EULER Project
Path-Vector Routing Stability Analysis

Florin Coras, Albert Lopez, Albert Cabellos

UPC

Dimitri Papadimitriou

Alcatel-Lucent

Introduction

BGP

- Inter-domain routing protocol used in the Internet
- AS-Path vector routing protocol
- Policing (without exchange of policies), traffic engineering

Affected by instability

- Policy- induced
- Protocol-induced

Effects

- Non-deterministic/unstable BGP states
- Long convergence time upon topology change/failure

Current Internet Growth Rate (1)

Traffic

- Traffic volume (per month): [8,9] Exabytes
- Traffic growth rate: 50% (+/- 5%) per year

Routing tables size

- Number of active Routing Table (RT) entries: 345k (Sep.2010)
- Growth rate: 15%-25% per year

Autonomous Systems (AS)

- Number of advertized AS: 35k (Sep.2010)
- Growth rate: 10% per year

Characteristic AS-path length

- Steady ~3.7

AS transit interconnection degree: growing (2.56 – 2.60)

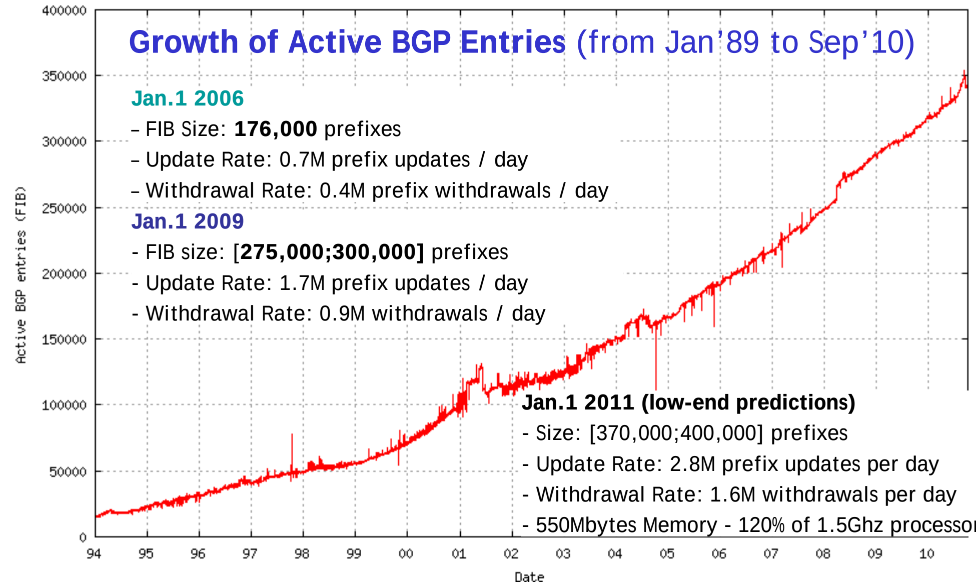

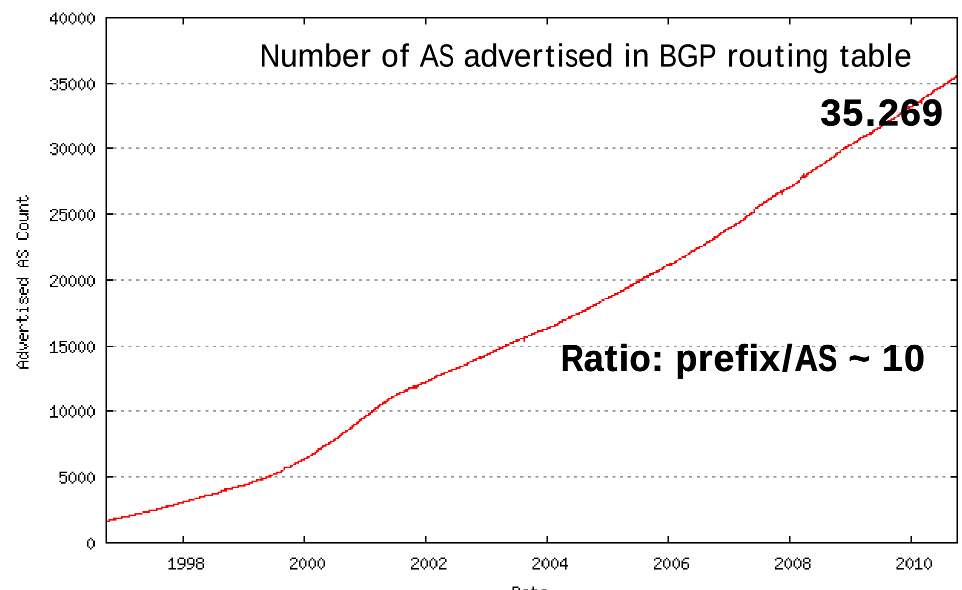

Source: BGP Routing Table Analysis Reports - http://bgp.potaroo.net

Current Internet Growth Rate (2)

