Stable Internet Route Selection

Brighten Godfrey
Matthew Caesar
Ian Haken
Scott Shenker
Ion Stoica
UC Berkeley

pbg@cs.berkeley.edu

NANOG 40
June 6, 2007

BGP instability: trouble

control plane	CPU cycles update processing uses majority of cycles on some core routers
data plane	degraded path quality BGP causes majority of packet loss bursts

Stable Route Selection:
simple technique to
significantly improve stability

What about Route Flap Damping?

- Introduces pathologies
- Impacts availability

  "...the application of flap damping in ISP networks is NOT recommended."

  --RIPE Route Working Group, May 2006

- Only helps for very unstable routes

Stable Route Selection

Given a choice between routes, select routes that are less likely to fail.

RFD philosophy	shut off bad routes
SRS philosophy	always pick a route if possible, but prefer more stable routes

Challenges

- Inferring stability of paths, locally
- Dependence: does one ISP's benefit require others' participation?
- Flexibility required

Outline

- Design
- Evaluation
  - Improvement in stability
  - Dependence
  - Flexibility
- Conclusion

Design

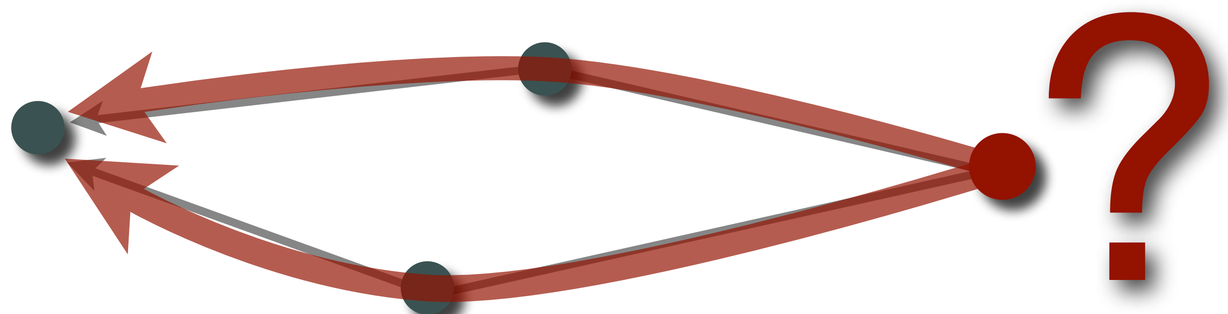

BGP decision process

1. Highest local pref
2. ~~Shortest path length~~
3. Lowest origin type
4. Lowest MED
5. eBGP- over iBGP-learned
6. Lowest IGP cost
7. Lowest router ID

SRS heuristic

2. Current route
3. Shortest path length
4. Longest uptime

Simplified processes

- Simulator has one router per AS, at most one link between AS's
- So...

Standard BGP	SRS
1. Highest local pref	1. Highest local pref
2. Shortest path length	2. Current route
3. Lowest router ID	3. Shortest path length
	4. Longest uptime

Outline

- Design
- Evaluation
    - Improvement in stability
    - Dependence
    - Flexibility
- Conclusion

Evaluation methodology

- Event-based BGP simulator

Topology	Internet AS-level
Local prefs	cust./prov./peer
AS-adjacency failures	inferred from RouteViews

