Convergence of Distributed Routing with Heterogeneous Population Dynamics

Walid Krichene Alexandre Bayen
Syrine Krichene Benjamin Drighes Roy Dong
Electrical Engineering and Computer Sciences, UC Berkeley

Routing Game

- Directed graph (V, E)
- Population k: paths $\mathcal{P}_k$
- Population distribution over paths $x_{\mathcal{P}_k} \in \Delta^{\mathcal{P}_k}$
- Loss on path i of population k: $\ell_i^k(x)$

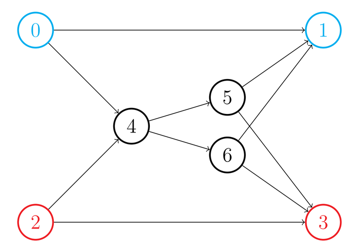

Figure: Example network

Nash equilibrium

- $x^* = (x^*_{\mathcal{P}_1}, \dots, x^*_{\mathcal{P}_k})$ is an equilibrium if $\forall k$,

$$\left\langle \ell_{\mathcal{P}_k}(x^*), x^*_{\mathcal{P}_k} \right\rangle \leq \left\langle \ell_{\mathcal{P}_k}(x^*), x_{\mathcal{P}_k} \right\rangle$$

Losses are minimal on the support of $x^*_{\mathcal{P}_k}$

- Rosenthal potential

$$f(x) \text{ Convex}$$
$$\nabla_{x^k} f(x) = \ell^k(x)$$
$$\mathcal{N} = \arg\min_{x \in \Delta^{\mathcal{P}_1} \times \dots \times \Delta^{\mathcal{P}_K}} f(x)$$

Optimality conditions:

$$\langle \ell(x^*), x - x^* \rangle \geq 0 \quad \forall x \qquad \Leftrightarrow \qquad \forall k,\ \forall x_{\mathcal{P}_k}, \left\langle \ell_{\mathcal{P}_k}(x^*_{\mathcal{P}_k}), x_{\mathcal{P}_k} - x^*_{\mathcal{P}_k} \right\rangle$$

Online Learning

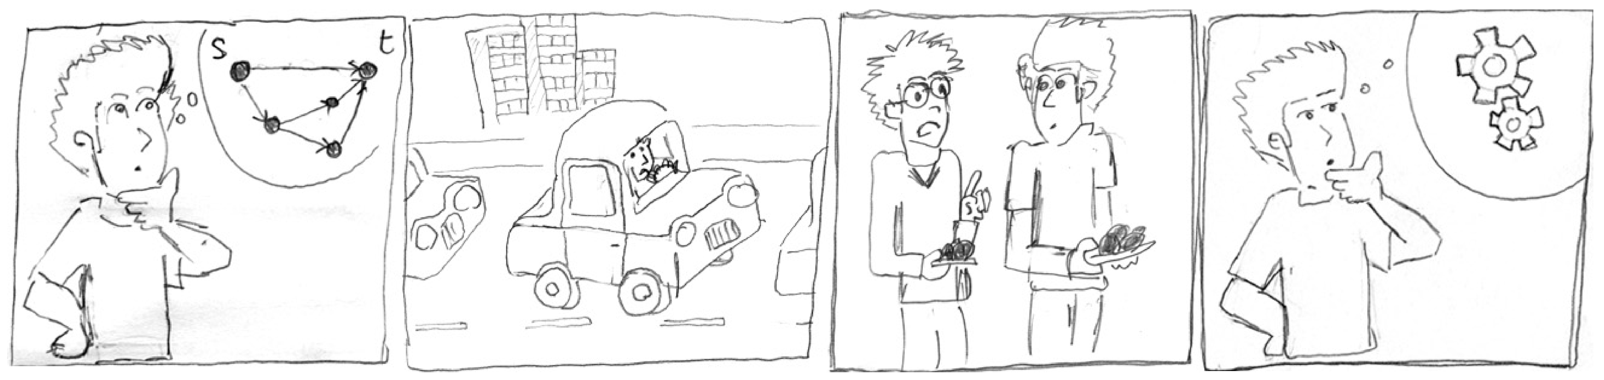

$x_{\mathcal{P}_k}^{(t)}$ | sample path $p \sim x_{\mathcal{P}_k}^{(t)}$ | $\ell_{\mathcal{P}_k}(x^{(t)})$ is revealed | $x_{\mathcal{P}_k}^{(t+1)} = u_k(x_{\mathcal{P}_k}^{(t)}, \text{history})$