Dynamics BGP updates (routing convergence)

- Between Jan.2006 and Jan.2009: prefix update and withdrawal rates per day increased by a factor of about 2.25-2.5 [Huston07]
  - Average: 2-3 per sec. – Peak: O(1000) per sec.
- BGP suffers from churn which increases load on routers due to topological failures and traffic engineering (prefix de-aggregation)
- BGP's path vector amplifies these problems (path exploration)

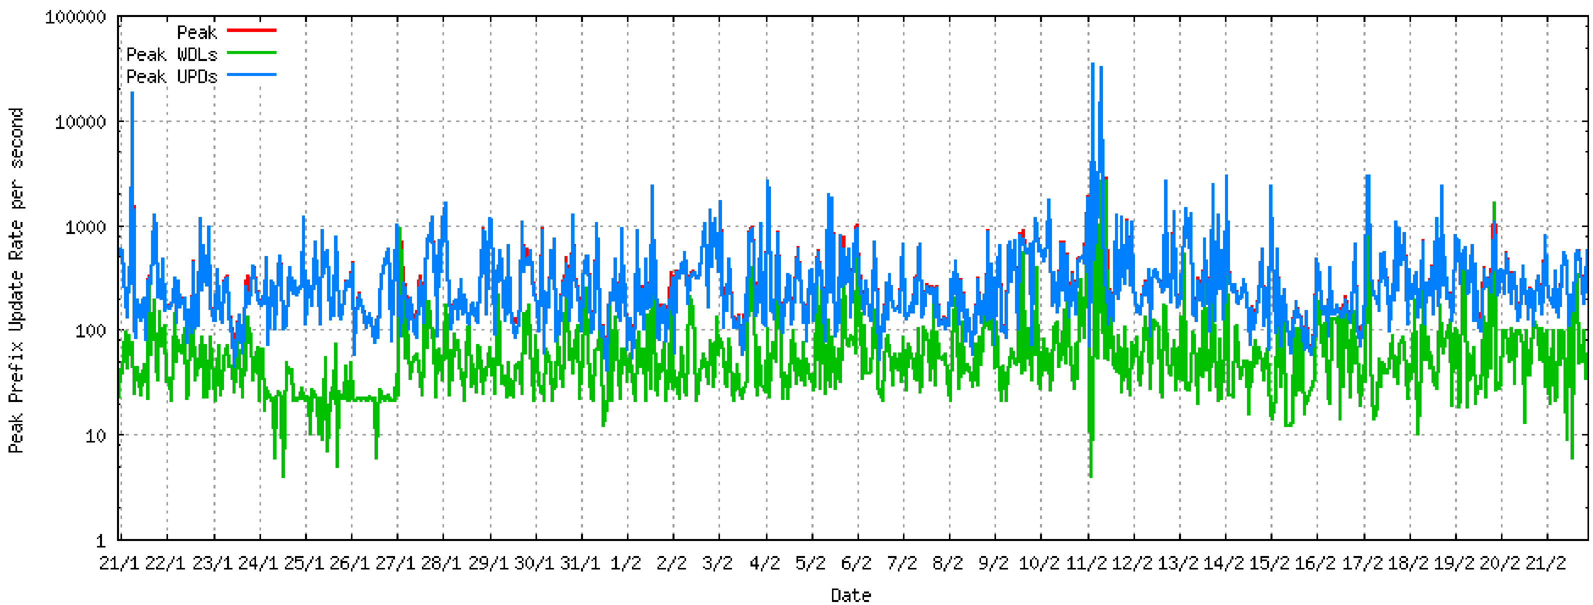

Objective

Overall

- Rigorous characterization of instability phenomena
- Robust framework for a straight forward and improved description of instability events

Specific objectives

- Methodology to process and interpret routing table data
- Set of stability metrics and means to apply them
- Ways of improving current and future routing systems' stability