- Measuring interruptions: route changes/ withdrawals

The bottom line

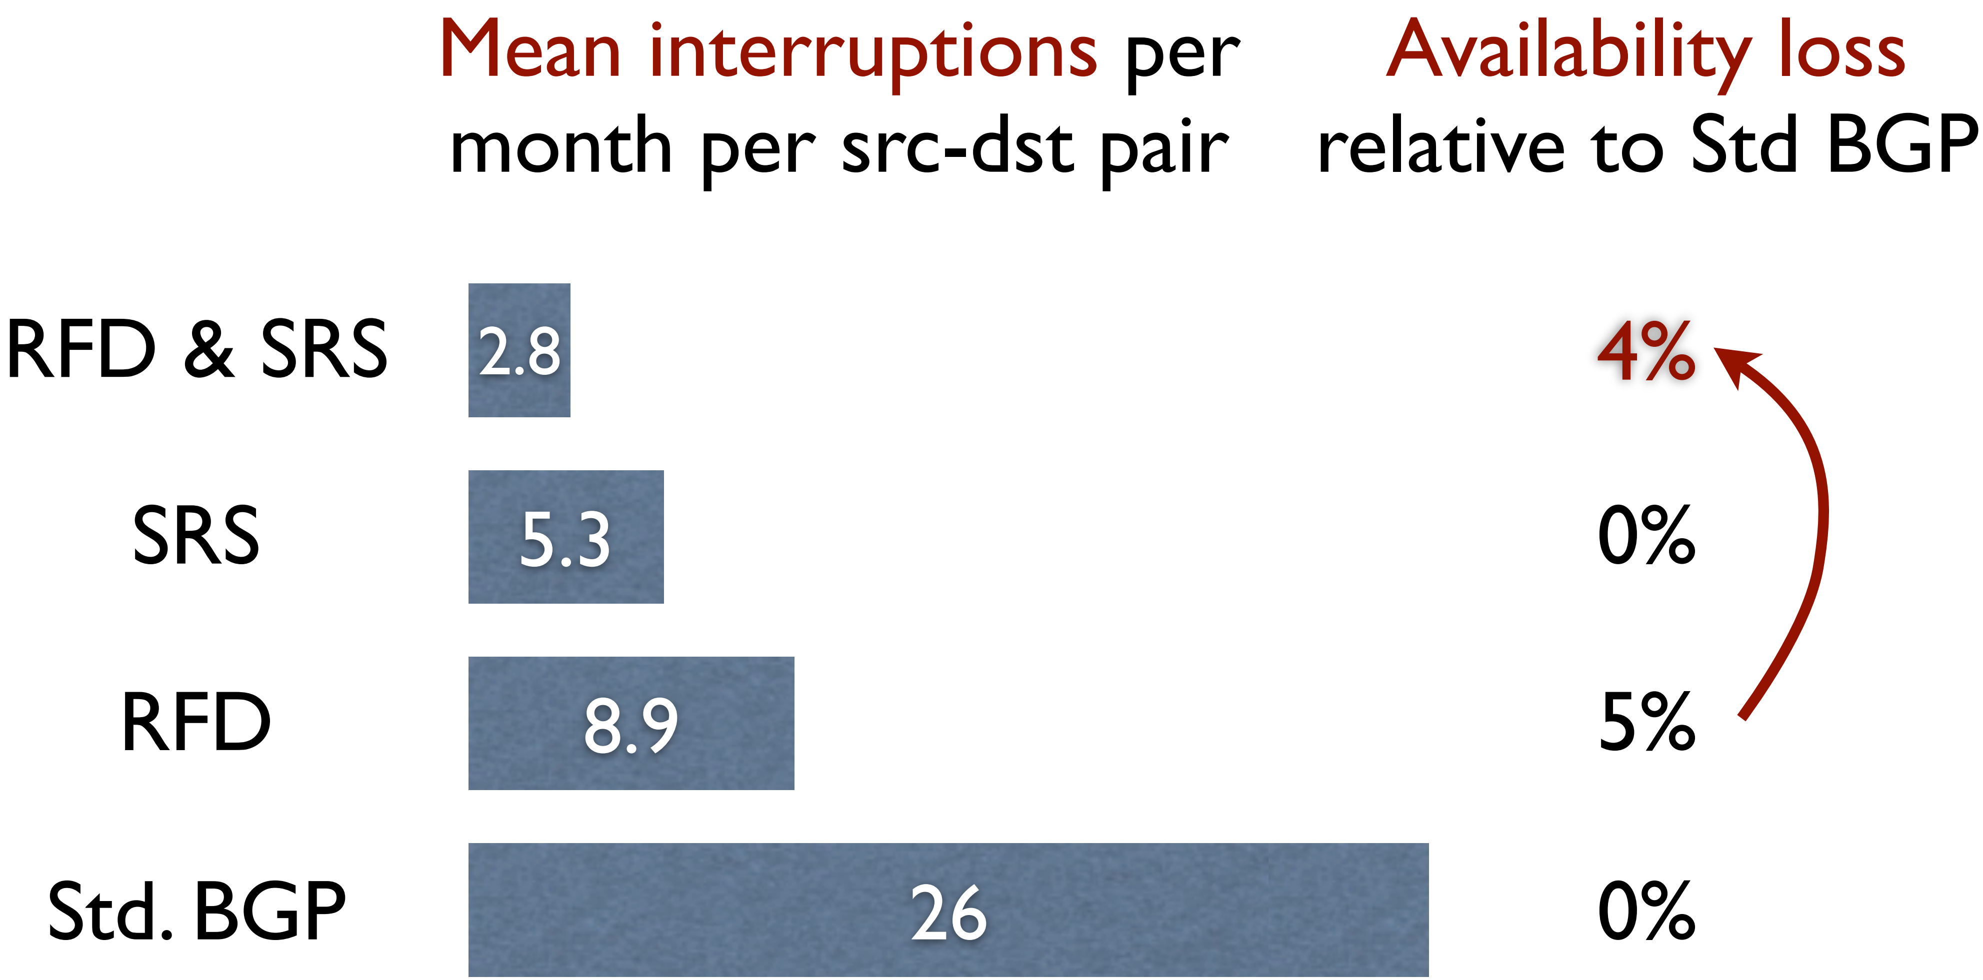

Dependence between ISPs

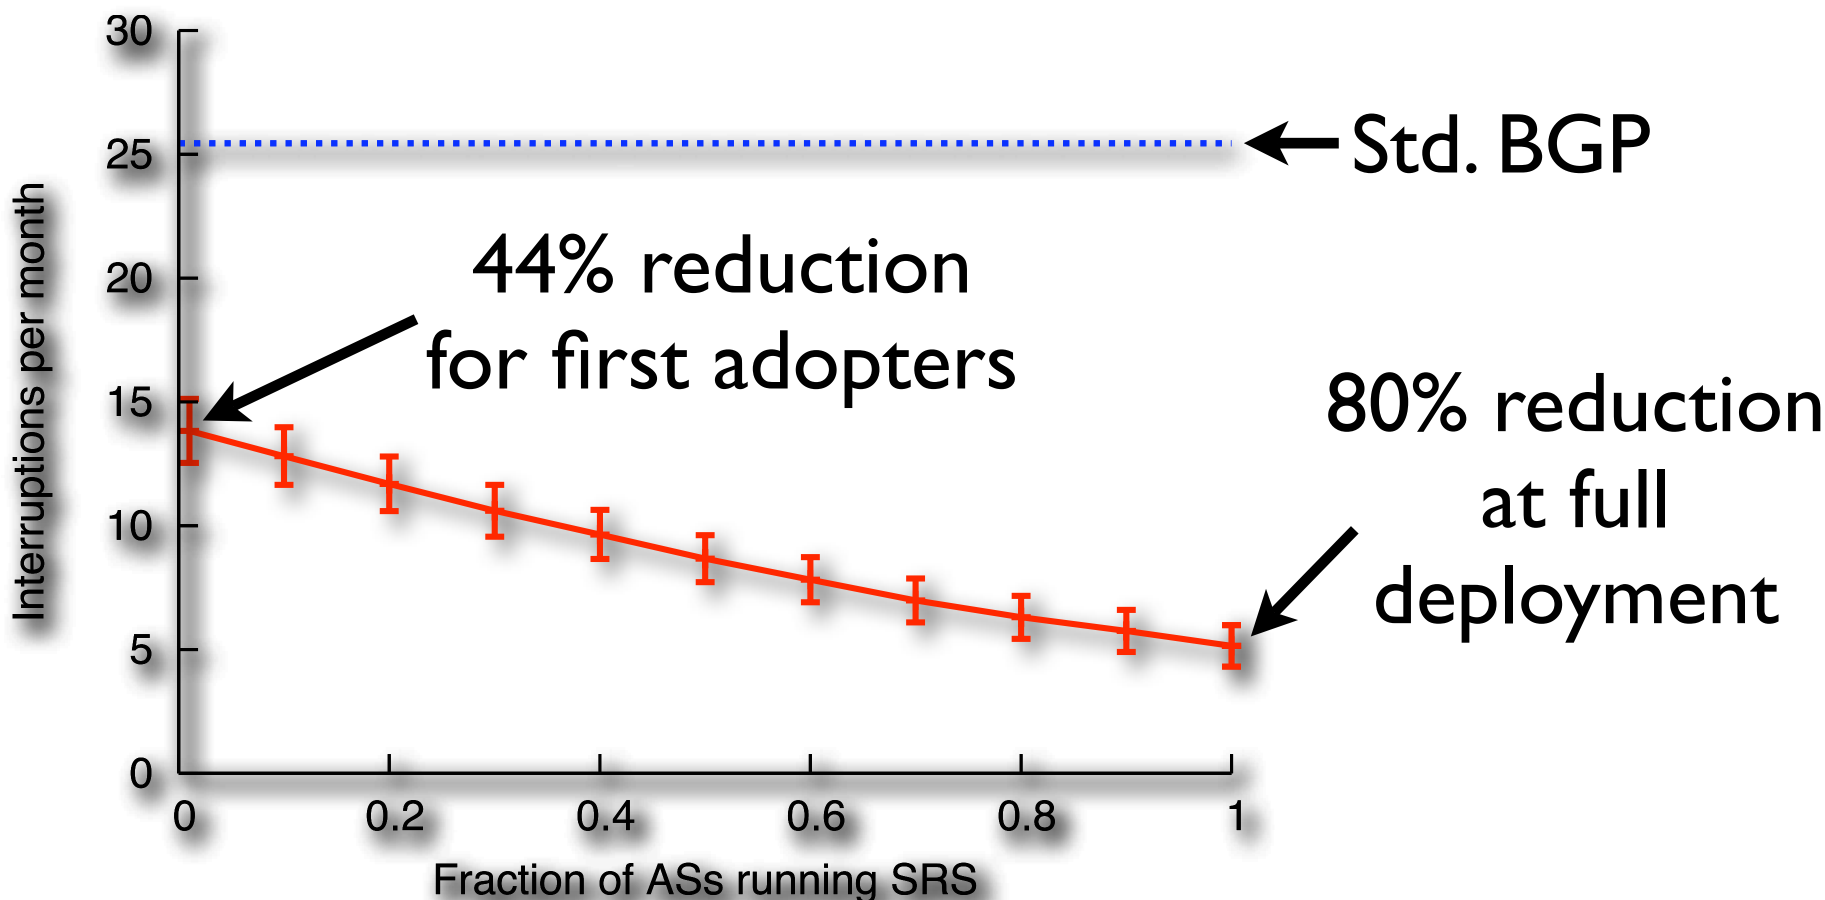

