International Graduate School on Control, Stuttgart, May 2024

Dissipation Inequalities and Quadratic Constraints for Control, Optimization, and Learning

Lesson 6: Applications to Optimization and Games

Murat Arcak[1] and Peter Seiler[2]

[1] University of California, Berkeley

[2] University of Michigan, Ann Arbor

Learning Objectives

In this lesson you will learn to:

- Apply dissipation inequality and IQC techniques to analyze the stability and convergence rates of optimization algorithms
- Identify Nash equilibria in a class of games, called population games
- Apply dissipation and IQC techniques to analyze convergence to Nash equilibria when agents continually revise their strategies

Outline

- Part I: Optimization
  - Review of first-order methods for convex optimization
  - Feedback perspective for optimization algorithms
  - ρ-hard IQCs
  - Convergence rate bounds using dissipation inequalities and IQCs
  - Example analysis of the heavy-ball algorithm
  - Extensions to stochastic gradient algorithms.

- Part 2: Population Games

Optimization

Goal: Minimize $f: \mathbb{R}^n \to \mathbb{R}$:

$$\min_{x \in \mathbb{R}^n} f(x)$$

Let $S(m, L)$ with $0 < m < L < \infty$ denote functions f such that:

- f is continuously differentiable so the gradient $\nabla f(x)$ exists for all $x \in \mathbb{R}^n$.
- f is m-strongly convex with L-Lipschitz gradients:

$$m\|x_1 - x_2\|^2 \le (\nabla f(x_1) - \nabla f(x_2))^\top (x_1 - x_2) \le L\|x_1 - x_2\|^2 \quad \forall x_1,\, x_2 \in \mathbb{R}^n$$

See the following for details:

- Boyd, Vandenberghe, Convex Optimization. Cambridge University Press, 2004.
- Nesterov, Introductory lectures on convex optimization: A basic course. Springer, 2013.
- Lessard, Recht, Packard, “Analysis and design of optimization algorithms via integral quadratic constraints,” SIAM Journal on Optimization, 2016.

Optimization

Goal: Minimize $f: \mathbb{R}^n \to \mathbb{R}$:

$$\min_{x \in \mathbb{R}^n} f(x)$$

If $f \in S(m, L)$ then there is a unique minimizer $x_* \in \mathbb{R}^n$.

We'll consider first-order methods that seek to iterate toward the minimizer from an initial condition $x_0 \in \mathbb{R}^n$, e.g.

- Gradient descent: $$x_{k+1} = x_k - \alpha \nabla f(x_k)$$
- Heavy ball: $$x_{k+1} = x_k - \alpha \nabla f(x_k) + \beta(x_k - x_{k-1}),$$
- Nesterov's method: $$x_{k+1} = y_k - \alpha \nabla f(y_k)$$ $$y_k = (1 + \beta) x_k - \beta x_{k-1}$$

Optimization

Goal: Minimize $f : \mathbb{R}^n \to \mathbb{R}$:

$$\min_{x \in \mathbb{R}^n} f(x)$$

If $f \in S(m, L)$ then there is a unique minimizer $x_* \in \mathbb{R}^n$.

Definition: Given a function $f \in S(m, L)$, the iterates converge to x_* with rate $\rho \in (0,1)$ if there exists a constant $c > 0$ such that

$$||x_k - x_*|| \leq c \, ||x_0 - x_*|| \, \rho^k \quad \forall x_0 \in \mathbb{R}^n$$

Questions: Do the iterates converge for all functions $f \in S(m, L)$? If yes, then what is the worst-case convergence rate:

$$\sup_{f \in S(m,L)} \rho(f)$$

where $\rho(f)$ denotes the converge rate for a specific $f \in S(m, L)$?

Optimization

Goal: Minimize $f: \mathbb{R}^n \to \mathbb{R}$:

$$\min_{x \in \mathbb{R}^n} f(x)$$

If $f \in S(m, L)$ then there is a unique minimizer $x_* \in \mathbb{R}^n$.

Definition: Given a function $f \in S(m, L)$, the iterates converge to x_* with rate $\rho \in (0,1)$ if there exists a constant $c > 0$ such that

$$||x_k - x_*|| \leq c\, ||x_0 - x_*||\, \rho^k \quad \forall x_0 \in \mathbb{R}^n$$

Questions: Do the iterates converge for all functions $f \in S(m, L)$? If yes, then what is the worst-case convergence rate:

$$\sup_{f \in S(m,L)} \rho(f)$$

where $\rho(f)$ denotes the converge rate for a specific $f \in S(m, L)$?

To simplify notation, we'll assume $x_* = 0$ going forward. This assumption is satisfied using an (unknown) coordinate translation.

Feedback Perspective [1]

Separate the gradient computation from the algorithm update:

$$G \begin{cases} \eta_{k+1} = A\eta_k + Bu_k \\ y_k = C\eta_k \end{cases}$$
$$u_k = \nabla f(y_k),$$

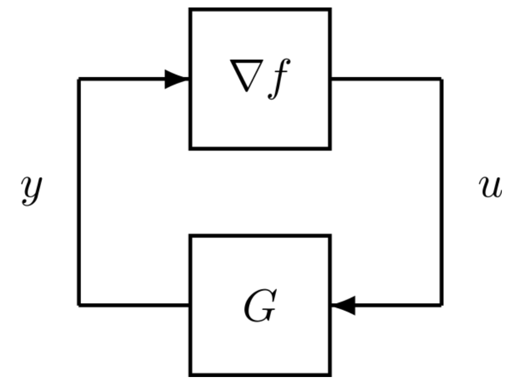

Example: Gradient Descent

$$x_{k+1} = x_k - \alpha \nabla f(x_k)$$

$$\eta_k := x_k \quad \Longrightarrow \quad \begin{aligned} \eta_{k+1} &= \eta_k - \alpha u_k \\ y_k &= \eta_k \\ u_k &= \nabla f(y_k), \end{aligned}$$

[1] Lessard, Recht, Packard, Analysis and Design of Optimization Algorithms via IQCs, SIAM, 2015

Feedback Perspective [1]

Separate the gradient computation from the algorithm update:

$$G \begin{cases} \eta_{k+1} = A\eta_k + Bu_k \\ y_k = C\eta_k \end{cases}$$

$$u_k = \nabla f(y_k),$$

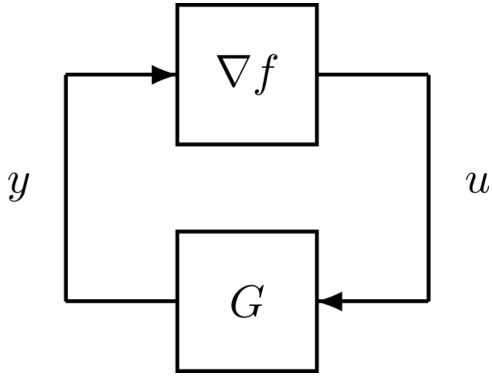

Example: Heavy-ball

$$x_{k+1} = x_k - \alpha \nabla f(x_k) + \beta(x_k - x_{k-1}),$$

$$\eta_k := \begin{bmatrix} x_k \\ x_{k-1} \end{bmatrix} \Rightarrow \begin{aligned} \eta_{k+1} &= \begin{bmatrix} 1+\beta & -\beta \\ 1 & 0 \end{bmatrix} \eta_k + \begin{bmatrix} -\alpha \\ 0 \end{bmatrix} u_k \\ y_k &= \begin{bmatrix} 1 & 0 \end{bmatrix} \eta_k \\ u_k &= \nabla f(y_k), \end{aligned}$$

[1] Lessard, Recht, Packard, Analysis and Design of Optimization Algorithms via IQCs, SIAM, 2015

Feedback Perspective for Optimization

Wake-up Problem

Nesterov's method is:

$$x_{k+1} = y_k - \alpha \nabla f(y_k)$$
$$y_k = (1+\beta)x_k - \beta x_{k-1}$$

Write this algorithm with the gradient computation separated from the algorithm update:

$$G \begin{cases} \eta_{k+1} = A\eta_k + Bu_k \\ y_k = C\eta_k \end{cases}$$
$$u_k = \nabla f(y_k),$$

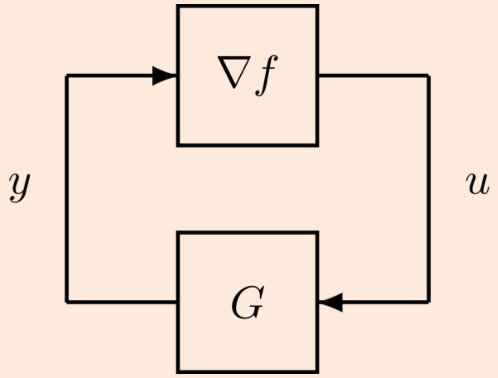

ρ-Hard IQCs

We can adapt IQCs to analyze convergence rates as in [1]. We'll focus on functions $f: \mathbb{R}^n \to \mathbb{R}$ with in $f \in S(m, L)$.

If $y \in \ell_2$ then define $u = \nabla f(y)$ by the sequence $u_k = \nabla f(y_k)$.

Definition: The gradient of $f \in S(m, L)$ satisfies the ρ-Hard IQC defined by a stable filter Ψ and a matrix $J = J^\top \in \mathbb{R}^{(2n)\times(2n)}$ if every $y \in \ell_2$ and $u = \nabla f(y)$ satisfies:

$$\sum_{k=0}^{T} \rho^{-2k} z_k^T J z_k \geq 0 \quad \forall T \geq 0$$

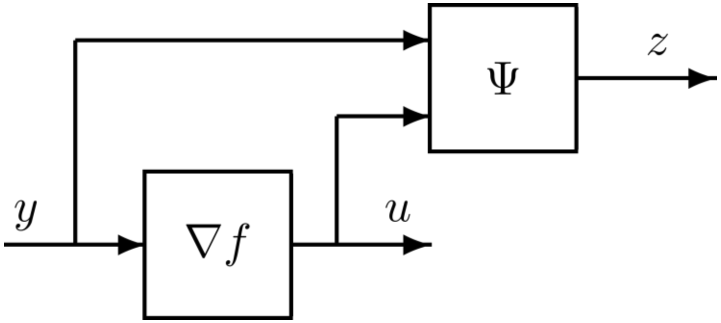

[1] Lessard, Recht, Packard, Analysis and Design of Optimization Algorithms via IQCs, SIAM, 2015

Sector-bound on Gradient

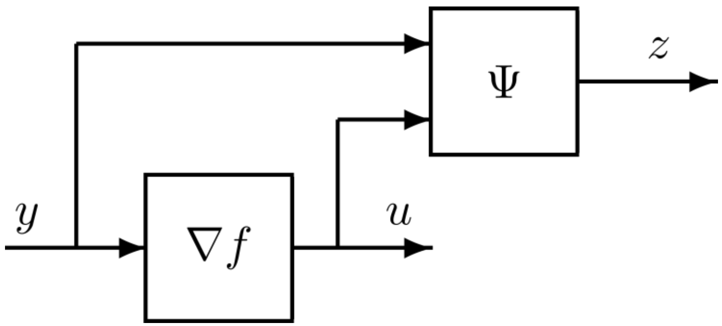

Assume $f \in S(m, L)$. Scalar graph below for illustration.

