Benefits of negotiated interdomain traffic engineering

Ratul Mahajan

David Wetherall

Tom Anderson

University of Washington

Problem

- Interdomain routing (BGP) decisions are based on very little information about other networks
  - poor performance
  - instability, oscillations
  - tedious, error-prone management

Example of poor performance

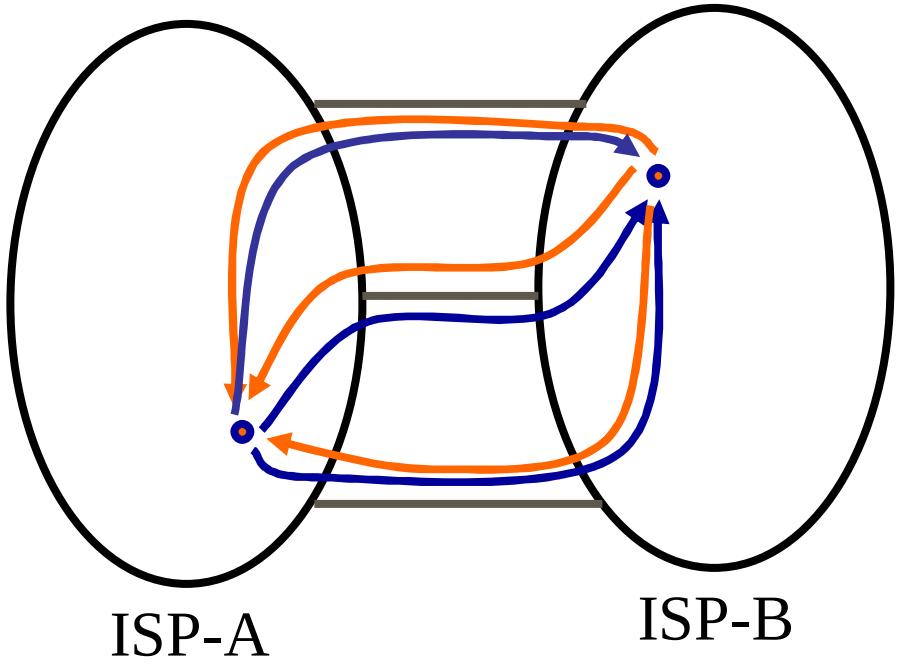

Example of oscillation

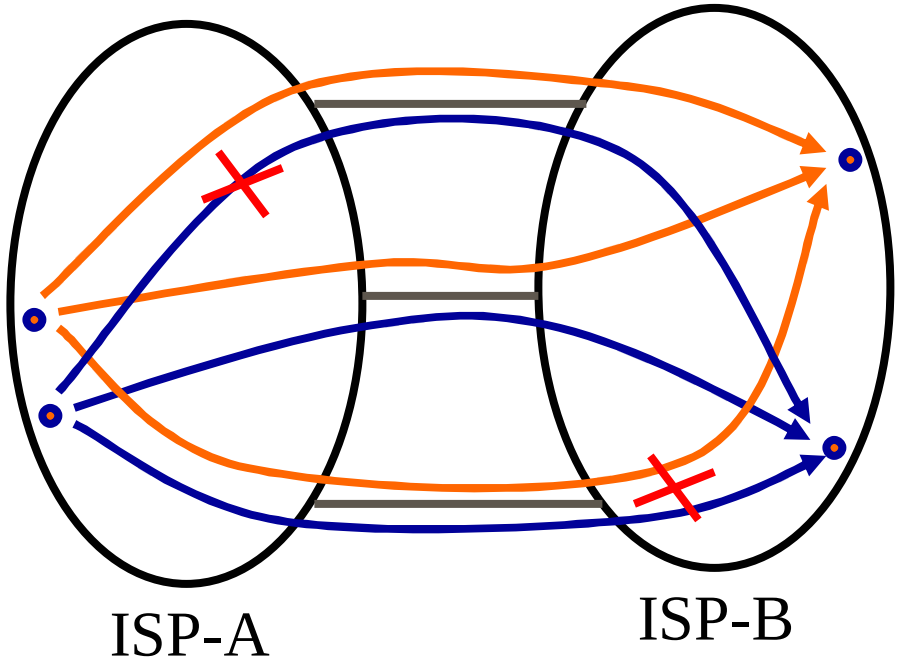

Current methodology

- Whenever interdomain routing changes need to be made
  - tweak-n-pray
  - call ahead
  - manually negotiate a mutually agreeable solution

Can we automate inter-ISP negotiation?

