Doing Don’ts: Modifying BGP Attributes within an Autonomous System

Luca Cittadini, **Stefano Vissicchio**, Giuseppe Di Battista
Università degli Studi RomaTre

IEEE/IFIP Network Operations and Management Symposium (NOMS) - April 2010

The Point of View of an AS

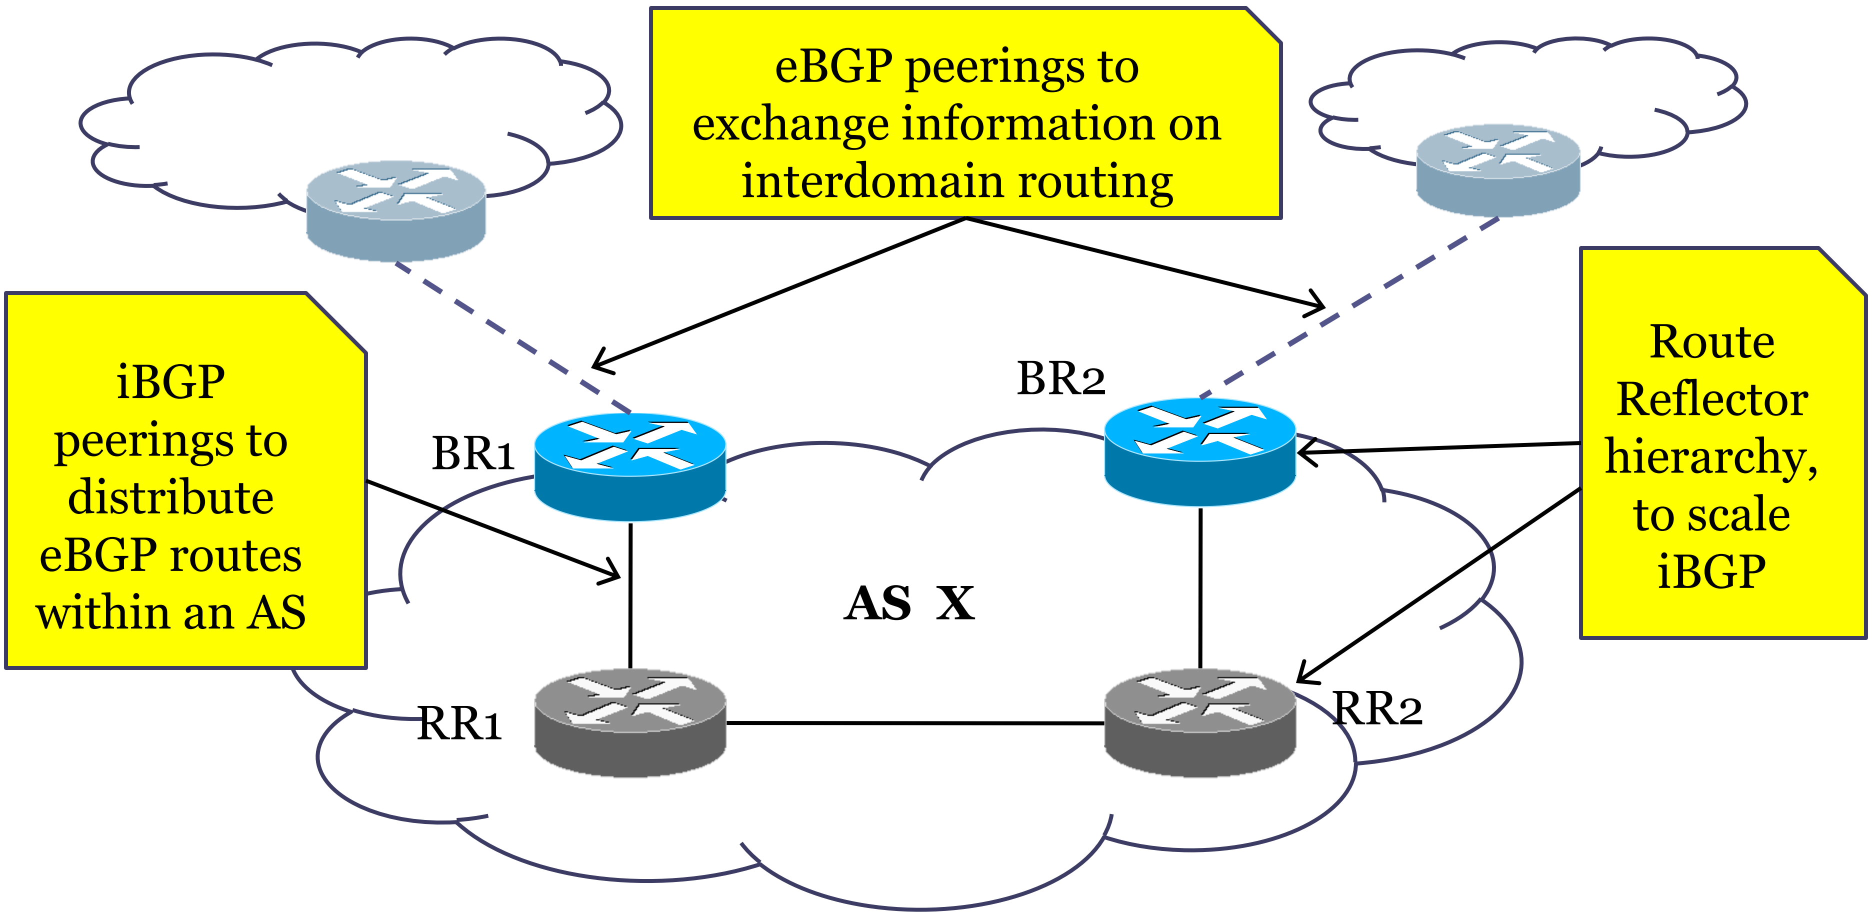

About The Role of iBGP

- iBGP just propagates eBGP routes across the AS
- Typically, no policy is applied in iBGP
  - Policies are applied only at the border
  - Best practices do not recommend applying policies in iBGP
  - Research community often assumes that iBGP messages are left untouched ([Griffin02],[Flavel08])