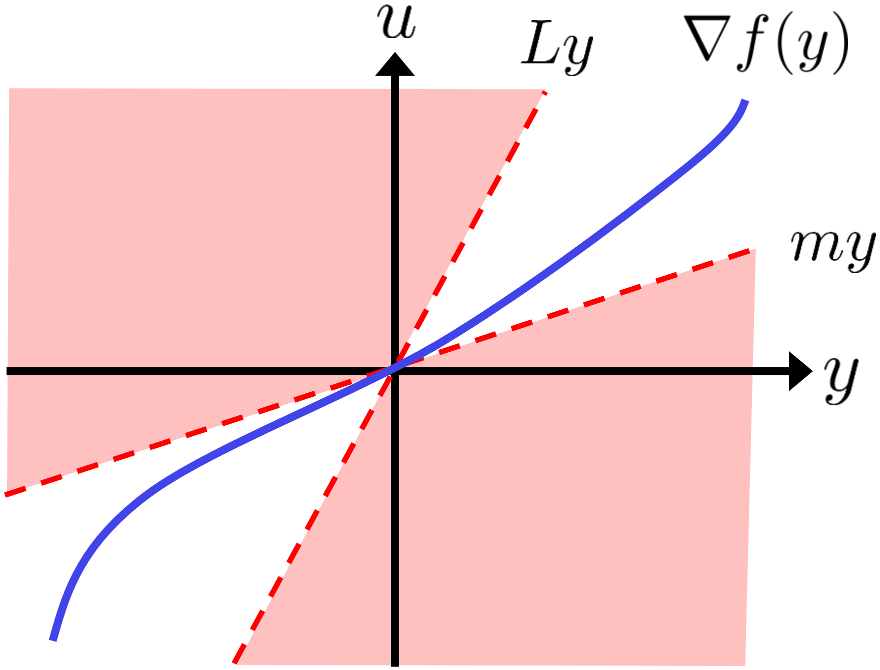

$$m\|y\|^2 \leq u^\top y \leq L\|y\|^2$$
$$\forall y \in \mathbb{R}^n$$

$$\begin{bmatrix} y \\ u \end{bmatrix}^\top \underbrace{\begin{bmatrix} -2mL & m+L \\ m+L & -2 \end{bmatrix}}_{:=J} \begin{bmatrix} y \\ u \end{bmatrix} \geq 0$$

$$\sum_{k=0}^{T} \rho^{-2k} \begin{bmatrix} y_k \\ u_k \end{bmatrix}^\top J \begin{bmatrix} y_k \\ u_k \end{bmatrix} \geq 0$$
$$\forall T \geq 0$$

∇f satisfies the ρ-Hard IQC defined by J and $\Psi = I$.

Dissipation/IQC Condition

Suppose there is a $\lambda \geq 0$ and a storage function $V(\eta) = \eta^\top P \eta$ with $P > 0$ such that the dissipation inequality (DI) holds along trajectories:

$$V(\eta_{k+1}) - \rho^2 V(\eta_k) + \lambda \begin{bmatrix} y_k \\ u_k \end{bmatrix}^\top \begin{bmatrix} -2mL & m+L \\ m+L & -2 \end{bmatrix} \begin{bmatrix} y_k \\ u_k \end{bmatrix} \leq 0$$

Then $\eta_k \to 0$ with rate $\leq \rho$ for all $f \in S(m, L)$.

$$G \begin{cases} \eta_{k+1} = A\eta_k + Bu_k \\ y_k = C\eta_k \\ u_k = \nabla f(y_k), \end{cases}$$

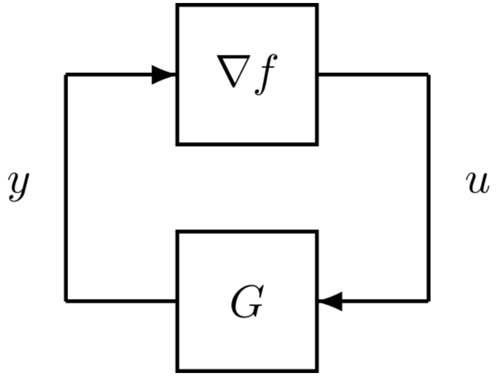

Dissipation/IQC Condition

Suppose there is a $\lambda \geq 0$ and a storage function $V(\eta) = \eta^\top P \eta$ with $P > 0$ such that the dissipation inequality (DI) holds along trajectories:

$$V(\eta_{k+1}) - \rho^2 V(\eta_k) + \lambda \begin{bmatrix} y_k \\ u_k \end{bmatrix}^\top \begin{bmatrix} -2mL & m+L \\ m+L & -2 \end{bmatrix} \begin{bmatrix} y_k \\ u_k \end{bmatrix} \leq 0$$

Then $\eta_k \rightarrow 0$ with rate $\leq \rho$ for all $f \in S(m, L)$.

Proof Sketch:

$$V(\eta_1) - \rho^2 V(\eta_0) + \lambda \begin{bmatrix} y_0 \\ u_0 \end{bmatrix}^\top \begin{bmatrix} -2mL & m+L \\ m+L & -2 \end{bmatrix} \begin{bmatrix} y_0 \\ u_0 \end{bmatrix} \leq 0$$

$$\rho^{-2} V(\eta_2) - V(\eta_1) + \lambda \rho^{-2} \begin{bmatrix} y_1 \\ u_1 \end{bmatrix}^\top \begin{bmatrix} -2mL & m+L \\ m+L & -2 \end{bmatrix} \begin{bmatrix} y_1 \\ u_1 \end{bmatrix} \leq 0$$

$$\vdots$$

$$\rho^{-2T} V(\eta_{T+1}) - \rho^{-2(T-1)} V(\eta_T) + \lambda \rho^{-2T} \begin{bmatrix} y_T \\ u_T \end{bmatrix}^\top \begin{bmatrix} -2mL & m+L \\ m+L & -2 \end{bmatrix} \begin{bmatrix} y_T \\ u_T \end{bmatrix} \leq 0$$

Dissipation/IQC Condition

Suppose there is a $\lambda \geq 0$ and a storage function $V(\eta) = \eta^\top P\, \eta$ with $P > 0$ such that the dissipation inequality (DI) holds along trajectories:

$$V(\eta_{k+1}) - \rho^2 V(\eta_k) + \lambda \begin{bmatrix} y_k \\ u_k \end{bmatrix}^\top \begin{bmatrix} -2mL & m+L \\ m+L & -2 \end{bmatrix} \begin{bmatrix} y_k \\ u_k \end{bmatrix} \leq 0$$

Then $\eta_k \to 0$ with rate $\leq \rho$ for all $f \in S(m, L)$.

Proof Sketch:

$$V(\eta_1) - \rho^2 V(\eta_0) + \lambda \begin{bmatrix} y_0 \\ u_0 \end{bmatrix}^\top \begin{bmatrix} -2mL & m+L \\ m+L & -2 \end{bmatrix} \begin{bmatrix} y_0 \\ u_0 \end{bmatrix} \leq 0$$

$$\rho^{-2} V(\eta_2) - V(\eta_1) + \lambda \rho^{-2} \begin{bmatrix} y_1 \\ u_1 \end{bmatrix}^\top \begin{bmatrix} -2mL & m+L \\ m+L & -2 \end{bmatrix} \begin{bmatrix} y_1 \\ u_1 \end{bmatrix} \leq 0$$

$$\vdots$$

$$\rho^{-2T} V(\eta_{T+1}) - \rho^{-2(T-1)} V(\eta_T) + \lambda \rho^{-2T} \begin{bmatrix} y_T \\ u_T \end{bmatrix}^\top \begin{bmatrix} -2mL & m+L \\ m+L & -2 \end{bmatrix} \begin{bmatrix} y_T \\ u_T \end{bmatrix} \leq 0$$

$$\rho^{-2T} V(\eta_{T+1}) - \rho^2 V(\eta_0) + \lambda \sum_{k=0}^{T} \rho^{-2k} \begin{bmatrix} y_k \\ u_k \end{bmatrix}^\top \begin{bmatrix} -2mL & m+L \\ m+L & -2 \end{bmatrix} \begin{bmatrix} y_k \\ u_k \end{bmatrix} \leq 0$$

Dissipation/IQC Condition

Suppose there is a $\lambda \geq 0$ and a storage function $V(\eta) = \eta^\top P \eta$ with $P > 0$ such that the dissipation inequality (DI) holds along trajectories:

$$V(\eta_{k+1}) - \rho^2 V(\eta_k) + \lambda \begin{bmatrix} y_k \\ u_k \end{bmatrix}^\top \begin{bmatrix} -2mL & m+L \\ m+L & -2 \end{bmatrix} \begin{bmatrix} y_k \\ u_k \end{bmatrix} \leq 0$$

Then $\eta_k \to 0$ with rate $\leq \rho$ for all $f \in S(m, L)$.

Proof Sketch:

$$\rho^{-2T} V(\eta_{T+1}) - \rho^2 V(\eta_0) + \underbrace{\lambda \sum_{k=0}^{T} \rho^{-2k} \begin{bmatrix} y_k \\ u_k \end{bmatrix}^\top \begin{bmatrix} -2mL & m+L \\ m+L & -2 \end{bmatrix} \begin{bmatrix} y_k \\ u_k \end{bmatrix}}_{\geq 0} \leq 0$$

$$\implies V(\eta_{T+1}) \leq \rho^{2T+2} V(\eta_0)$$

Dissipation/IQC Condition

Suppose there is a $\lambda \geq 0$ and a storage function $V(\eta) = \eta^\top P \eta$ with $P > 0$ such that the dissipation inequality (DI) holds along trajectories:

$$V(\eta_{k+1}) - \rho^2 V(\eta_k) + \lambda \begin{bmatrix} y_k \\ u_k \end{bmatrix}^\top \begin{bmatrix} -2mL & m+L \\ m+L & -2 \end{bmatrix} \begin{bmatrix} y_k \\ u_k \end{bmatrix} \leq 0$$

Then $\eta_k \to 0$ with rate $\leq \rho$ for all $f \in S(m, L)$.

Proof Sketch:

$$\rho^{-2T} V(\eta_{T+1}) - \rho^2 V(\eta_0) + \underbrace{\lambda \sum_{k=0}^{T} \rho^{-2k} \begin{bmatrix} y_k \\ u_k \end{bmatrix}^\top \begin{bmatrix} -2mL & m+L \\ m+L & -2 \end{bmatrix} \begin{bmatrix} y_k \\ u_k \end{bmatrix}}_{\geq 0} \leq 0$$

$$\implies V(\eta_{T+1}) \leq \rho^{2T+2} V(\eta_0)$$

Finally, use $\lambda_{min}(P)||\eta||^2 \leq \eta^\top P \eta \leq \lambda_{max}(P)||\eta||^2$ to show:

$$\|\eta_{T+1}\| \leq c\|\eta_0\|\, \rho^{T+1} \text{ where } c := \sqrt{\frac{\lambda_{max}(P)}{\lambda_{min}(P)}}$$

Dissipation/IQC Condition

Wake-up Problem

The Disspation/IQC condition with the sector bound is:

$$V(\eta_{k+1}) - \rho^2 V(\eta_k) + \lambda \begin{bmatrix} y_k \\ u_k \end{bmatrix}^\top \begin{bmatrix} -2mL & m+L \\ m+L & -2 \end{bmatrix} \begin{bmatrix} y_k \\ u_k \end{bmatrix} \leq 0$$

Write the LMI corresponding to this condition. How would you find the best (smallest) convergence rate bound using this condition?