Model (1)

Preliminaries

- Stability of the routing system is characterized by the system's response to inputs of finite amplitude
- Inputs may be
  - Internal: routing protocol configuration change
  - External: topological changes
- Formalism for characterizing route and routing table stability defined in draft-ietf-grow-rss-00.txt [ID-grow-rss09]

Model (2)

A route is defined as:

$$r_i(t) = \{\text{destination, path, attributes}\}$$

We define function φ_i to measure the stability of route *i*:

$$\varphi_i(\mathrm{t}+1) = \begin{cases} \varphi_i(t)+1, & update \\ \varphi_i(t)-1, & no\,update \end{cases}, \qquad \varphi_i(t) \geq 0$$

And $\Delta r_i(t+1)$ to quantify the change in stability of a route

$$\Delta r_i(t+1) = \begin{cases} [\varphi_i(t)+1]/[\varphi_i(t+1)+1], \; update \\ [\varphi_i(t+1)]/[\varphi_i(t)], \; no\ update \end{cases}, \qquad \varphi_i(t) > 0$$

Model (3)

BGP router comprises 3 logical routing information bases:

- **Adj-RIBs-In**: stores routing information learned through updates from all the peers
- **Loc-RIB**: local routing data obtained after the BGP selection process
- **Adj-RIBs-Out**: contains routes that are advertised to peers

Routing table stability defined as the average of the changes in stability of the routes composing it

Stability Metrics

Set of stability metrics for quantifying

- BGP selected routes stability (draft-ietf-rss)
- Relative stability of routes with identical destinations (most stable)

$$\Delta\varphi_{i,j}(\mathrm{t}+1) = [\varphi_{i,j}(\mathrm{t}+1)+1]/[\varphi_{i,stable}(\mathrm{t})+1]$$

- Relative stability between learned routes and the one selected by BGP (best selectable)

$$\Delta\varphi_{i,j}(\mathrm{t}+1) = [\varphi_{i,j}(\mathrm{t}+1)+1]/[\varphi_{i,selected}(\mathrm{t})+1]$$

- Differential stability between most stable route and that selected by BGP

$$\delta\varphi_{i}(\mathrm{t}) = \varphi_{i,selected}(\mathrm{t}) - \varphi_{i,to\,be\,selected}(\mathrm{t})$$

Experiment Tool

BGP feed processing tool that

- Parses significant volumes of real BGP data from routers under the control of RouteViews
- Computes the values of the proposed metrics when applied to the obtained BGP data

Proposed stability metrics provide a framework for methodological interpretation of BGP routing information

Experiment: Dataset

Obtained from RouteViews [rv] - a project that comprises archives containing BGP feeds from a set of worldwide distributed Linux PCs running Zebra

Our experiments use one of the Oregon based Zebra router's BGP feeds

- The router currently keeps 46 peering sessions with routers belonging to 38 ASNs
- It contains around **11M** routes (Adj_RIBs_In)
- Note: we implement the Zebra BGP decision process to infer the Loc_RIB

Experiment: Dataset (2)

RouteViews routers use multi-hop BGP sessions to peer with transit providers and other ASes at selected locations around the globe

- Dumps are in either Cisco or Zebra format
- RIBs are collected every two hours
- Update files are rotated every 15 minutes
- Archives run to 1997 for some routers