Our Contributions: Doing Don’ts

- Consequences of changing iBGP messages are poorly understood
- We investigate pros and cons
  - **why** an ISP should (not) modify iBGP messages?
  - **who** currently change iBGP attributes in the real world?
  - **what** are the consequences on iBGP stability?
    - technique and tool for detecting instabilities
  - **how** can an ISP change iBGP attributes without affecting iBGP stability?
    - configuration guidelines

Why Should an AS Change iBGP attributes?

- attributes influence the *BGP decision process*

Step	Criterion
1	Prefer routes with higher LOCAL-PREFERENCE
2	Prefer routes with lower AS-PATH length
3	Prefer routes with lower ORIGIN
…	…

- An ISP can change iBGP attributes as an additional knob to better engineer its traffic
  - recent Cisco IOS releases allow operators to partially modify the decision process

Simple Scenario

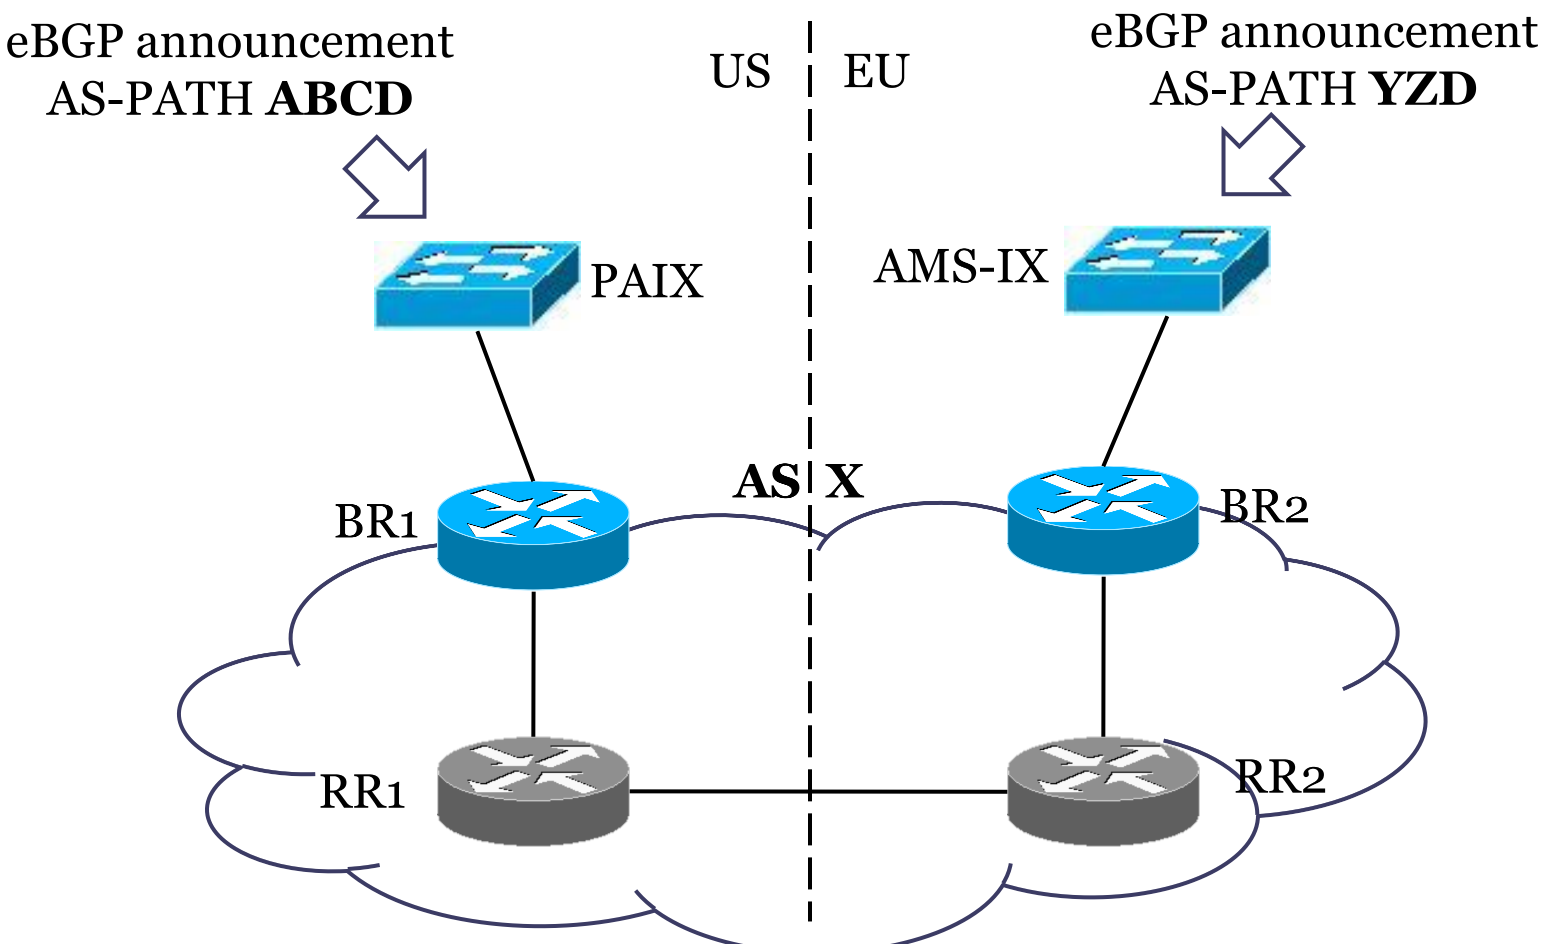

Simple Scenario

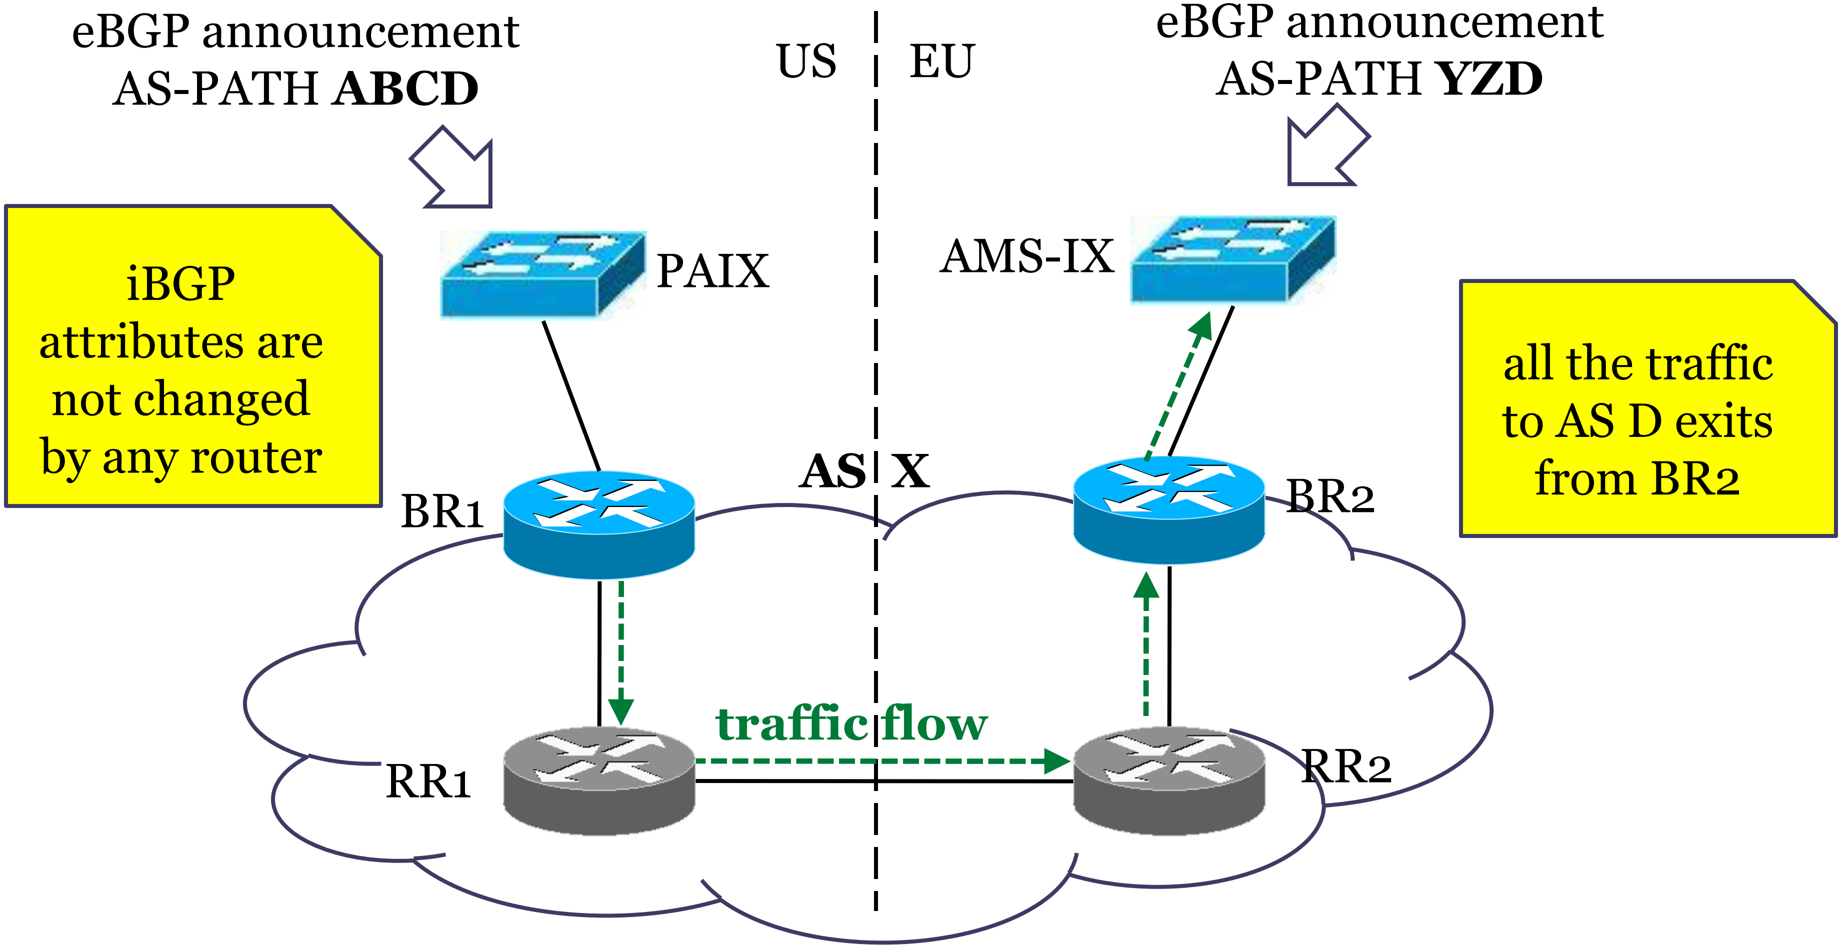

Simple Scenario

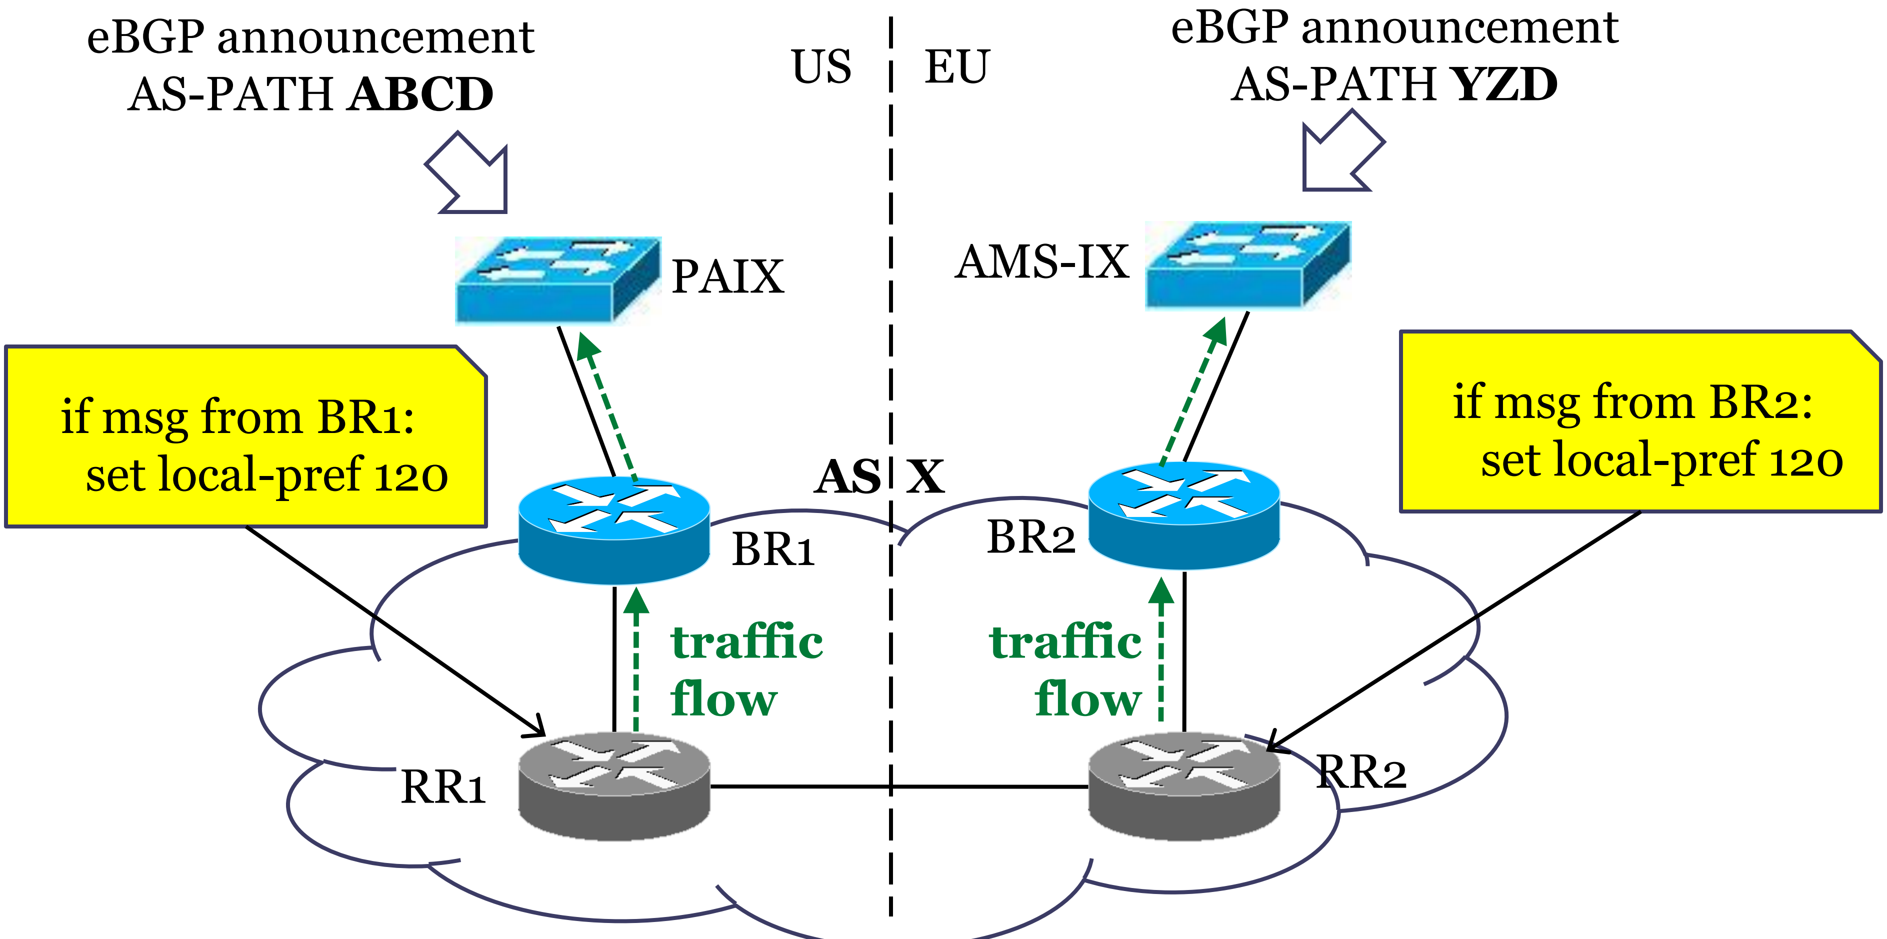

Who engineers iBGP in real world?

route-server.phx1>sh ip bgp 189.90.12.0/24

Paths: (4 available, best #1)

13878 15180 28189 28189 28189 28189, (received & used)

67.17.64.89 from 67.17.80.210 (67.17.80.210)

Origin IGP, metric 0, localpref 300, valid, internal, best

Community: 3549:4471 3549:30840

Originator: 67.17.81.221, Cluster list: 0.0.0.92

…

28189 28189 28189 28189 28189 28189 28189, (received &

67.17.64.89 from 67.17.82.40 (67.17.82.40)

Origin IGP, metric 0, localpref 300, valid, internal

Community: 3549:4950 3549:34076

Originator: 200.186.0.67, Cluster list: 0.0.2.109, 0.0.5.2

routes with different AS-PATH length

internal router

iBGP Attribute Changing in Internet

- We estimated the number of ASes which exhibit simultaneously active routes having different AS-PATH length
  - conservative approach
  - dataset: RIBs from RIS and Routeviews (May 2009)
- We found that 1,838 ASes out of 32,066 (5.7%) change iBGP attributes

What Impact on Routing Stability?

- We use a custom extension of the Stable Path Problem (SPP) framework as the model
  - an undirected graph represents (i)BGP peerings