Main problem

Define a class $\mathcal{C}$ of algorithms (update rules) such that

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

Extension: Losses are noisy $\hat{\ell}_{\mathcal{P}_k}(x^{(t)})$ (e.g. loss is not observed, but estimated) with

$$\mathbb{E}[\hat{\ell}_{\mathcal{P}_k}(x^{(t)}) | x^{(t)}] = \ell_{\mathcal{P}_k}(x^{(t)})$$

Regret

Instantaneous regret:

$$r^{(t)}(x) = \left\langle \ell(x^{(t)}), x^{(t)} - x \right\rangle$$
$$x^{(t)} \to \mathcal{N} \Leftrightarrow \limsup_t \sup_x r^{(t)}(x) \leq 0$$

Average cumulative regret

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

Convergence of average strategies

$$\bar{x}^{(t)} = \frac{1}{t} \sum_{\tau \leq t} x^{(\tau)} \to \mathcal{N} \Leftrightarrow \limsup_t \sup_x R^{(t)}(x) \leq 0$$

By convexity of f,

$$f\left(\frac{1}{t}\sum_{\tau \leq t} x^{(\tau)}\right) - f(x) \leq \frac{1}{t}\sum_{\tau \leq t} f(x^{(\tau)}) - f(x) \leq \frac{1}{t}\sum_{\tau \leq t} \left\langle \ell(x^{(t)}), x^{(t)} - x \right\rangle = R^{(t)}(x)$$

When does $x^{(t)}$ converge?

Observation

If $f(x^{(t)})$ is eventually monotone, then $f(x^{(t)}) \to f^*$.

Class 1: Approximate Replicator dynamics

Replicator equation

$$\forall p \in \mathcal{P}_k, \frac{dx_p^k}{dt} = x_p^k \left(\left\langle \ell_{\mathcal{P}_k}(x), x_{\mathcal{P}_k} \right\rangle - \ell_p^k(x) \right) \tag{1}$$

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

Discretization of the continuous-time replicator dynamics

$$x_p^{(t+1)} - x_p^{(t)} = \eta_t x_p^{(t)} \left(\left\langle \ell^k(x^{(t)}), x^{(t)}_{\mathcal{P}_k} \right\rangle - \ell^k(x^{(t)}) \right) + \eta_t U_p^{(t+1)}$$

Convergence of AREP updates

If $\eta_t \downarrow 0$ and $\sum \eta_t = \infty$, then

$$x^{(t)} \to \mathcal{N}$$

Proof:

- f is a Lyapunov function for Nash equilibria in the continuous system.
- Affine interpolation of $x^{(t)}$ is an asymptotic pseudo trajectory.

