FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Distributed Online Learning in Multi-Agent Systems

Alex Bayen

Dept. of Electrical Engineering & Computer Sciences,
UC Berkeley, CA, USA

Berkeley University of California
Massachusetts Institute of Technology
NSF
Michigan
Vanderbilt University

Project status within FORCES scope

Framework for cybersecurity of physical infrastructure

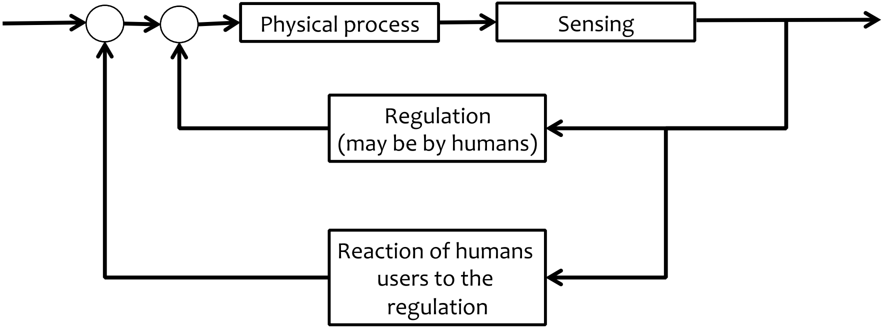

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Roughly one attack a month on the traffic management infrastructure

The *Italian Job* (2003)

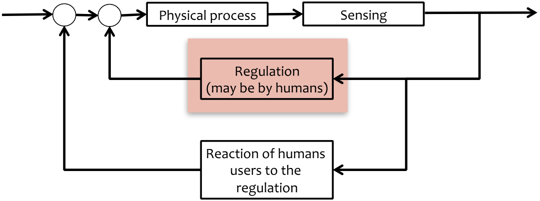

Roughly one attack a month on the traffic management infrastructure

The *Italian Job* (2003)

The “real” *Italian Job* (2007)

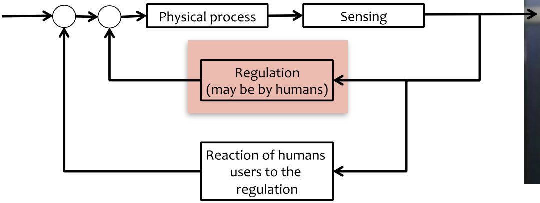


Los Angeles Times LOCAL SPORTS ENTERTAINMENT NATION WORLD BUSINESS OPINION LIFESTYLE MORE

YOU ARE HERE: LAT Home → Collections → Los Angeles

Key signals targeted, officials say

Two accused of hacking into L.A.'s traffic light system plead not guilty. They allegedly chose intersections they knew would cause major jams.

January 09, 2007 | Sharon Bernstein and Andrew Blankstein | Times Staff Writers

Roughly one attack a month on the traffic management infrastructure

The *Italian Job* (2003)

The "real" *Italian Job* (2007)

NC DOT signs hacked (2014)

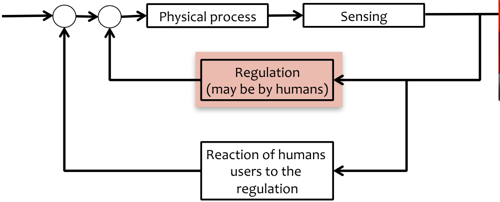

MEMBERS: Register | Login

WWAY 3 abc NewsChannel Celebrating 50 YEARS 1964-2014

LIVE. LOCAL. INT

SOUTHEASTERN NORTH CAROLINA'S #

Home News Weather Sports Videos Community Features Programming Ab

FBI investigating hacked NCDOT digital road signs

Submitted by WWAY on Sat, 05/31/2014 - 9:55am.

READ MORE: News New Hanover County News Crime Cybercrime FBI Hacking N.C. NCDOT Transportation

Like 17 Tweet 5 Share 6

WILMINGTON, NC (WWAY) -- The North Carolina Department of Transportation says the FBI is looking into a group that hacked into at least five digital road signs yesterday, including one in New Hanover County.

The DOT says it is also evaluation the security measures in place for its digital road signs after a group changed the intended transportation-related messages on the signs to an advertisement for its Twitter account. According to a news released, the DOT corrected the messages as soon as it discovered the hackings.

The DOT says the hacked message boards are on Carolina Beach Road in New Hanover county, I-40 and I-240 in Asheville, US 421 in Winston-Salem and I-77 near the North Carolina/Virginia state line.

riously," NCDOT Chief Information Officer David Ulmer said in

Roughly one attack a month on the traffic management infrastructure

The *Italian Job* (2003)

The "real" *Italian Job* (2007)

NC DOT signs hacked (2014)

Snail operations (2014)

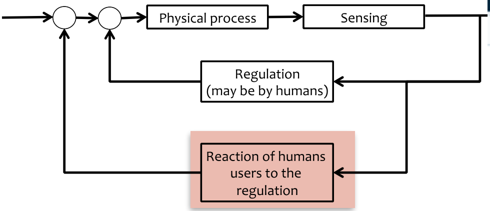

Roughly one attack a month on the traffic management infrastructure

The *Italian Job* (2003)

The "real" *Italian Job* (2007)

NC DOT signs hacked (2014)

Snail operations (2014)

Waze / Google hacked (2014)

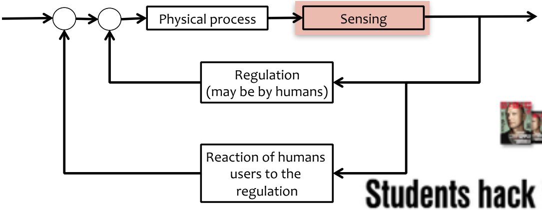

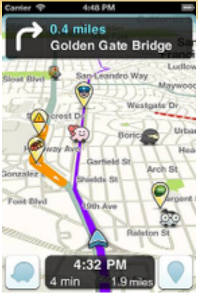

WIRED

Students hack Waze, send in army of traffic bots

TECHNOLOGY / 25 MARCH 14 / by NICHOLAS TUFNELL

Two Israeli students have successfully hacked popular

AND DIGITAL EDITIONS

Roughly one attack a month on the traffic management infrastructure

The *Italian Job* (2003)
The “real” *Italian Job* (2007)
NC DOT signs hacked (2014)
Snail operations (2014)
Waze / Google hacked (2014)
Sensys Attack (2014)

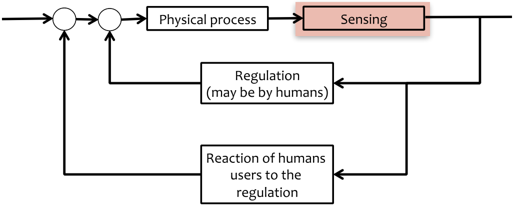

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Roughly one attack a month on the traffic management infrastructure

The *Italian Job* (2003)

The "real" *Italian Job* (2007)

NC DOT signs hacked (2014)

Snail operations (2014)

Waze / Google hacked (2014)

Sensys Attack (2014)

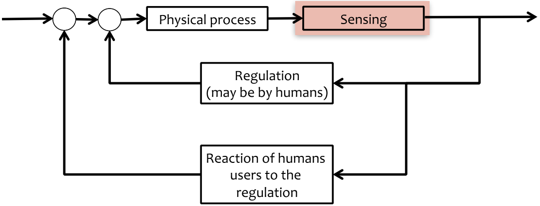

Cesar Cerrudo in downtown New York City, conducting field test of vulnerable traffic sensors. Photo: Courtesy of Cesar Cerrudo

Roughly one attack a month on the traffic management infrastructure

The *Italian Job* (2003)
The “real” *Italian Job* (2007)
NC DOT signs hacked (2014)
Snail operations (2014)
Waze / Google hacked (2014)
Sensys Attack (2014)

