

Mitigating Exploits of the Current Interdomain Routing Infrastructure

Aaron D. Jaggard, Vijay Ramachandran, and Rebecca N. Wright

<http://dimacs.rutgers.edu/~adj/mitigating/>



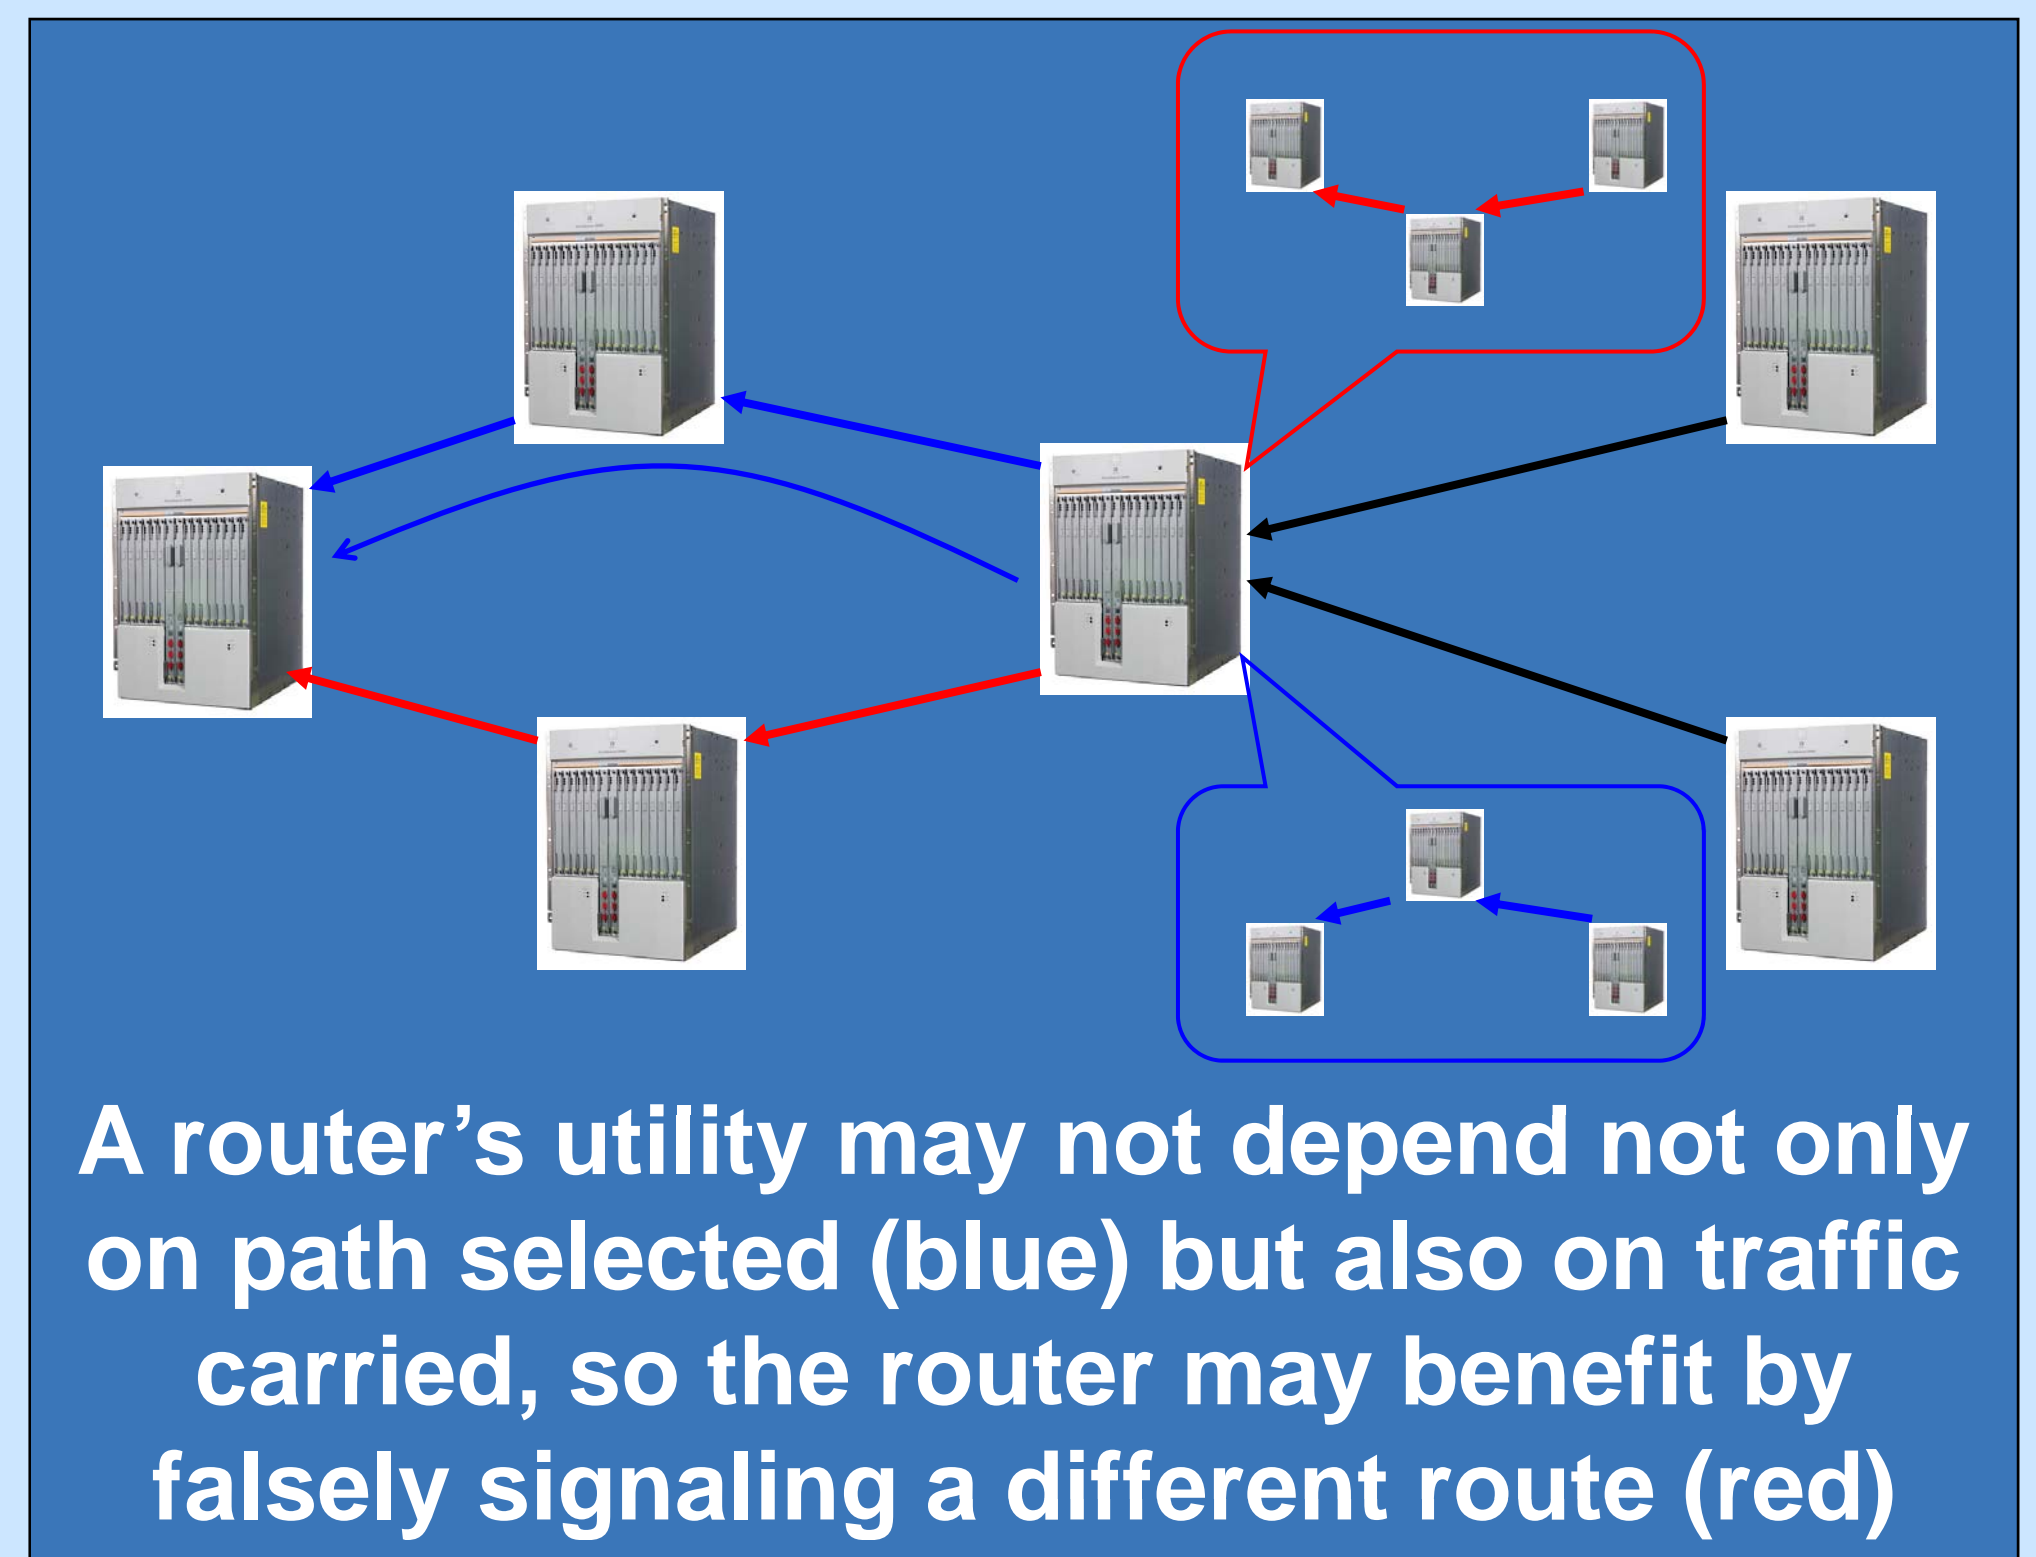
A more realistic model of incentives for routing