G $\left\{ \begin{aligned} \eta_{k+1} &= A\eta_k + Bu_k \\ y_k &= C\eta_k \end{aligned} \right.$

$$u_k = \nabla f(y_k),$$

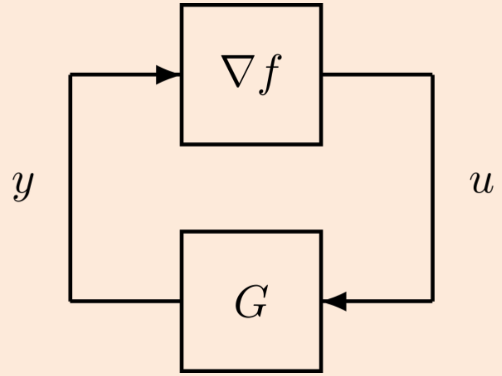

Analysis of Heavy-ball Algorithm

The Heavy-ball algorithm is:

$$x_{k+1} = x_k - \alpha \nabla f(x_k) + \beta(x_k - x_{k-1}),$$

The Heavy-ball parameters (α, β) can be tuned to achieve the optimal rate on quadratic functions in $S(m, L)$:

$$\alpha_0 := \frac{4}{(\sqrt{L} + \sqrt{m})^2} \text{ and } \beta_0 := \left(\frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1}\right)^2,$$

where $\kappa := \frac{L}{m}$ is the condition ratio. These parameters give the following rate on <u>*quadratic functions*</u> in $S(m, L)$*:*

$$\rho = \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1}$$

This is Nesterov's lower bound for the general class $S(m, L)$.

Analysis of Heavy-ball Algorithm

Use the dissipation/IQC condition to find upper bound on the worst-case rate for Heavy-ball with (α_0, β_0) over all functions in $S(m, L)$.

a) We can prove Heavy-ball is convergent up to $\kappa \approx 6$. See [1] for details.

b) Nestorov's lower bound is equal to the rate achieved by Heavy-ball on quadratic functions in $S(m, L)$.

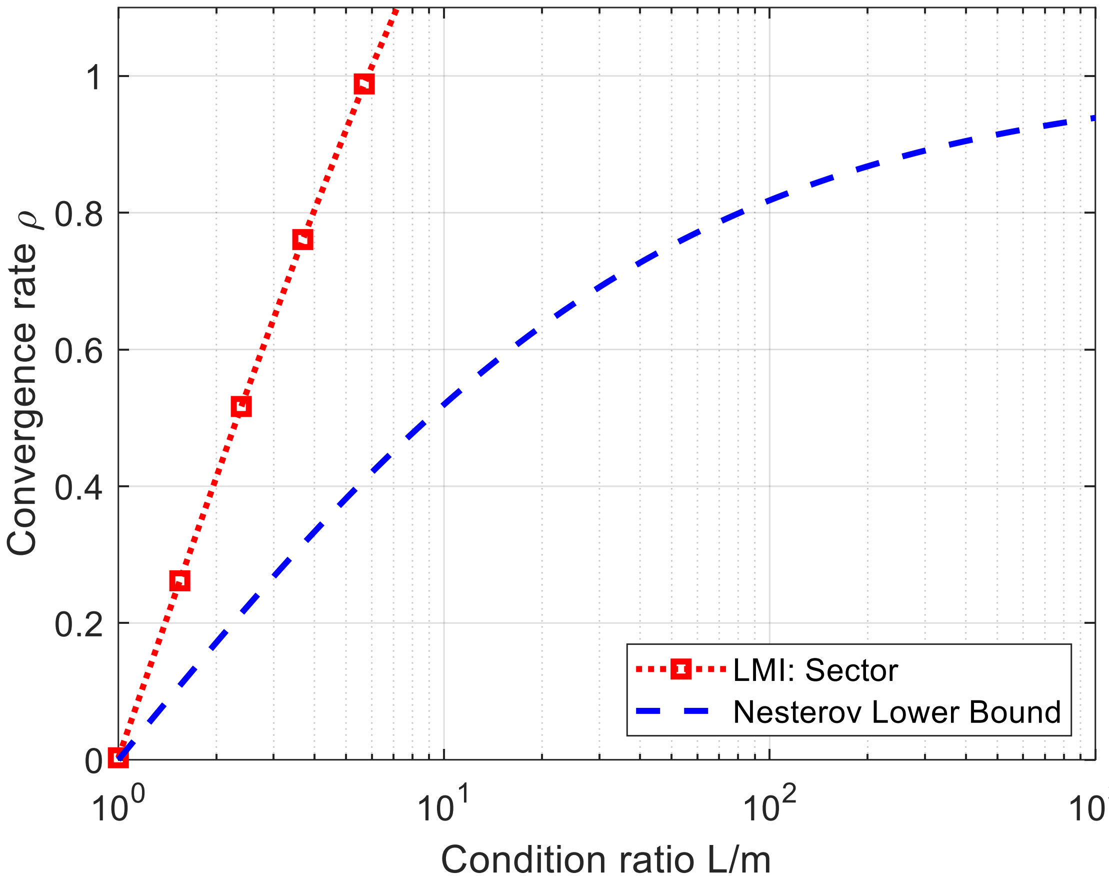

[1] Lessard, Recht, Packard, Analysis and Design of Optimization Algorithms via IQCs, SIAM, 2015

Off-by-1 ρ-Hard IQC [1]

For any $h_1 \in [0, \rho]$ the gradient of $f \in S(m, L)$ satisfies the ρ-Hard IQC defined by

$$\Psi \text{ with } (A_\Psi, B_\Psi, C_\Psi, D_\Psi) = \left(0, \begin{bmatrix} -L & 1 \end{bmatrix}, \begin{bmatrix} h_1 \\ 0 \end{bmatrix}, \begin{bmatrix} L & -1 \\ -m & 1 \end{bmatrix}\right)$$

$$J := \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

Comments:

- This can be combined with the sector ρ-hard IQC in the LMI condition.
- The sector and off-by-1 IQCs are special cases of a larger family of Zames-Falb ρ-hard IQCs.

[1] Lessard, Recht, Packard, Analysis and Design of Optimization Algorithms via IQCs, SIAM, 2015

Analysis of Heavy-ball Algorithm

Use the dissipation/IQC condition to find upper bound on the worst-case rate for Heavy-ball with (α_0, β_0) over all functions in $S(m, L)$.

a) The Weighted off-by-1 IQC improves the bound compared to the sector IQC.

b) We can prove Heavy-ball is convergent up to $\kappa \approx 18$. See [1] for details.

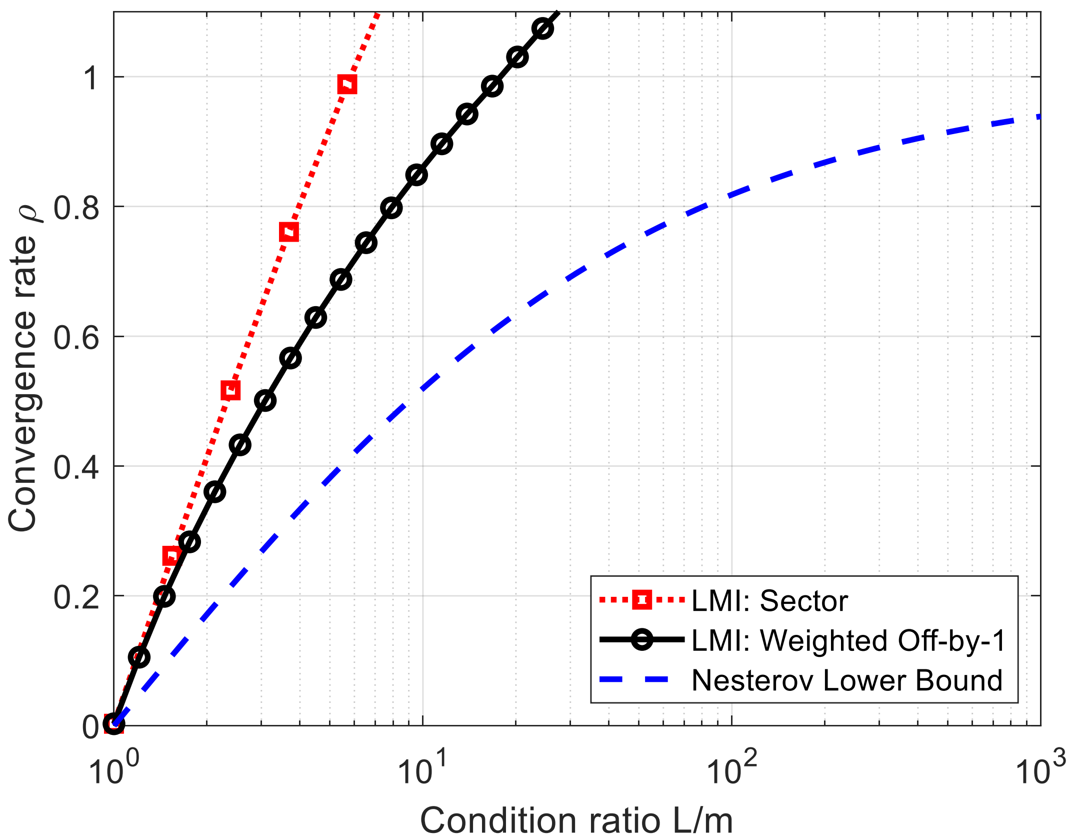

[1] Lessard, Recht, Packard, Analysis and Design of Optimization Algorithms via IQCs, SIAM, 2015

Analysis of Heavy-ball Algorithm

The LMI is feasible iff $\kappa < 9 + 4\sqrt{5} \approx 18$ **and** $\exists f \in S\left(m, \left(9 + 4\sqrt{5}\right) m\right)$ such that Heavy-ball with (α_0, β_0) has a limit cycle [1].

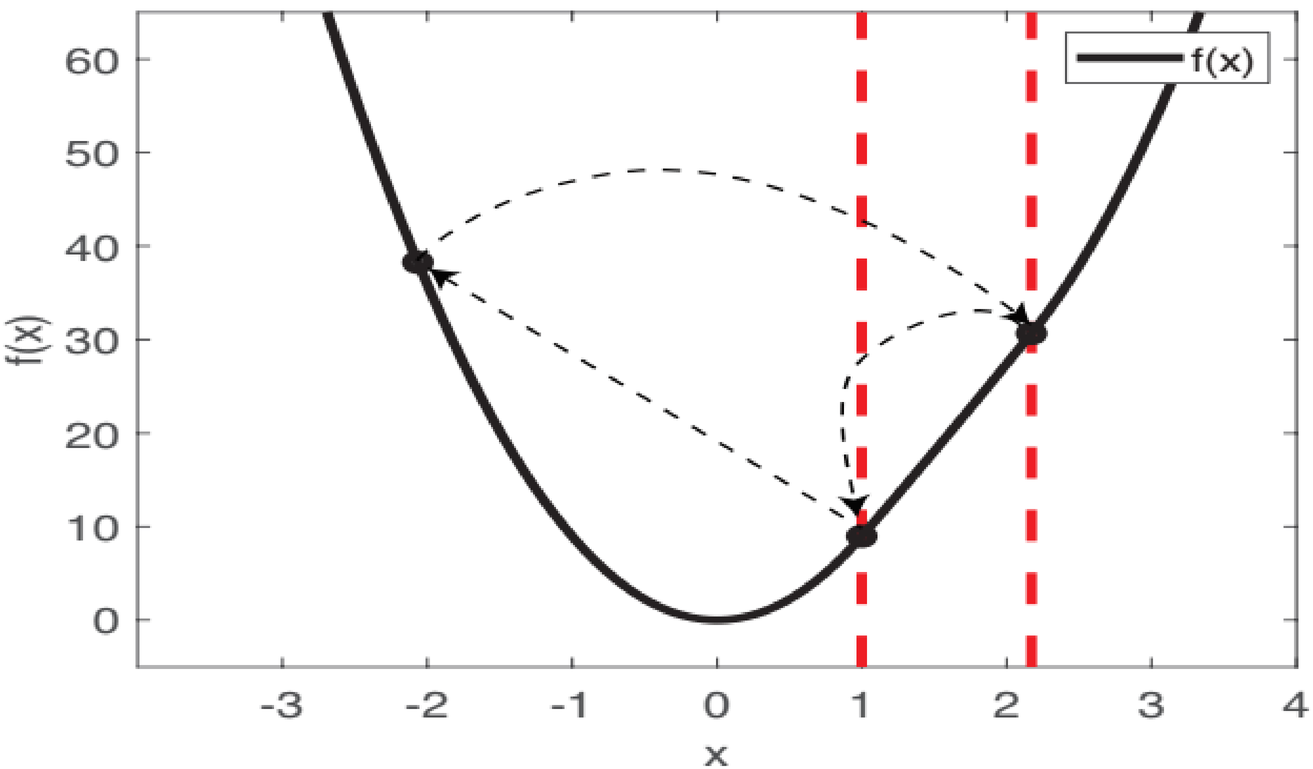

[1] Badithela, Seiler, Analysis of the Heavy-ball Algorithm using IQCs, ACC 2019.

