Towards Coordinated Interdomain Traffic Engineering

Ratul Mahajan
David Wetherall
Tom Anderson

University of Washington

ISPs need interdomain TE

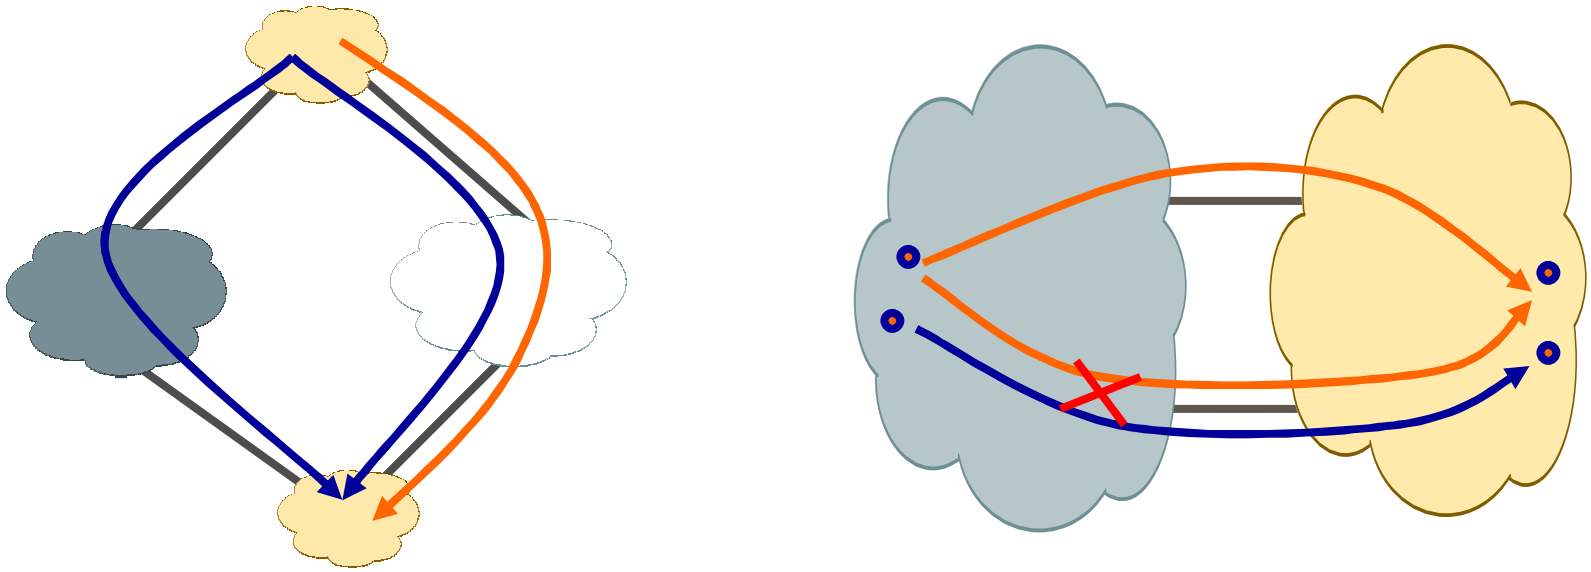

- Change how traffic enters or exits an ISP
  - to reduce resource usage, improve performance, balance load
  - deal with unforeseen events, e.g., failure and overload

Interdomain TE today

- The original design of BGP did not support TE
- A plethora of *post hoc* techniques
  - MEDs, communities, AS-path prepending, smart routing, prefix splitting, selective announcements, …
- Unilateral, uncoordinated actions but non-local impact
  - instabilities due to conflicting control
  - unpredictable traffic control