Dependence between ISPs

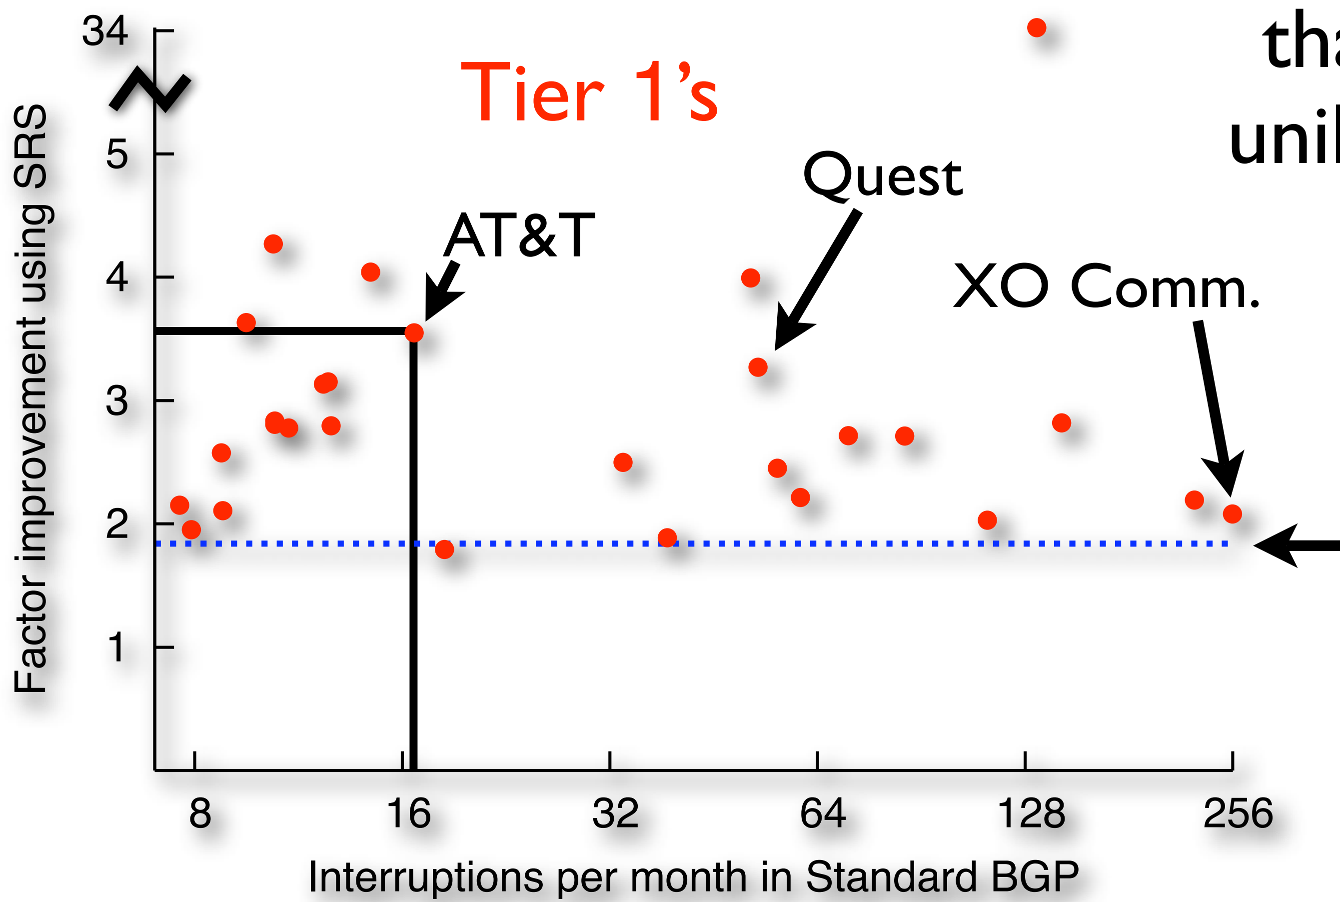

Tier 1s benefit more than other ISPs when unilaterally running SRS

Route flexibility

Flexibility also needed for

- load balancing
- business relationships
- path length

What is the tradeoff with these other objectives?

How much flexibility?