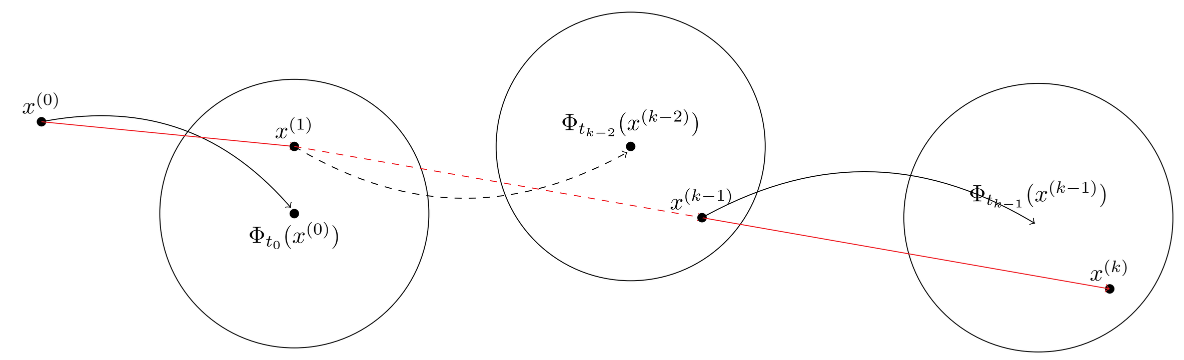

- However, unknown convergence rate

Class 1: Distributed Stochastic Mirror Descent (DSMD) Dynamics

Given, for each population,

- A Bregman divergence D_{ψ^k}
- A sequence of learning rates η_t^k

Algorithm 1 DSMD Method

1: **for** $t \in \mathbb{N}$ **do**
2: $\hat{\ell}^{(t)}$, such that $\mathbb{E}[\hat{\ell}^{(t)} | \mathcal{F}_{t+1}] = \ell^{(t)} = \nabla f(x^{(t)})$
3: $$x_{\mathcal{P}_k}^{(t+1)} = \arg\min_{x_{\mathcal{P}_k} \in \mathcal{X}_k} \left\langle \hat{\ell}^{(t)}_{\mathcal{P}_k}, x_{\mathcal{P}_k} - x_{\mathcal{P}_k}^{(t)} \right\rangle + \frac{1}{\eta_t^k} D_{\psi^k}(x_{\mathcal{P}_k}, x_{\mathcal{P}_k}^{(t)})$$
4: **end for**

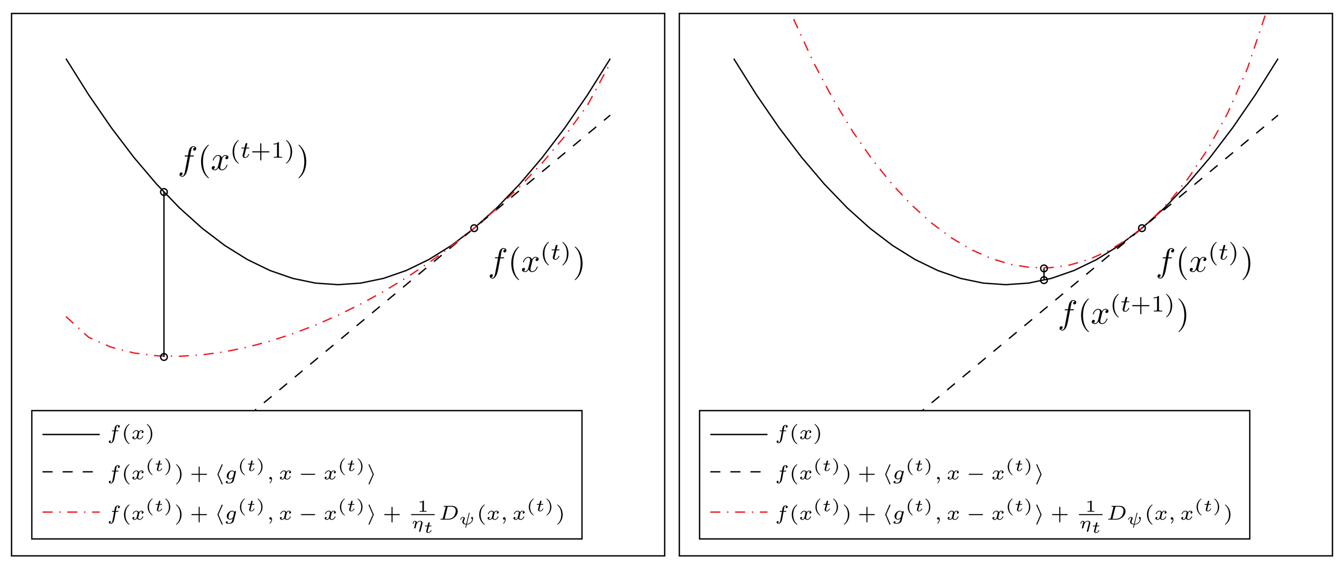

Convergence rates of MD regret (non stochastic)