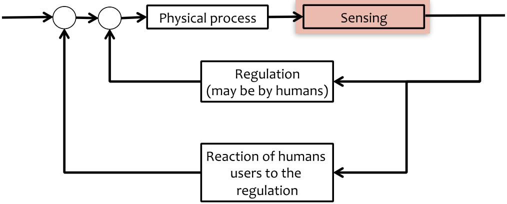

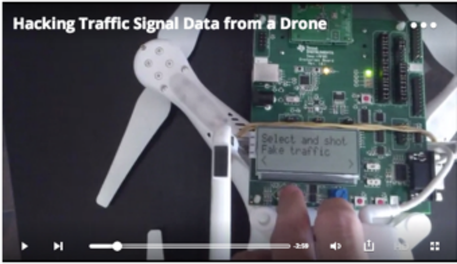

Presented at the previous FORCES all hands

Framework for cybersecurity of physical infrastructure

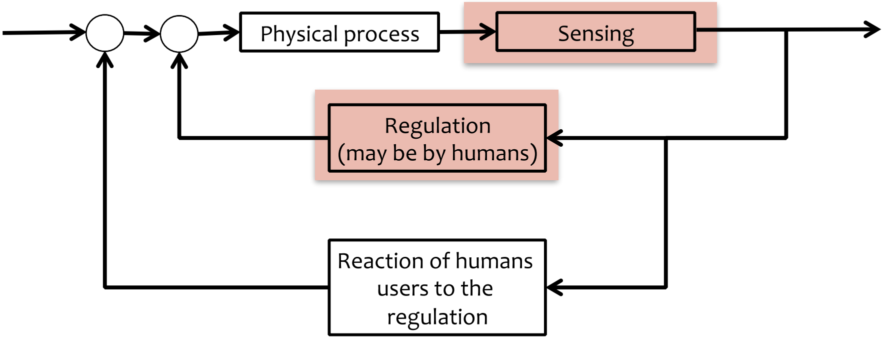

Presented at the previous FORCES all hands

Attack on the sensing infrastructure, and the control infrastructure

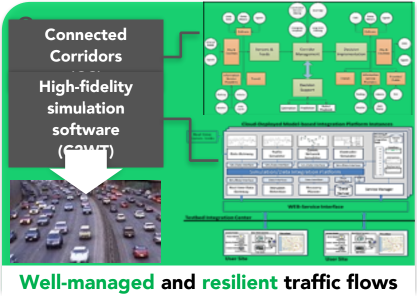

Presented at the this FORCES all hands

How do humans learn from attacks / ops / to protect / attack performance of the system

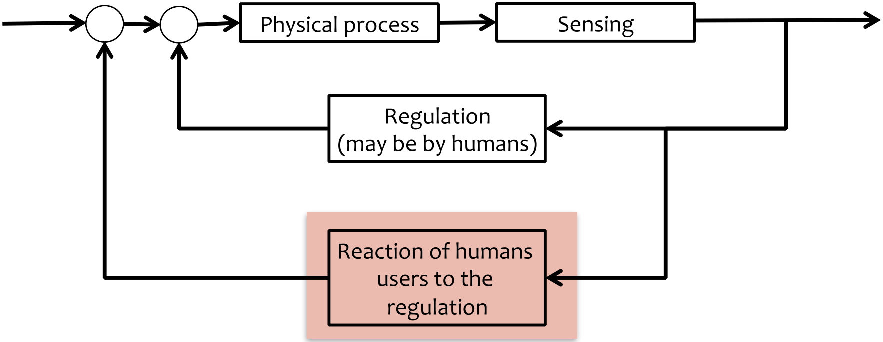

Context: (non-)collaborative players

Commonly spread ansatz about decision making in routing: if you reduce your own travel time, you help the public good.

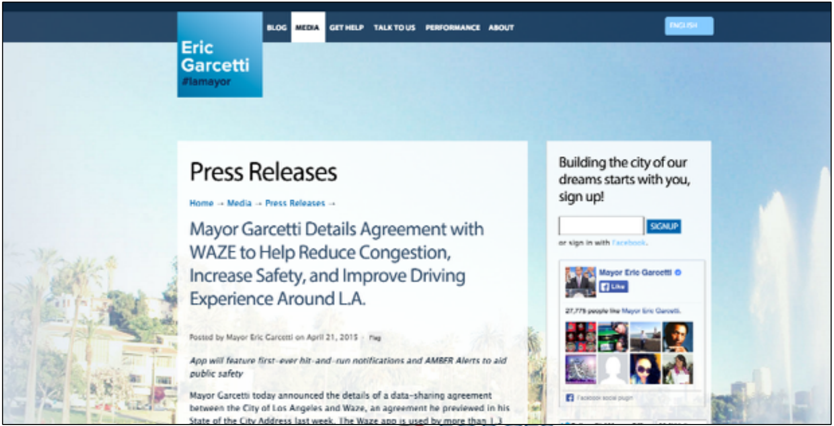

Context: (non-)collaborative players

Commonly spread ansatz about decision making in routing: if you reduce your own travel time, you help the public good.

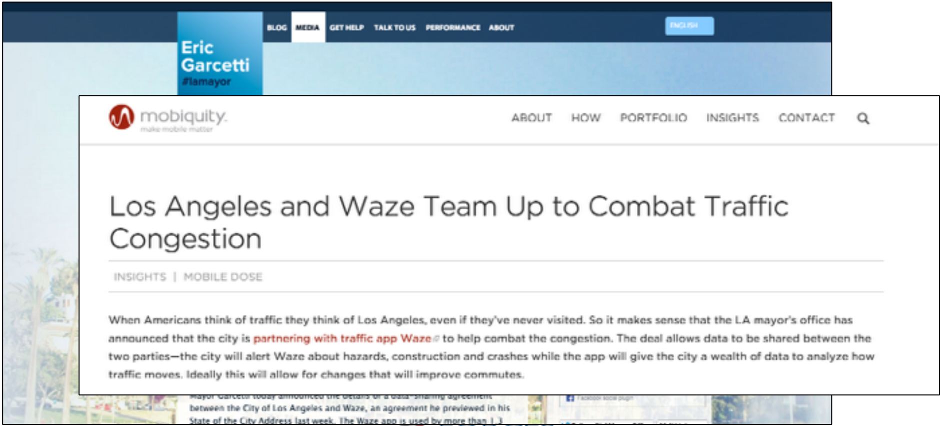
BLOG MEDIA GET HELP TALK TO US PERFORMANCE ABOUT

Eric Garcetti
#lamayor

mobiquity
make mobile matter

ABOUT HOW PORTFOLIO INSIGHTS CONTACT

Los Angeles and Waze Team Up to Combat Traffic Congestion

INSIGHTS | MOBILE DOSE

When Americans think of traffic they think of Los Angeles, even if they've never visited. So it makes sense that the LA mayor's office has announced that the city is partnering with traffic app Waze to help combat the congestion. The deal allows data to be shared between the two parties—the city will alert Waze about hazards, construction and crashes while the app will give the city a wealth of data to analyze how traffic moves. Ideally this will allow for changes that will improve commutes.

between the City of Los Angeles and Waze, an agreement he previewed in his
State of the City Address last week. The Waze app is used by more than 1.3

Context: (non-)collaborative players

Commonly spread ansatz about decision making in routing: if you reduce your own travel time, you help the public good.

