Disabling External Influence in Social Networks via Edge Recommendation

Victor Amelkin, Ambuj K. Singh

University of California, Santa Barbara

Introduction

- **Goal:** Defend opinion distribution against external influence.
- Large directed strongly connected social network of n agents.
- Row-stochastic interpersonal influence matrix $W \in [0,1]^{n \times n}$.
- $x(t) \in [0,1]^n$ – agents' opinions at time t.
- $\pi \in \mathbb{R}^n$ — network nodes' eigenvector centralities.
- DeGroot model ($\approx$ Markov chain): $x(t+1) = Wx(t) \to \langle \pi, x(0) \rangle \mathbb{1}$.
- In what follows, we use notation $x = x(0)$.
- $\langle \boldsymbol{\pi}, \boldsymbol{x} \rangle$ – **asymptotic consensus value**; "opinion attractor".

Problem Statement by Example

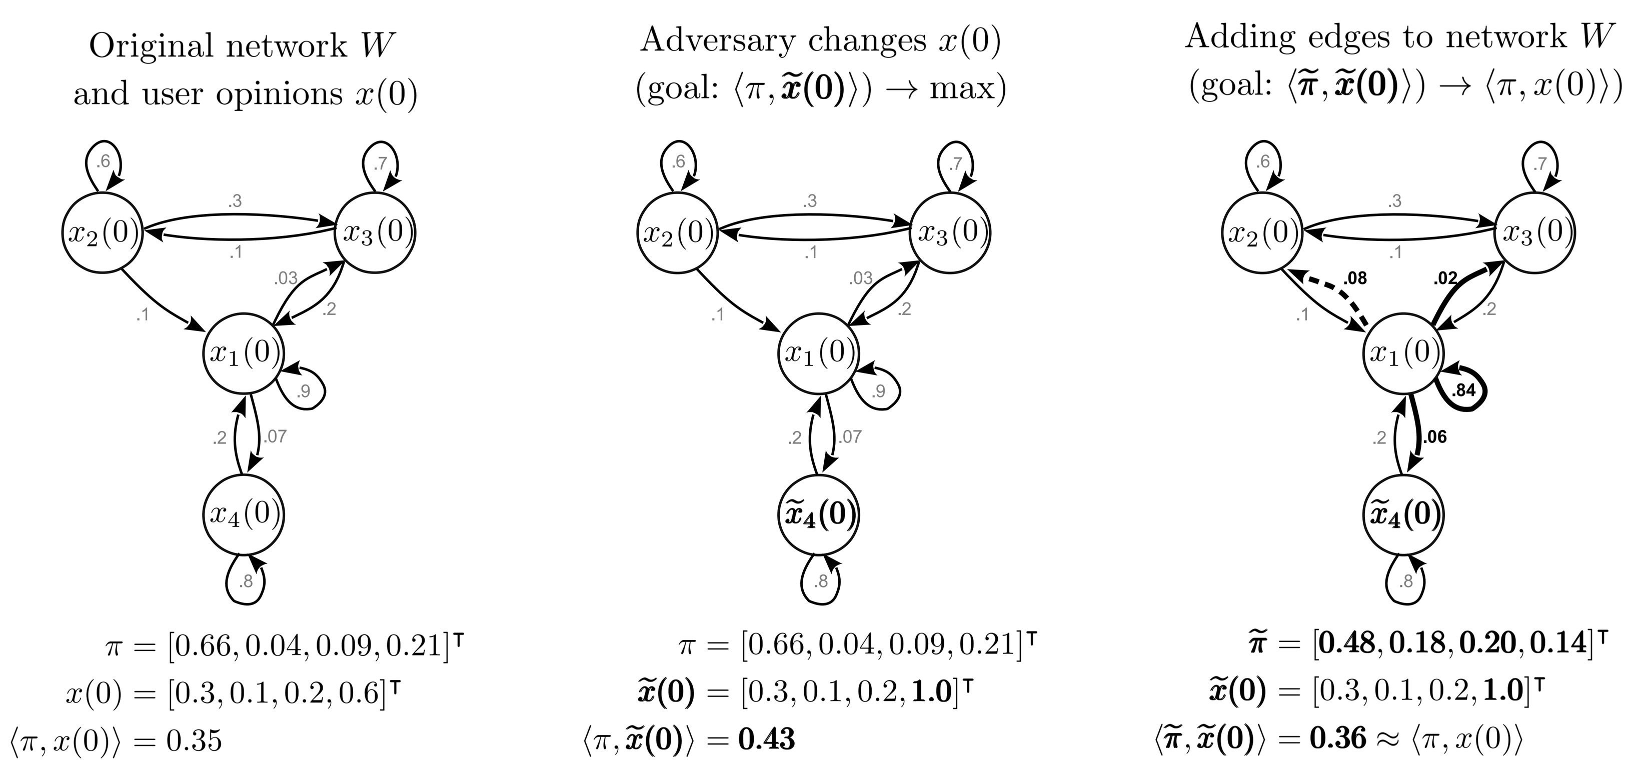

Problem and Its Hardness

- **Adversary's Goal:** Maximize $\langle \pi, \widetilde{x} \rangle$ via altering $x \to \widetilde{x}$
- **Our Goal:** Return $\langle \widetilde{\pi}, \widetilde{x} \rangle$ back to $\langle \pi, x \rangle$ via altering $\pi \to \widetilde{\pi}$ by the means of recommending new links to users. For a single-edge (r, c) perturbation, $\widetilde{W} = W - \theta_{rc}\,\mathrm{diag}(e_r)W + \theta_{rc} e_r e_c^\intercal$.

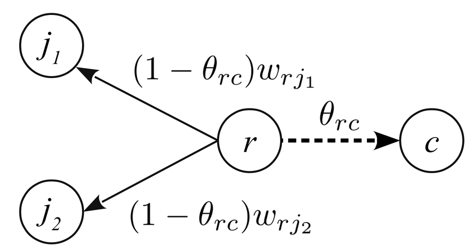

- **Problem:** $\mathrm{DIVER}(W, k, x, \widetilde{x}) = \arg\min_{\widetilde{W}} |\langle \widetilde{\pi}(\widetilde{W}), \widetilde{x} \rangle - \langle \pi, x \rangle|$,

  where the perturbed $\widetilde{W}$ differs from W by k new edges, with weight θ_{ij} of an added edge (i, j) being predefined (by user i). DIVER is NP-hard.

General Solution for DIVER

$$\mathrm{DIVER}(W, k, x, \widetilde{x}) = \arg\min_{\widetilde{W}} |\langle \widetilde{\pi}, \widetilde{x} \rangle - \langle \pi, x \rangle|$$