Regret bound

$$\sup_{x^k \in m_k(S_k)\Delta^{\mathcal{P}_k}} R^{k(t)}(x^k) \leq \frac{L_k^2}{2\ell_{\psi^k}} \sum_{\tau=1}^{t} \eta_\tau^k + \frac{D_k}{\eta_t^k} \tag{2}$$

In particular, If $\eta_t^k \to 0$ and $\sum_{\tau \leq t} \eta_\tau^k = \infty$, then

- $\limsup_t \sup_{x^k} \frac{R^{k(t)}}{t} \to 0$ at a rate $O\left(\frac{\sum_{\tau=1}^t \eta_\tau^k}{t} + \frac{1}{t\eta_t^k}\right)$
- For example, if $\eta_t^k = \theta(t^{-\alpha_k})$, $\alpha_k \in (0, 1)$, then $O(t^{-\min(\alpha_k, 1-\alpha_k)})$.
- If ℓ is Lipschitz, the potentials $f(x^{(t)})$ are eventually decreasing, so $f(x^{(t)}) \to \mathcal{N}$ at the same rate.

Numerical simulation: MD method

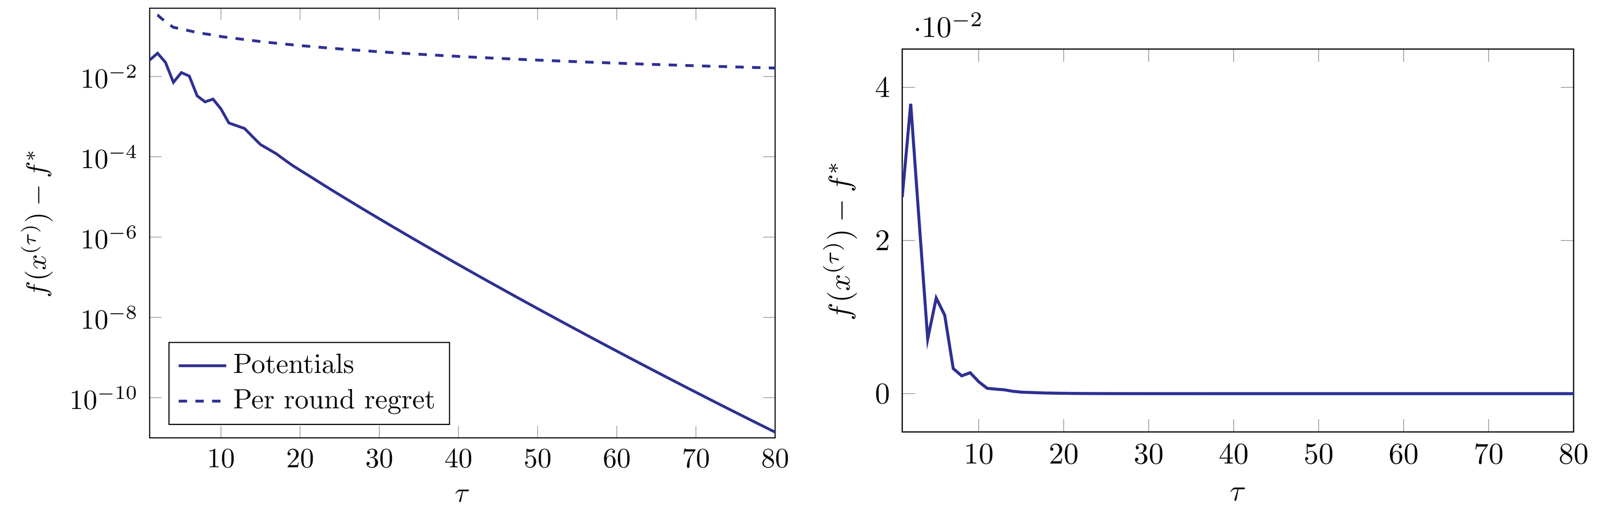

Figure: Potential values $f(x^{(t)}) - f^*$.

