QoS Enhancements to BGP in Support of Multiple Classes of Service

A. Terzis[†], M. Liang[†], L. Benmohamed[*], E. Naber[*]
Computer Science Department[†]
Applied Physics Lab[*]
Johns Hopkins University
ID: *draft-liang-bgp-qos-00.txt*

Motivation

- Applications of emerging IP networks require network paths with diverse QoS characteristics
  - Global Information Grid (GIG)
    - Large-scale internet for the US government
    - Diverse link characteristics
      - OC-192 backbone, satellite links, tactical wireless networks, etc.
    - Diverse applications
      - VoIP, bulk data transfer, etc.

Current Limitations

- BGP design limits its ability to provide multiple QoS paths
  - *Reachability* under policy constraints is BGP's focus.
  - Single route advertised → Limited alternate route visibility
  - Path selection logic does not consider QoS characteristics
    - Number of AS hops is the only rough indication of path quality

Goal

- Expose multiple network paths to applications with different QoS requirements
  - Paths can have multiple QoS attributes (bandwidth, delay, loss, etc.)
  - Paths span multiple administrative domains

Proposed BGP Changes

1. Maintain multiple QoS metrics for each path
2. Exchange multiple paths per destination
3. Prune the set of known paths to a dominant set while maintaining optimality
4. Choose a particular path from this dominant set for the unique QoS requirements of a traffic class

Maintaining QoS metrics

- Need to accumulate QoS parameters across E2E path
- Different accumulation rules for different QoS metrics.
  - *Additive metrics*: *e.g.* latency
  - *Multiplicative metrics*: *e.g.* packet loss rate
  - *"Min" metrics*: *e.g.* bandwidth.

Dominant Path Selection Algorithm (DPSA)

- BGP routers are allowed to advertise multiple paths for a given destination prefix
- DPSA reduces the number of paths exchanged while exposing “best” paths (optimality)
  - Path *P* dominates a set *S* of paths if it can provide better QoS than any path in S for all QoS metrics of interest
- If more than one dominant paths have identical QoS metric values, the path with lower AS_PATH hop count and lower next-hop IP is preferred.