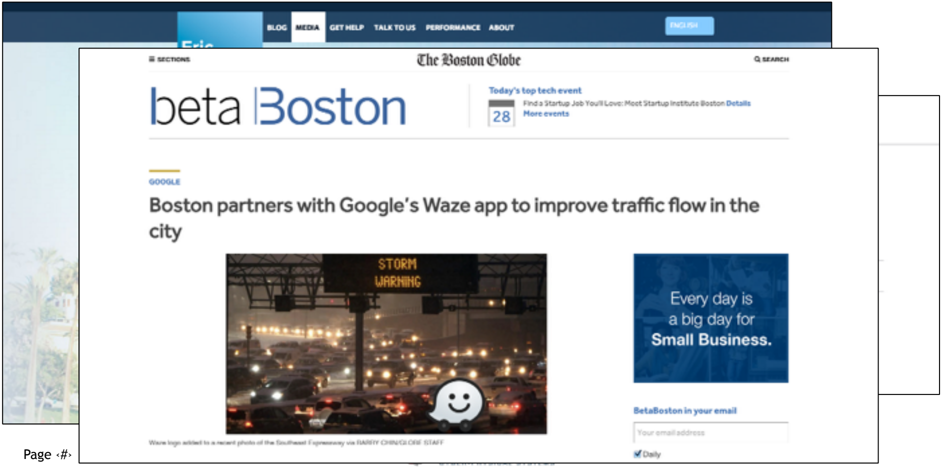

Page ‹#›

Context: (non-)collaborative players

Reality: without optimal coordination, users experience the difference between social optimal and Nash

Context: (non-)collaborative players

Reality: without optimal coordination, users experience the difference between social optimal and Nash

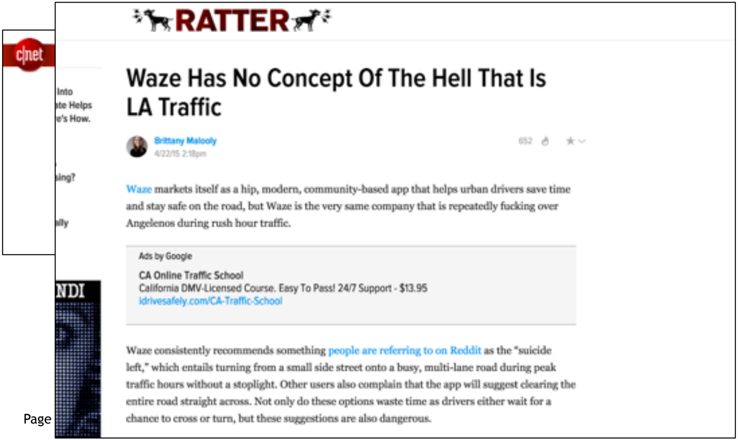
RATTER

Waze Has No Concept Of The Hell That Is LA Traffic

Brittany Malooly
4/22/15 2:18pm

Waze markets itself as a hip, modern, community-based app that helps urban drivers save time and stay safe on the road, but Waze is the very same company that is repeatedly fucking over Angelenos during rush hour traffic.

Waze consistently recommends something people are referring to on Reddit as the "suicide left," which entails turning from a small side street onto a busy, multi-lane road during peak traffic hours without a stoplight. Other users also complain that the app will suggest clearing the entire road straight across. Not only do these options waste time as drivers either wait for a chance to cross or turn, but these suggestions are also dangerous.

Page

Context: (non-)collaborative players

Reality: without optimal coordination, users experience the difference between social optimal and Nash

Context: (non-)collaborative players

Reality: without optimal coordination, users experience the difference between social optimal and Nash

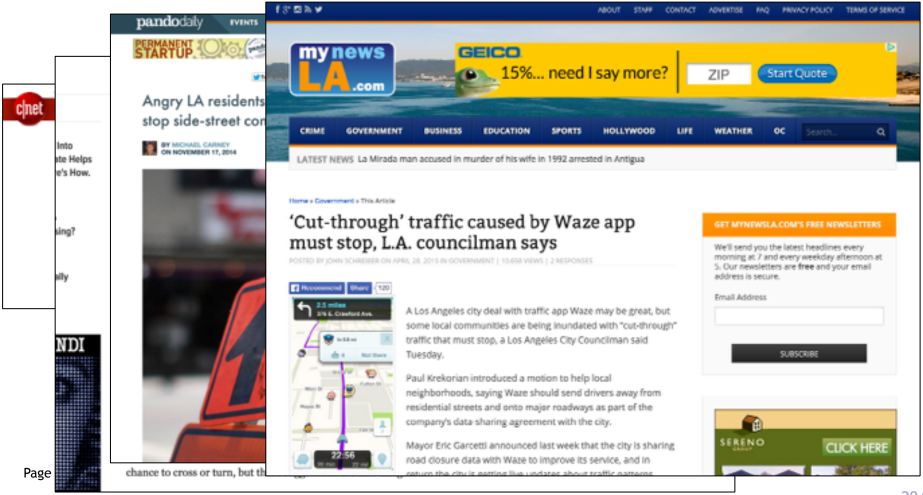

$20.8 Billion dollar question

Annual cost of congestion: $20.8B (TTI Urban Mob. Rep. 2012)

Problem set up: one shot / repeated game

N players are routing traffic

- Private sector apps (Google, Waze, Apple, INRIX etc.)
- Some public sector apps (511)

Except some specific public agencies, none of these players are solving for social optimal solutions.

At best, all of these are providing Nash solutions, i.e. routes in which each user has no incentive to change his/her trajectory.

What if companies “learned” from the past?

FORCES FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Problem set up: one shot / repeated game

Structure of a learning game:

1) Make choice (route)
2) Perform action (drive the route)
3) Compare outcome with external information
4) Learn (for next action)

Outline

1. Introduction

2. Convergence of agent dynamics
   - Background
   - Approximate replicator dynamics (AREP)
   - Distributed stochastic mirror descent dynamics (DSMD)

3. Application to routing

Outline

1. Introduction
2. Convergence of agent dynamics
   - Background
   - Approximate replicator dynamics (AREP)
   - Distributed stochastic mirror descent dynamics (DSMD)
3. Application to routing

Interaction of K decision makers

Decision maker k faces a sequential decision problem

At iteration t

(1) chooses probability distribution $x_{\mathcal{A}_k}^{(t)}$ over action set $\mathcal{A}_k$

(2) discovers a loss function $\ell_{\mathcal{A}_k}^{(t)} : \mathcal{A}_k \to [0, 1]$

(3) updates distribution

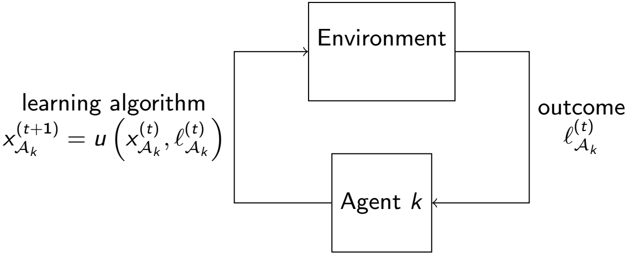

Figure: Sequential decision problem.

Loss of agent k affected by strategies of other agents.
Does not know this function, only observes its value.

Interaction of K decision makers

Decision maker k faces a sequential decision problem

At iteration t

(1) chooses probability distribution $x^{(t)}_{\mathcal{A}_k}$ over action set $\mathcal{A}_k$

(2) discovers a loss function $\ell^{(t)}_{\mathcal{A}_k} : \mathcal{A}_k \to [0,1]$

(3) updates distribution

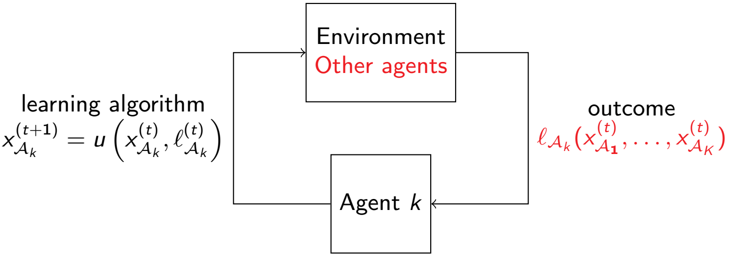

Figure: Sequential decision problem.