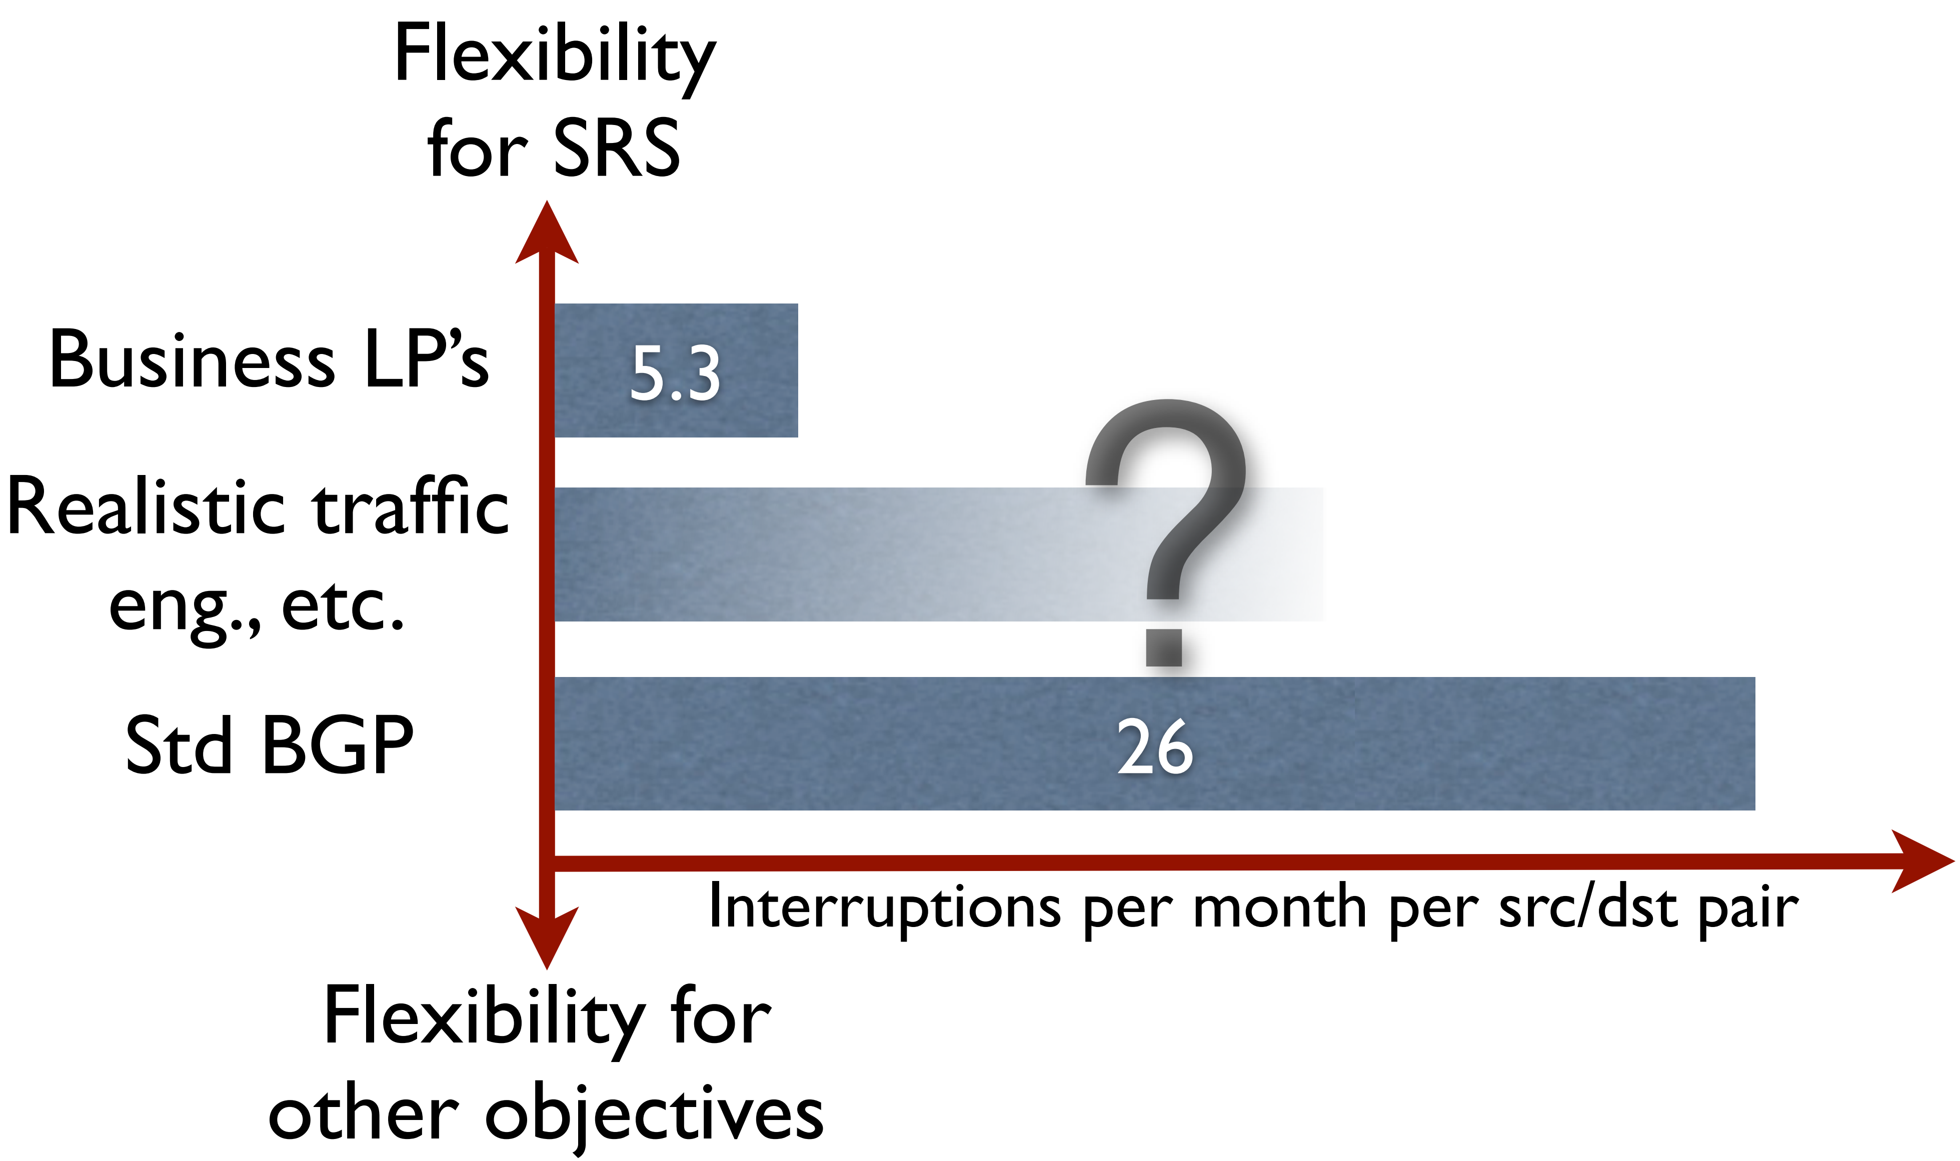

Tradeoffs: path length

SRS only
4% longer!

(Hypothetical "Longest paths" strategy:
32% longer)

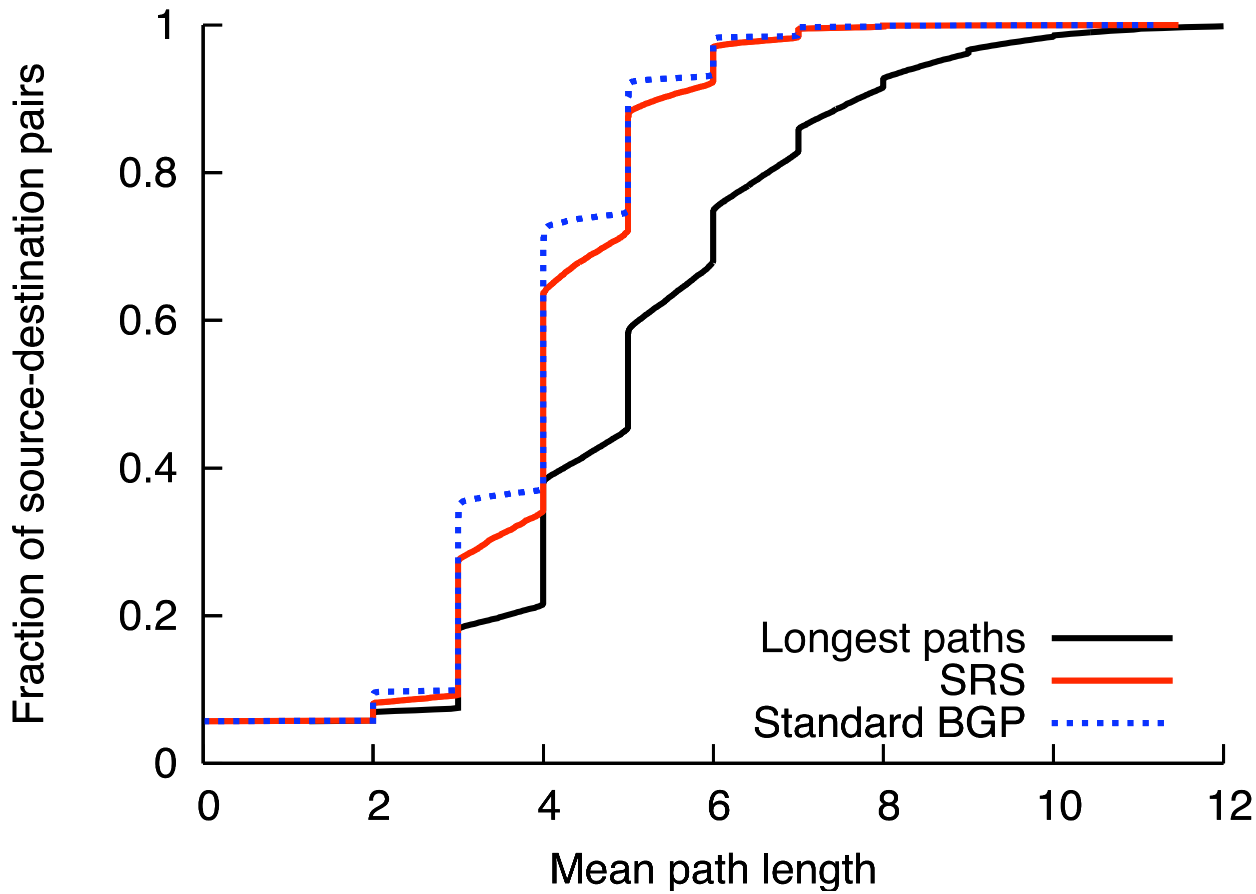

Outline

- Design
- Evaluation
  - Improvement in stability
  - Dependence
  - Flexibility
- Conclusion

Summary

- Stable Route Selection: use flexibility in path selection to optimize for stability
  - Significantly more stable
  - No impact on availability
  - Very low stretch
- Ongoing work: implementation & deployment

Questions for operators

1. How useful is stability?
2. What are the barriers?
   (nondeterminism, traffic engineering...)
3. How much flexibility would
   be available to SRS?