Extension to Stochastic Optimization

Finite Sum Minimization

$$\min_{x \in \mathbf{R}^n} \frac{1}{n} \sum_{i=1}^{n} f^i(x)$$

- Certain convexity/Lipschitz assumptions
- Application to empirical risk minimization

Stochastic Gradient (SG) is widely used

- Fixed stepsize: Convergence to tolerance of optimal
- Decreasing stepsize: Sublinear convergence

Many recent methods (SAGA, Finito, SDCA) with linear convergence and similar iteration cost as SG.

SAGA

Randomly sample i_k at each step

$$x_{k+1} = x_k - \alpha \left(\nabla f^{i_k}(x_k) - y_k^{i_k} + \tfrac{1}{n} \textstyle\sum_{i=1}^{n} y_k^i \right)$$

$$y_{k+1}^i := \begin{cases} \nabla f^i(x_k) & \text{if } i = i_k \\ y_k^i & \text{else} \end{cases}$$

Extension to Stochastic Optimization [1,2]

Express stochastic optimization algorithms with:

- Gradient, ∇f
- Markov Jump System representation for optimization algorithm

$$y_k = \begin{bmatrix} \nabla f^1(u_k) \\ \vdots \\ \nabla f^n(u_k) \end{bmatrix}$$

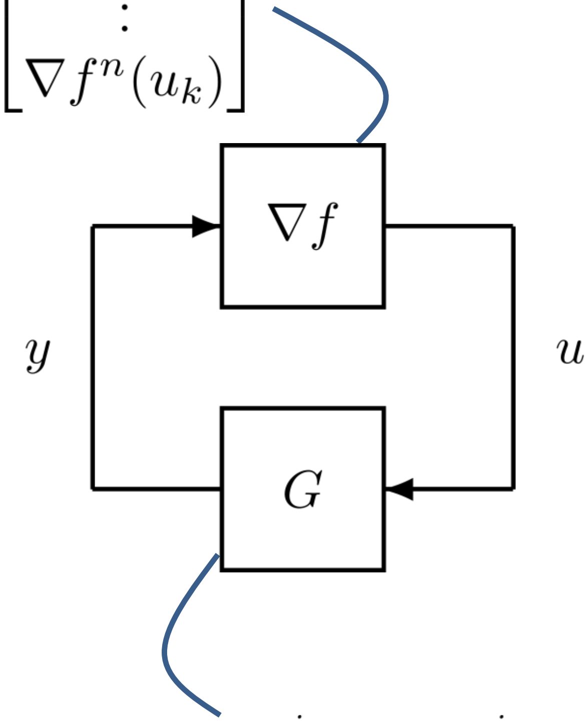

$$\eta_{k+1} = A^{i_k}\eta_k + B^{i_k} y_k$$
$$u_k = C\eta_k$$

Automated Analysis with IQC/SDP

- Characterize ∇f with IQCs
- "Small" SDPs to certify convergence-rate
- Analytical proofs guided by SDP solutions.

[1] Hu, Seiler, Rantzer, A Unified Analysis of Stochastic Optimization Methods Using Jump System Theory and Quadratic Constraints, COLT 2017.

[2] Hu, A Robust Control Perspective on Optimization of Strongly-Convex Functions, Ph.D. , 2016.

Additional Extensions

There are a variety of additional extensions including:

- Van Scoy, Freeman, Lynch, The fastest known globally convergent first-order method for minimizing strongly convex functions," IEEE Control Systems Letters, 2018.
- Drori, Teboulle, Performance of first-order methods for smooth convex minimization: A novel approach, Mathematical Programming, 2014.
- Taylor, Hendrickx, Glineur, Smooth strongly convex interpolation and exact worst-case performance of first-order methods, Mathematical Programming, 2017.
- Mohammadi, Razaviyayn, Jovanovic, Robustness of accelerated first-order algorithms for strongly convex optimization problems, IEEE TAC, 2021.
- Scherer, Ebenbauer, Holicki. Optimization Algorithm Synthesis based on Integral Quadratic Constraints: A Tutorial, CDC, 2023.

Outline

- Part I: Optimization
- Part 2: Population Games
  - What are population games?
  - Best response and Nash equilibria
  - Monotone games
  - Learning rules and evolutionary dynamics
  - Recalling relevant dissipativity notions
  - Convergence to Nash equilibria from dissipativity

What are Population Games?

Continuum model for large numbers of anonymous agents, playing one of a finite number of strategies.

Example:

Congestion Games

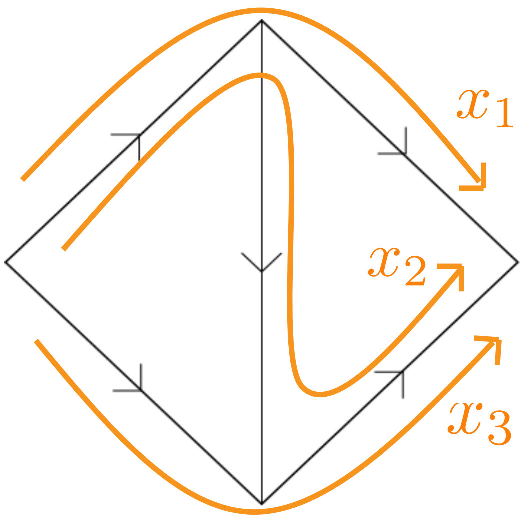

x_i: flow using route

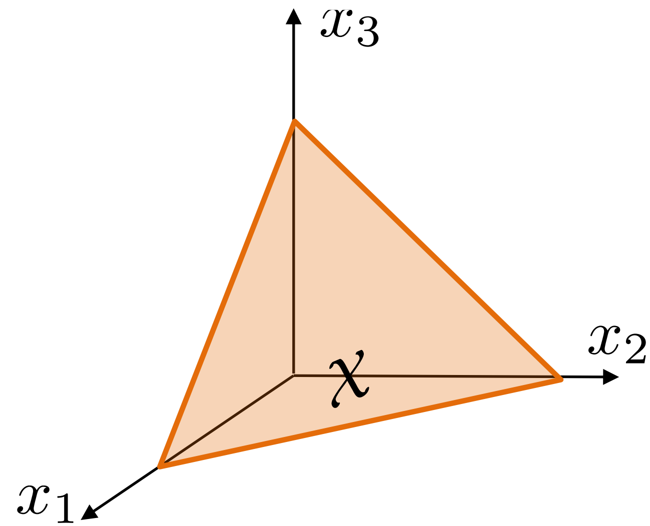

state space $\mathcal{X}$ is a simplex

For $p > 1$ populations (e.g., different origin-destination pairs):

$$\mathcal{X} = \mathcal{X}^1 \times \cdots \times \mathcal{X}^p, \quad \mathcal{X}^r \subset \mathbb{R}^{n^r}$$

n^r: # of strategies available to population r, $n = \sum_{r=1}^{p} n^r$

$x = \{x^r\}_{r=1}^{p} \in \mathcal{X}$ called "social state" and $m^r = \mathbf{1}^\top x^r$ "demand"

What are Population Games?

Example: Congestion games with multiple origin-destination pairs

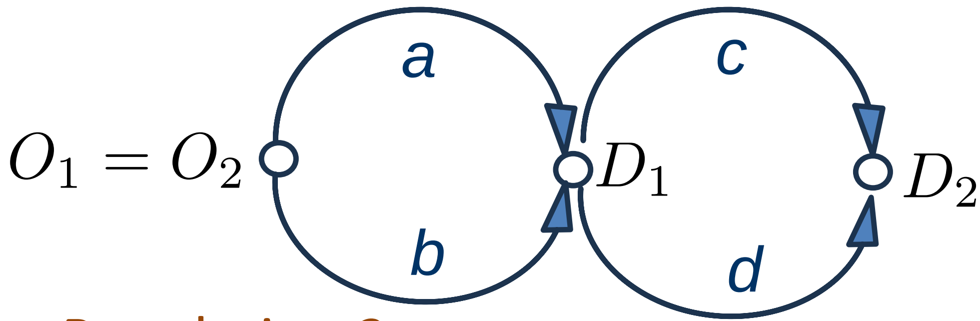

Population 1:

O_1/D_1 :

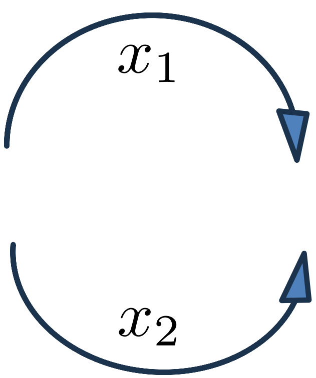

$$x_1 + x_2 = m^1$$

Population 2:

O_2/D_2 :

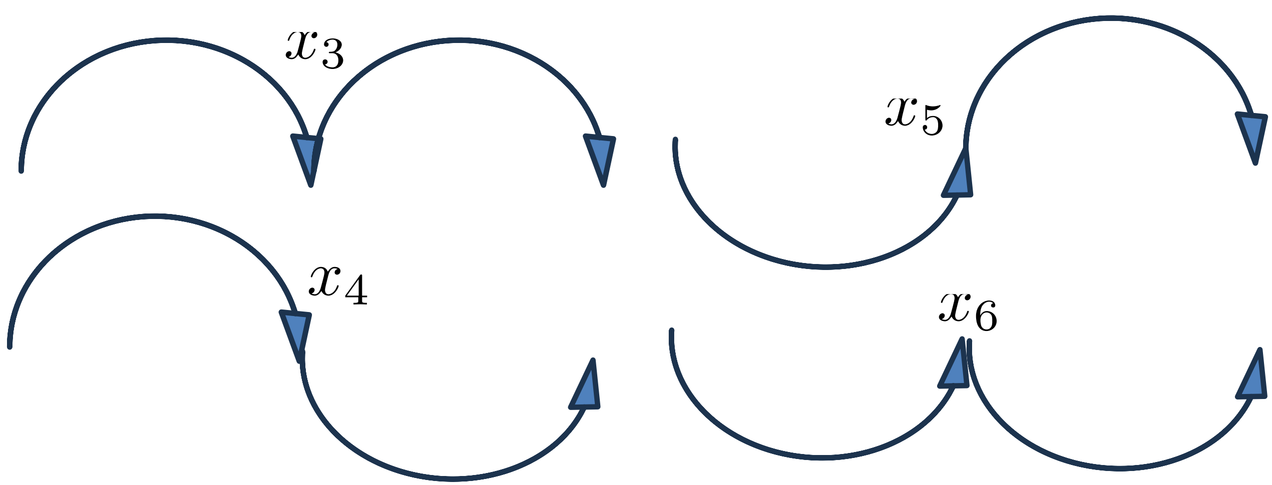

$$x_3 + x_4 + x_5 + x_6 = m^2$$

What are Population Games?