Routers share routes (*signaling*) and then use the best route they know to send data to a destination (*forwarding*). These actions are not coupled, allowing routers to falsely signal routes if this benefits them.

Approach

We formally model the decoupling of forwarding from signaling and show that natural sufficient conditions for good signaling behavior also guarantee acyclic forwarding. However, these do not guarantee *truthful* signaling.

We then explore conditions that guarantee routers do not have an incentive to lie in signaling and the role of other conditions that have proved important for related problems.



Approach and Impact

New approach

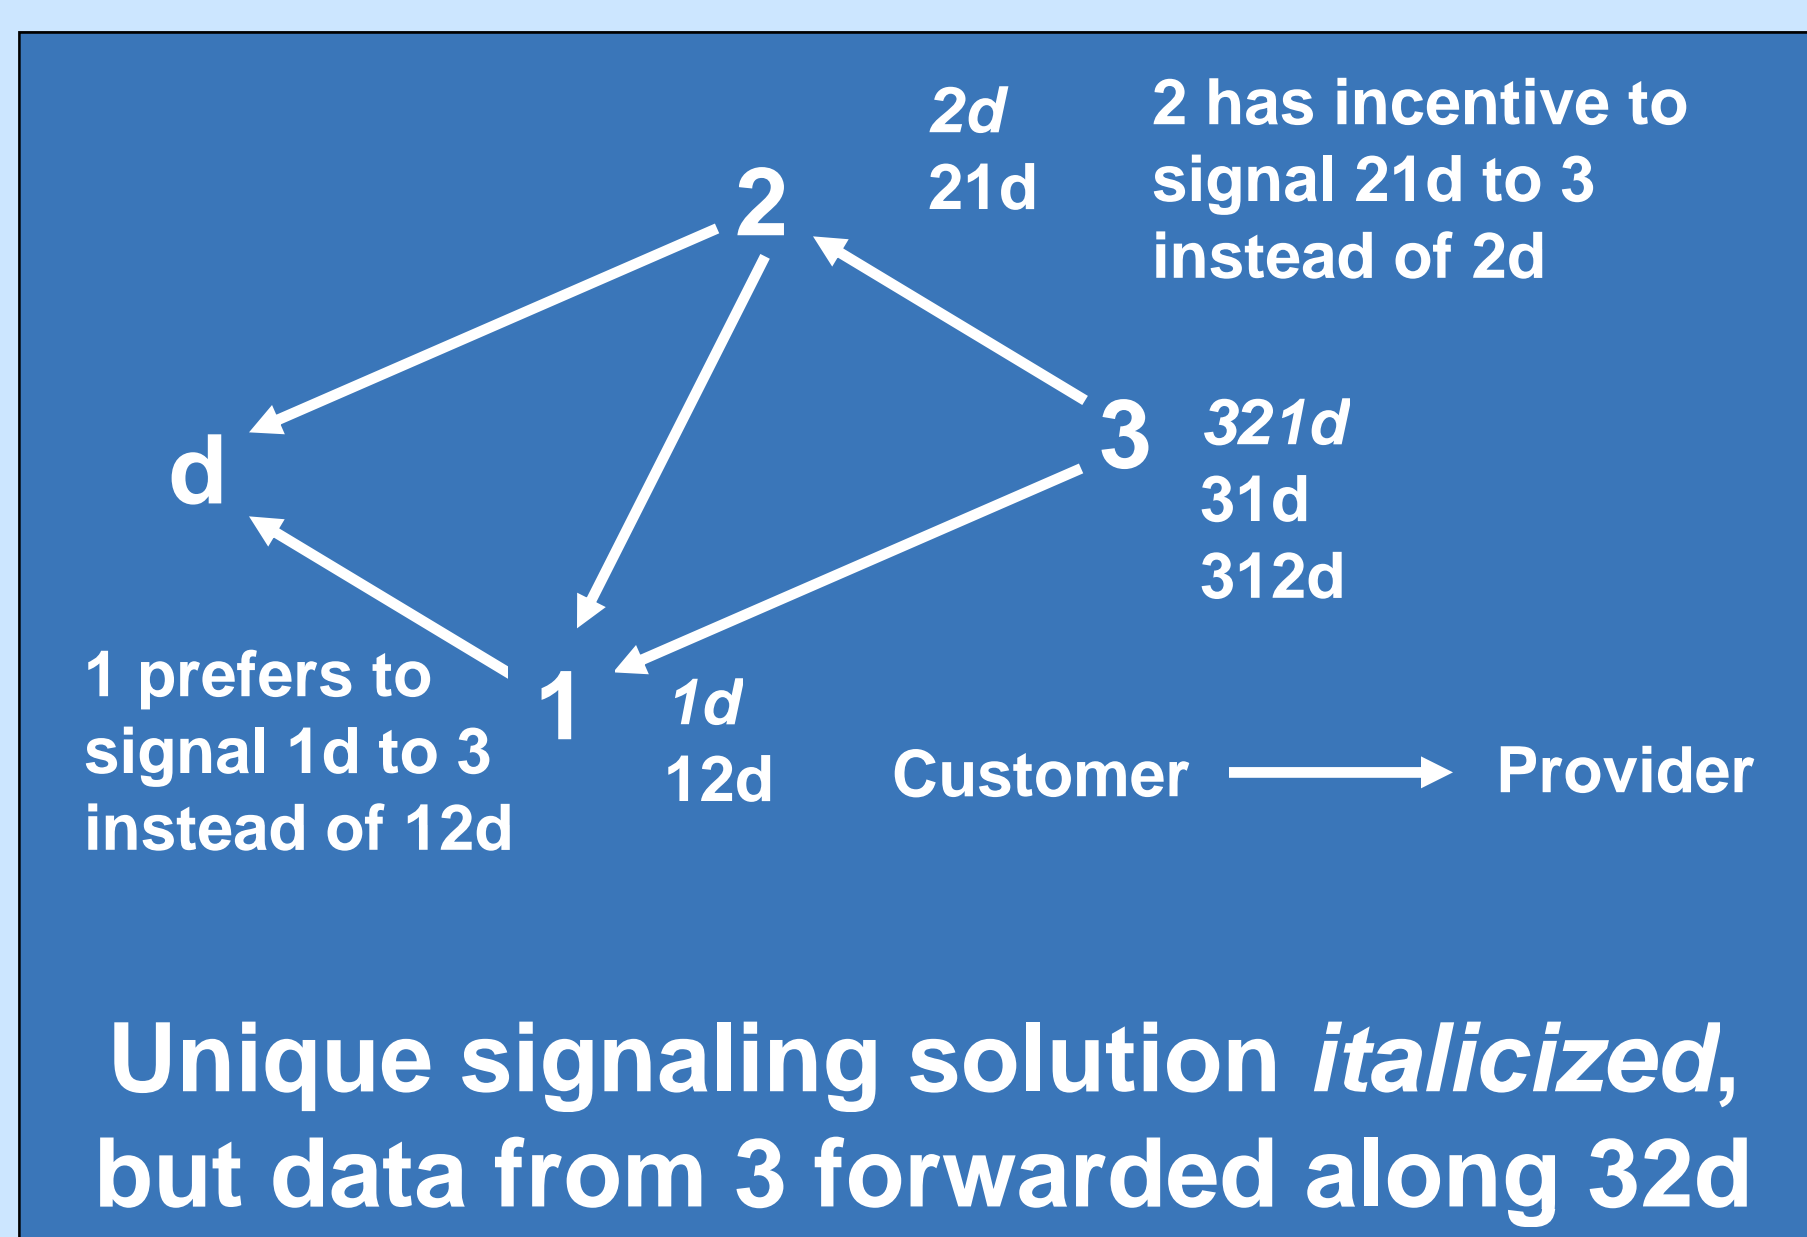
- Model real-world decoupling of forwarding from signaling
- Study utility functions that depend on signaling and not just forwarding

Research Impact

- Shifts static and dynamic models of routing to more realistic perspective
- Positive and negative results about dishonest signaling

Technical Description

- We extend the classic Stable Paths Problem (SPP) to a framework (FS-SPP) that decouples forwarding from signaling. Nodes (Autonomous Systems) have forwarding preferences and separate signaling preferences for each of their neighbors.
- We also extend a networking game so that each node's utility depends not only on the routes it selects for forwarding, but also (in a variety of ways) on the other nodes that route through it. We include the potential for filtering to be a non-strategic action (as in real-world routing) and study the effects of different filtering restrictions on this game.



Results

- Modified Gao-Rexford constraints guarantee stable signaling and acyclic forwarding, but not that these match
- In a network with only next-hop routing policies, routers cannot unilaterally act strategically to force the network to switch from one solution to a more preferred solution
- Examples sharpen boundary between incentive-compatible systems and those that aren't