Defense Against Shortest Path Attacks

Benjamin A. Miller[1], Zohair Shafi[1], Wheeler Ruml[2], Yevgeniy Vorobeychik[3], Tina Eliassi-Rad[1], and Scott Alfeld[4]

[1]Northeastern University, Boston, MA, USA
[2]University of New Hampshire, Durham, NH, USA
[3]Washington University in St. Louis, St. Louis, MO, USA
[4]Amherst College, Amherst, MA, USA

Problem Statement

Scenario: An adversary can divert traffic to specific parts of a graph by removing edges

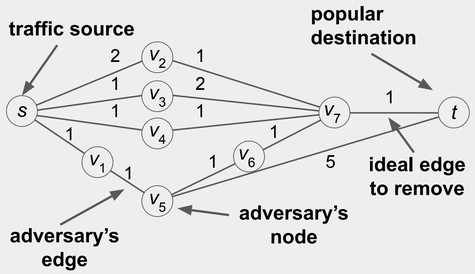

***Goal*: Increase attacker's budget required to successfully attack**

Assumed Context:

- We have a graph and know the true weights
- But we publish *approximate* weights
- Desired outcome: Weight uncertainty will increase cost of attack
- Constraint: Minimize negative impact on legitimate users

Research Questions:

- How much does adding noise to weights increase the attacker's budget?
- What is the impact on legitimate users of the network?

Attack: PATHATTACK*

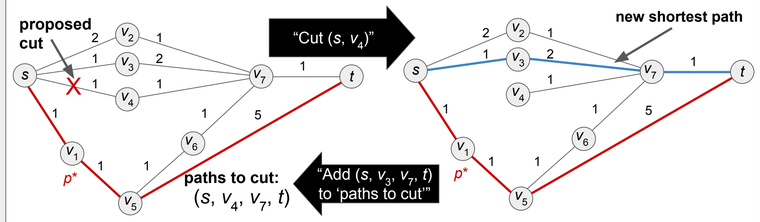

- Algorithm to approximate optimal path cut cost
- Iteratively add constraints until *p** is shortest

*B. A. Miller et al., "PATHATTACK: Attacking Shortest Paths in Complex Networks," in *ECML PKDD*, 2021

Defense: Noisy Weights†

- Prior work used noisy weights to achieve differential privacy
- We use the same technique to increase adversarial uncertainty, increasing cost

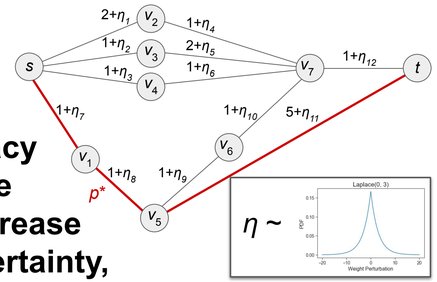

†A. Sealfon, "Shortest paths and distances with differential privacy," in *PODS*, 2016

Adversarial Adaptations

Informed Attacker

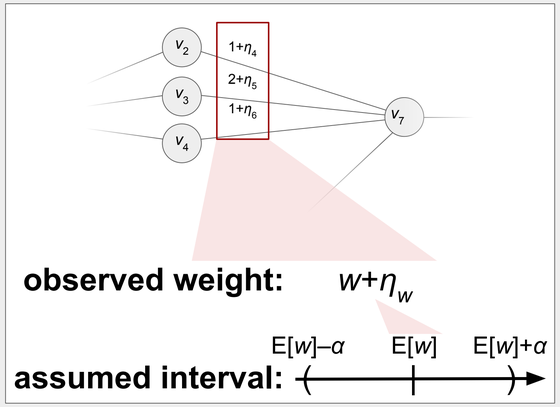

- Considers an interval around the advertised value
- Runs a modified PATHATTACK, which cuts all paths that *might* not be longer than *p**
- Lower bound for alternative paths must be *shorter than upper bound* for *p**

Oracle-Enhanced Attacker

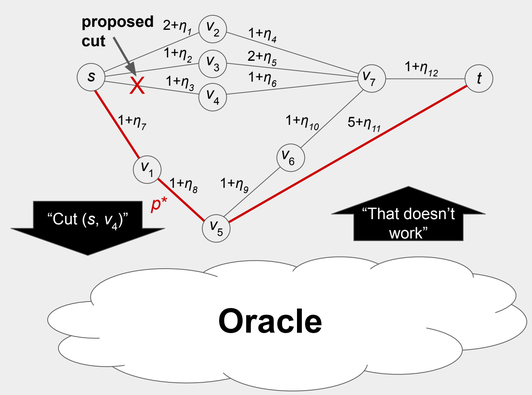

- Oracle: answers whether a proposed attack will be successful
- If not, use the standard PATHATTACK oracle to find the next shortest path (in expectation)
- Continue until oracle returns positive result

Results

- Graphs with weights drawn from a Poisson distribution
- Edge removal costs are equal to weights
- Randomly selected *s* and *t*
- 50th shortest path is *p**
- Noise distribution: uniform based on proportion of true weight
- *α*-informed attacker assumes true weight falls within *α* of expectation

- Can increase attacker's cost by over an order of magnitude
- Results vary based on topology
- Higher attacker cost and/or lower success probability come at expense of less reliable weights

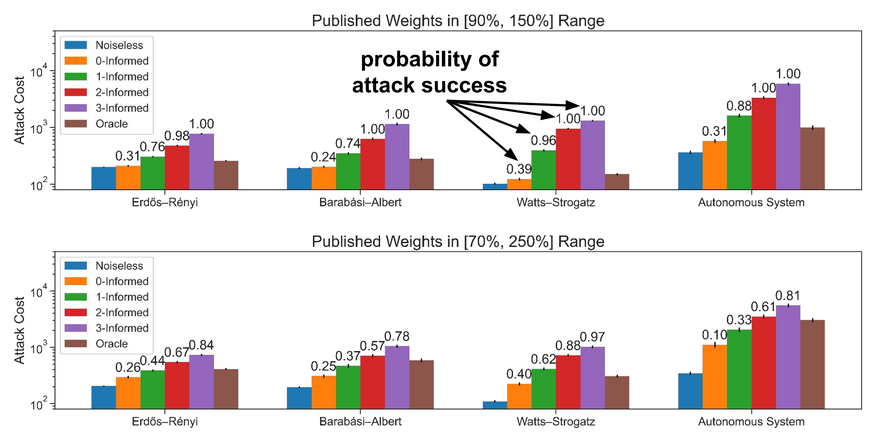

This material is based upon work supported by the United States Air Force under Air Force Contract No. FA8702-15-D-0001 and the Combat Capabilities Development Command Army Research Laboratory (under Cooperative Agreement Number W911NF-13-2-0045). Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the United States Air Force or Army Research Laboratory.