

Communication in Distributed Systems - KIVS II

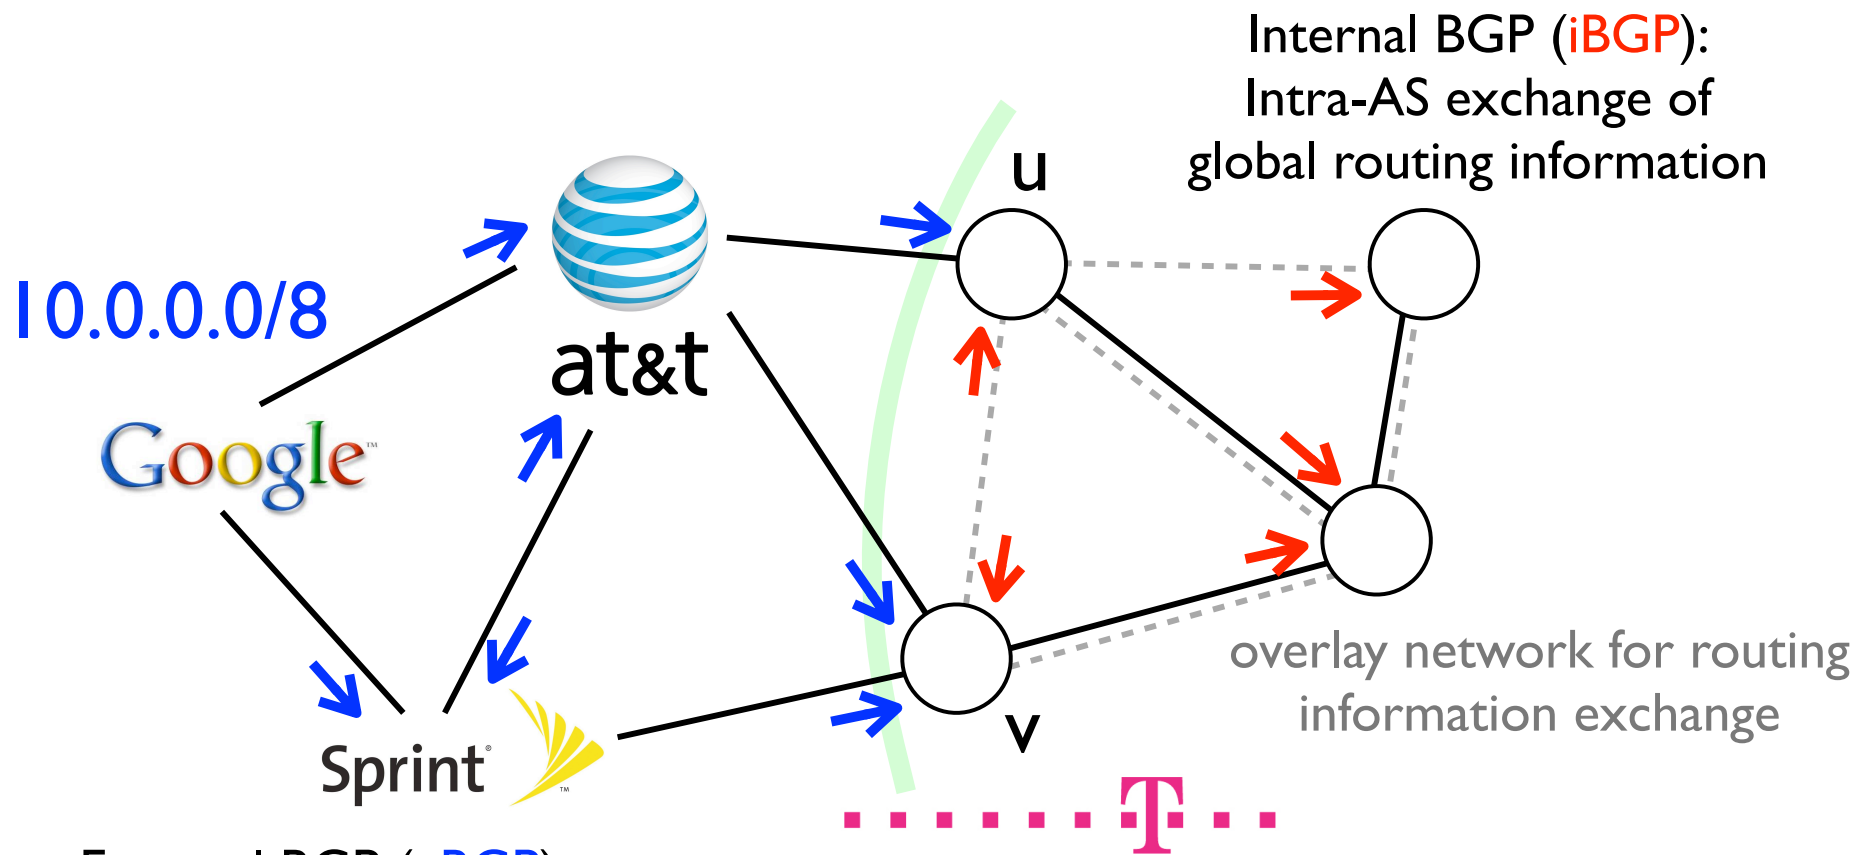
The Scope of the iBGP Routing Anomaly Problem