Convergence rate of DSMD

Convergence rate of DSMD

Distributed SMD such that $\eta_t^k = \frac{\theta_k}{t^{\alpha_k}}$ with $\alpha_k \in (0, 1)$. Then

$$\mathbb{E}[f(x^{(t)})] - f(x^*) \leq \left(1 + \sum_{\tau=1}^{t} \frac{1}{\tau}\right) \sum_{k=1}^{K} \left(\frac{1}{t^{1-\alpha_k}} \frac{D}{\theta_k} + \frac{\theta_k G}{2\ell_\psi (1 - \alpha_k)} \frac{1}{t^{\alpha_k}} \right)$$

Numerical simulation: DSMD method

- Centered Gaussian noise on edges.
- Population 1: Hedge with $\eta_t^1 = t^{-0.1}$
- Population 2: Hedge with $\eta_t^2 = \frac{1}{2} t^{-0.5}$

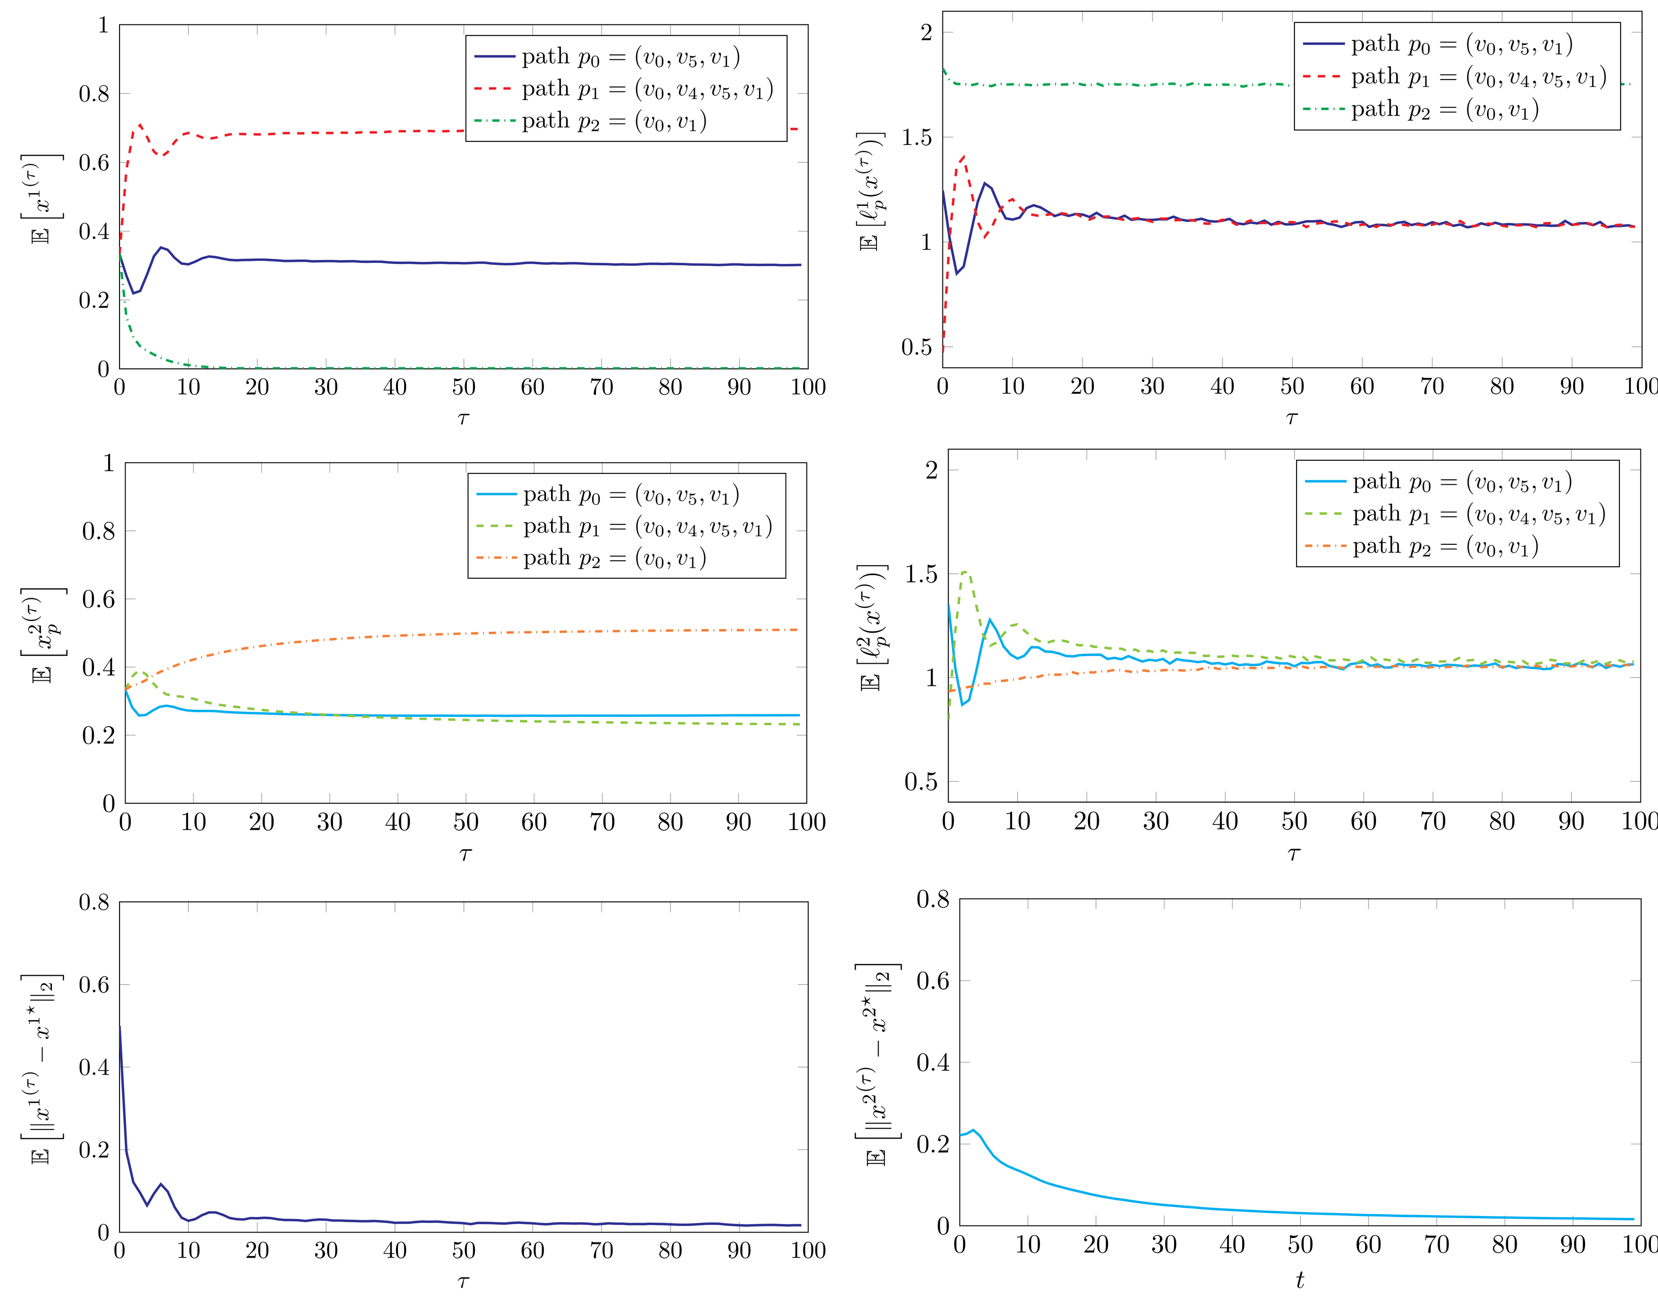

Figure: Population distributions, path losses, and expected distance to equilibrium

Summary

- Convergence of two classes of dynamics: AREP and DSMD.
- AREP is larger, but no convergence rates.
- Convergence of DSMD is for heterogeneous populations (each uses a different update rule)
- Robust to noise (e.g. losses are estimated)

This work was supported in part by FORCES (Foundations Of Resilient CybEr-physical Systems), which receives support from the National Science Foundation (NSFaward numbers CNS-1238959, CNS-1238962, CNS-1239054, CNS-1239166).