Loss of agent k affected by strategies of other agents.
Does not know this function, only observes its value.

Convergence to equilibria

- Can we guarantee $x^{(t)} \to \mathcal{X}^*$? Can players arrive to equilibrium?
  - $x^{(t)} = (x^{(t)}_{\mathcal{A}_1}, \dots, x^{(t)}_{\mathcal{A}_K})$
  - $\mathcal{X}^*$ set of equilibria.
- Convergence rates?
- Robustness to stochastic perturbations?

Examples of decentralized decision makers

Routing game

- Player drives from source to destination node
- Chooses path from $\mathcal{A}_k$
- Mass of players on each edge determines cost on that edge.

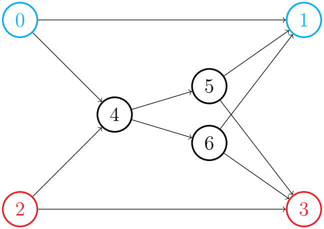

Figure: Routing game

Examples of decentralized decision makers

Routing game

- Player drives from source to destination node
- Chooses path from $\mathcal{A}_k$
- Mass of players on each edge determines cost on that edge.

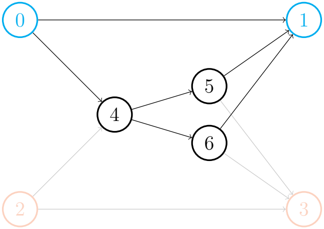

Figure: Routing game

Online learning model: illustration

$x^{(t)}_{\mathcal{A}_1} \in \Delta^{\mathcal{A}_1}$ Sample $a \sim x^{(t)}$ Discover $\ell^{(t)}_{\mathcal{A}_1} \in [0, 1]^{\mathcal{A}_1}$ Update $x^{(t+1)}_{\mathcal{A}_1}$

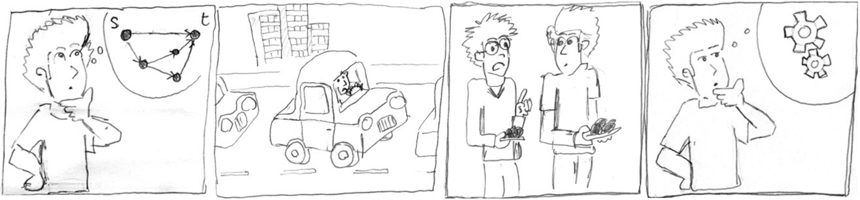

Convergence of distributed learning

for $t \in \mathbb{N}$ **do**
  Each agent k plays $x_{\mathcal{A}_k}^{(t)}$ (independently)
  Reveal loss vector $\ell_{\mathcal{A}_k}(x^{(t)}) \in [0,1]^{\mathcal{A}_k}$
  Update

$$x_{\mathcal{A}_k}^{(t+1)} = u_k \left(x_{\mathcal{A}_k}^{(t)}, \ell_{\mathcal{A}_k}(x^{(t)}) \right)$$

end for

Main problem

Define class of dynamics $\mathcal{C}$ such that

$$u_k \in \mathcal{C} \ \forall k \Rightarrow x^{(t)} \to \mathcal{X}^*$$

A brief review

Discrete time: uses regret analysis

- Hannan consistency: [5]
- Hedge algorithm for two-player games: [4]
- Online learning in games: [2]

Most studies prove the convergence of time-averaged strategies

$$\bar{x}^{(t)} = \frac{1}{t} \sum_{\tau \leq t} x^{(\tau)}$$

[5] James Hannan. Approximation to Bayes risk in repeated plays. *Contributions to the Theory of Games*, 3:97–139, 1957

[4] Yoav Freund and Robert E Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1):79–103, 1999

[2] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games.* Cambridge University Press, 2006

Convergence of $\bar{x}^{(t)}$ Vs convergence of $x^{(t)}$

Routing game example

Population 1

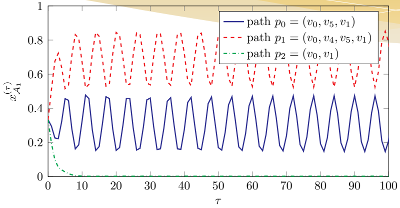

Population 2

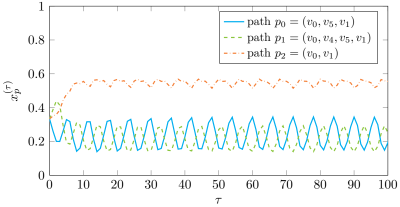

Figure: Population distributions

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Convergence of $\bar{x}^{(t)}$ Vs convergence of $x^{(t)}$

Routing game example

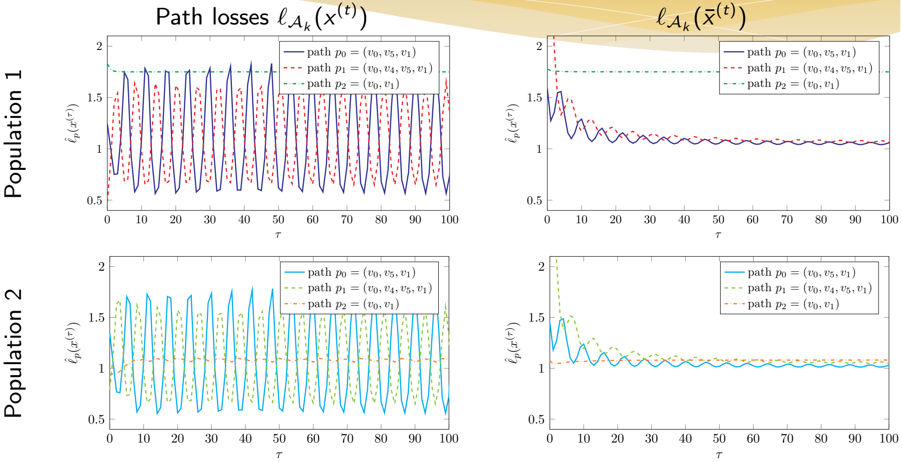

Figure: Path losses

FORCES FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Contributions

For the class of convex potential games.
Discrete time:

- $x^{(t)} \to \mathcal{X}^*$ for class of approximate replicator (AREP) dynamics [7]
- $x^{(t)} \to \mathcal{X}^*$ for class of distributed mirror descent (DMD) dynamics [9]
- $x^{(t)} \to \mathcal{X}^*$ for class of distributed stochastic mirror descent (DSMD) dynamics [6]

---

[7] Walid Krichene, Benjamin Drighès, and Alexandre Bayen. On the convergence of no-regret learning in selfish routing.
In *31st International Conference on Machine Learning (ICML)*. JMLR, 2014

[9] Walid Krichene, Syrine Krichene, and Alexandre Bayen. Convergence of mirror descent dynamics.
In *European Control Conference (ECC), accepted*, 2015

[6] Syrine Krichene, Walid Krichene, Roy Dong, and Alexandre Bayen. Convergence of stochastic mirror descent and applications to distributed routing.
In *Allerton Conference on Communication, Control and Computing, in preparation*, 2015

Outline

1. Introduction
2. Convergence of agent dynamics
   - Background
   - Approximate replicator dynamics (AREP)
   - Distributed stochastic mirror descent dynamics (DSMD)
3. Application to routing

Convex potential

Convex potential

Assume that $\exists f$ convex on $\mathcal{X} = \Delta^{\mathcal{A}_1} \times \cdots \times \Delta^{\mathcal{A}_K}$ such that

$$\nabla_{x_{\mathcal{A}_k}} f(x) = \ell_{\mathcal{A}_k}(x)$$

Then

$$\mathcal{X}^* = \arg\min_{x \in \Delta^{\mathcal{A}_1} \times \cdots \times \Delta^{\mathcal{A}_K}} f(x)$$

is the set of Nash equilibria.