Introduction - Internal BGP · The Scope of the Problem · Conclusion

Table of Content

Introduction - Internal BGP

The Scope of the iBGP Routing Anomaly Problem



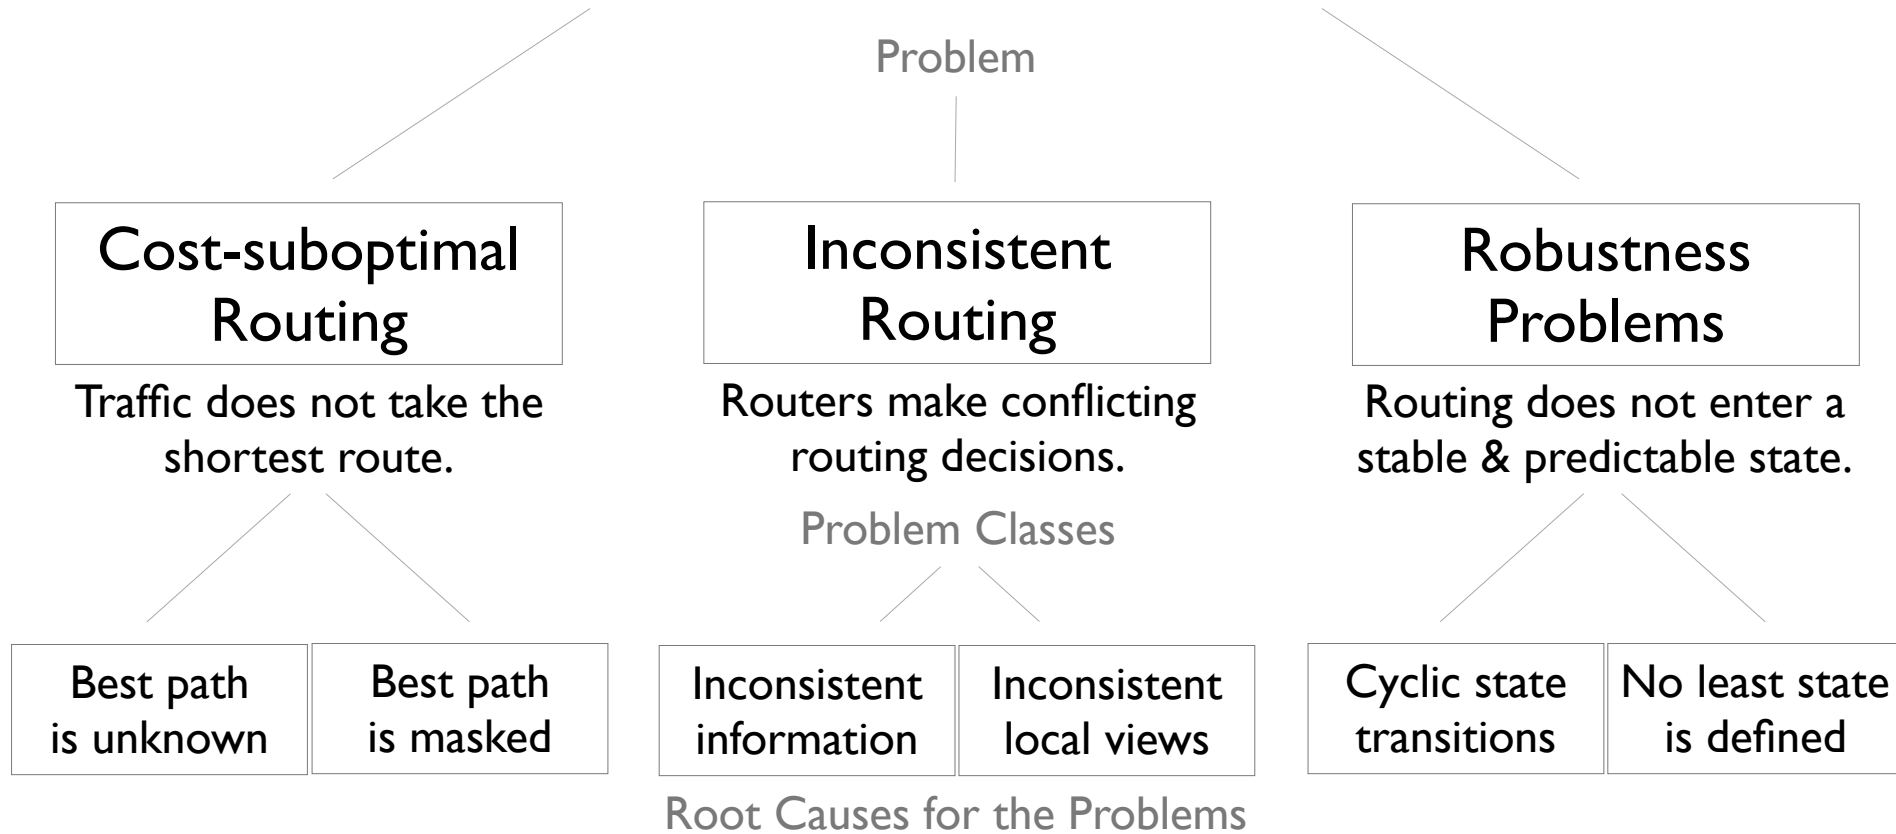
External BGP (eBGP):
Inter-AS exchange of
global routing information

