


# A simple coordination mechanism for interdomain routing

Ratul Mahajan\*  
David Wetherall  
Tom Anderson

*University of Washington*  
*(\*now @ Microsoft Research)*

# The nature of Internet routing today

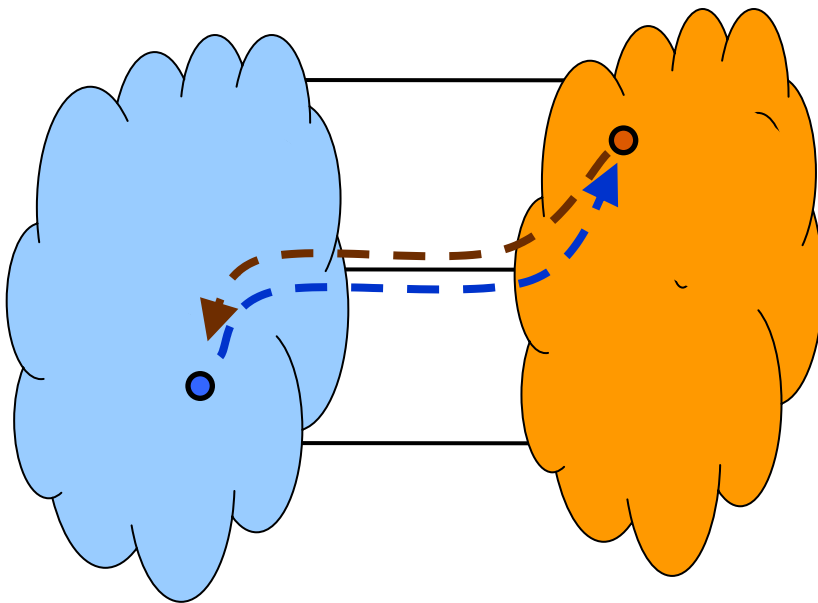


Within a contractual framework,  
ISPs select paths that are best  
for themselves

## Potential downsides

- q higher BW provisioning
- q requires manual tweaking to fix egregious problems
- q inefficient end-to-end paths

# An alternative approach: coordinated routing



ISPs have joint control

- q path selection is based on the preferences of both ISPs


Potential benefits

- q lower BW provisioning
- q no egregious cases that need manual tweaking
- q efficient end-to-end paths
  - basis for interdomain QoS

# Existing mechanisms cannot implement coordinated routing

Route optimization boxes help (stub) ISPs pick better routes from those available

MEDs implement receiver's preferences



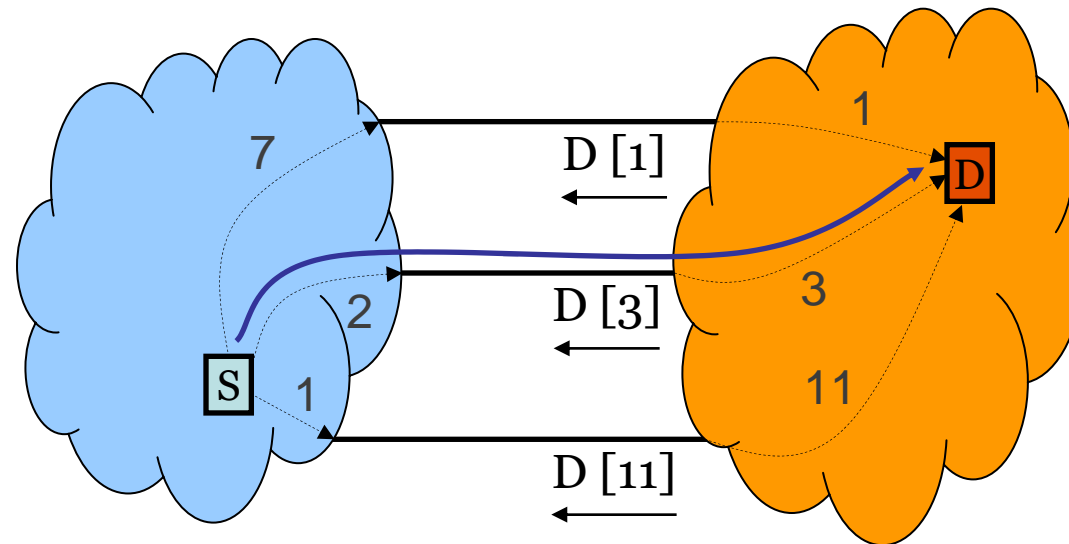
# What makes for a good coordination mechanism?