- Minimize manual firefighting
- As good or better than manual negotiation
- Honors real-world constraints
  1. controlled information disclosure
  2. independent of optimization criteria
  3. flexible outcomes

Current status

- Negotiation between two neighboring ISPs
- High-level methodology
- Evaluation of the potential benefit

Simplified negotiation methodology

1. Assign a numeric preference (like MEDs) to each routing option for each flow
   - opaque & independent of optimization criterion
2. Exchange preference lists
3. Take turns to propose routing options
   - find good compromises
   - reassign prefs if needed (load dependent)
4. Stop when one of the ISP wants to

Example of negotiation

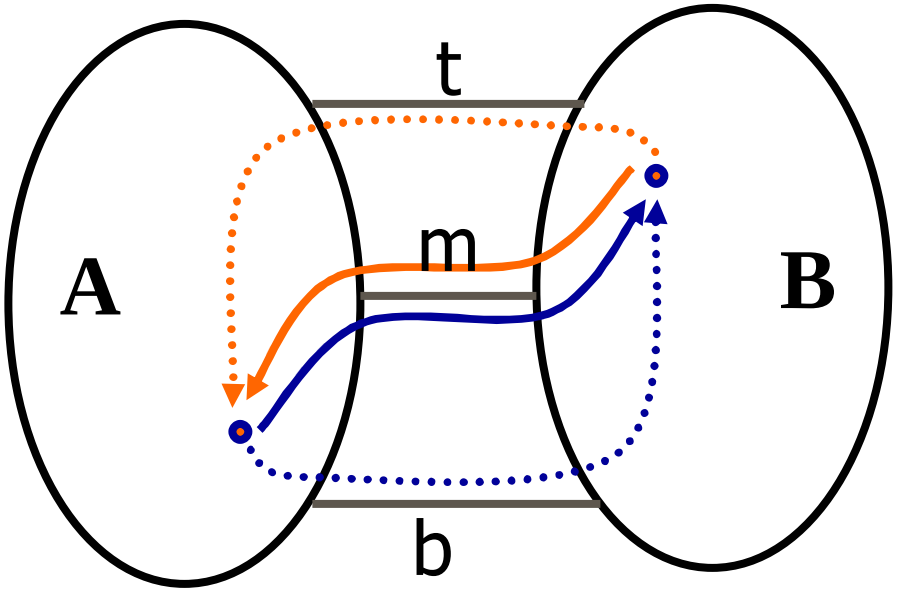

- A and B negotiate for 2 flows

	A→B	B→A
t	(-6, 6)	(0, 0)
m	(-1, 4)	(4, -1)
b	(0, 0)	(6, -6)

- Trade small sacrifices for bigger gains such that both ISPs win

Evaluation

- Compare three routing methodologies
  1. default: early-exit, selfish
  2. optimal: globally best across the two ISPs
  3. negotiated
- Dataset: 65 measured PoP-level ISP topologies; synthetic traffic models
- Evaluate latency reduction and hotspot avoidance

Experiment 1: Latency reduction

- Higher latency
  - ⇒ poorer performance
  - ⇒ more resource usage ⇒ costlier
- Measure latency of traffic when routed using the three routing mechanisms
  - default, optimal, negotiated