Payoff of each strategy depends on social state through the function

$$F : \mathcal{X} \to \mathbb{R}^n$$

Example: payoff for a route in congestion games is the negation of travel time on that route:

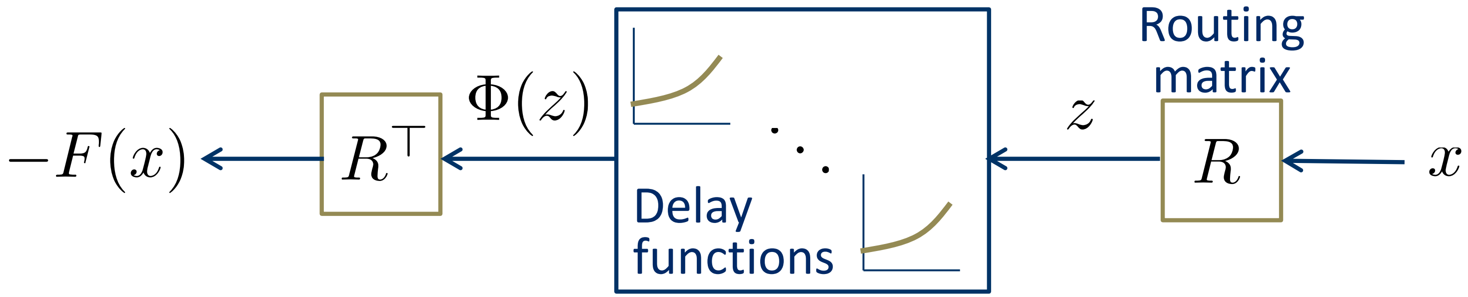

For the network on the previous slide, the payoff for, say, route 3 is:

$$\begin{aligned} F_3(x) &= -\Phi_a(z_a) - \Phi_c(z_c) \\ &= -\Phi_a(x_1 + x_3 + x_4) \\ &\quad - \Phi_c(x_3 + x_5) \end{aligned}$$

$$\begin{bmatrix} z_a \\ z_b \\ z_c \\ z_d \end{bmatrix} = \begin{bmatrix} 1 & 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \end{bmatrix}$$

What are Population Games?

Wake-up Problems

1) Consider the simple network below and draw the state space on the (x_1, x_2) plane, assuming the demand is normalized to $m = 1$.

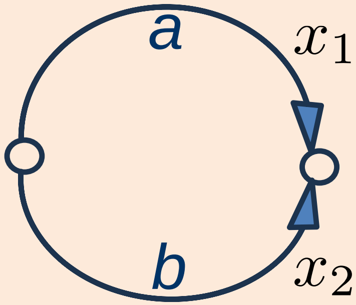

2) Next, write the payoff function $F(x)$ assuming delay functions

$$\Phi_a(x_1) = 1 \qquad \Phi_b(x_2) = x_2$$

i.e., link a is a wide road where travel time doesn't depend on the flow, but link b takes longer when there is more flow on it.

Note each link is a route, so $R = I$ in this simple example.

3) Which road would you take?

Best Response and Nash Equilibria

Best Response: given vector $\pi \in \mathbb{R}^n$ of payoffs for each strategy, $\bar{x} \in \mathcal{X}$ is said to be a best response to π if

$$x^\top \pi \leq \bar{x}^\top \pi \quad \forall x \in \mathcal{X}$$

Example: Suppose $n = 3$ and $\pi_1 = \pi_2 > \pi_3$. Then, every point on this edge is a best response to π.

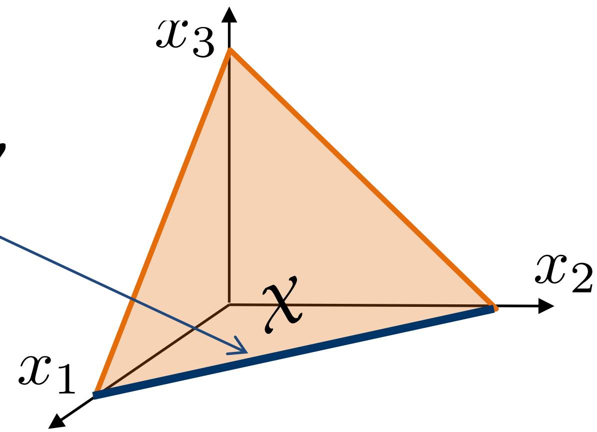

State $\bar{x} \in \mathcal{X}$ is a **Nash Equilibrium** if it is a best response to the payoff at that state:

$$x^\top F(\bar{x}) \leq \bar{x}^\top F(\bar{x}) \quad \forall x \in \mathcal{X}$$

If F is continuous, the set of Nash Equilibria, $NE(F)$, is nonempty by Kakutani's Fixed Point Theorem for set-valued maps.

Best Response and Nash Equilibria

Equivalent characterizations[1] of Nash Equilibria:

1) $\mathrm{NE}(F) = \{x \in \mathcal{X} : \ \exists j \text{ s.t. } F_i(x) < F_j(x) \ \Rightarrow \ x_i = 0\}$

"At equilibrium, a strategy with inferior payoff can't be in use."

2) $\mathrm{NE}(F) = \{x \in \mathcal{X} : \ x_i > 0 \ \Rightarrow \ F_i(x) \geq F_j(x) \ \forall j\}$

"Every strategy in use must earn maximal payoff."

Corollary to Characterization #2:

A point $x \in \mathrm{int}(\mathcal{X})$ is Nash iff $F_1(x) = F_2(x) = \cdots = F_n(x)$

More generally, if we define $\mathcal{S}(x) = \{i \in \{1, \cdots, n\} : \ x_i > 0\}$

$$F_i(x) = F_j(x) \quad \forall i, j \in \mathcal{S}(x)$$

[1]Assuming single population for simplicity

Best Response and Nash Equilibria

Example: Pigou network

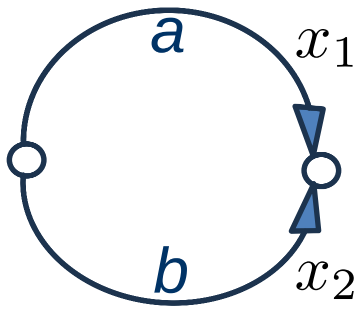

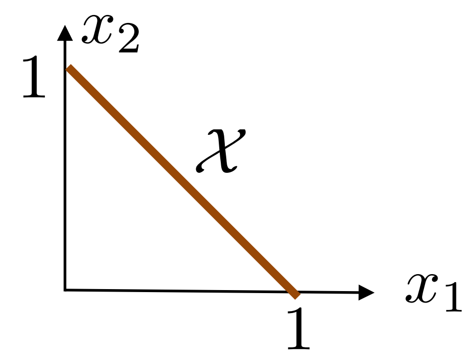

$$\Phi_a(x_1) = 1, \Phi_b(x_2) = x_2 \;\Rightarrow\; F(x) = \begin{bmatrix} -1 & -x_2 \end{bmatrix}^\top$$

$x = (0, 1)$ is a best response to $F(x) = \begin{bmatrix} -1 & -1 \end{bmatrix}^\top$ → Nash ✓

No interior Nash, because $F_1(x) = F_2(x) \Rightarrow x_2 = 1, x_1 = 0$

$x = (1, 0)$ not best response to $F(x) = \begin{bmatrix} -1 & 0 \end{bmatrix}^\top$ → not Nash

This example shows that Nash equilibria may be inefficient: the social optimum is $x = (0.5, 0.5)$, which reduces average travel time to 0.75, but this is not a Nash equilibrium.

Best Response and Nash Equilibria

Example: Rock-Paper-Scissors

$$F(x) = Ax = \begin{bmatrix} 0 & -l & w \\ w & 0 & -l \\ -l & w & 0 \end{bmatrix} \begin{bmatrix} x_R \\ x_P \\ x_S \end{bmatrix} \; w, l > 0, \;\; x_R + x_P + x_S = 1$$

First, look for Nash equilibria in the interior of the simplex.

By Characterization 2, this means $F(x) = Ax = \alpha \mathbf{1}$ for some $\alpha \in \mathbb{R}$

$$\Rightarrow \begin{bmatrix} A & -\mathbf{1} \end{bmatrix} \begin{bmatrix} x \\ \alpha \end{bmatrix} = 0$$

Null space of $\begin{bmatrix} A & -\mathbf{1} \end{bmatrix}$ spanned by $\begin{bmatrix} \mathbf{1} \\ w - \ell \end{bmatrix}$, *i.e.*, $x_R = x_P = x_S$

Since $x_R + x_P + x_S = 1$, unique solution $x_R = x_P = x_S = \dfrac{1}{3}$

Can also show no Nash on the boundary of the simplex. Thus, unique Nash equilibrium.

Best Response and Nash Equilibria

Caption Contest (from the New Yorker magazine)

- Listen you guys, I just went through this with Eenie, Meenie, Minie and Moe.
- This place was so much nicer when you were a priest, a rabbi and an Irishman.
- Pay up front. I know how this ends.

Monotone Games

Population game defined by payoff $F : \mathcal{X} \to \mathbb{R}^n$ is **monotone** if

$$(y - x)^\top (F(y) - F(x)) \leq 0 \quad \forall x, y \in \mathcal{X}$$

and **strictly monotone** if the inequality is strict when $x \neq y$.

Theorem: $\mathrm{NE}(F)$ is a convex set when F is monotone, and a single point when F is strictly monotone.

Suppose F is defined on $\mathbb{R}^n_+$ and continuously differentiable, so the Jacobian matrix $DF(x)$ defined as $DF_{ij}(x) = \dfrac{\partial F_i(x)}{\partial x_j}$ exists.

Then F is monotone iff

$$z^\top (DF(x) + DF(x)^\top) z \leq 0 \quad \forall x \in \mathcal{X}, z \in T\mathcal{X}$$

where $T\mathcal{X}$ is tangent space of $\mathcal{X}$. Strictly monotone if $< 0 \ \forall z \neq 0.$

Monotone Games

Example: Rock-Paper-Scissors

$$DF(x) + DF(x)^\top = (w-l) \underbrace{\begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix}}_{\lambda_1=\lambda_2=-1,\, \lambda_3=2}$$

Schur decomposition:

$$DF(x) + DF(x)^\top = (w-l)(-v_1 v_1^\top - v_2 v_2^\top + 2 v_3 v_3^\top)$$

where orthonormal eigenvectors v_1, v_2 span $T\mathcal{X}$ and $v_3 \perp T\mathcal{X}$.

Thus, for $z \in T\mathcal{X}$,

$$z^\top (DF(x) + DF(x)^\top) z = -(w-l)\|z\|^2$$

Monotone for $w \geq l$; strictly monotone for $w > l$.