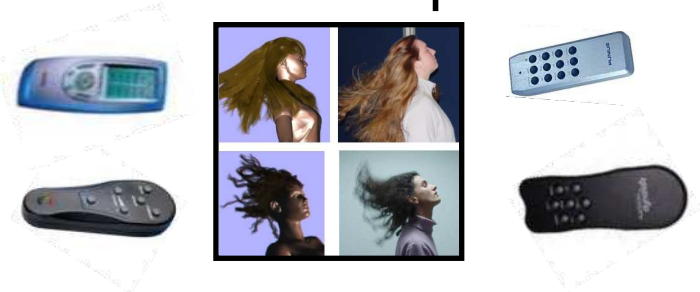

- Complex network operations, much manual intervention and coordination

Example of conflicting control

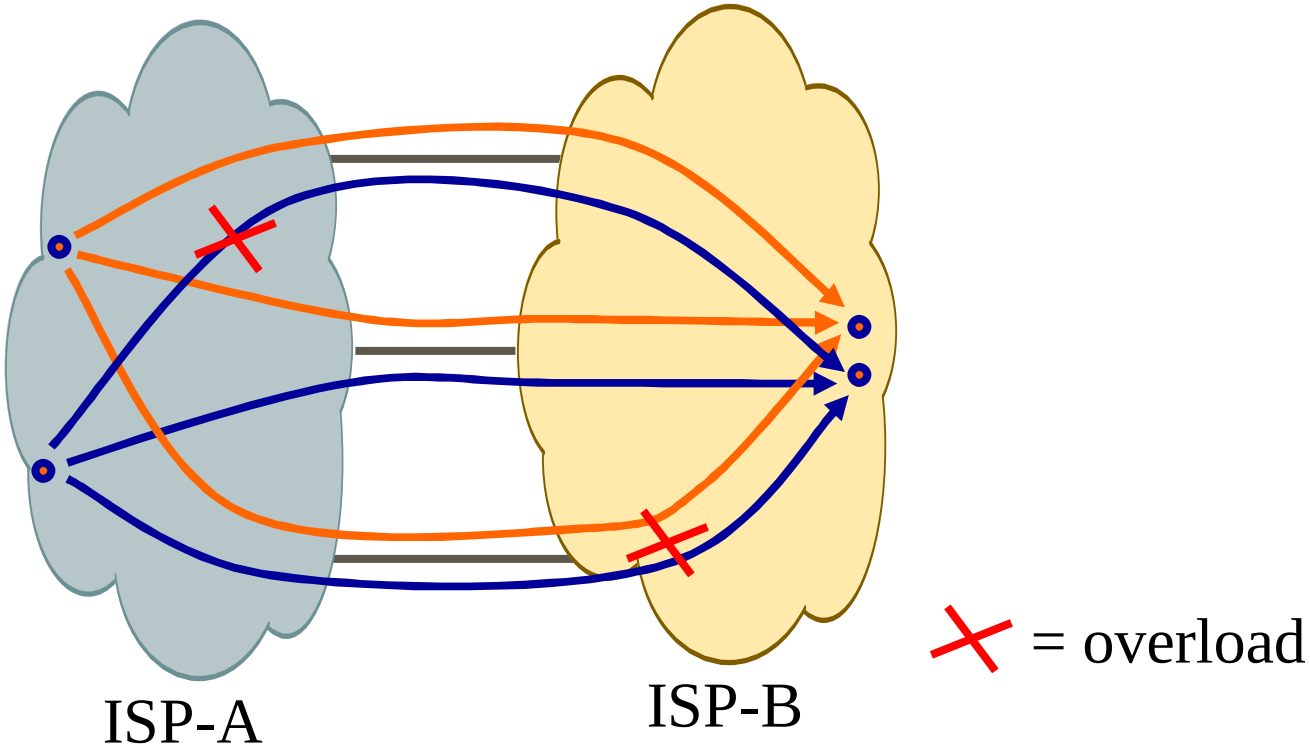

Example of unpredictable control

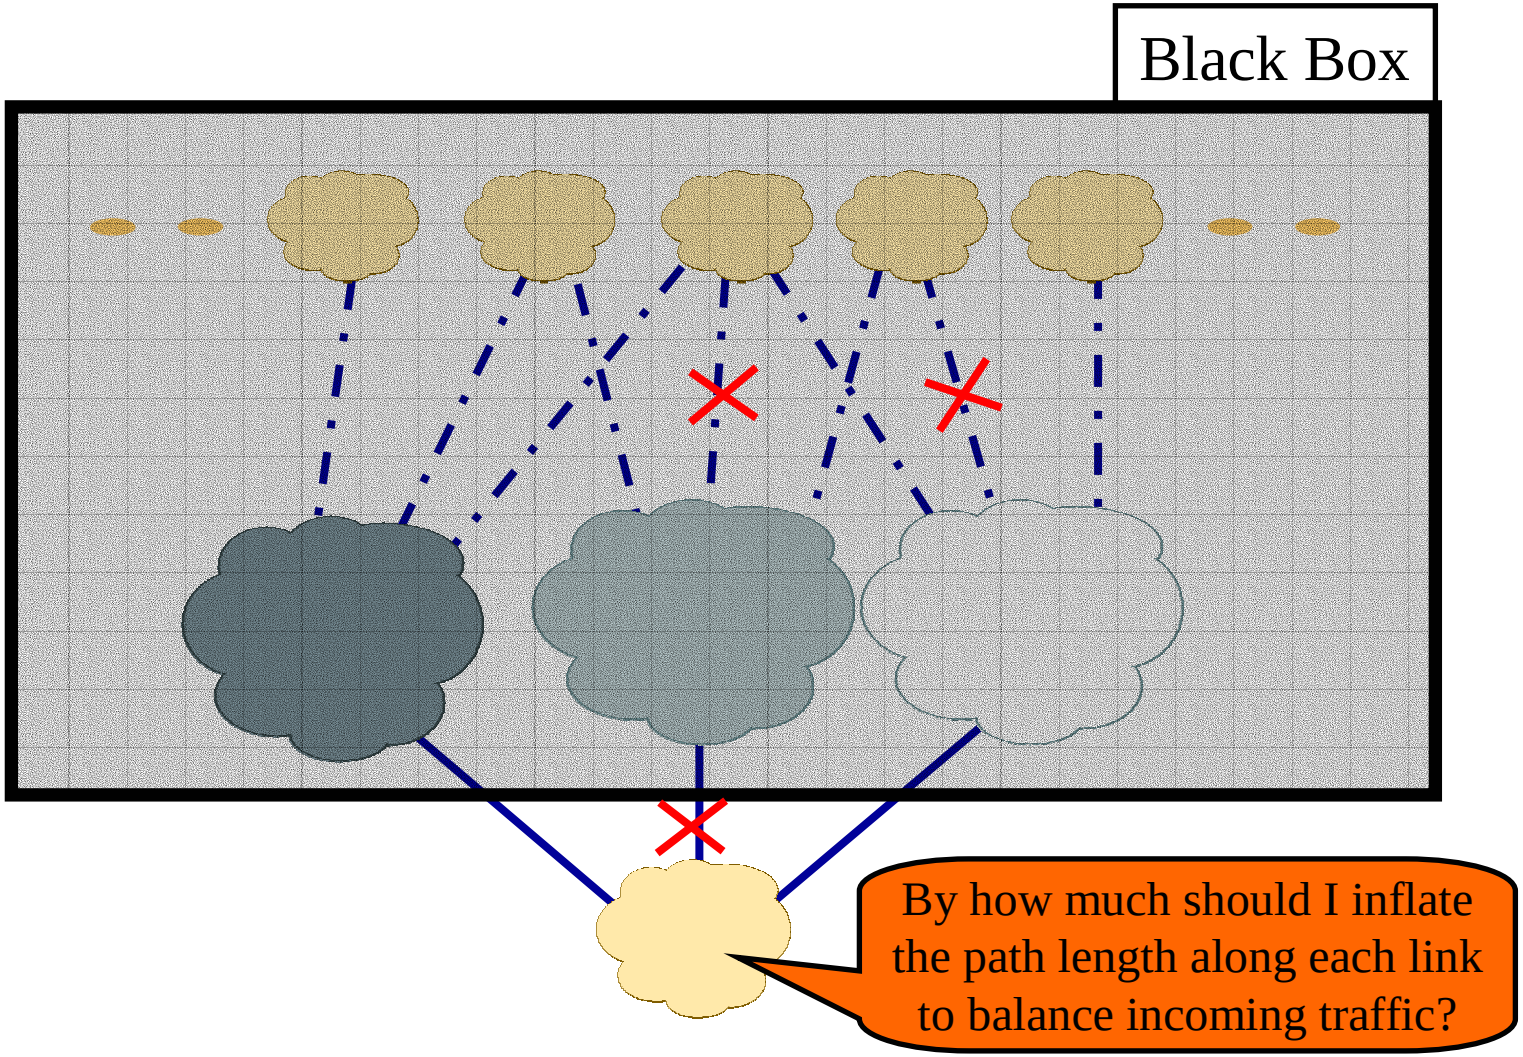

Coordinated TE

- Implement routing changes with the cooperation of impacted ISPs
  - precludes instability and unpredictability
- Properties of the coordination mechanism
  - provide predictable control over traffic
  - accommodate different optimization objectives
    - e.g., latency vs. utilization
  - accommodate different interests (“tussle”)
  - disclose little information