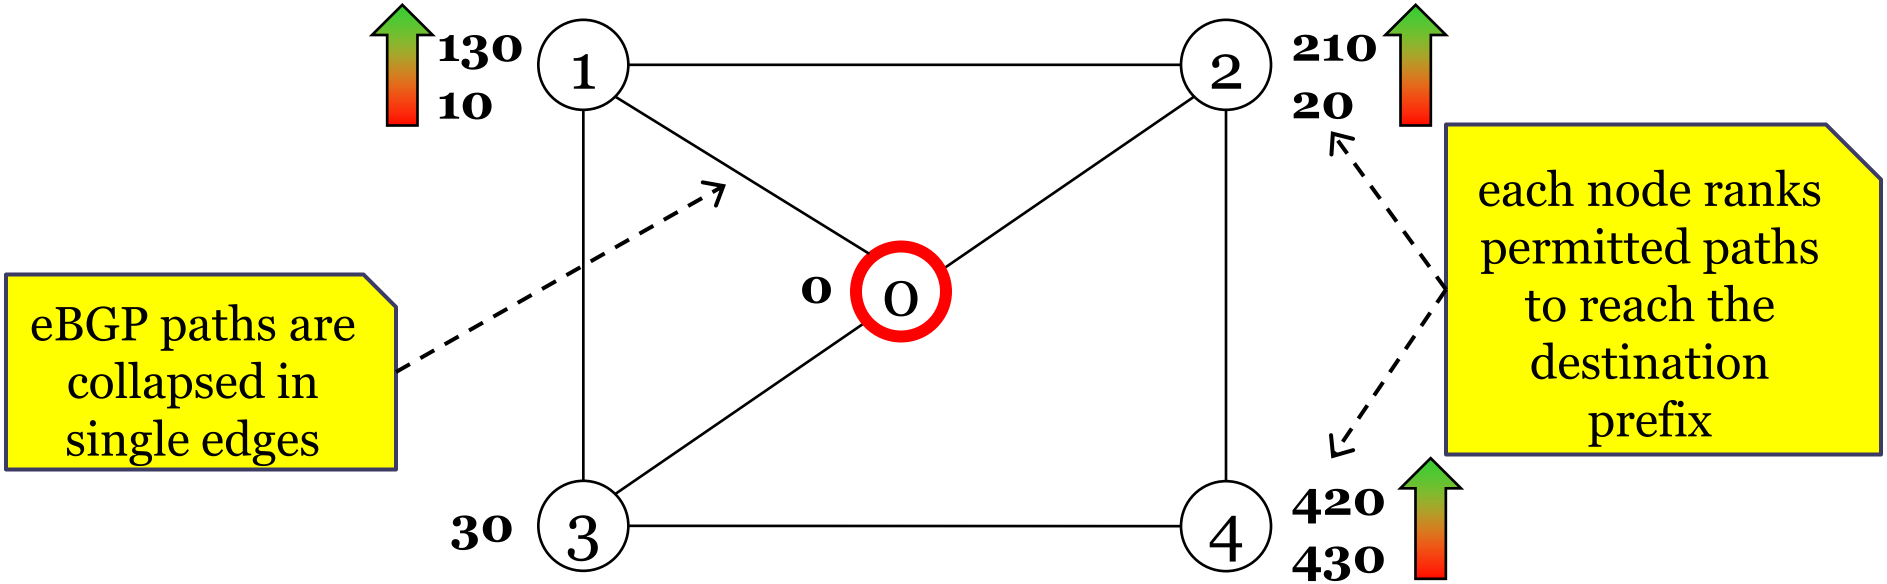

More Flexibility = More Instability

- Theo: BGP configurations that allow iBGP attribute changing can generate a larger set of oscillations than BGP configurations where iBGP attributes are not modified.

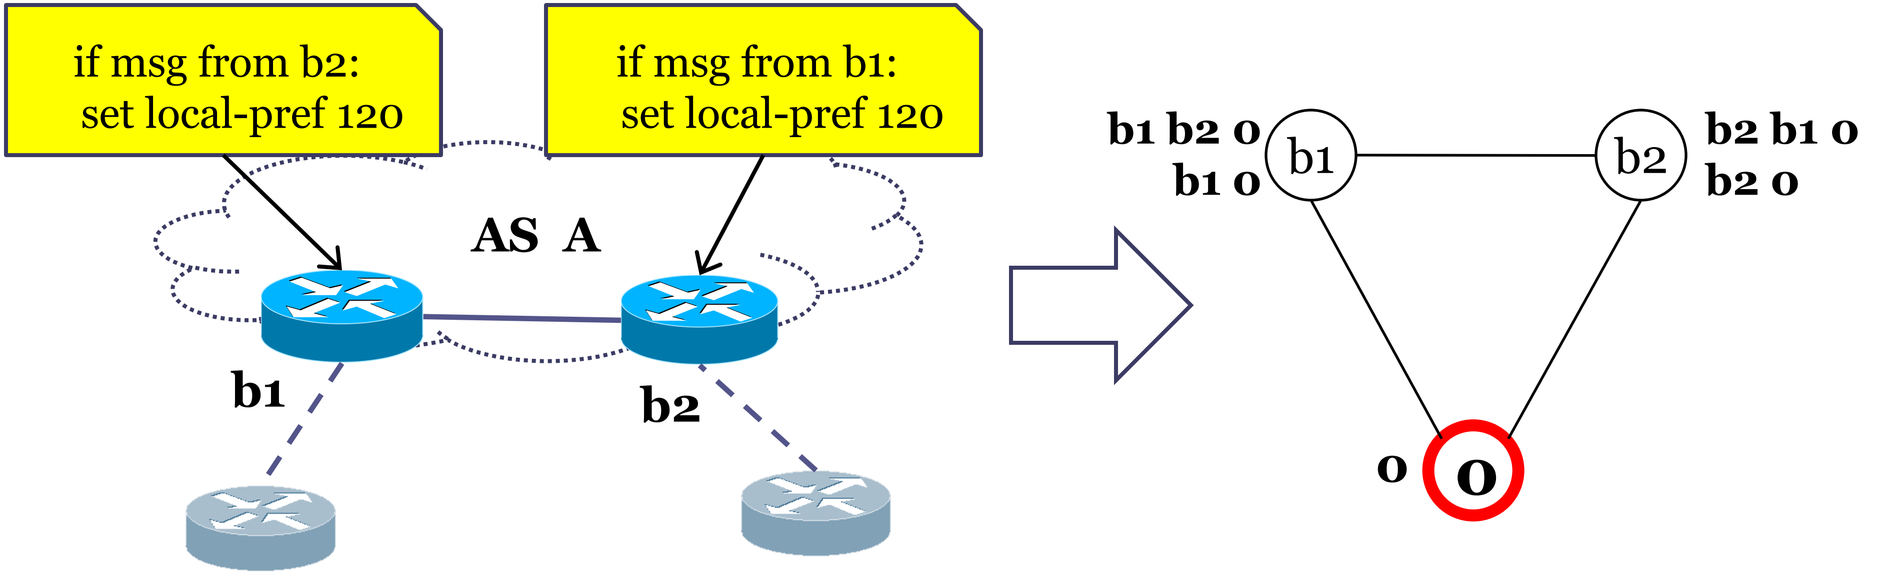

Detecting Oscillations

- We propose a technique to automatically check iBGP configuration for routing stability
  - Known detection techniques assume that iBGP attributes are left unchanged [Flavel08]
  - Conveniently translating an iBGP configuration into an SPP instance allows us to check for stability even when iBGP attributes are changed
  - The SPP instance is then checked with a known polynomial heuristic algorithm [Cittadini08]

A Stability Checker Tool

- We built and evaluated a prototype tool able to check iBGP configurations for stability