Results: Stability of selected route

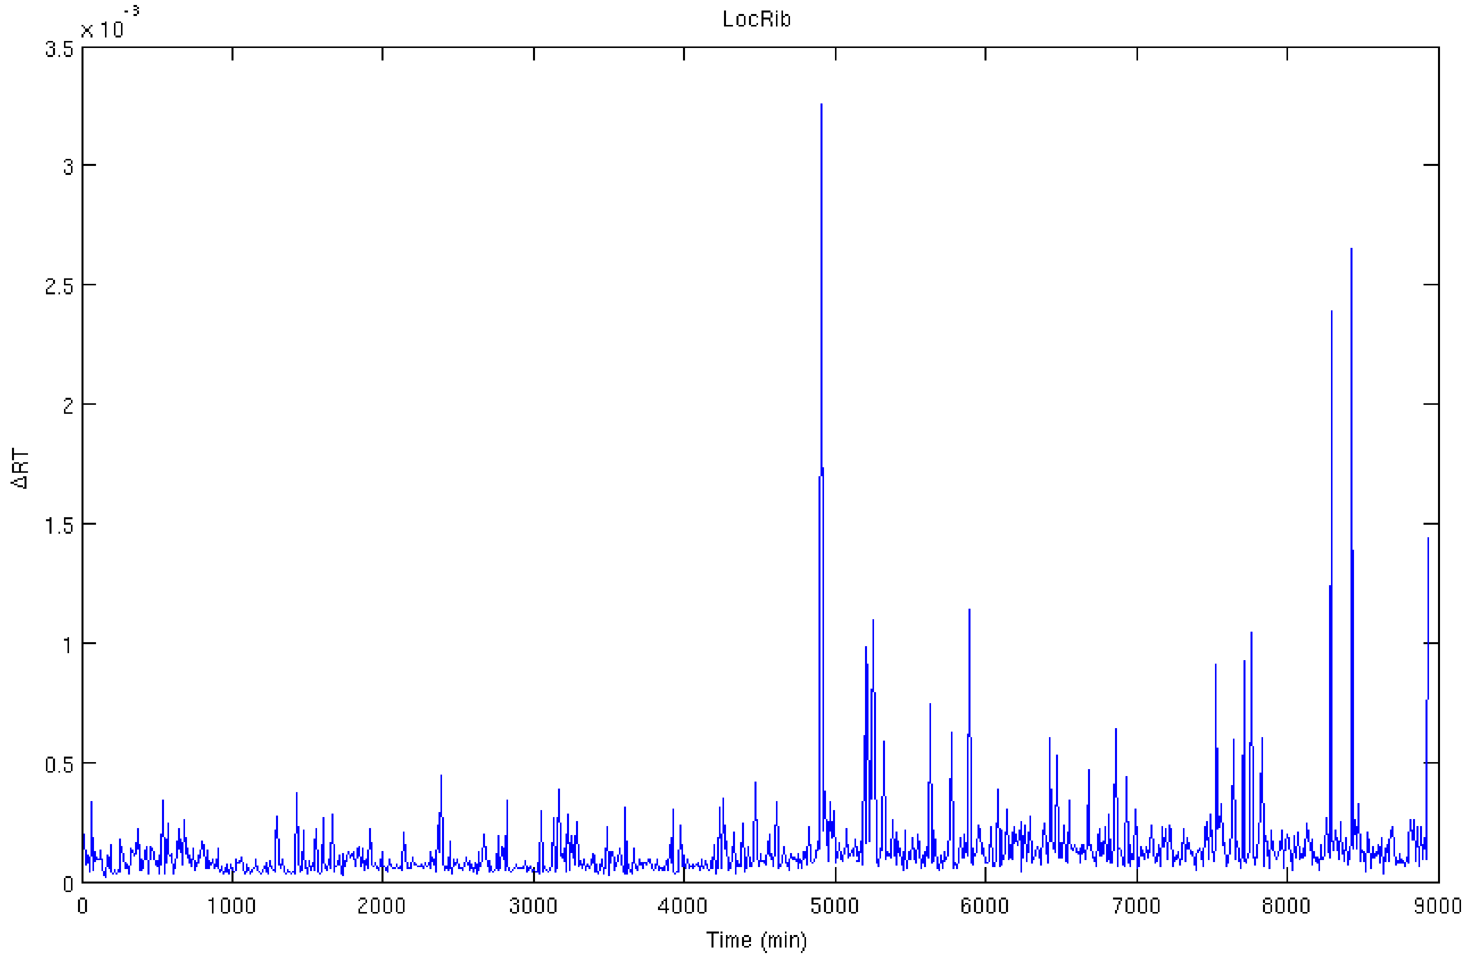

Results: best selectable and most stable metric (2)