MEDs have some nice properties

- q ISPs can express their own metrics
- q ISPs are not required to disclose sensitive info.
- q lightweight
- q requires only pairwise contracts

Provides joint control and benefits all ISPs

# Our solution: *Wiser*



Operates in a lowest-cost routing framework

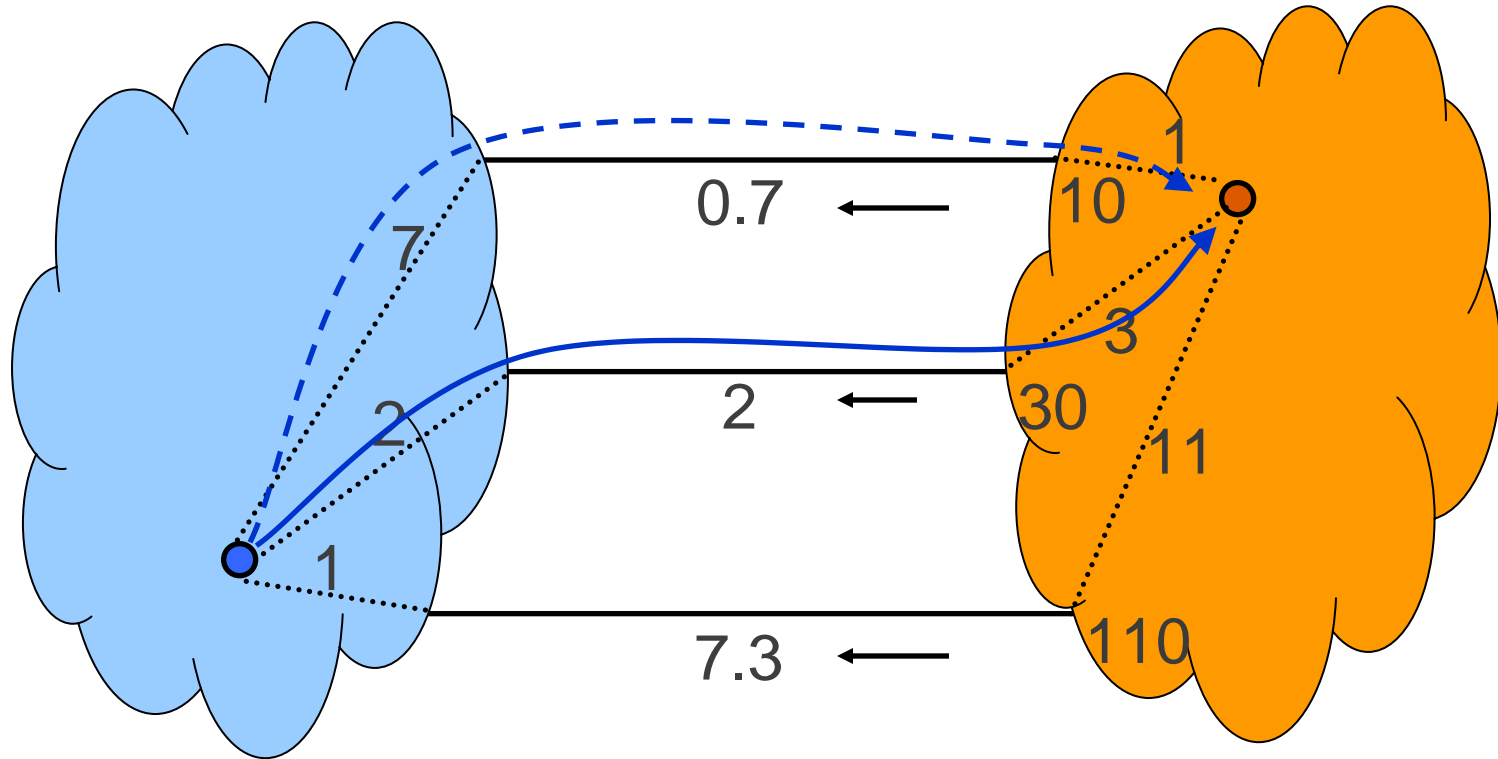
- q downstream ISPs advertise their cost
- q upstream ISPs select paths based on both their own and received costs