Write $\ell(x) = \nabla f(x) = (\ell_{\mathcal{A}_1}(x), \dots, \ell_{\mathcal{A}_K}(x))$

First order optimality conditions

$$\begin{array}{ccc} \text{Nash condition} & \Leftrightarrow & \text{first order optimality} \\ \forall k, \forall x_{\mathcal{A}_k}, \left\langle \ell_{\mathcal{A}_k}(x^*), x^*_{\mathcal{A}_k} - x_{\mathcal{A}_k} \right\rangle \leq 0 & \Leftrightarrow & \forall x \in \mathcal{X}, \langle \ell(x^*), x^* - x \rangle \leq 0 \end{array}$$

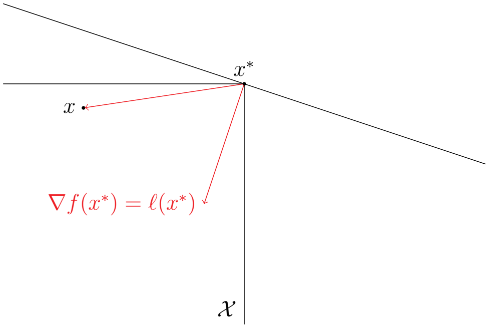

Figure: First order optimality conditions of the potential f

Convergence of $\bar{x}^{(t)}$

Cumulative regret

$$R^{(t)}_{\mathcal{A}_k} = \sup_{x_{\mathcal{A}_k} \in \Delta^{\mathcal{A}_k}} \sum_{\tau \le t} \left\langle x^{(\tau)}_{\mathcal{A}_k} - x_{\mathcal{A}_k}, \ell_{\mathcal{A}_k}(x^{(\tau)}) \right\rangle$$

Convergence of averages

$$\left[\forall k, \ \limsup_t \frac{R^{(t)}_{\mathcal{A}_k}}{t} \le 0 \right] \Rightarrow \bar{x}^{(t)} \to \mathcal{X}^*$$

Sufficient condition for $(x^{(t)})_t \to \mathcal{X}^*$

$f(x^{(t)})$ eventually decreasing

$$\Downarrow$$

$$f(x^{(t)}) \to f^*$$

$$\Downarrow$$

$$x^{(t)} \to \mathcal{X}^*$$

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Convergence of $\bar{x}^{(t)}$

Cumulative regret

$$R^{(t)}_{\mathcal{A}_k} = \sup_{x_{\mathcal{A}_k} \in \Delta^{\mathcal{A}_k}} \sum_{\tau \le t} \left\langle x^{(\tau)}_{\mathcal{A}_k} - x_{\mathcal{A}_k}, \ell_{\mathcal{A}_k}(x^{(\tau)}) \right\rangle$$

Convergence of averages

$$\left[\forall k,\ \limsup_t \frac{R^{(t)}_{\mathcal{A}_k}}{t} \le 0\right] \Rightarrow \bar{x}^{(t)} \to \mathcal{X}^*$$

Sufficient condition for $(x^{(t)})_t \to \mathcal{X}^*$

$f(x^{(t)})$ eventually decreasing

$$\Downarrow$$

$$f(x^{(t)}) \to f^*$$

$$\Downarrow$$

$$x^{(t)} \to \mathcal{X}^*$$

Outline

Replicator dynamics

Replicator equation [11]

$$\forall a \in \mathcal{A}_k, \frac{dx_a}{dt} = x_a \left(\langle \ell_{\mathcal{A}_k}(x), x_{\mathcal{A}_k} \rangle - \ell_a(x) \right) \quad (1)$$

Theorem: [3]

Every solution of the ODE (1) converges to the set of its stationary points.

[11] Jörgen W Weibull. *Evolutionary game theory*.
MIT press, 1997

[3] Simon Fischer and Berthold Vöcking. On the evolution of selfish routing.
In *Algorithms–ESA 2004*, pages 323–334. Springer, 2004

Replicator dynamics

Replicator equation [11]

$$\forall a \in \mathcal{A}_k, \frac{dx_a}{dt} = x_a \left(\langle \ell_{\mathcal{A}_k}(x), x_{\mathcal{A}_k} \rangle - \ell_a(x) \right) \tag{1}$$

Theorem: [3]

Every solution of the ODE (1) converges to the set of its stationary points.

---

[11] Jörgen W Weibull. *Evolutionary game theory.*
MIT press, 1997

[3] Simon Fischer and Berthold Vöcking. On the evolution of selfish routing.
In *Algorithms–ESA 2004*, pages 323–334. Springer, 2004

Approximate REPlicator update

Discretization of the continuous-time replicator dynamics

$$x_a^{(t+1)} - x_a^{(t)} = \eta_t x_a^{(t)} \left(\left\langle \ell_{\mathcal{A}_k}(x^{(t)}), x_{\mathcal{A}_k}^{(t)} \right\rangle - \ell_a(x^{(t)}) \right) + \eta_t U_a^{(t+1)}$$

- $(U^{(t)})_{t \geq 1}$ perturbations that satisfy for all $T > 0$,

$$\lim_{\tau_1 \to \infty} \max_{\tau_2 : \sum_{t=\tau_1}^{\tau_2} \eta_t < T} \left\| \sum_{t=\tau_1}^{\tau_2} \eta_t U^{(t+1)} \right\| = 0$$

- η_t discretization time steps.

[1] Michel Benaïm. Dynamics of stochastic approximation algorithms.
In *Séminaire de probabilités XXXIII*, pages 1–68. Springer, 1999

Convergence to Nash equilibria

Theorem [8]

Under AREP updates, if $\eta_t \downarrow 0$ and $\sum \eta_t = \infty$, then

$$x^{(t)} \to \mathcal{X}^*$$

- Affine interpolation of $x^{(t)}$ is an asymptotic pseudo trajectory.

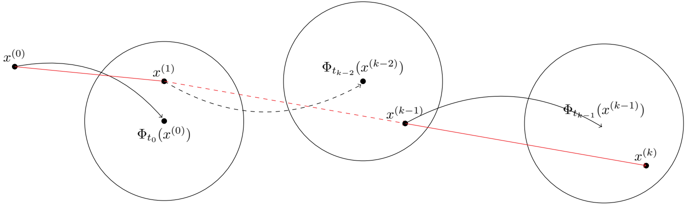

- Use f as a Lyapunov function.

However, No convergence rates.

[8] Walid Krichene, Benjamin Drighès, and Alexandre Bayen. Learning nash equilibria in congestion games.
SIAM Journal on Control and Optimization (SICON), to appear, 2014

Convergence to Nash equilibria

Theorem [8]

Under AREP updates, if $\eta_t \downarrow 0$ and $\sum \eta_t = \infty$, then

$$x^{(t)} \to \mathcal{X}^*$$

- Affine interpolation of $x^{(t)}$ is an asymptotic pseudo trajectory.

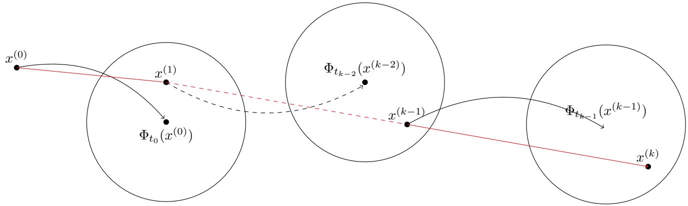

- Use f as a Lyapunov function.

However, No convergence rates.