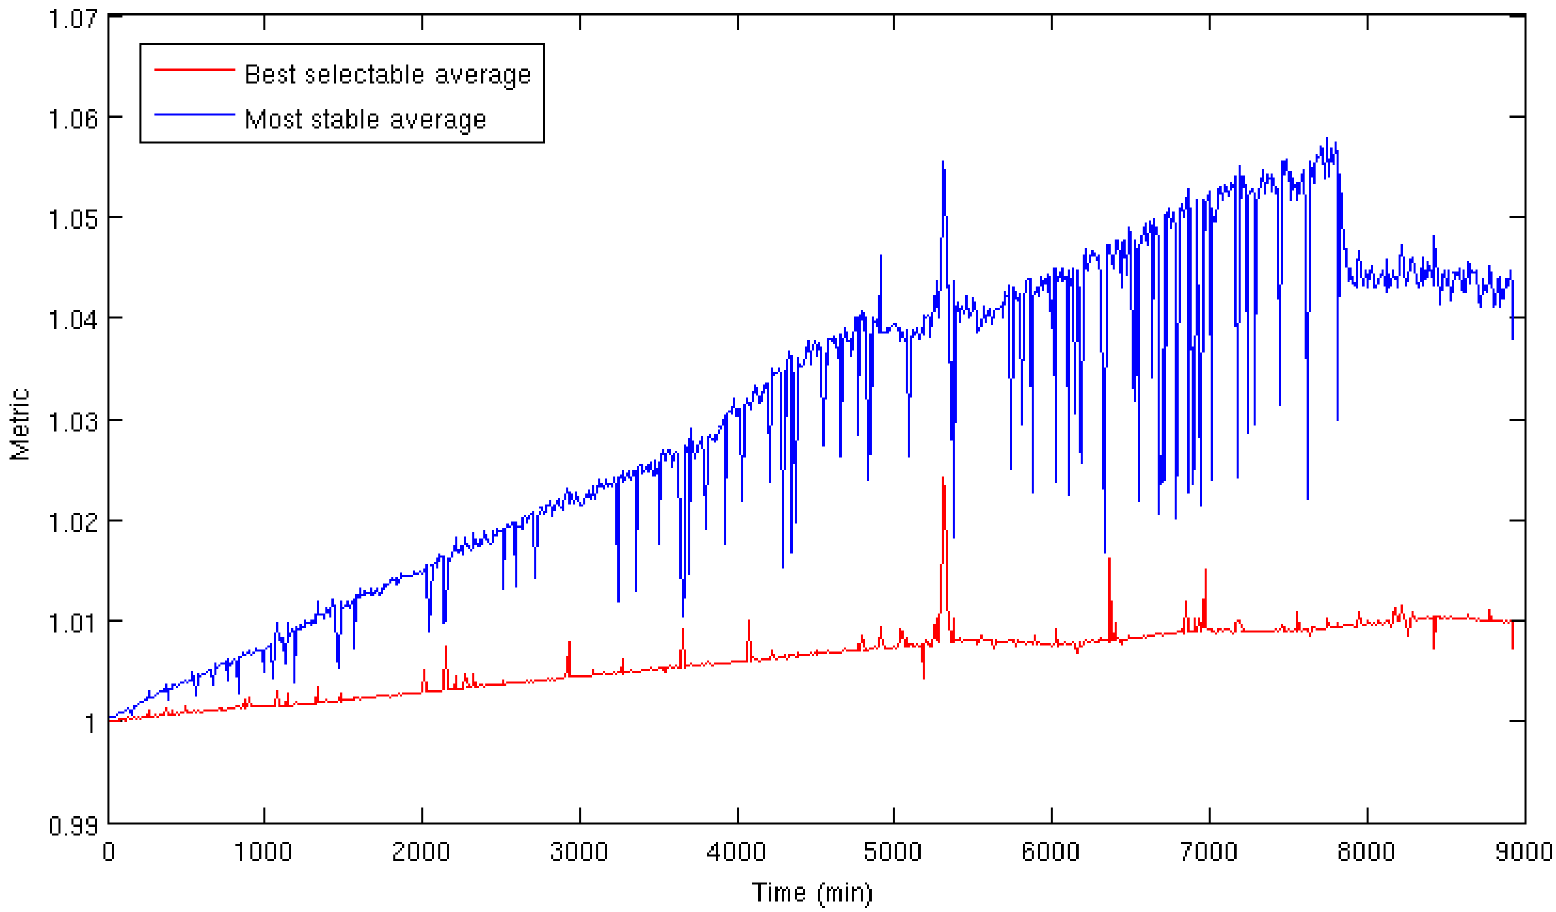

Conclusions/Ongoing work

We have defined a set of stability metrics that characterize BGP's instabilities

Initial results suggest that they can be used to both identify and quantify the impact of instability events on local routing tables

Ongoing work includes determining if our stability function $\delta\varphi_i(\mathrm{t})$ can be integrated within the BGP decision process and the possible trade-offs

References

[Huston07] G. Huston, *http://www.potaroo.net/ispcol/2007-06/dampbgp.html* , Damping BGP, 2007.

[ID-grow-rss09]