Internal BGP (iBGP):
Intra-AS exchange of
global routing information

overlay network for routing
information exchange

iBGP: Exchange of Global Routing Information within ASs

iBGP Correctness Problems: „unwanted behavior“



iBGP Routing Anomalies and Their Root Causes



The problems are well-known.

The root causes are well-known.

The scope of the problem is unknown!

Which effect may appear under which scheme?

iBGP Routing Anomalies: State of the Art

The Scope of the Problem

The Scope of the iBGP Routing Anomaly Problem

Standardized iBGP

Full-mesh iBGP

Each router peers with every other $O(n^2)$.

Route Reflection

Paths are aggregated, best paths are reflected.

AS Confederations

AS is divided into smaller full-meshed sub-ASs.


For each scheme and root cause:



- Verify (by example) that root cause may appear (identify the circumstances).
- Verify that root cause cannot appear.


Determine the Scope of the Problem: What Did We Do?

Router v

	MED	External	IGP	BGP ID
r_v	Sprint:1	Yes	0	x.001 
q_v	AT&T:1	Yes	0	x.010


Router u

	MED	External	IGP	BGP ID
r_u	Sprint:1	No	X	x. 0 20
p_u	AT&T:2	Yes	0	x.011



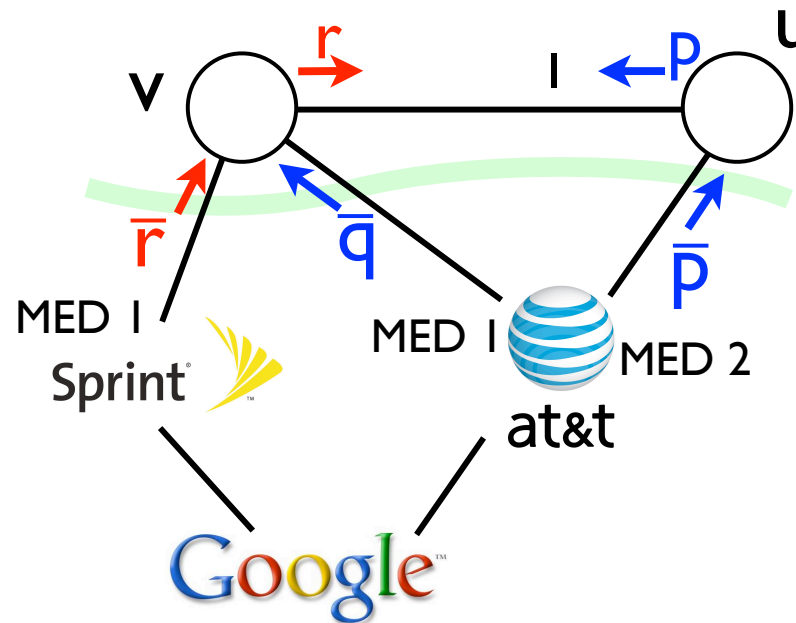
Example: **Path Masking** may Appear in Full-meshed Configurations.