Towards a solution

- Essential building blocks of the coordination mechanism
  - two-way routing information exchange for predictable control
  - route negotiation to effectively accommodate different interests

(both are outside the confines of “BGP model” of routing)

- Coordination example: peering point selection

Two-way information exchange

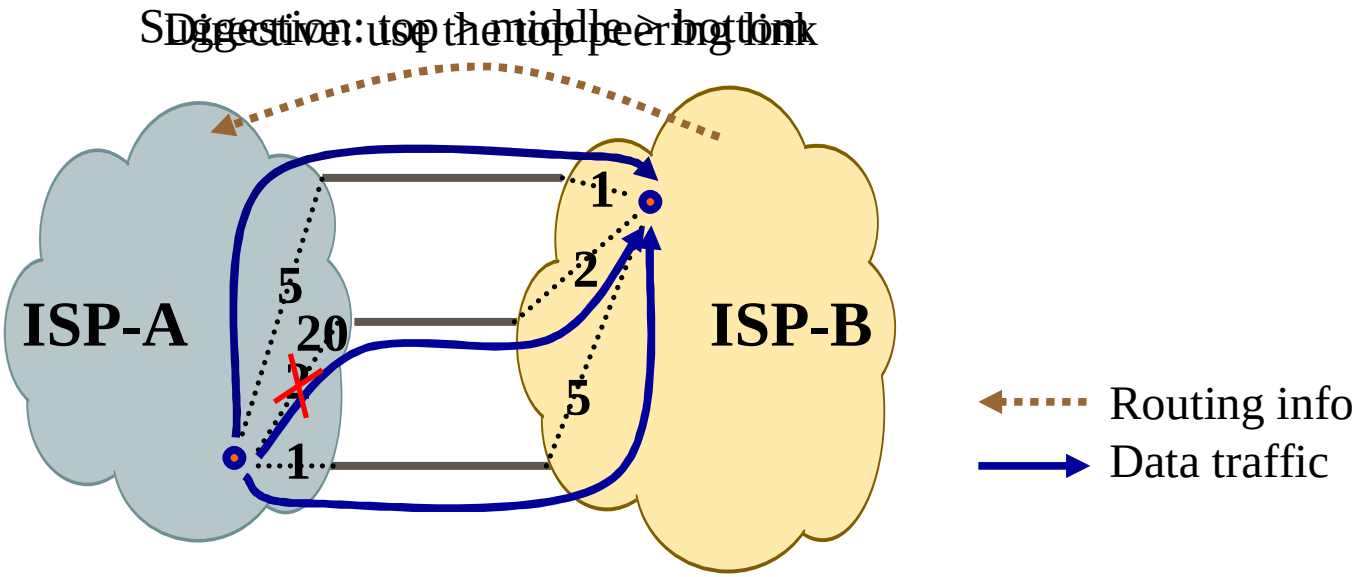

- Limitations of one-way routing info exchange
  - *directives* disable upstream control
  - *suggestions* have an unpredictable impact
- Two-way exchange is essential for predictable, joint control

Route negotiation

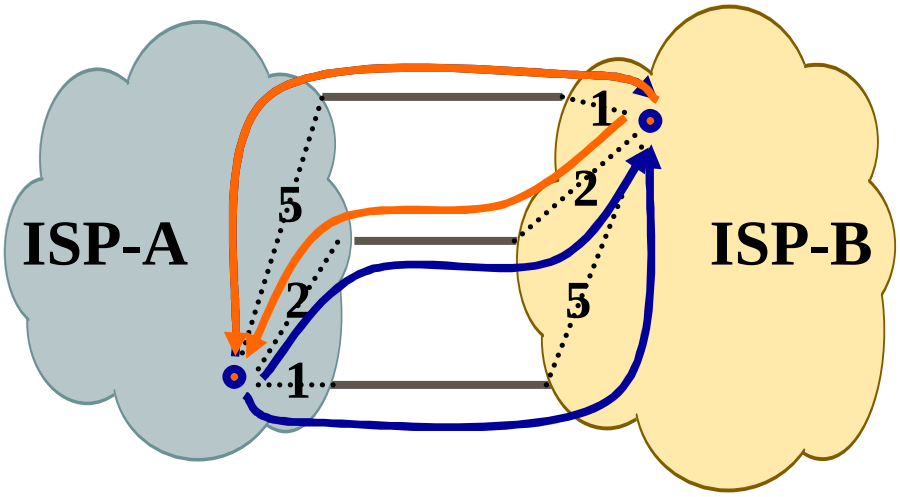

- Flow movement might present a conflict
  - one ISP loses and the other gains relative to unilateral routing
- Negotiate across flows and time
  - trade small losses for bigger gains
  - overall gain for all ISPs (“win-win”)