---

[8] Walid Krichene, Benjamin Drighès, and Alexandre Bayen. Learning nash equilibria in congestion games.
SIAM Journal on Control and Optimization (SICON), to appear, 2014

Outline

1. Introduction
2. Convergence of agent dynamics
    - Background
    - Approximate replicator dynamics (AREP)
    - Distributed stochastic mirror descent dynamics (DSMD)
3. Application to routing

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Mirror Descent [10]

minimize	$f(x)$	convex function
subject to	$x \in \mathcal{X} \subset \mathbb{R}^d$	convex, compact set

Algorithm 1 MD Method with learning rates (η_t)

1: **for** $t \in \mathbb{N}$ **do**
2: $\quad g^{(t)} \in \partial f(x^{(t)})$
3: $\quad x^{(t+1)} = \arg\min_{x \in \mathcal{X}} \left\langle g^{(t)}, x \right\rangle + \frac{1}{\eta_t} D_{\psi_t}(x, x^{(t)})$
4: **end for**

η_t: learning rate
D_ψ: Bregman divergence

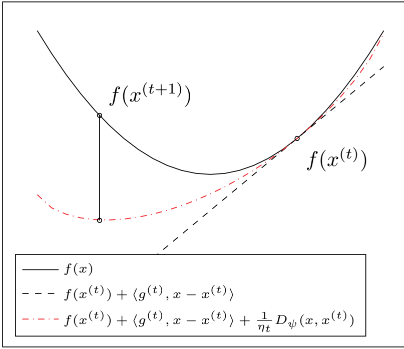

[10] A. S. Nemirovsky and D. B. Yudin. *Problem complexity and method efficiency in optimization.* Wiley-Interscience series in discrete mathematics. Wiley, 1983

Mirror Descent [10]

minimize	$f(x)$	convex function
subject to	$x \in \mathcal{X} \subset \mathbb{R}^d$	convex, compact set

Algorithm 2 MD Method with learning rates (η_t)

for $t \in \mathbb{N}$ **do**
  $g^{(t)} \in \partial f(x^{(t)})$
  $x^{(t+1)} = \arg\min_{x \in \mathcal{X}} \left\langle g^{(t)}, x \right\rangle + \frac{1}{\eta_t} D_{\psi_t}(x, x^{(t)})$
end for

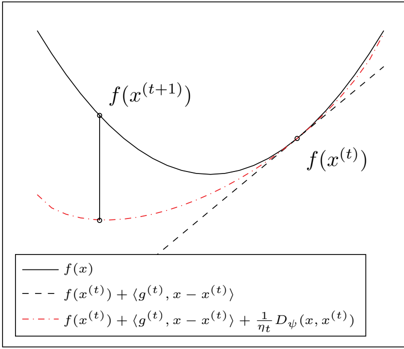

η_t: learning rate
D_ψ: Bregman divergence

[10] A. S. Nemirovsky and D. B. Yudin. *Problem complexity and method efficiency in optimization.*
Wiley-Interscience series in discrete mathematics. Wiley, 1983

Distributed Mirror Descent with heterogeneous agents

Distributed convex optimization

$$\begin{array}{ll} \text{minimize} & f(x) \\ \text{subject to} & x \in \mathcal{X} = \mathcal{X}_1 \times \cdots \times \mathcal{X}_K \end{array}$$

Algorithm 3 Distributed MD

for $t \in \mathbb{N}$ **do**

  Agent k observes $g_{\mathcal{A}_k}^{(t)}$

  Update

$$x_{\mathcal{A}_k}^{(t+1)} = \arg\min_{x_{\mathcal{A}_k} \in \mathcal{X}_k} \left\langle g_{\mathcal{A}_k}^{(t)}, x_{\mathcal{A}_k} - x_{\mathcal{A}_k}^{(t)} \right\rangle + \frac{1}{\eta_t^k} D_{\psi^k}(x_{\mathcal{A}_k}, x_{\mathcal{A}_k}^{(t)})$$

end for

- D_{ψ^k} and η_t^k depends on k

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

A true descent

Under mirror descent, $f(\bar{x}^{(t)}) \to f^*$.

A true descent [9]

If f has L-Lipschitz gradient, and $\eta_t \downarrow 0$, then eventually,

$$f(x^{(t+1)}) \leq f(x^{(t)})$$

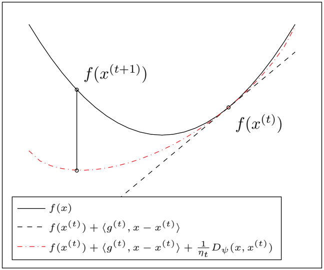

Figure: Mirror Descent iteration with decreasing η_t

[9] Walid Krichene, Syrine Krichene, and Alexandre Bayen. Convergence of mirror descent dynamics.
In *European Control Conference (ECC)*, accepted, 2015.

A true descent

Under mirror descent, $f(\bar{x}^{(t)}) \to f^*$.

A true descent [9]

If f has L-Lipschitz gradient, and $\eta_t \downarrow 0$, then eventually,

$$f(x^{(t+1)}) \leq f(x^{(t)})$$

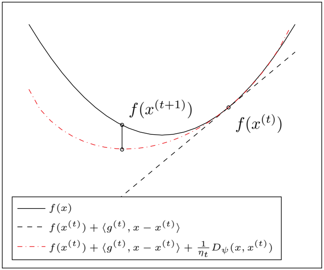

Figure: Mirror Descent iteration with decreasing η_t

[9] Walid Krichene, Syrine Krichene, and Alexandre Bayen. Convergence of mirror descent dynamics.
In *European Control Conference (ECC)*, accepted, 2015.

A true descent

Under mirror descent, $f(\bar{x}^{(t)}) \to f^*$.

A true descent [9]

If f has L-Lipschitz gradient, and $\eta_t \downarrow 0$, then eventually,

$$f(x^{(t+1)}) \leq f(x^{(t)})$$

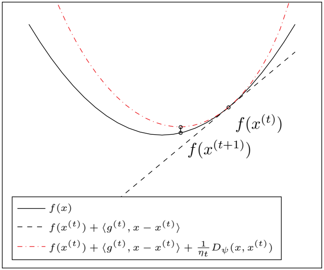

Figure: Mirror Descent iteration with decreasing η_t

[9] Walid Krichene, Syrine Krichene, and Alexandre Bayen. Convergence of mirror descent dynamics.
In *European Control Conference (ECC)*, accepted, 2015

FORCES FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

A true descent

Consequence:

Theorem: Convergence of DMD [9]

Suppose f has L Lipschitz gradient. Then under the DMD class with $\eta_t \downarrow 0$ and $\sum \eta_t = \infty$,

$$f(x^{(t)}) - f^* = O\left(\frac{\sum_{\tau \leq t} \eta_\tau}{t} + \frac{1}{t\eta_t} + \frac{1}{t}\right)$$

How robust is the convergence if losses are stochastic?
Players do not observe the true loss, but have an estimate $\hat{\ell}(x^{(t)})$

[9] Walid Krichene, Syrine Krichene, and Alexandre Bayen. Convergence of mirror descent dynamics.
In *European Control Conference (ECC), accepted*, 2015

A true descent

Consequence:

Theorem: Convergence of DMD [9]

Suppose f has L Lipschitz gradient. Then under the DMD class with $\eta_t \downarrow 0$ and $\sum \eta_t = \infty$,

$$f(x^{(t)}) - f^* = O\left(\frac{\sum_{\tau \le t} \eta_\tau}{t} + \frac{1}{t\eta_t} + \frac{1}{t}\right)$$

How robust is the convergence if losses are stochastic?
Players do not observe the true loss, but have an estimate $\hat{\ell}(x^{(t)})$

[9] Walid Krichene, Syrine Krichene, and Alexandre Bayen. Convergence of mirror descent dynamics.
In *European Control Conference (ECC), accepted*, 2015