Router v

	MED	External	IGP	BGP ID
r_v	Sprint:1	Yes	0	x.001 
q_v	AT&T:1	Yes	0	x.010
p_v	AT&T:2	No	X	x. 0 12

Router u

	MED	External	IGP	BGP ID
r_u	Sprint:1	No	X	x. 0 20
p_u	AT&T:2	Yes	0	x.011



u routes via an exit point with suboptimal MED.

Example: **Path Masking** may Appear in Full-meshed Configurations.

If router u would know all paths....

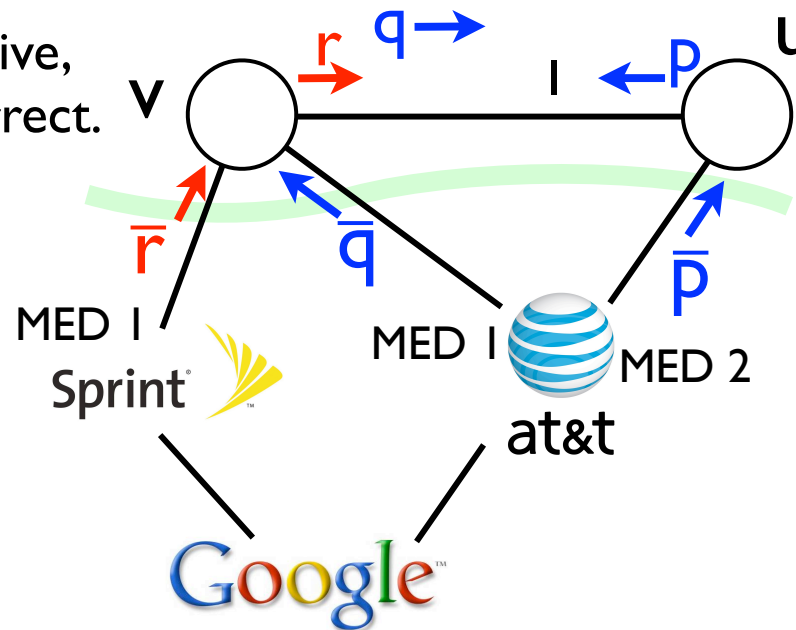
	MED	External	IGP	BGP ID
r_u	Sprint:1	No	1	x.020
q_u	AT&T:1	No	1	x.020
p_u	AT&T:2	Yes	X	x.012

Router u

	MED	External	IGP	BGP ID
r_u	Sprint:1	No	X	x.020
p_u	AT&T:2	Yes	0	x.011

In any case: NOT p!

from a global perspective,
the chosen path is incorrect.



Example: **Path Masking** may Appear in Full-meshed Configurations.

	Expressiveness Problems		Consistency Problems		Robustness Problems	
	Best path is unknown	Best path is masked	Inconsistent information	Inconsistent local views	Cyclic state transitions	No least state
Full-mesh iBGP	No	possible	No	possible	No	No
Route Reflection AS Confederations	trivial	possible	trivial	possible	possible	possible

Anomalies in Common iBGP Architectures

Advertisement of receiver-specific routing information.

Additional iBGP sessions.

Network designs that ensure a common view within clusters.

Modification of the BGP Path Selection Process.

Design Restrictions.

Resulting Workarounds and Solutions

Conclusion

The Scope of the iBGP Routing Anomaly Problem

IBGP anomalies may even appear if full-meshed iBGP is applied.



IBGP anomalies are no direct result of information reduction.

Using AS Confederations or Route Reflection, all classes of problems may appear.

Being interested in the details, take a look at the paper!

Conclusion and Take Aways




Use the results to solve the problems in new drafts.

Develop workarounds for the existing standards.

Specify new, anomaly-free iBGP schemes.

Main Aspects for Future Work



<http://cs.bonn.edu/IV/ub/>

Questions?