# Problems with vanilla lowest-cost routing

ISP costs are incomparable

Can be easily gamed

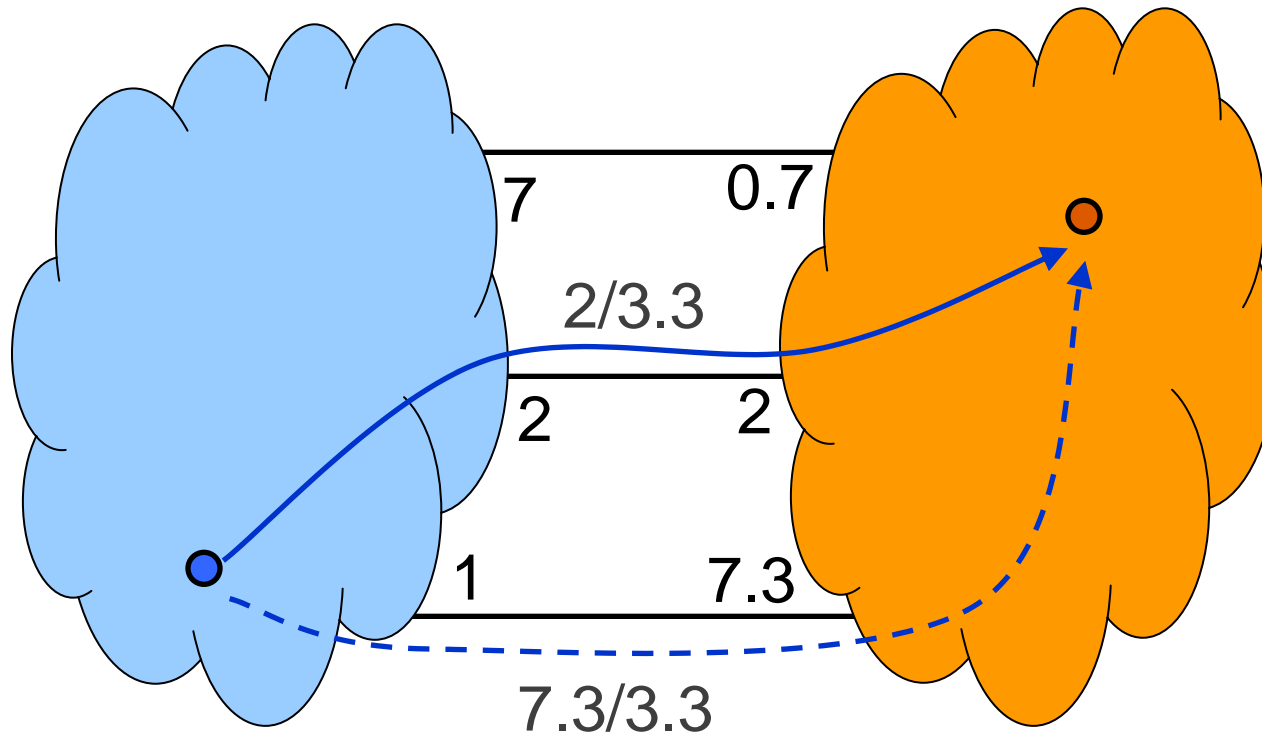
# Cost normalization



Normalize costs such that both ISPs have “equal say”