- **Method:** reduce $\langle \widetilde{\pi}, \widetilde{x} \rangle$ through iterative edge addition until it gets close enough to $\langle \pi, x \rangle$

$$\langle \pi, \widetilde{x} \rangle - \langle \widetilde{\pi}, \widetilde{x} \rangle \to \max$$

- **Central Question:** How does $\langle \widetilde{\pi}, \widetilde{x} \rangle$ change when a single edge (r, c) with weight θ_{rc} is added to network W?

Network Perturbation Analysis

- Single-edge perturbation: $\widetilde{W} = W - \theta_{rc}\,\mathrm{diag}(e_r)W + \theta_{rc} e_r e_c^\intercal$

Theorem 1. *Under single-edge perturbation of W with edge (r, c) having weight θ_{rc}, the eigenvector centrality changes as follows:*

$$\widetilde{\pi}_j = \pi_j \left[1 - \frac{\theta_{rc}(m_{cj} \cdot (1 - \delta\{j, c\}) - m_{rj} + 1)}{m_{rr} + \theta_{rc}(m_{cr} - m_{rr} + 1)} \right],$$

where m_{ij} is the mean first passage time (MFPT) from state i to state j of Markov chain W, and $\delta\{\cdot,\cdot\}$ is Kronecker's delta.

Theorem 2. *Under single-edge perturbation of W with edge (r, c) having weight θ_{rc}, the asymptotic consensus value changes as follows:*

$$f_\pi(r, c) = \langle \pi, \widetilde{x} \rangle - \langle \widetilde{\pi}, \widetilde{x} \rangle = \theta_{rc} \frac{\sum_{j=1}^{n} \pi_j (m_{cj} \cdot (1 - \delta\{j, c\}) - m_{rj} + 1)\widetilde{x}_j}{m_{rr} + \theta_{rc}(m_{rc} - m_{rr} + 1)}.$$

How to efficiently solve DIVER in large networks?

- Approach to DIVER: iteratively adding edges (r, c) with top values $f_\pi(r, c)$ until satisfied with the value of $\langle \widetilde{\pi}, \widetilde{x} \rangle$.
- **Issue 1:** There are $\mathcal{O}(n^2)$ candidate edges in a sparse network.
- **Issue 2:** How to efficiently compute $f_\pi(r, c)$?
  - Evaluation of a single $f_\pi(r, c)$ involves summation over $\mathcal{O}(n)$ terms.
  - Direct computation of MFPTs m_{ij} would cost at least $\mathcal{O}(n^3)$.