Distributed Stochastic Mirror Descent (DSMD)

At iteration t

- Have a stochastic vector $\hat{g}^{(t)}$
- $\hat{g}^{(t)}$ unbiased: $\mathbb{E}\left[\hat{g}^{(t)}|\mathcal{F}_{t-1}\right] \in \partial f(x^{(t)})$ a.s. ($\mathcal{F}_t$ natural filtration of $(\hat{g}^{(t)})$)

Algorithm 4 DSMD dynamics

1: **for** $t \in \mathbb{N}$ **do**
2: Agent k observes $\hat{g}^{(t)}_{\mathcal{A}_k}$
3: Update

$$x^{(t+1)}_{\mathcal{A}_k} = \arg\min_{x_{\mathcal{A}_k} \in \mathcal{X}_k} \left\langle \hat{g}^{(t)}_{\mathcal{A}_k}, x_{\mathcal{A}_k} - x^{(t)}_{\mathcal{A}_k} \right\rangle + \frac{1}{\eta^k_t} D_{\psi^k}(x_{\mathcal{A}_k}, x^{(t)}_{\mathcal{A}_k})$$

4: **end for**

Assume

- $\exists G > 0$ s.t. $\mathbb{E}\left[\|\hat{g}^{(t)}\|_*^2\right] \leq G \; \forall t$

FORCES FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Distributed Stochastic Mirror Descent (DSMD)

At iteration t

- Have a stochastic vector $\hat{g}^{(t)}$
- $\hat{g}^{(t)}$ unbiased: $\mathbb{E}\left[\hat{g}^{(t)} | \mathcal{F}_{t-1}\right] \in \partial f(x^{(t)})$ a.s.

  ($\mathcal{F}_t$ natural filtration of $(\hat{g}^{(t)})$)

Algorithm 5 DSMD dynamics

1: **for** $t \in \mathbb{N}$ **do**
2: Agent k observes $\hat{g}^{(t)}_{\mathcal{A}_k}$
3: Update

$$x^{(t+1)}_{\mathcal{A}_k} = \arg\min_{x_{\mathcal{A}_k} \in \mathcal{X}_k} \left\langle \hat{g}^{(t)}_{\mathcal{A}_k}, x_{\mathcal{A}_k} - x^{(t)}_{\mathcal{A}_k} \right\rangle + \frac{1}{\eta^k_t} D_{\psi^k}(x_{\mathcal{A}_k}, x^{(t)}_{\mathcal{A}_k})$$

4: **end for**

Assume

- $\exists G > 0$ s.t. $\mathbb{E}\left[\|\hat{g}^{(t)}\|_*^2\right] \leq G \ \forall t$

Convergence of DSMD

Existing result: $\mathbb{E}[f(\bar{x}^{(t)})] \to f^*$

Our results

- All convex functions (including non-smooth)

$$x^{(t)} \overset{a.s.}{\to} \mathcal{X}^*$$

$$\mathbb{E}\left[f(x^{(t)})\right] - f^\star = O\left(\sum_k \frac{\log t}{t^{\min(\alpha_k, 1-\alpha_k)}}\right)$$

- Strongly convex functions

$$\mathbb{E}\left[D_\psi(x^\star, x^{(t)})\right] = O(\sum_k t^{-\alpha_k})$$

Rates are for $\eta_t^k = \frac{\theta_k}{t^{\alpha_k}}$, $\alpha_k \in (0, 1)$

[6] Syrine Krichene, Walid Krichene, Roy Dong, and Alexandre Bayen. Convergence of stochastic mirror descent and applications to distributed routing.
In *Allerton Conference on Communication, Control and Computing, in preparation*, 2015

Outline

1. Introduction

2. Convergence of agent dynamics
   - Background
   - Approximate replicator dynamics (AREP)
   - Distributed stochastic mirror descent dynamics (DSMD)

3. Application to routing

Application to the routing game

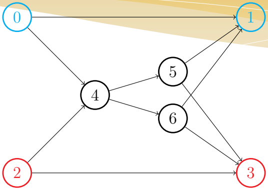

Figure: A strongly convex example.

- Centered Gaussian noise on edges.
- Population 1: Hedge with $\eta_t^1 = t^{-1}$
- Population 2: Hedge with $\eta_t^2 = t^{-1}$

Hedge algorithm

$$x_{\mathcal{A}_k}^{(t+1)} \propto x_a^{(t)} e^{-\eta_t \ell_a^{(t)}}$$

Routing game with strongly convex potential

Mass distributions $x_{\mathcal{A}_k}^{(t)}$ Path losses $\ell_{\mathcal{A}_k}(x^{(t)})$

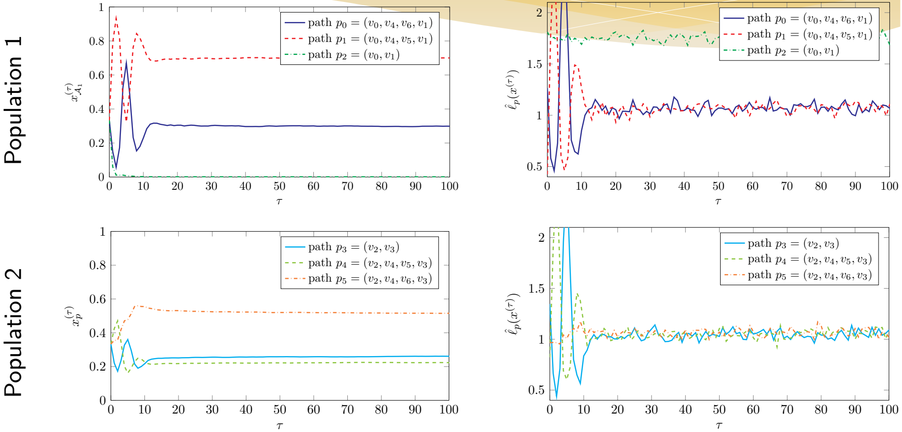

Figure: Population distributions and noisy path losses

Routing game with strongly convex potential

Mass distributions $\mathbb{E}\, x_{\mathcal{A}_k}^{(t)}$ Path losses $\mathbb{E}\, \ell_{\mathcal{A}_k}(x^{(t)})$

Population 1

Population 2

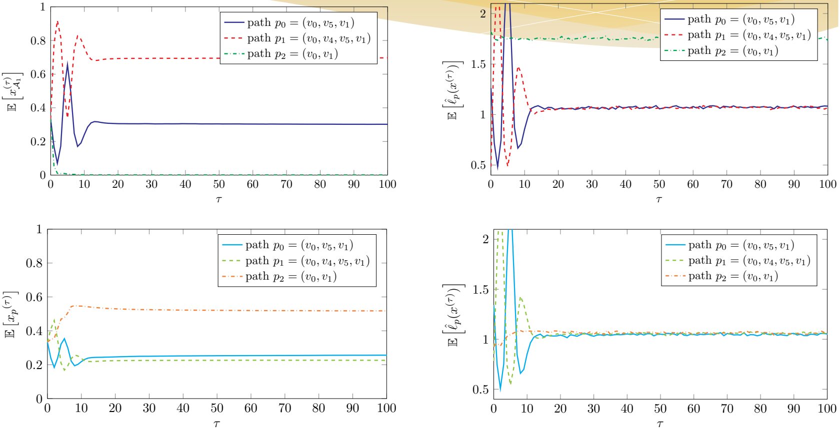

Figure: Population distributions and noisy path losses

Routing game with strongly convex potential

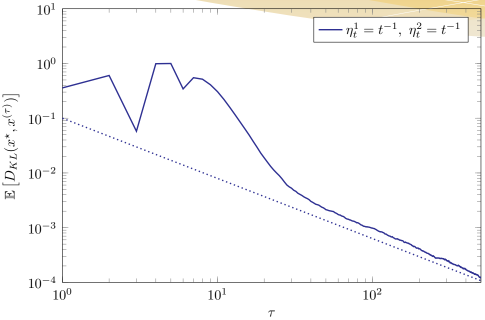

Figure: Distance to equilibrium.

