Analysis of Path-vector Routing Stability

D.Papadimitriou (Alcatel-Lucent Bell Labs)
dimitri.papadimitriou@alcatel-lucent.com
A.Cabellos (Technical University of Catalonia, Barcelona)

ALGOTEL 2012

May 29 – June 1, 2012

Introduction

- Border Gateway Protocol (BGP)
  - Inter-domain (AS) routing protocol of the Internet routing system
  - (AS-)Path vector routing algorithm
  - Capabilities
    - Policing (without exchange of policies)
    - Prefix and community-based traffic engineering
- Affected by instability
  - **Policy-induced instability**: conflicting policy interactions
  - **Protocol-induced instability**: path exploration
- Effects
  - Non-deterministic unstable states (dispute wheels)
  - Delayed BGP convergence (long convergence time)

Routing system/state stability

- Characterized by its response (in terms of processing of routing information) to inputs of finite amplitude
- Inputs may be classified as follows
  - **Internal events**
    - Routing protocol configuration change
    - Software changes/updates
  - **External events**
    - Topological changes
    - Policy changes

  Both types of events lead to exchange of routing updates that may result in routing states changes
- Note: BGP does not differentiate routing updates with respect to their root cause, their identification (origin), etc. during its selection process

Objectives

1) Develop a method to systematically process and interpret data part of BGP Routing Information Bases (RIB) to identify and characterize occurrences of BGP routing system instability

2) Determine a set of stability metrics to measure local stability properties of path-vector routing

3) Investigate how path-vector routing behavior and network dynamics mutually influence each other

By means of these metrics

- Develop a method to analyze effects/impacts of BGP policy- and protocol-induced instability on local routers.
- Derive a stability decision criterion that can be applied as part of the BGP route selection process
- Study applicability of this decision criterion using real BGP datasets

BGP and Stability Metrics

- Stability of selected route r_i at time t (Loc_RIB): $\varphi_i(t)$
- Most stable in Adj_RIB_In: relative stability between learned routes with identical dest. d at time t+1 and most stable learned route for dest d. at time t: $\Delta\varphi_{i,j}(t+1) = [\varphi_{i,j}(t+1)+1]/[\varphi_{i,stable}(t)+1]$
- Best selectable route in Adj_RIB_In: relative stability between learned routes with identical dest. d at time t+1 and route selected by BGP for dest d. at time t: $\Delta\varphi_{i,j}(t+1) = [\varphi_{i,j}(t+1)+1]/[\varphi_{i,selected}(t)+1]$
- **Differential stability** between most stable candidate route and the route currently selected by BGP: $\delta\varphi_i(t) = \varphi_{i,selected}(t) - \varphi_{i,candidate}(t)$

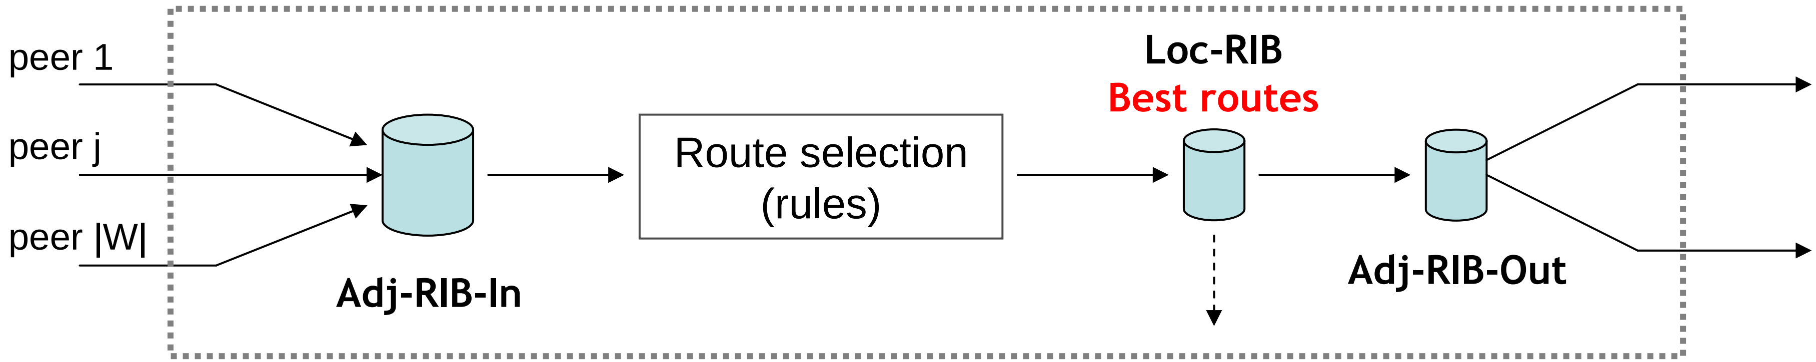

Differential stability

- Differential stability between selected route at time t for dest. d (stored in Loc_RIB) and newly selected route at time t+1 for same dest. d: $\delta\varphi_i$ ($i \in [1,|D|]$)
  - Characterizes stability of selected route $r_i(t)$ at time t (active route for dest.d) against stability of newly selected route $r^*_i(t)$ at time t for the same destination that would replace $r_i(t)$ at time t+1:

    $\delta\varphi_i(t) = \varphi_i(t) - \varphi^*_i(t)$

    where, $\varphi_i(t) = \varphi_{i,selected}(t)$ and $\varphi^*_i(t) = \varphi_{i,candidate}(t)$
  - If $\delta\varphi_i(t) > 0$ then replacement of route $r_i(t)$ by $r^*_i(t)$ increases stability

    Otherwise, safest decision is to keep currently selected route $r_i(t)$
- Application of metric $\delta\varphi_i$ in BGP route selection
  - Prevents replacement of more stable routes by less stable one
  - Enables selection of more stable routes than currently selected routes
- Local Proof of consistency of stability-based selection with preferential-based selection (path ranking) [SIGMETRICS11]

Measurement-based Results

- Data set: routeviews BGP data set containing around 11M routes (Adj_RIBs_In)
- Stability decision criteria ($\delta\varphi_i(t) > 0$) leads to AS-Path length/stretch decrease ($\delta\rho_i(t) > 0$) for more than 25% of the routes
- Using this decision criteria no negative/detrimental effects for 90% of the routes
- Only 10% of the routes would be stretch increasing ($\delta\rho_i(t) < 0$ with $\delta\varphi_i(t) > 0$) but stretch increase would be limited to 1 or 2 AS-hops

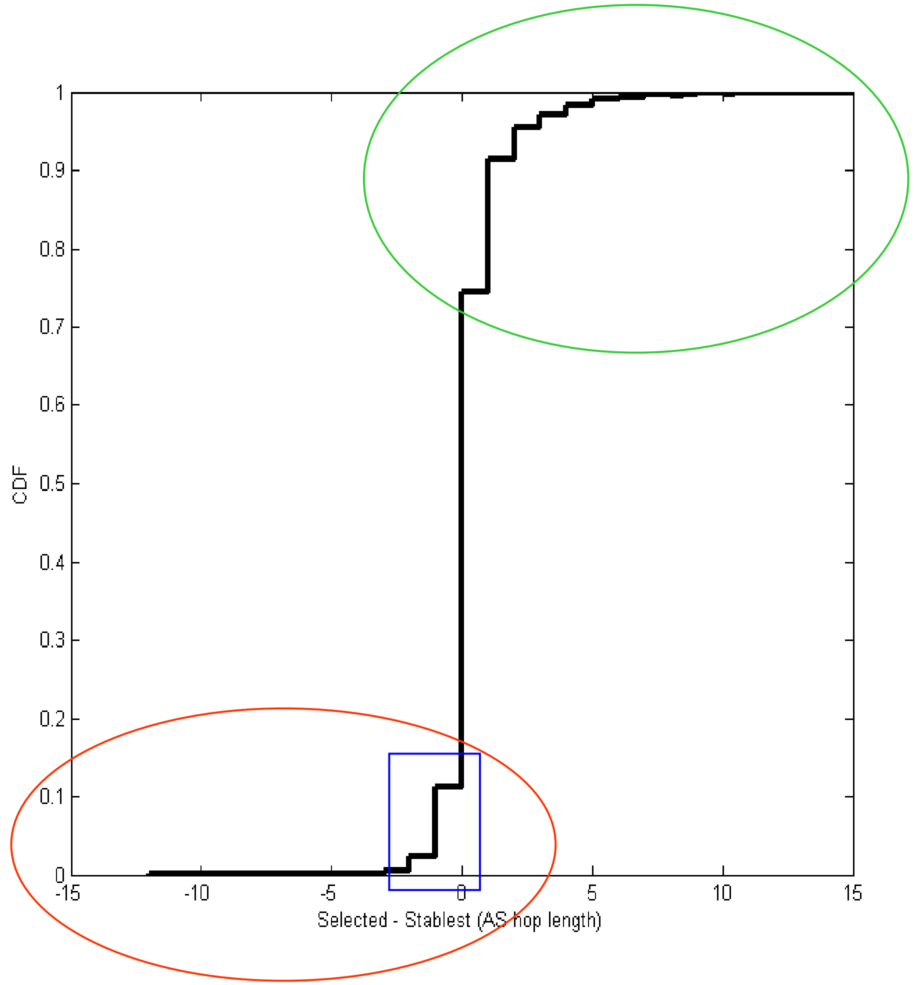

Conclusion

- Differential stability-based decision criterion that can be taken into account as part of the BGP route selection process.
- A significant fraction of the routes (90%) selected by means of this process is not stretch increasing.
  - If one would admit an AS_path length increase of one AS-hop, only a minor fraction of the routes (about 2%) would be penalized by a higher stretch increase (two AS-hops and above)
- Future/ongoing work includes
  - Verify general trade-offs between stability-based route selection and resulting stretch increase/decrease factor on selected routing paths
  - Determine necessary but sufficient conditions for preventing potential oscillations to occur (as local action of selecting a more stable route shall not induce unwanted perturbation(s) on neighboring routing states)
  - Generalize formulation of stability function (and variation increments)
  - Model extended to discriminate between policy vs. protocol instability