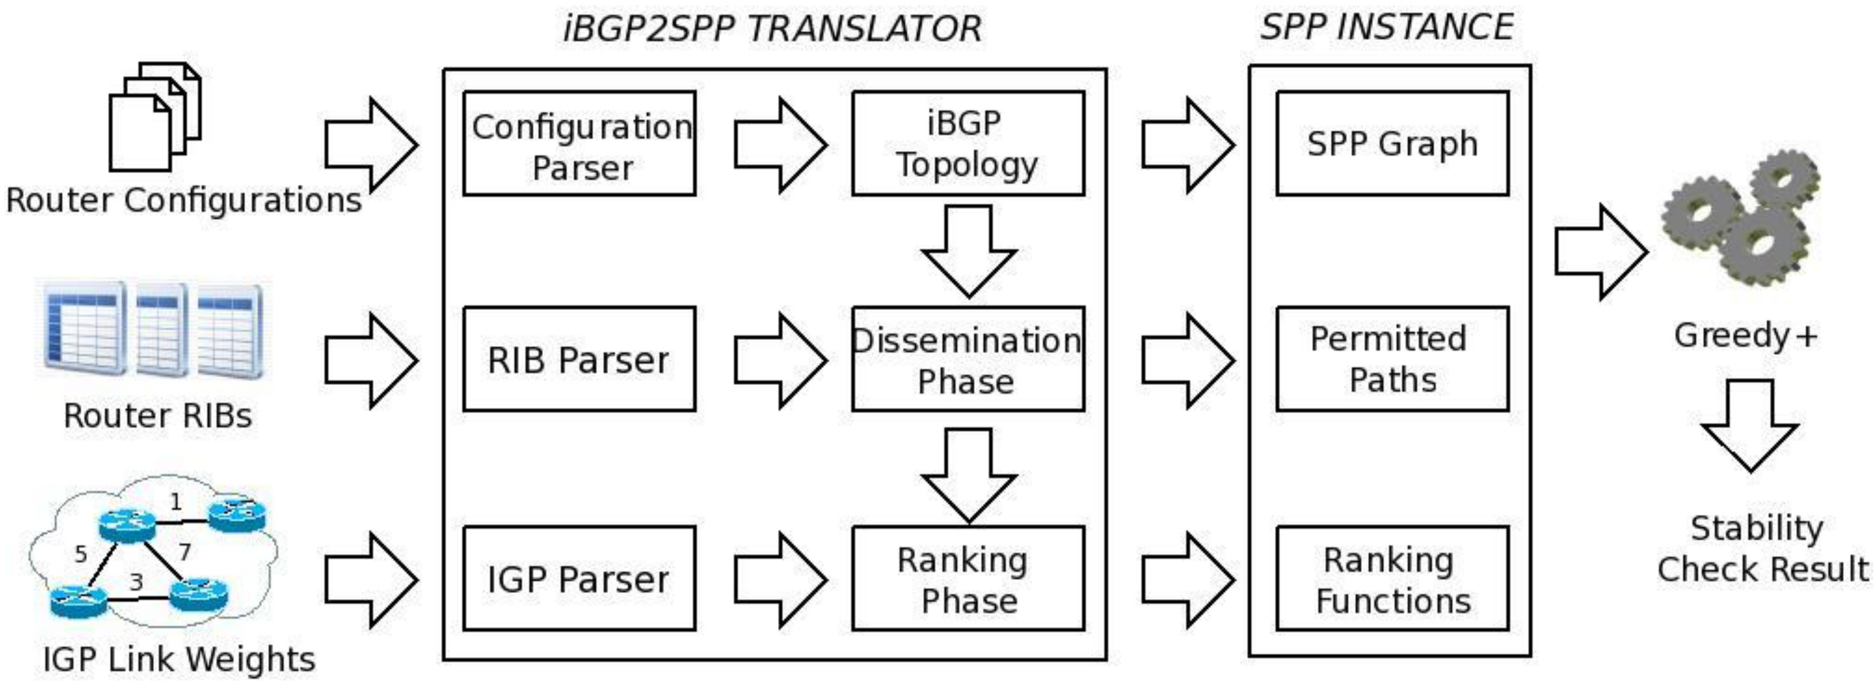

Evaluation of the Tool

- We evaluated the stability checker on both synthetic and real-world configurations
  - in-vitro configurations with up to 1100 iBGP speaking routers
    - three route reflection levels
    - 20 eBGP routes injected
  - real-world configuration of a medium-sized Italian ISP for all the prefixes in the full RIB
    - the entire test takes only few minutes