Coordination example

- Goal: select peering links between two ISPs
  - to improve the performance and stability of traffic exchanged between them
  - "base case" of the overall problem

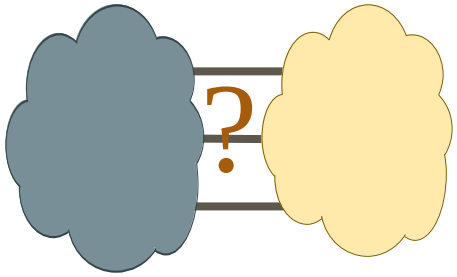

- High-level methodology:
  - ISPs share coarse preferences for receiving *and* sending each flow over each peering link
    - opaque (like MEDs); derived from respective optimization metrics
  - negotiate routing patterns that lead to mutual gain
    - take turns to propose better routing paths for flows

Coordination example (2)

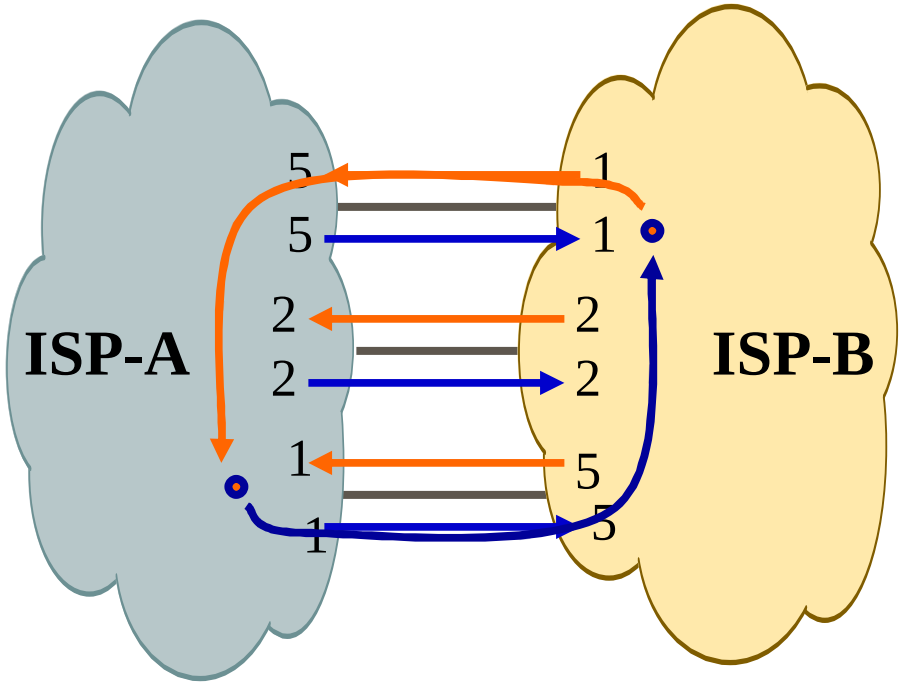

	A→B	B→A
t	(5, 1)	(5, 1)
m	(2, 2)	(2, 2)
b	(1, 5)	(1, 5)

4 [2]	4 [2]

- Predictable traffic paths
- Better paths for both ISPs

Evaluation

- Compare three routing methodologies
  - unilateral, globally optimal, negotiated
  - two metrics: latency reduction, avoiding overload after failures
- Using inferred topologies and synthetic traffic demands
- Results:
  - negotiated routing closely approximates the globally optimal
  - negotiation is win-win; globally optimal can be win-lose

Summary

- Internet needs a principled interdomain TE architecture
  - replace the collection of ad hoc, unilateral techniques
- ISP coordination is essential
  - predictable control over traffic I/O
  - prevents inadvertent resource policy violations
  - win-win solutions provide an incentive to negotiate
- Future work:
  - the nature of Internet-wide ISP negotiation
    - scalability and gaming issues
  - the role of negotiation in the overall commercial relationship