Efficient candidate edge selection

- Focus on $\mathcal{O}(n)$ candidate edges, outgoing from $n_{src} \ll n$ nodes.
- (Most good candidate edges emanate from a small number of nodes.)
- In hierarchical networks, these edge sources are top-centrality nodes.

Efficient computation of $f_\pi(r, c)$

- In hierarchical networks, f_π is largely determined by a small (n_{src}) number of top-centrality nodes.

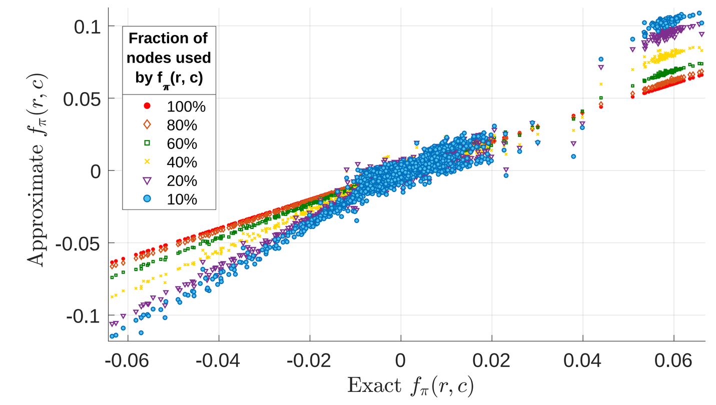

- We can estimate MFPTs via finite-time random walks; all the MFPTs to and from n_{src} top-centrality nodes will converge in $\mathcal{O}(n)$ time.

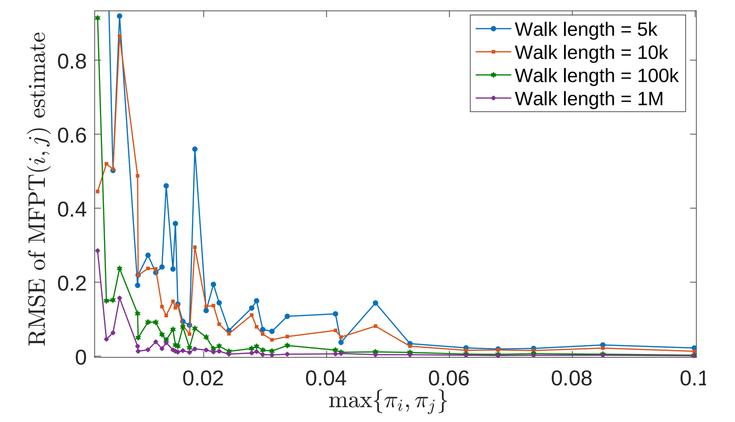

- **Outcome:** $\mathcal{O}(n)$-time heuristic for DIVER for hierarchical networks.

Result

- Solving DIVER via iterative edge addition using $n_{src} = 0.2n$ of network nodes:

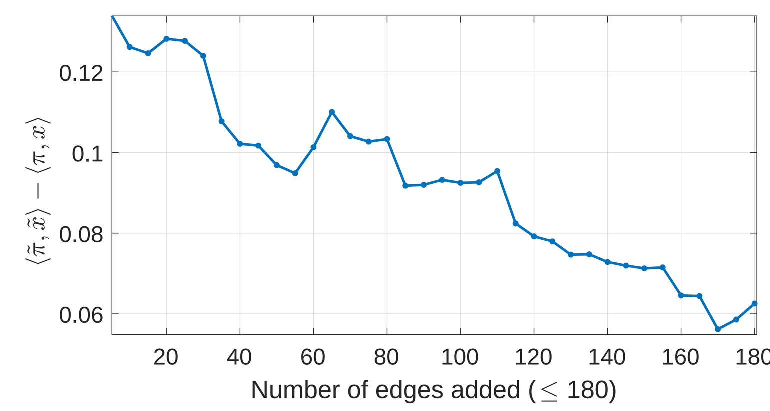

- [1] Amelkin V., Singh A.K., *"Disabling external influence in social networks via edge recommendation"*, in submission (2018).

Quantitative Network-based Models of Adaptive Team Behavior
Milestone: 1.4

This material is based upon work supported by, or in part by, the U.S. Army Research Laboratory and the U.S. Army Research Office under grant number W911NF-15-1-0577.

https://muriteams.cs.ucsb.edu