For $\eta_t^k = \frac{\theta_k}{\ell_f t^{\alpha_k}},\ \alpha_k \in (0,1]$, $\mathbb{E}\left[D_\psi(x^\star, x^{(t)})\right] = O(\sum_k t^{-\alpha_k})$

Routing game with weakly convex potential

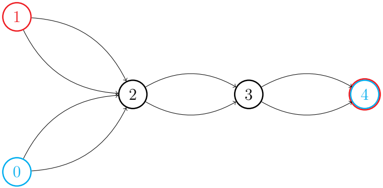

Figure: A weakly convex example.

Routing game with weakly convex potential

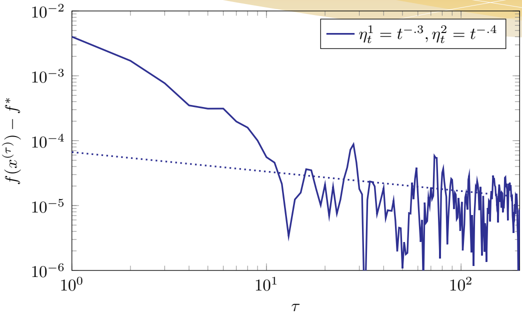

Figure: Potential values.

For $\frac{\theta_k}{t^{\alpha_k}}$, $\alpha_k \in (0,1)$, $\mathbb{E}\left[f(x^{(t)})\right] - f^\star = O\left(\sum_k \frac{\log t}{t^{\min(\alpha_k, 1-\alpha_k)}}\right)$

FORCES FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Routing game with weakly convex potential

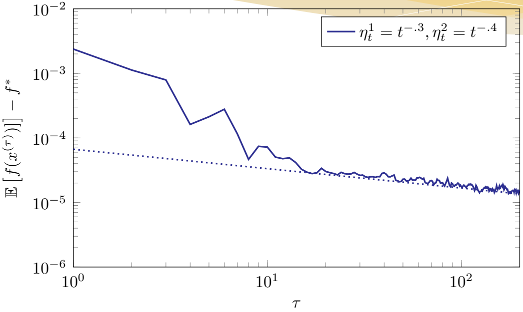

Figure: Potential values.

For $\frac{\theta_k}{t^{\alpha_k}}$, $\alpha_k \in (0,1)$, $\mathbb{E}\left[f(x^{(t)})\right] - f^\star = O\left(\sum_k \frac{\log t}{t^{\min(\alpha_k, 1-\alpha_k)}}\right)$

FORCES FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Routing game with weakly convex potential

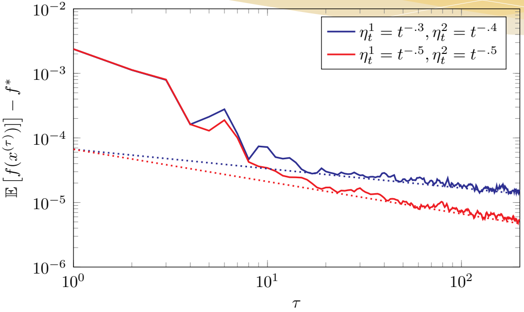

Figure: Potential values.

For $\frac{\theta_k}{t^{\alpha_k}}$, $\alpha_k \in (0,1)$, $\mathbb{E}\left[f(x^{(t)})\right] - f^\star = O\left(\sum_k \frac{\log t}{t^{\min(\alpha_k, 1-\alpha_k)}}\right)$

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Routing game with weakly convex potential

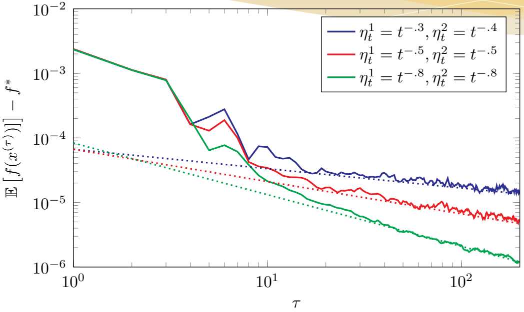

Figure: Potential values.

For $\frac{\theta_k}{t^{\alpha_k}}$, $\alpha_k \in (0,1)$, $\mathbb{E}\left[f(x^{(t)})\right] - f^\star = O\left(\sum_k \frac{\log t}{t^{\min(\alpha_k, 1-\alpha_k)}}\right)$

Summary

Convergence guarantees for distributed dynamics

- Under no-regret learning, $\bar{x}^{(t)} \to \mathcal{N}$
- Under AREP dynamics, $x^{(t)} \xrightarrow{a.s.} \mathcal{N}$
- Under DMD dynamics, $x^{(t)} \to \mathcal{N}$ with rate $O\left(\frac{\sum_{\tau \le t} \eta_\tau}{t} + \frac{1}{t\eta_t} + \frac{1}{t}\right)$
- Under Stochastic MD, $x^{(t)} \xrightarrow{a.s.} \mathcal{N}$, and
  $\mathbb{E}\, f(x^{(t)}) \to f^*$ with rate $O\left(\sum_k \frac{\log t}{t^{\min(\alpha_k, 1-\alpha_k)}}\right)$
- if potential is strongly convex, $\mathbb{E}\, D_\psi(x^*, x^{(t)}) \to 0$ with rate $O(\sum_k t^{-\alpha_k})$

Applications

- Distributed machine learning
- Used as a model for optimal control of potential games

Ongoing and future work

- Efficient Bregman projections (CDC 2015)
- Learning on a Continuum (NIPS 2015)
- Fitting of the learning model to observed dynamics

Applications

- Distributed machine learning
- Used as a model for optimal control of potential games

Ongoing and future work

- Efficient Bregman projections (CDC 2015)
- Learning on a Continuum (NIPS 2015)
- Fitting of the learning model to observed dynamics

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Thank you

References I

[1] Michel Benaïm. Dynamics of stochastic approximation algorithms. In *Séminaire de probabilités XXXIII*, pages 1–68. Springer, 1999.

[2] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games.* Cambridge University Press, 2006.

[3] Simon Fischer and Berthold Vöcking. On the evolution of selfish routing. In *Algorithms–ESA 2004*, pages 323–334. Springer, 2004.

[4] Yoav Freund and Robert E Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1):79–103, 1999.

[5] James Hannan. Approximation to Bayes risk in repeated plays. *Contributions to the Theory of Games*, 3:97–139, 1957.

[6] Syrine Krichene, Walid Krichene, Roy Dong, and Alexandre Bayen. Convergence of stochastic mirror descent and applications to distributed routing. In *Allerton Conference on Communication, Control and Computing, in preparation*, 2015.

References II

[7] Walid Krichene, Benjamin Drighès, and Alexandre Bayen. On the convergence of no-regret learning in selfish routing. In *31st International Conference on Machine Learning (ICML)*. JMLR, 2014.

[8] Walid Krichene, Benjamin Drighès, and Alexandre Bayen. Learning nash equilibria in congestion games. *SIAM Journal on Control and Optimization (SICON), to appear*, 2014.

[9] Walid Krichene, Syrine Krichene, and Alexandre Bayen. Convergence of mirror descent dynamics. In *European Control Conference (ECC), accepted*, 2015.

[10] A. S. Nemirovsky and D. B. Yudin. *Problem complexity and method efficiency in optimization*. Wiley-Interscience series in discrete mathematics. Wiley, 1983.

[11] Jörgen W Weibull. *Evolutionary game theory*. MIT press, 1997.