Results: Latency reduction

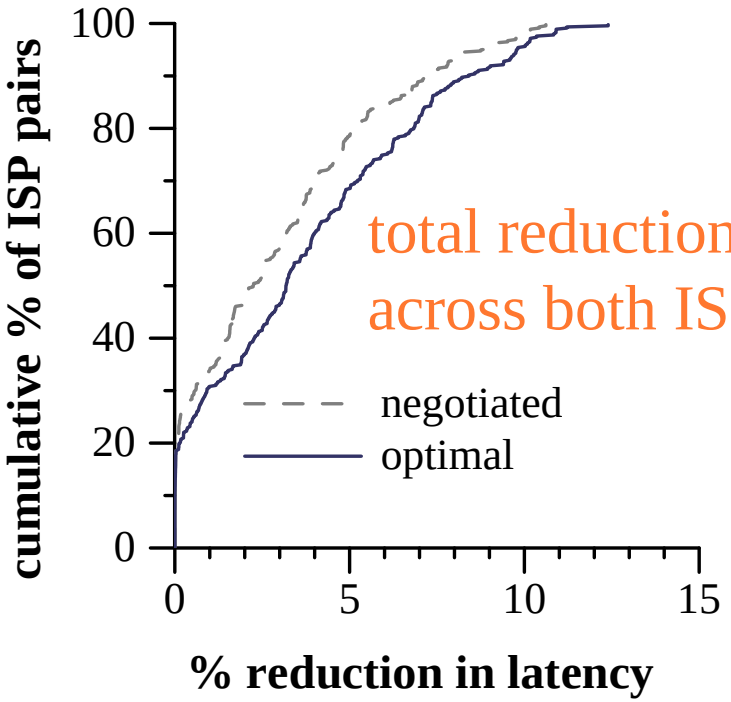

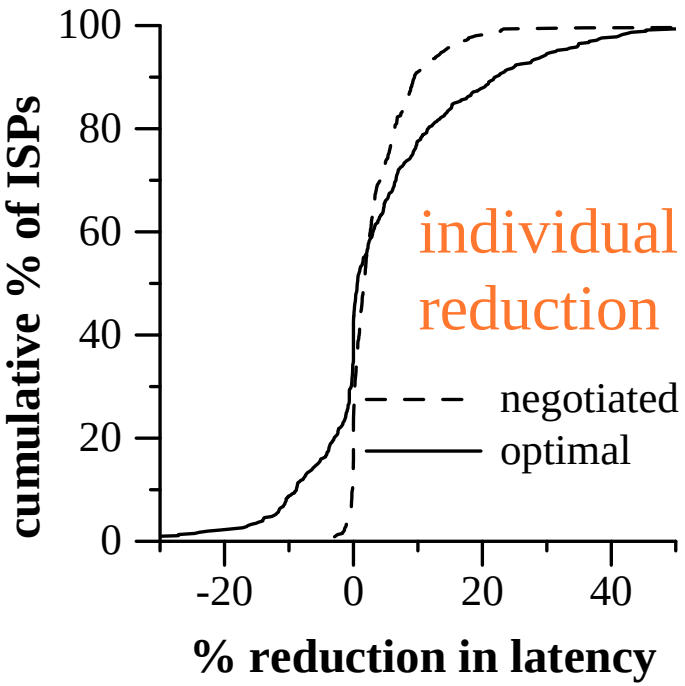

- Small aggregate latency reduction
  - some flows gain big
  - is this valuable?
- Individual ISPs can lose with the optimal
- Negotiation is win-win

Experiment 2: Hotspot avoidance

- Sudden changes (failures, DoS attacks) can cause short-term overload
  - fighting these is a major time sink

1. Assume that a peering link failed
2. Reroute flows traversing the failed link
3. Measure the potential for overload using multiplicative increase in link load

Results: Hotspot avoidance

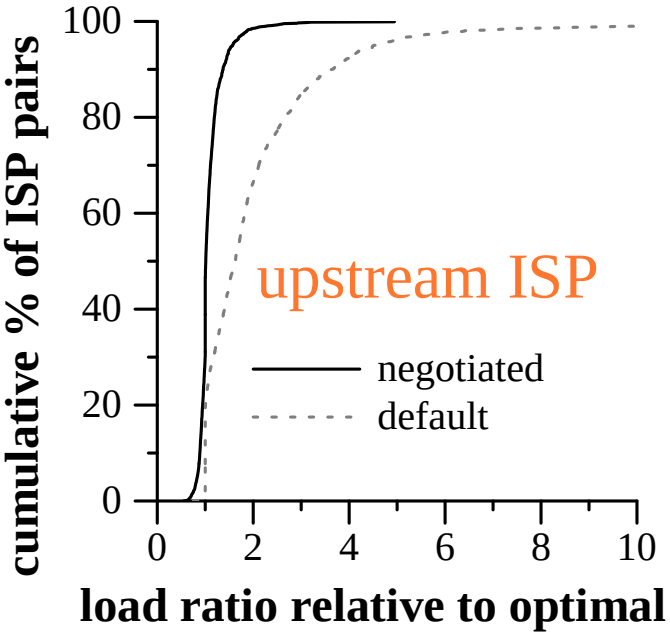

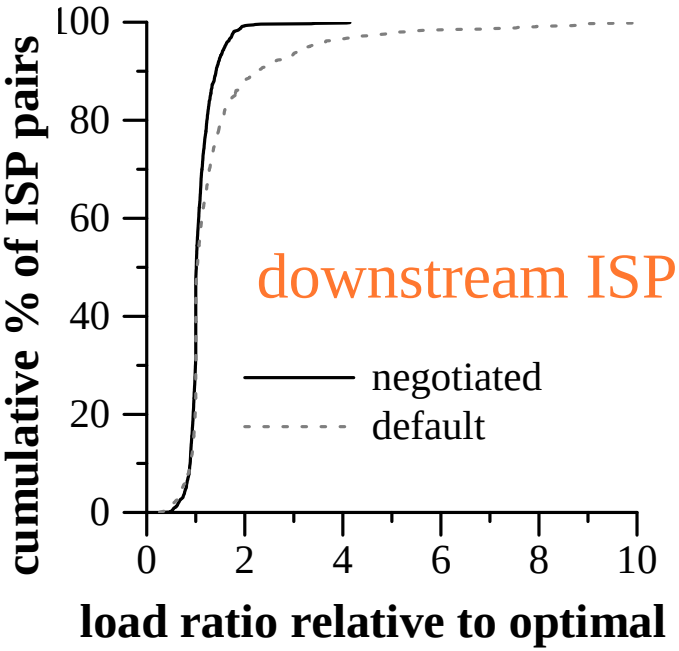

- Default routing tends to overload certain links
- Negotiation reduces the possibility of hotspots
  - fewer problems for the operators to resolve

Summary

- Interdomain routing decisions are based largely on local information
  - poor performance, instability
  - tedious, error-prone management
- Automated negotiation can help
- Feedback:
  - would you use it to talk to your neighbors?
  - http://www.cs.washington.edu/research/networking/negotiation