DPSA Example

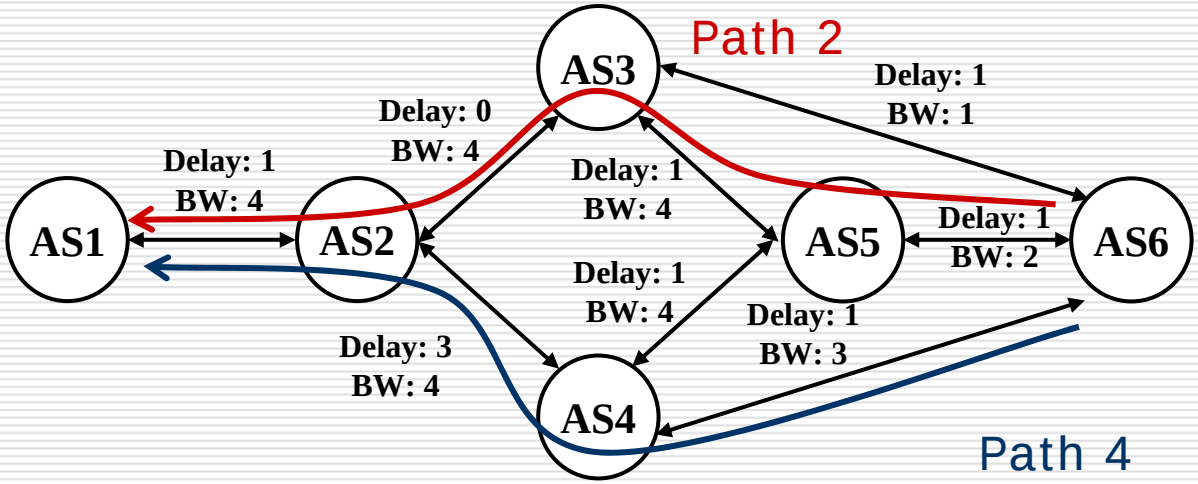

Dominant Paths

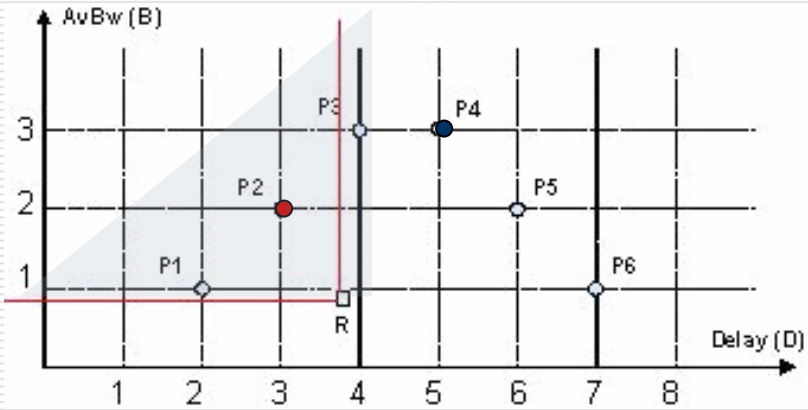

Route Pinning

- Routers have multiple paths to a destination prefix, packets should follow QoS compliant path
- A set of network-wide traffic classes with different QoS requirements is predefined
- Forwarding decision is based on packet destination address and class identifier stored in fields
  - DS Field in IP header

Route Pinning – Our Approach

- For each destination prefix in routing table, every traffic class is assigned to at most one path.
  - Class-assignment information at each border router is also injected into IGP routers
- Forwarding decision is based on packet destination address and class identifier stored in fields
  - DS Field in IP header

Class Assignment Algorithm

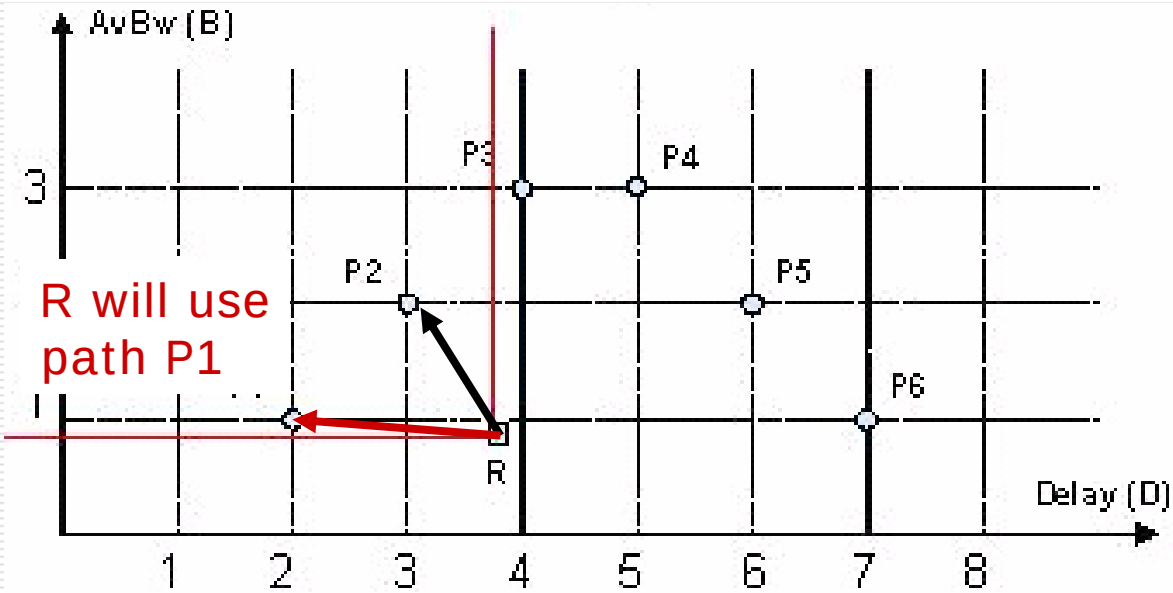

Changes to BGP route decision process

- QoS routing needs routes from all neighbors to ideally be enabled.
- All enabled routes in Loc-RIB undergo DPSA before Output Policy Engine.
- All enabled routes in Loc-RIB undergo DPSA and class-assignment algorithm before FIB.