Performance of the Tool

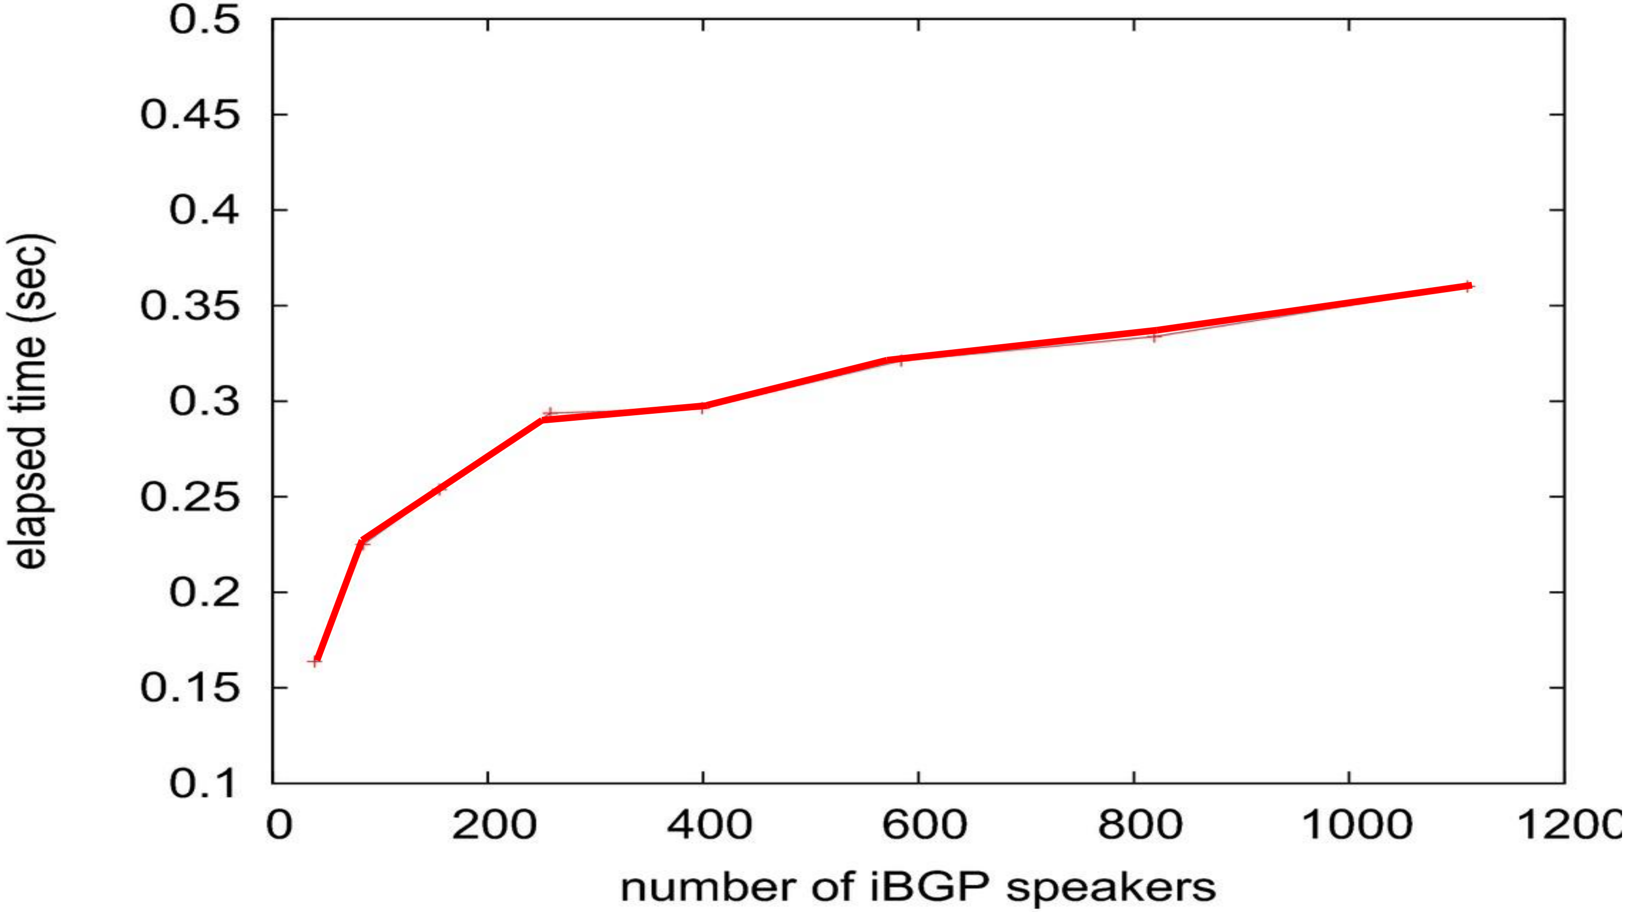

How Can We Prevent Oscillations?

- The more flexibility comes at the cost of increased instability
  - Is it possible to profitably change iBGP attributes without affecting routing stability and requiring low additional management complexity?
- We provide some configuration guidelines
  - two main high-level requirements
    - routers should be ranked according to revenues and costs [GaoRexford00]
    - internal transit cost should be minimized

Configuration Guidelines

Guideline A. Every iBGP speaker assigns to each route a local-preference value such that

- $LP_{cust} > LP_{peer} > LP_{prov}$

Guideline B. Route reflectors prefer the routes propagated from their own clients

- $LP_{cust\text{-}client} > LP_{cust\text{-}non\text{-}client} > LP_{peer\text{-}client} > LP_{peer\text{-}non\text{-}client}$ ..

We also formally proved that these guidelines guarantee routing stability

Conclusions

- We explored the possibility of changing attributes in iBGP, evaluating pros

	Aspect	Contributions
Pro	Better traffic engineering	We showed a simple scenario in which an ISP is allowed to easily perform fine-grained TE
Pro	Complete control of traffic flow	We showed that iBGP attribute changing can prevent traffic shift and modification of load balancing
Pro	Added flexibility	We showed that some real-world ASes already change iBGP attributes

Conclusions

- We explored the possibility of changing attributes in iBGP, evaluating pros and cons

	Aspect	Contributions
Con	Stability Problem is exacerbated	• We defined configuration guidelines which ensure stability • We proposed a technique for detecting routing instabilities • We realized a prototype tool able to efficiently detect instabilities
Con	Management complexity	Our guidelines do not add significant configuration complexity