Very interested in feedback and collaborations

pbg@cs.berkeley.edu

Backup slides

AS adjacency mean session time distribution

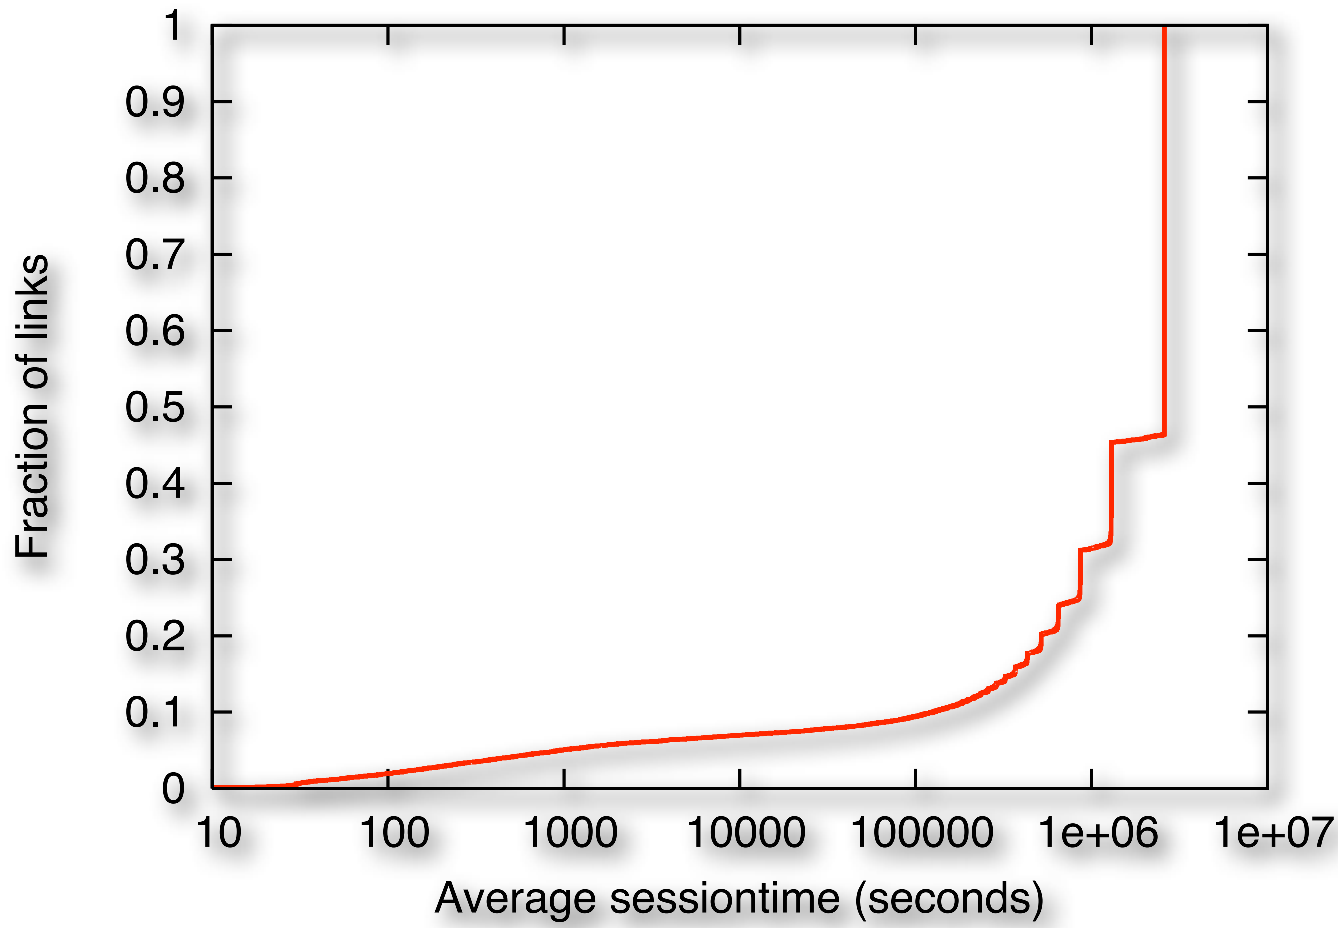

Attribution of improvement

instability = interruptions per event x # events

speed convergence

~8% better

avoid failures

Majority of the improvement

SRS vs. flap damping

SRS: better mean improvement

RFD: better for worst ~0.5% of src-dst pairs

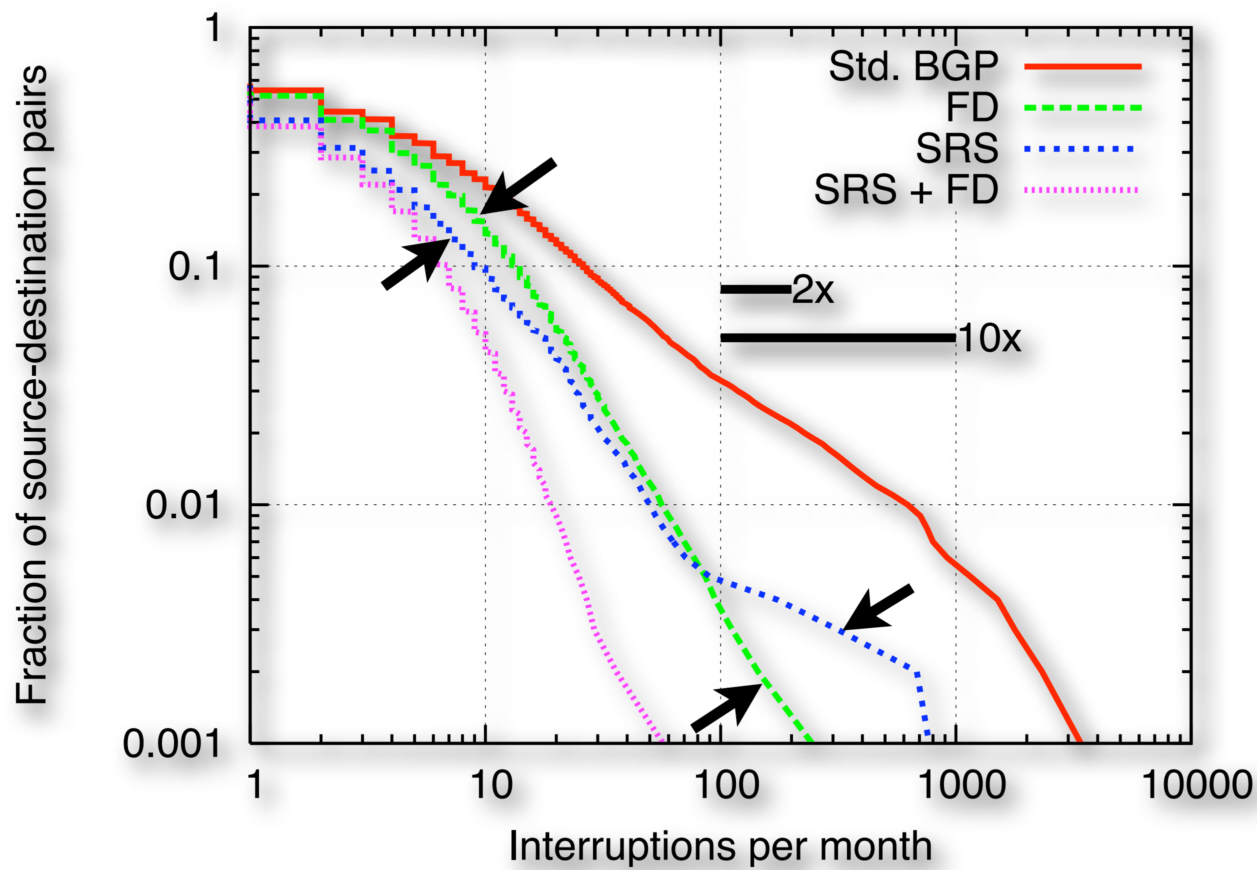

SRS vs. flap damping

SRS is more conservative	always pick a route if one is advertised
SRS is more aggressive	use any flexibility available for stability