Monotone Games

Example: Congestion Games

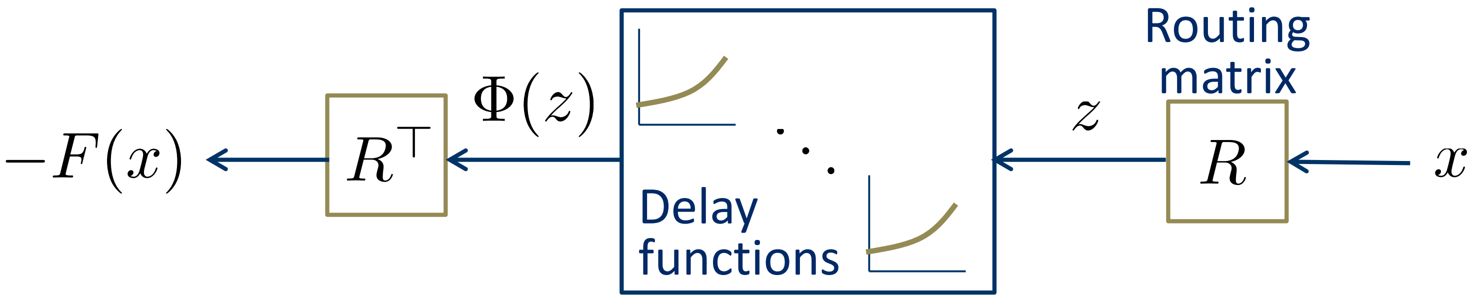

$$R_{ij} = \begin{cases} 1 & \text{if route } j \text{ uses link } i \\ 0 & \text{otherwise} \end{cases}$$

If $Q \succeq 0$ then $R^\top Q R \succeq 0$ for any R, because $z^\top (R^\top Q R) z = (Rz)^\top Q (Rz)$

$$DF(x) = -R^\top \begin{bmatrix} \Phi_1' & & \\ & \Phi_2' & \\ & & \ddots \end{bmatrix} R \preceq 0 \;\Rightarrow\; F \text{ monotone}$$

$DF(x) \prec 0 \; \forall x$ (strictly monotone) if $\Phi_i'(s) > 0 \; \forall s$ for each link i, and null space of R is $\{0\}$.

Monotone Games

Wake-up Problem

True or False? A congestion game where the routing matrix has trivial null space admits a unique Nash equilibrium.

Learning Rules and Evolutionary Dynamics

How do agents update their strategies?

Given current social state x and payoff π, the function

$$\pi, x \mapsto \rho_{ij}(\pi, x)$$

represents the switch rate from strategy i to strategy j.

The form of this function describes a "rule" by which agents switch to more favorable strategies with limited information.

Examples:

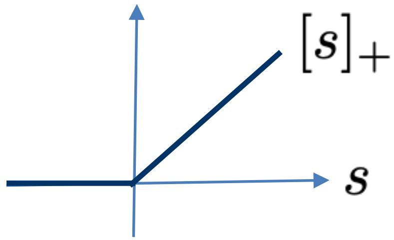

Learning by Imitation: $\rho_{ij}(\pi, x) = x_j[\pi_j - \pi_i]_+$

MJ Smith Rule: $\rho_{ij}(\pi, x) = [\pi_j - \pi_i]_+$

Brown-von Neumann-Nash (BNN) Rule:

$$\rho_{ij}(\pi, x) = [\pi_j - \pi_{\text{avg}}]_+ \quad \pi_{\text{avg}} := (x^\top \pi)/m$$

Learning Rules and Evolutionary Dynamics

Evolutionary Dynamics Model (EDM)

Given the switch rate defined by learning rule $\pi, x \mapsto \rho_{ij}(\pi, x)$

the mass of agents playing strategy i evolves according to

$$\dot{x}_i = \sum_{j \in \mathcal{S}} x_j \rho_{ji}(x, \pi) - x_i \sum_{j \in \mathcal{S}} \rho_{ij}(x, \pi)$$

or $\dot{x} = \nu(x, \pi)$ in concise notation.

Example: For imitation learning, where $\rho_{ij}(\pi, x) = x_j [\pi_j - \pi_i]_+$

$$\dot{x}_i = x_i \sum_{j \in \mathcal{S}} x_j (\pi_i - \pi_j)$$

Note: $\nu(x, \pi)$ lies in the tangent cone to $\mathcal{X}$ at x for each (x, π). Therefore, $\mathcal{X}$ is invariant under the EDM.

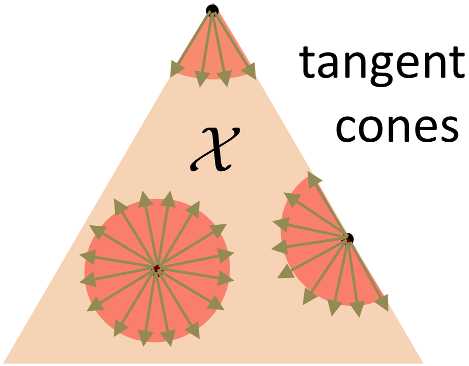

Learning Rules and Evolutionary Dynamics

Evolutionary Dynamics Model (EDM)

For multiple populations $r = 1, \ldots, p$:

$$\dot{x}_i^r = \sum_{j \in \mathcal{S}^r} x_j^r \rho_{ji}^r(x^r, \pi^r) - x_i^r \sum_{j \in \mathcal{S}^r} \rho_{ij}^r(x^r, \pi^r)$$

$$\dot{x}^r = \nu^r(x^r, \pi^r) \qquad (\mathrm{EDM}^r)$$

where $\pi = F(x)$:

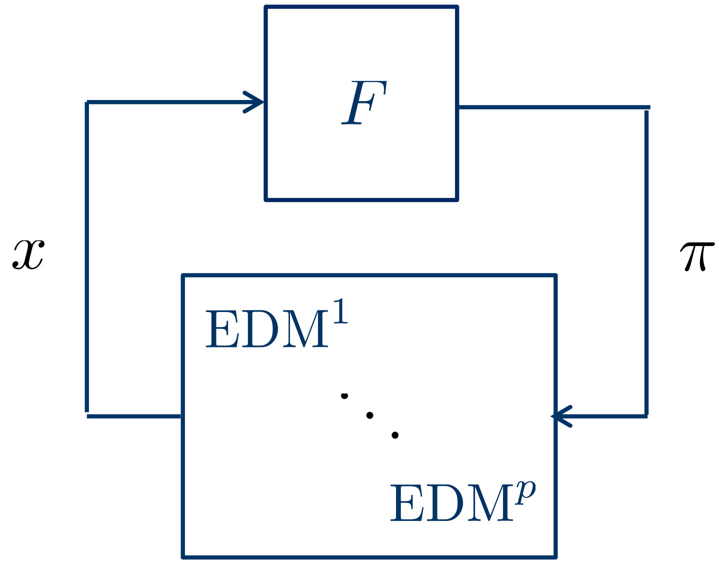

Learning Rules and Evolutionary Dynamics

Do trajectories of the EDM converge to Nash equilibria?

Not necessarily. Periodic orbits and chaotic attractors possible:

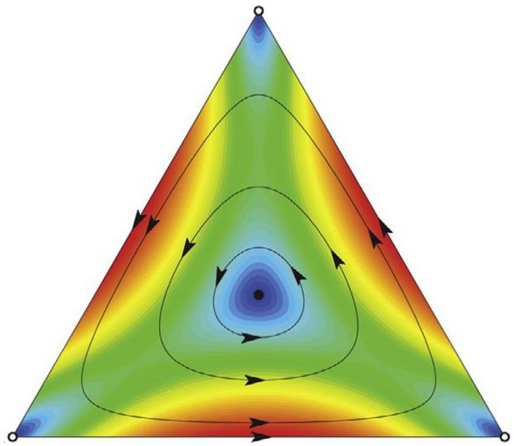

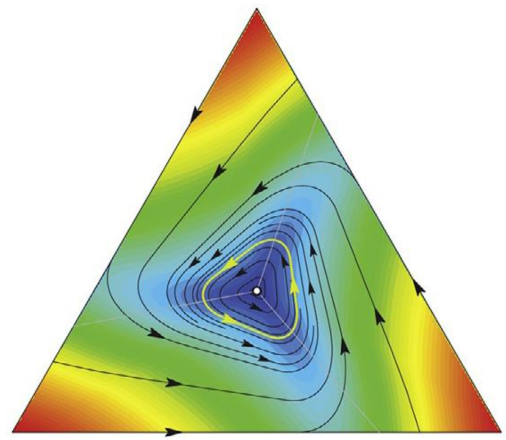

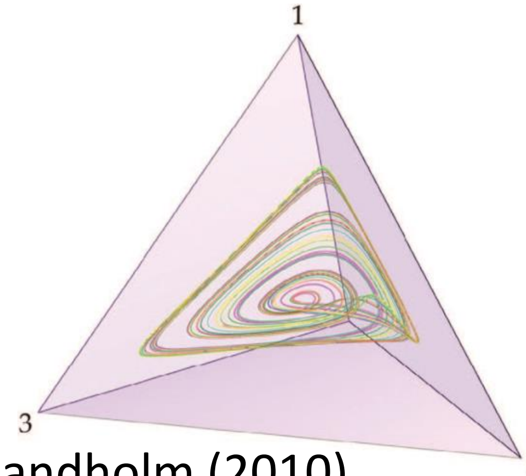

Sandholm (2010)

Convergence guarantees, when possible, justify the assumption that players are at equilibrium, common in game theory literature.

Convergence with myopic learning rules relaxes global information requirements associated with the equilibrium assumption.

We will prove convergence in monotone games for several learning rules using dissipativity.

Learning Rules and Evolutionary Dynamics

Origins of Congestion Games

Equilibrium concepts in route choices were explored by Wardrop (1952) and Beckmann et al. (1956). No explicit reference to game theory (which was just emerging then) but many game theoretic notions implicit in these publications. Connection to game theory and work of Nash was made later by others. Transportation literature still uses the term “Wardrop equilibrium” rather than “Nash equilibrium.”

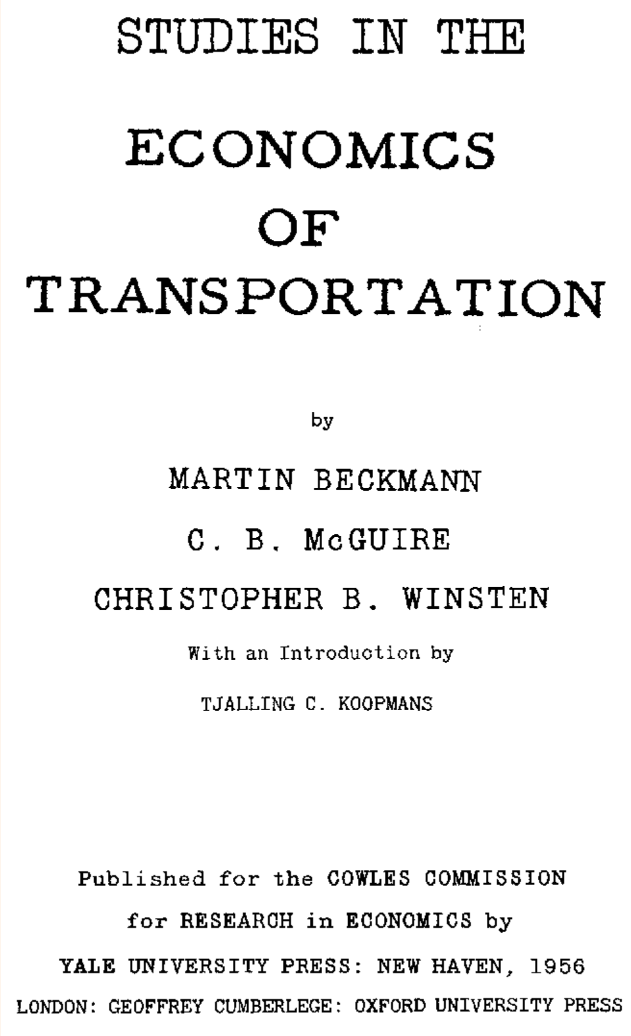
STUDIES IN THE

ECONOMICS
OF
TRANSPORTATION

by

MARTIN BECKMANN

C. B. McGUIRE

CHRISTOPHER B. WINSTEN

With an Introduction by

TJALLING C. KOOPMANS

Published for the COWLES COMMISSION
for RESEARCH in ECONOMICS by
YALE UNIVERSITY PRESS: NEW HAVEN, 1956
LONDON: GEOFFREY CUMBERLEGE: OXFORD UNIVERSITY PRESS

