings involve an interconnection of decision making elements whose collective behavior depends on a compilation of local decisions that are based on partial information about each other and the state of the world. Accordingly, there is extensive work in game theory that is relevant to the engineering agenda. Similarities notwithstanding, there remain important differences between the constraints and objectives in societal and engineered systems that require looking at game theoretic methods from a new perspective.

Please ask for more!

Summary: game theory describes interactions



Economics:

- Individuals (traders)
- Strategies (buy/sell)
- Outcome (profit/loss)

wired.co.uk

Mechanism design:

- Players (doctors and hospitals)
- Strategies (applications)
- Outcome (Matching)

NRMP

Game theory and (distributed) control...



Biology:

- Individuals (honeybees)
- Strategies (foraging nectar)
- Outcome (survival)

bee-care.bayer.com



Distributed control:

- Distributed agents (turbines)
- Actions (orientation)
- System performance (energy)

studyindenmark.dk

Broad agenda comparison

	Biology	Social systems	Mechanism design	Distributed control
Game structure	given	given	manipulable	manipulable
Actions	given	given	given	given
Payoffs	given	given	given	manipulable
Information	given	given	manipulable	given

Thanks & Acknowledgements



Bary Pradelski



Jeff Shamma



Peyton Young