# Bounds on cost usage

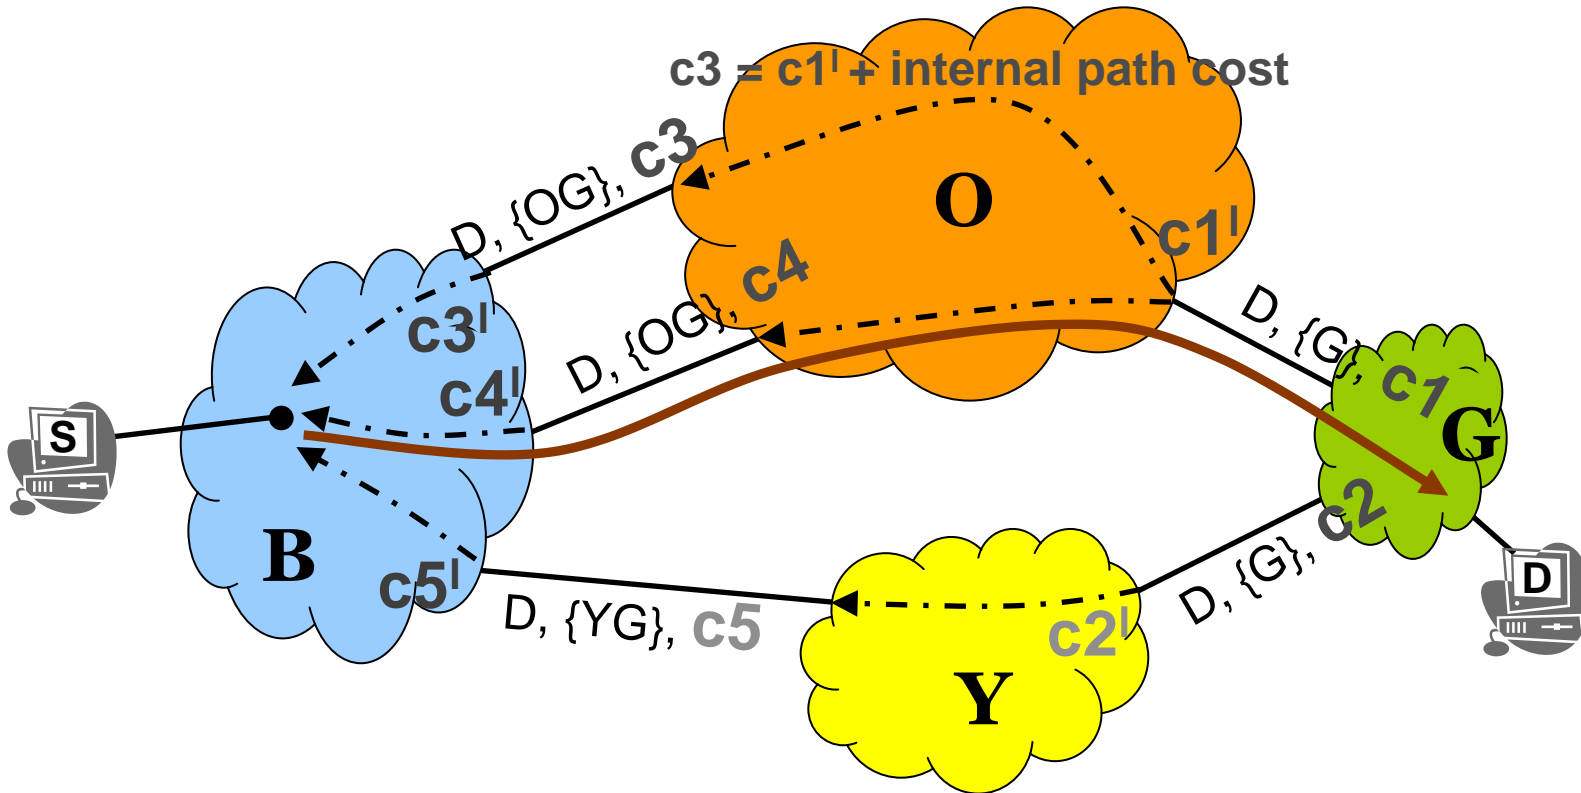


Downstreams log cost usage of the upstream ISPs

Compute the ratio of avg. cost of paths used and announced

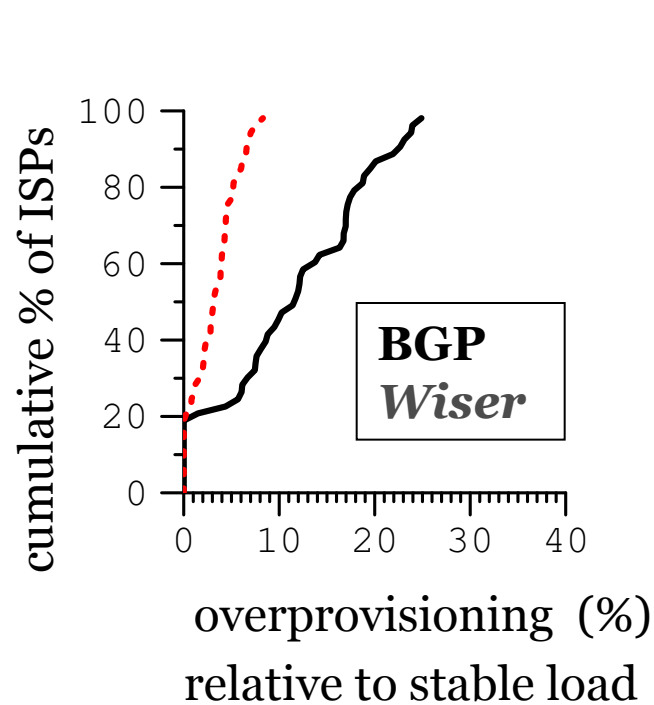
Contractually stipulate a bound on the ratio

# Wiser in action

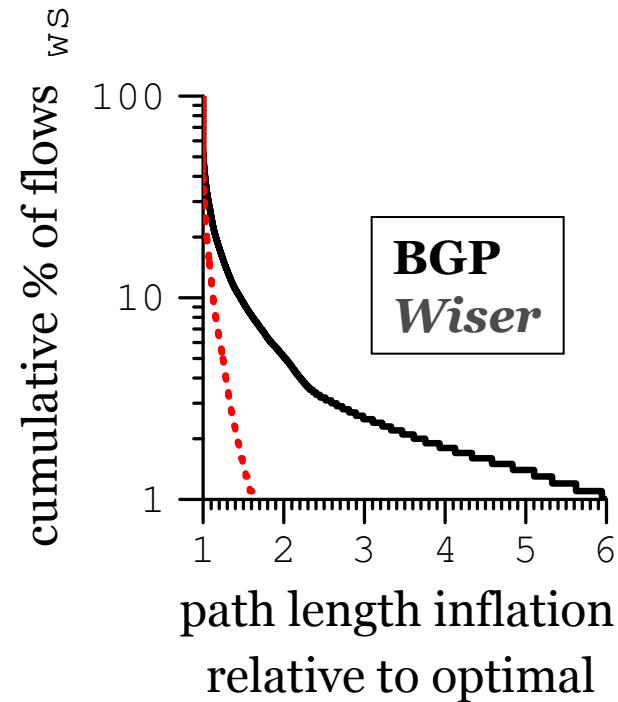


Converting all internal costs within a node into link costs

# Example results



*Wiser* requires lower bandwidth provisioning



*Wiser* produces shorter paths  
Significant benefit in the tail

# Implementation

XORP prototype

Simple, backward-compatible extensions to BGP

- embed costs in non-transitive BGP communities
- border routers jointly compute normalization factors and log cost usage
- a slightly modified BGP decision process

Benefits even the first two ISPs that deploy it

# Summary

*Wiser* is a simple mechanism to coordinate interdomain routing

- may lower provisioning, reduce manual tweaking, produce efficient paths, and help with interdomain QoS

Feedback: <[ratul@cs.washington.edu](mailto:ratul@cs.washington.edu)>

Details/code:

<http://www.cs.washington.edu/research/networking/negotiation/>