Recalling Relevant Dissipativity Notions

The system:

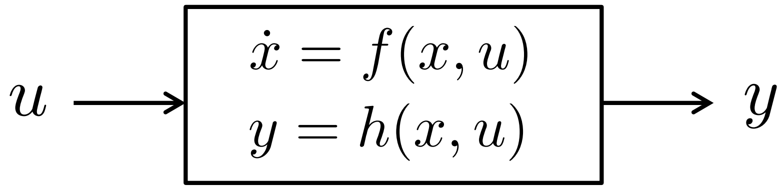

is **Equilibrium Independent Dissipative (EID)** if ∃ a storage function $V : \mathcal{X} \times \bar{\mathcal{X}} \to \mathbb{R}_{\geq 0}$ **s.t.** $\forall x \in \mathcal{X}, \bar{x} \in \bar{\mathcal{X}}, u \in \mathcal{U},$

$$V(\bar{x},\bar{x}) = 0, \quad \nabla_x V(x,\bar{x})^\top f(x,u) \leq s(u-\bar{u}, y-\bar{y})$$

where $\bar{u}, \bar{y}$ are functions of $\bar{x}$ through $f(\bar{x},\bar{u}) = 0,\ \bar{y} = h(\bar{x},\bar{u})$;

Delta Dissipative if ∃ a storage function $S : \mathcal{X} \times \mathcal{U} \to \mathbb{R}_{\geq 0}$ s.t.

$$S(x,u) = 0 \Leftrightarrow f(x,u) = 0$$

$$\nabla_x S(x,u)^\top f(x,u) + \nabla_u S(x,u)^\top v \leq s(v,w) \quad \forall x,u,v$$

where $w := \nabla_x h(x,u)^\top f(x,u) + \nabla_u h(x,u)^\top v.$

Recalling Relevant Dissipativity Notions

Stability Criteria

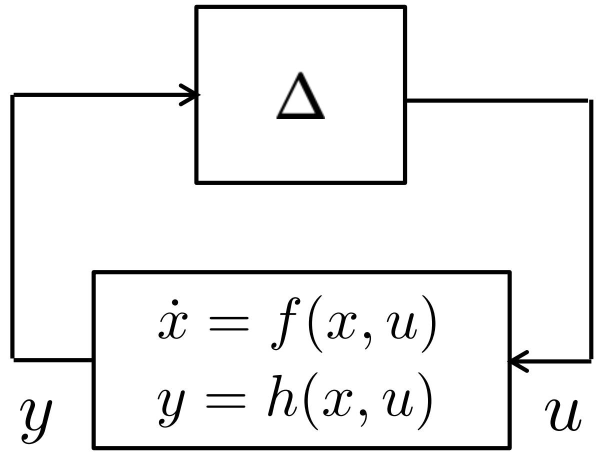

Assume: 1) Δ is a static map and the interconnection is well posed; that is, $u = \Delta(h(x,u))$ has sol'n $u = g(x)$;
2) interconnection has equilibrium x^*.

If the system $\dot{x} = f(x,u)$, $y = h(x,u)$ is EID with supply rate s and storage function V such that $V(x,\bar{x}) > 0\ x \neq \bar{x}$, and Δ satisfies

$$s(u-\bar{u}, y-\bar{y}) \leq 0$$

then x^* is stable and $V(\cdot, x^*)$ is a Lyapunov function.

If, instead, the system is delta dissipative with supply rate s and storage function S, and Δ satisfies

$$s(\dot{u}(t), \dot{y}(t)) \leq 0 \quad \forall t$$

then x^* is stable and a Lyapunov function is $V(x) = S(x, g(x))$.

Recalling Relevant Dissipativity Notions

Example: If the system

$$\dot{x} = f(x, u)$$
$$y = h(x, u)$$

is equilibrium independent *passive* (EIP), the complementary condition to be satisfied by $u = F(y)$ is:

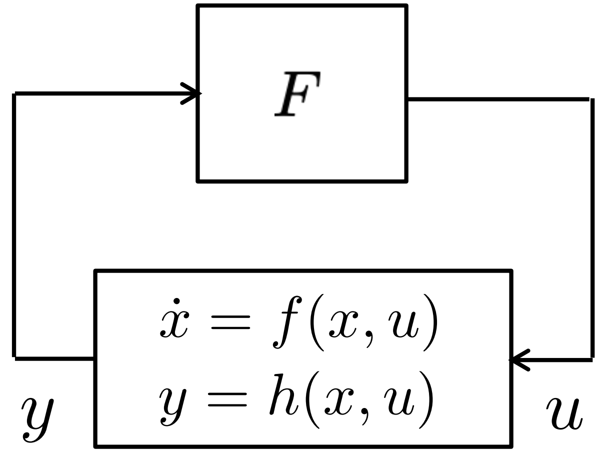

$$s(u - \bar{u}, y - \bar{y}) = (y - \bar{y})^\top (u - \bar{u}) = (y - \bar{y})^\top (F(y) - F(\bar{y})) \leq 0$$

i.e., F is "monotone."

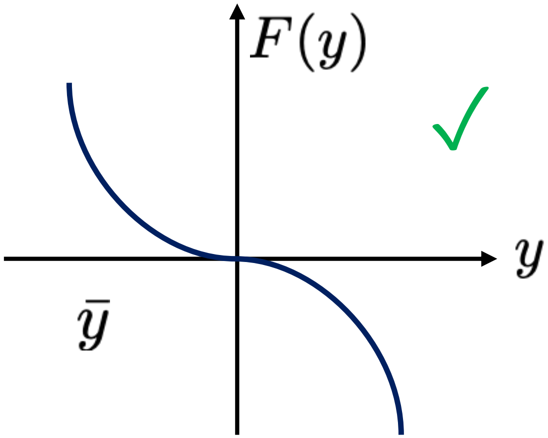

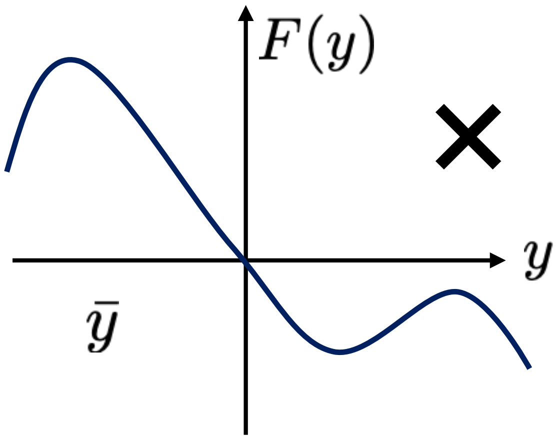

Recalling Relevant Dissipativity Notions

Example: If the system

$$\dot{x} = f(x,u)$$
$$y = h(x,u)$$

is delta *passive*, the new condition is:

$$s(\dot{u}(t), \dot{y}(t)) = \dot{y}(t)^\top \dot{u}(t) \leq 0$$

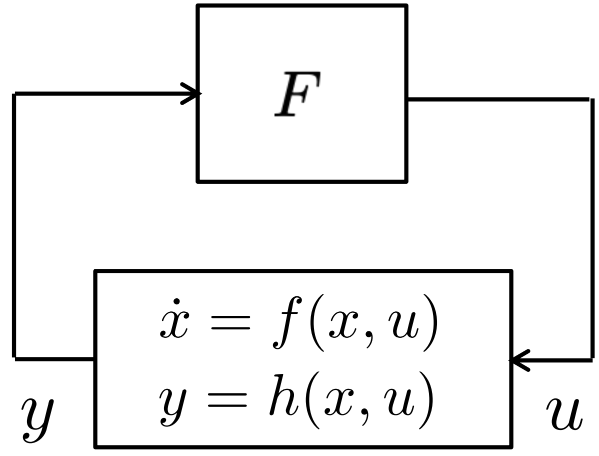

Differentiate $u(t) = F(y(t))$: $\dot{u}(t) = DF(y(t))\dot{y}(t)$

Then, $\dot{y}(t)^\top \dot{u}(t) = \frac{1}{2}\dot{y}(t)^\top (DF(y(t)) + DF(y(t))^\top)\dot{y}(t)^\top \leq 0$

is equivalent to: $$DF(y) + DF(y)^\top \preceq 0 \quad \forall y$$

$$\Updownarrow$$

$$(y - \bar{y})^\top (F(y) - F(\bar{y})) \leq 0 \quad \forall y, \bar{y}$$

Same complementary condition as the one for EIP: monotone F.

Recalling Relevant Dissipativity Notions

Wake-up Problem

Consider the system $\dot{x} = F(x)$, which can be decomposed as

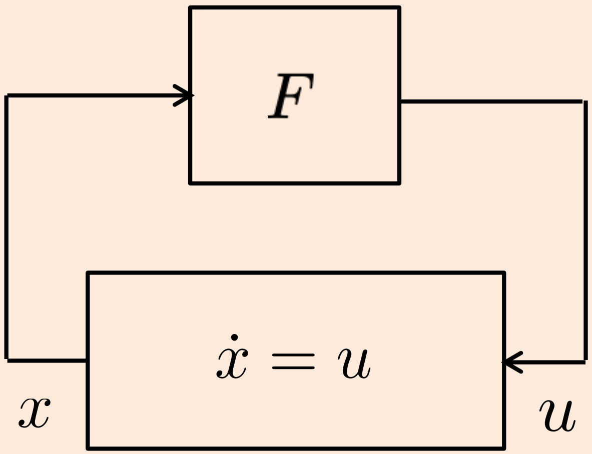

and $\dot{x} = u$ is delta passive with storage function:

$$S(x,u) = \frac{1}{2}u^\top u$$

Suppose F satisfies the monotonicity condition and an equilibrium exists. What is a Lyapunov function that proves its stability?

Convergence to Nash Equilibria

Case 1: Monotone payoff + Equilibrium Independent Passive EDM

Imitation learning leads to EIP (but not delta passive[1]) EDM:

$$\dot{x}_i = x_i \sum_j x_j(\pi_i - \pi_j) = x_i(m\pi_i - x^\top \pi)$$

$\bar{x} \in \operatorname{int}(X), \bar{\pi} = \alpha \mathbf{1}$ equilibrium candidates for any α.

$$S(x, \bar{x}) = \frac{1}{m} \sum_i \{x_i - \bar{x}_i - \bar{x}_i \ln(x_i / \bar{x}_i)\}$$

well defined in $\operatorname{int}(X)$, nonnegative, zero only when $x = \bar{x}$, and

$$\nabla_x S(x, \bar{x})^\top \nu(x, \pi) = (x - \bar{x})^\top (\pi - \bar{\pi}) \qquad \textbf{(EIP)}$$

$\bar{x} \in \operatorname{int}(X)$ is an equilibrium for EDM in feedback with $\pi = F(x)$ iff it is a Nash eq. Asymptotically stable if F is *strictly* monotone:

$$(x - \bar{x})^\top (F(x) - F(\bar{x})) < 0 \quad x \neq \bar{x}.$$

[1] shown in (Park, Shamma, Martins, 2018)

Convergence to Nash Equilibria

Case 2: Monotone payoff + Delta Passive EDM

EDM with BNN and Smith learning rules has the "Nash stationarity" property:

$$\nu(x,\pi) = 0 \;\Leftrightarrow\; x \text{ is a best response to } \pi$$

Since equilibria of EDM satisfy $\nu(x, F(x)) = 0$, they correspond to Nash equilibria of the game defined by F.