SRS with less flexibility

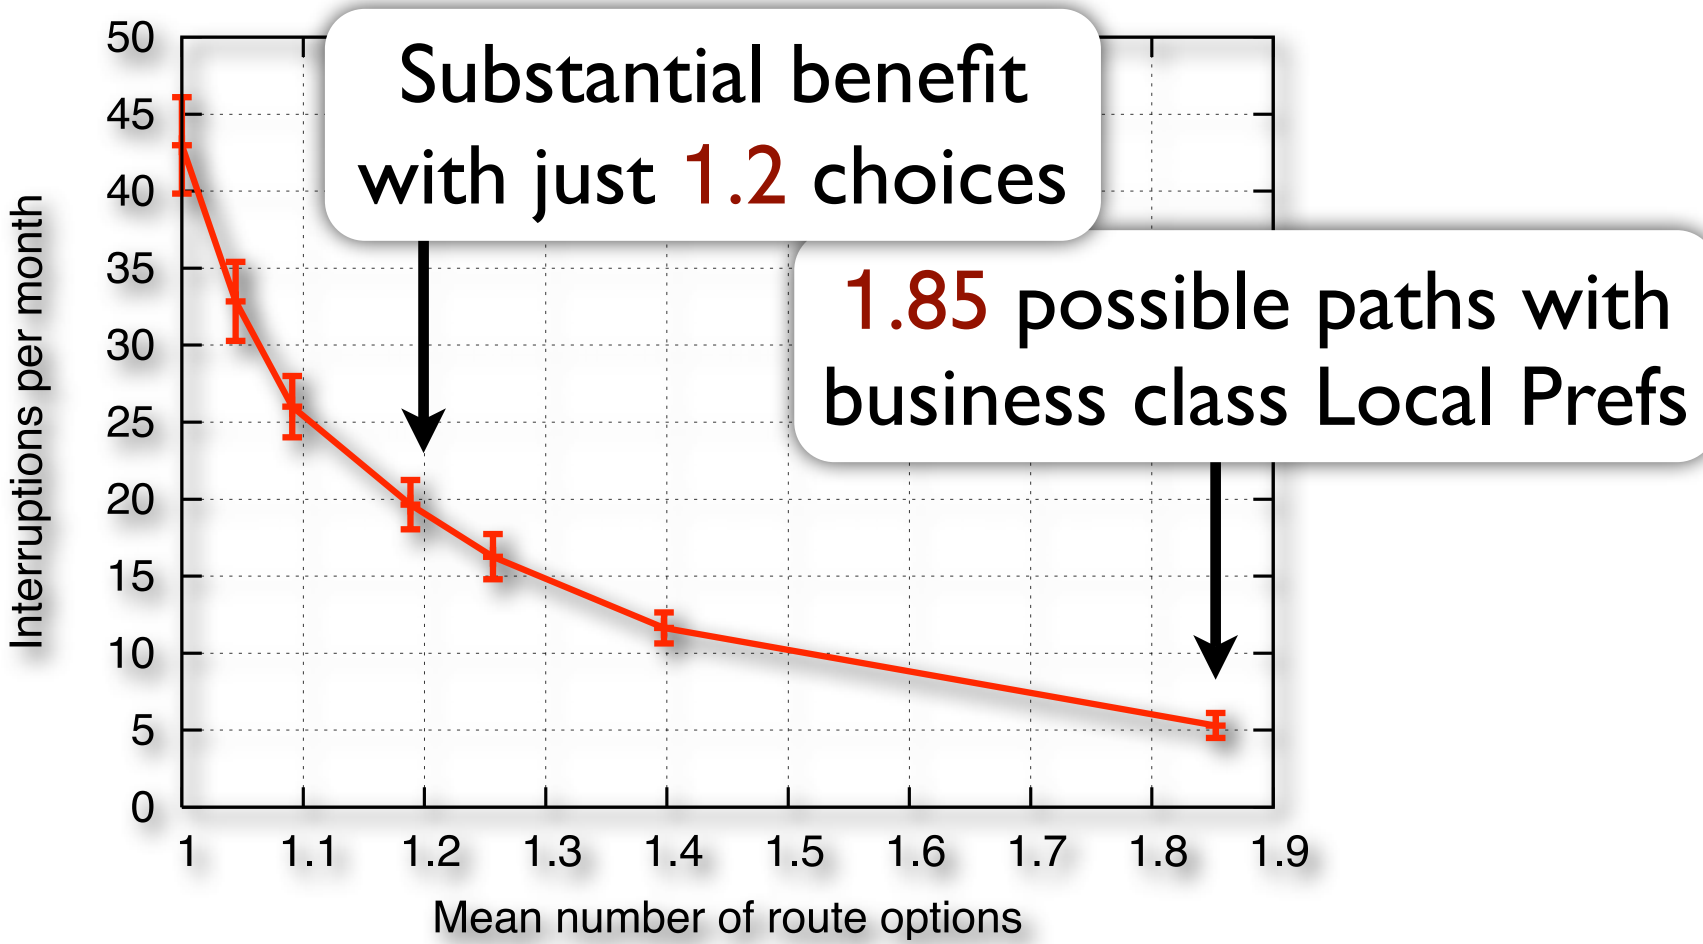

But how much flexibility in practice?

SRS convergence

- Convergence depends on decision process
- If heuristic is passive

SRS Heuristic
2. Current route
3. Lowest path length
4. Longest uptime

  - Any stable state for Std BGP is still stable
  - Gao-Rexford constraints still sufficient to guarantee convergence to stable state
- (Simulations: *slightly* faster convergence)

SRS can converge where standard BGP doesn't

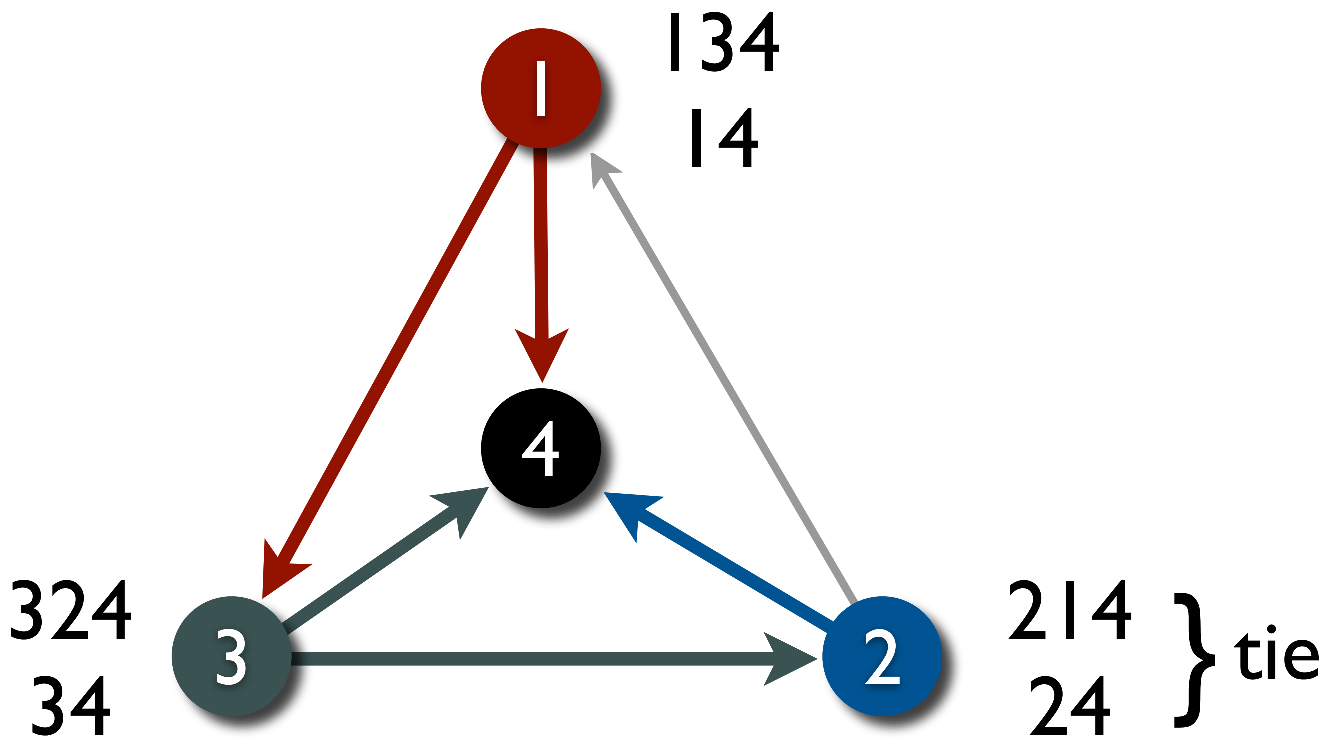

Dispute wheel

Acknowledgements

- This material is based upon work supported by a National Science Foundation Graduate Research Fellowship
- Traffic sign image from Manual of Traffic Signs, by Richard C. Moeur (http://www.trafficsign.us)