Changes to BGP packet

- Need to extend AS_PATH attribute to store QoS attributes
- Modeled after TLV (Type-Length-Value) model

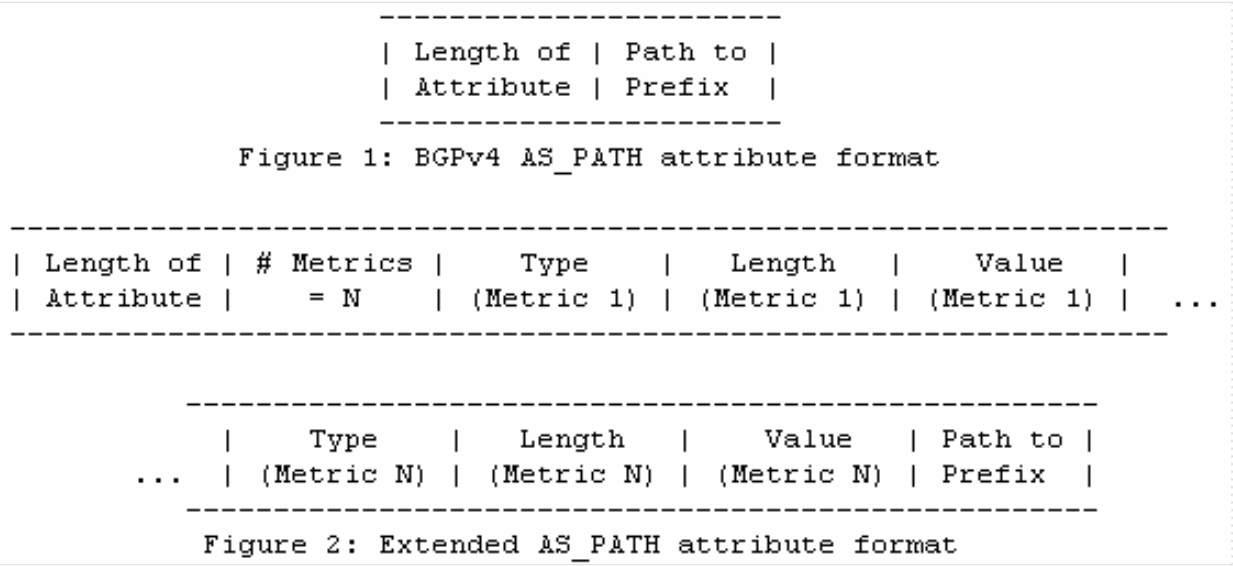

Preliminary simulation results

- ns-2 simulations of GIG-like topology
  - Vary network size
  - Link in dominant path is removed (phase 2), re-added (phase 3)
- Metrics
  - Convergence time
  - Number of updates

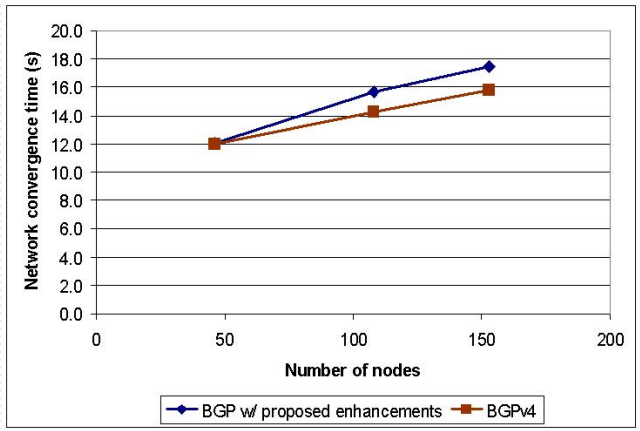

# Nodes	BGPv4	BGP w/ proposed enhancements		
	Phase 1 #msgs	Phase 1 #msgs	Phase 2 #msgs	Phase 3 #msgs
48	38174	38188	2790	1477
108	307469	335729	2384	1943
153	468398	484657	1081	1757

Questions?

Class Assignment Algorithm

- Class-assignment algorithm matches at most one path to each traffic class under each destination prefix in routing table.
- From the set of paths satisfying a traffic class' QoS requirement, the algorithm chooses the path offering the best QoS service.
  - In case of two QoS metrics, the chosen path would be the "furthest" path.