EDM with BNN and Smith learning rules is delta passive:

$$S(x,\pi) = \frac{1}{2}\sum_{j\in\mathcal{S}}(\max\{0, \pi_j - \pi_{\text{avg}}\})^2 \quad \text{(BNN)}$$

$$S(x,\pi) = \frac{1}{2}\sum_{i\in\mathcal{S}} x_i \sum_{j\in\mathcal{S}}(\max\{0, \pi_j - \pi_i\})^2 \quad \text{(Smith)}$$

In each case: $S(x,\pi) = 0 \Leftrightarrow \nu(x,\pi) = 0$

Convergence to Nash Equilibria

Case 2: Monotone payoff + Delta Passive EDM

and $\nabla_x S(x,\pi)^\top \nu(x,\pi) + \nabla_\pi S(x,\pi)^\top p \leq -\sigma(x,\pi) + \nu(x,\pi)^\top p$

$\sigma(x,\pi) \geq 0$ everywhere and $\sigma(x,\pi) = 0 \Leftrightarrow \nu(x,\pi) = 0.$

Strictness of passivity guarantees asymptotic stability when F is monotone, even if monotonicity is not strict.

Conclusion: For BNN and Smith learning rules, trajectories of the EDM in feedback with $\pi = F(x)$ converge to Nash equilibria if

$$(y-x)^\top (F(y) - F(x)) \leq 0 \qquad \forall x, y \in \mathcal{X}$$

Delta passivity shown for broader learning rules encompassing BNN and Smith (see tutorial by Park, Martins, Shamma, CDC'19).

Convergence to Nash Equilibria

Example: Rock-Paper-Scissors

$$F(x) = Ax = \begin{bmatrix} 0 & -l & w \\ w & 0 & -l \\ -l & w & 0 \end{bmatrix} \begin{bmatrix} x_R \\ x_P \\ x_S \end{bmatrix}$$

When $w \geq l, \ \ (y-x)^\top (F(y) - F(x)) \leq 0 \ x, y \in \mathcal{X},$ thus BNN and Smith learning rules guarantee convergence to Nash equilibria

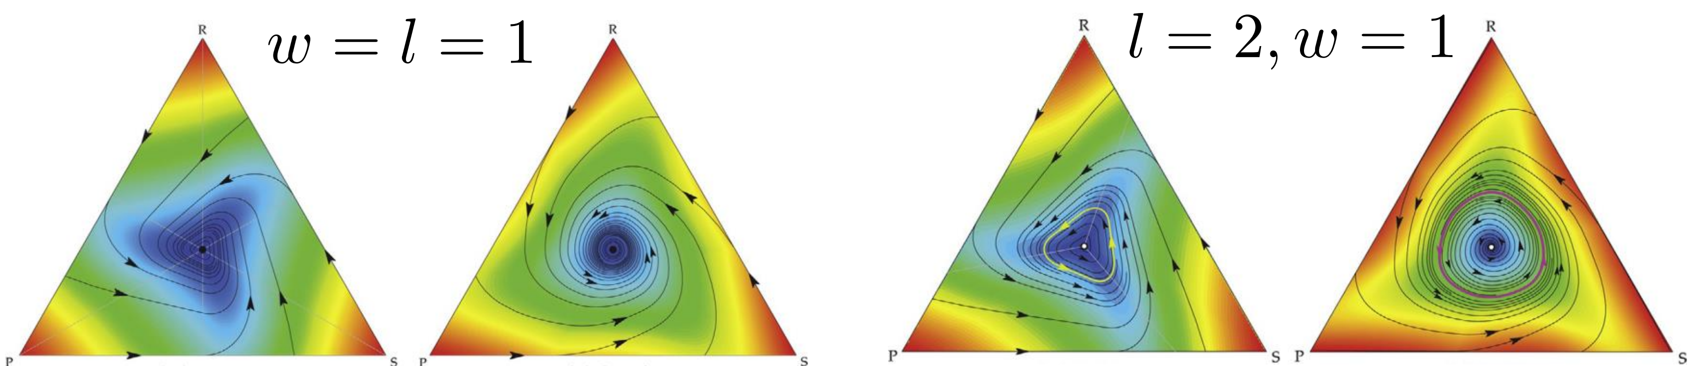

BNN Smith BNN Smith

Sandholm (2010)

Convergence to Nash Equilibria

Multiple Populations

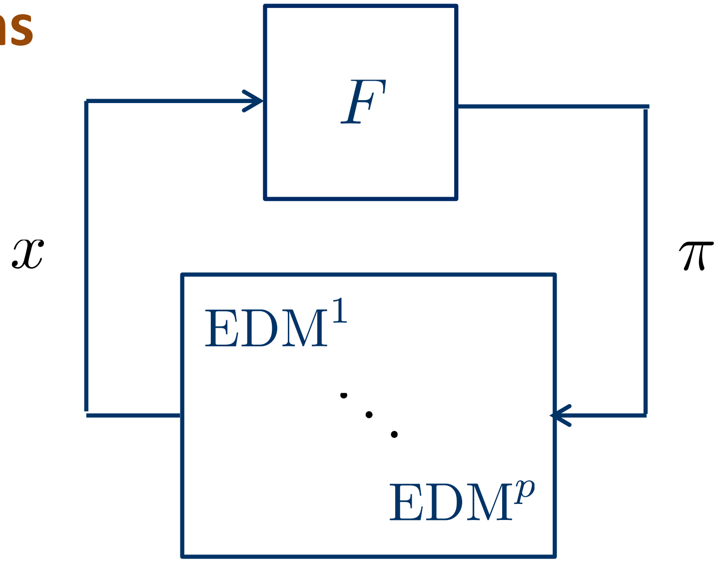

If each EDM has Nash Stationarity and delta passivity properties, then the monotonicity requirement on F can be relaxed as:

$$(y-x)^\top W(F(y)-F(x)) \leq 0 \quad \forall x, y \in \mathcal{X}$$

for some $W = \begin{bmatrix} w_1 I_{n^1} & & \\ & \ddots & \\ & & w_p I_{n^p} \end{bmatrix}, \; w_1, \cdots, w_p > 0.$

Convergence to Nash Equilibria

Example: Congestion Game with Mixed Autonomy

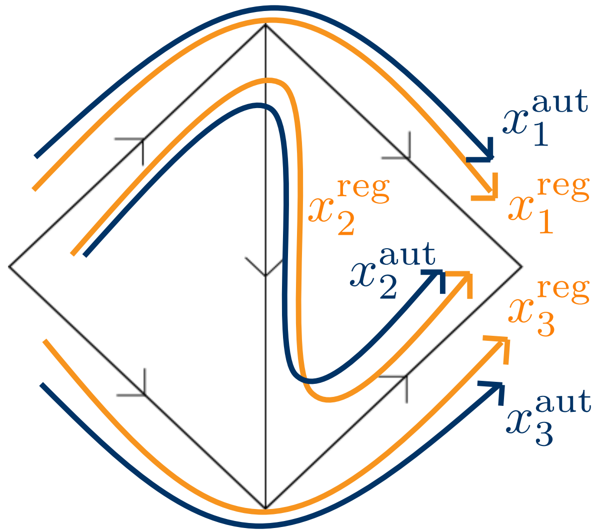

Autonomous vehicles can maintain shorter headway than regular vehicles. Payoff model below treats them as a separate population and discounts their externality by a factor of $\mu \in (0,1)$:

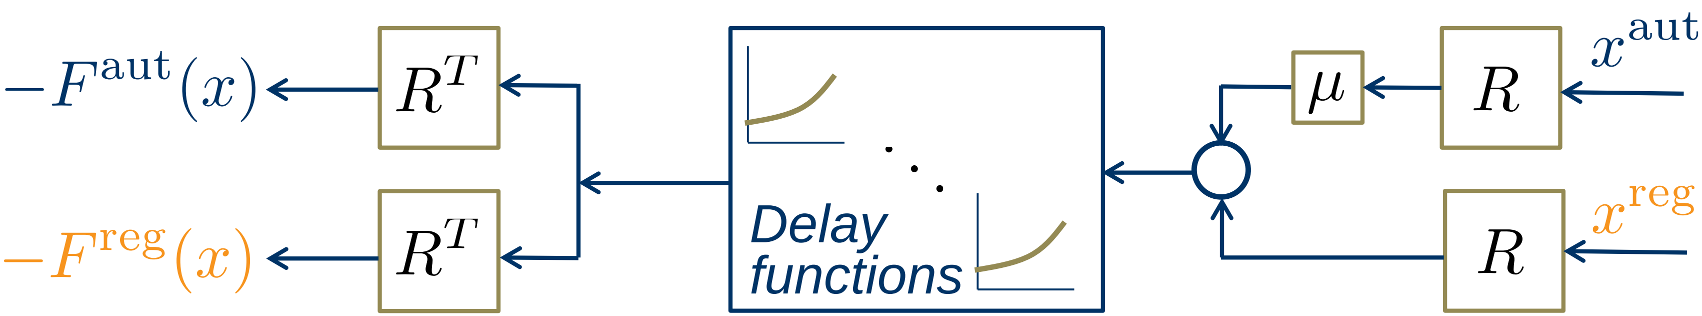

Convergence to Nash Equilibria

Example: Congestion Game with Mixed Autonomy

When $\mu \neq 1$ payoff loses the structure that ensures monotonicity:

$$DF(x) = -[R \quad R]^\top \begin{bmatrix} \Phi_1' & & \\ & \Phi_2' & \\ & & \ddots \end{bmatrix} [\mu R \quad R]$$

but WF with $W = \begin{bmatrix} \mu I & 0 \\ 0 & I \end{bmatrix}$ is monotone:

$$DWF(x) = -[\mu R \quad R]^\top \begin{bmatrix} \Phi_1' & & \\ & \Phi_2' & \\ & & \ddots \end{bmatrix} [\mu R \quad R]$$

This guarantees convergence to Nash equilibria for learning rules that have Nash Stationarity and delta passivity properties, such as BNN and Smith rules. See (Arcak, Martins, TCNS'21) for further results.

Convergence to Nash Equilibria

An Invitation to Population Games

Population games are harmonious with control theory and many research opportunities exist: combining evolutionary dynamics with physical dynamics, influencing equilibria and transients with feedback ("mechanism design"), applications to networks, etc.

Further Reading:

- Sandholm, Population Games and Evolutionary Dynamics, MIT Press, 2010
- Quijano et al., The role of population games and evolutionary dynamics in distributed control systems, Control Systems, pp. 70-97, February 2017
- Park, Shamma, Martins, Passivity and evolutionary game dynamics, IEEE CDC 2018
- Park, Martins, Shamma, From population games to payoff dynamics models: a passivity-based approach, IEEE CDC 2019
- Arcak and Martins, Dissipativity tools for convergence to Nash equilibria in population games, IEEE Trans. Control of Network Systs., pp. 39-50, vol.8, 2021

Summary

This lesson demonstrated applications of dissipativity in optimization and game theory, proving convergence to optima and Nash equilibria.

Part 1:

- Introduced a feedback systems perspective for a class of first-order optimization algorithms.
- Used dissipation inequality and IQC techniques to analyze the stability and convergence rates of these algorithms.

Part 2:

- Introduced population games, and notions of Nash equilibria and monotone games.
- Revealed classes of learning rules that ensure convergence to Nash equilibria in monotone games. Equilibrium independent and delta passivity played key roles.

Self-Study Problems

See Web site for problems and solutions.

sites.google.com/berkeley.edu/dissipation-iqc