Backbone Networks

Mike Freedman

COS 461: Computer Networks

http://www.cs.princeton.edu/courses/archive/spr14/cos461/

Networking Case Studies

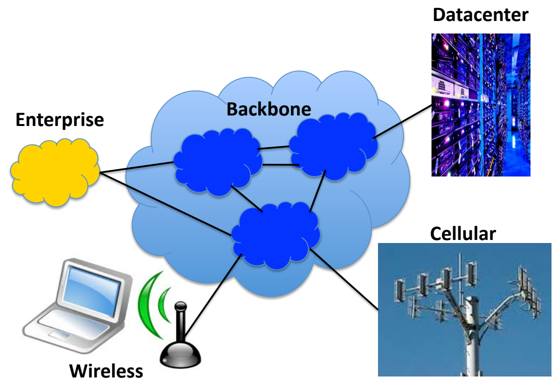

2

Backbone Topology

3

Backbone Networks

- Backbone networks
  - Multiple Points-of-Presence (PoPs)
  - Lots of communication between PoPs
  - Accommodate traffic demands and limit delay

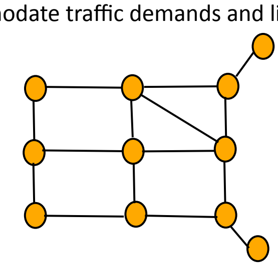

4

Abilene Internet2 Backbone

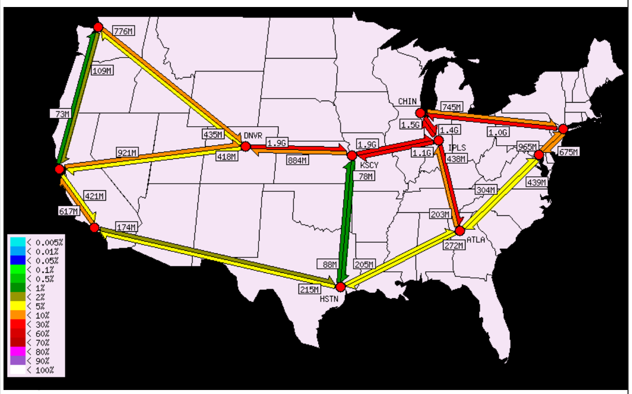

Points-of-Presence (PoPs)

- Inter-PoP links
  - Long distances
  - High bandwidth
- Intra-PoP links
  - Short cables between racks or floors
  - Aggregated bandwidth
- Links to other networks
  - Wide range of media and bandwidth

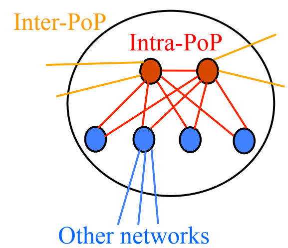

6

Where to Locate Nodes and Links

- Placing Points-of-Presence (PoPs)
  - Large population of potential customers
  - Other providers or exchange points
  - Cost and availability of real-estate
  - Mostly in major metropolitan areas (“NFL cities”)
- Placing links between PoPs
  - Already fiber in the ground
  - Needed to limit propagation delay
  - Needed to handle the traffic load

7

Peering

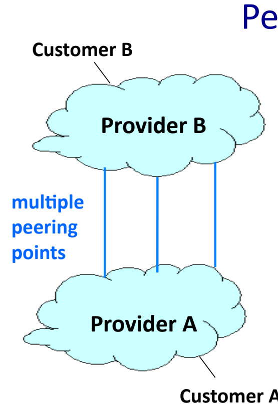

- Exchange traffic between customers
  - Settlement-free
- Diverse peering locations
  - Both coasts, and middle
- Comparable capacity at all peering points
  - Can handle even load

8

Combining Intradomain and Interdomain Routing

9

Intradomain Routing

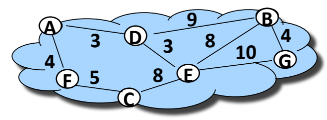

- Compute shortest paths between routers
  – Router C takes path C-F-A to router A
- Using link-state routing protocols
  – E.g., OSPF, IS-IS

10

Interdomain Routing

- Learn paths to remote destinations
  – AT&T learns two paths to Yale
- Applies local policies to select a best route

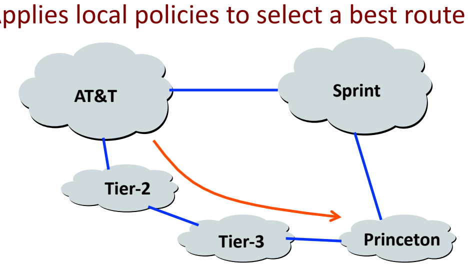

11

An AS is Not a Single Node

- Multiple routers in an AS
  – Need to distribute BGP information within the AS
  – Internal BGP (iBGP) sessions between routers

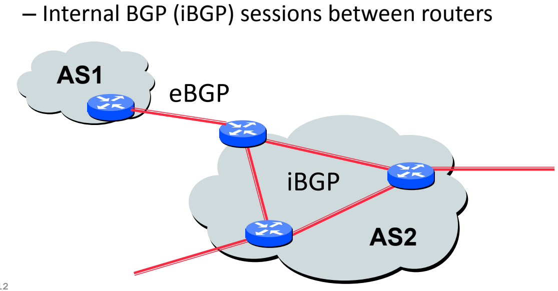

12

Internal BGP and Local Preference

- Both routers prefer path through AS 100
- … even though right router learns external path

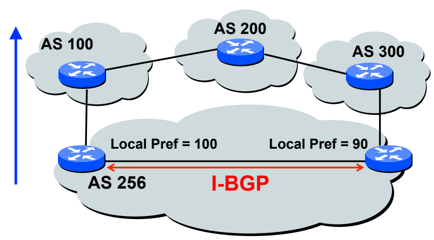

Hot-Potato (Early-Exit) Routing

- Hot-potato routing
  - Each router selects the closest egress point
  - … based on the path cost in intradomain protocol
- BGP decision process
  - Highest local preference
  - Shortest AS path
  - Closest egress point
  - Arbitrary tie break

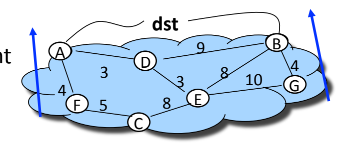

Hot-Potato Routing

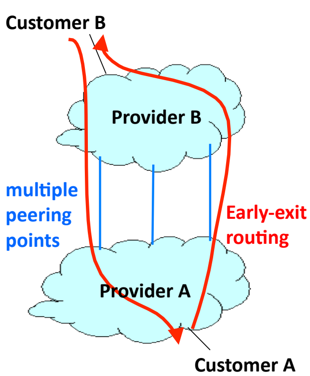

- Selfish routing
  - Each provider dumps traffic on the other
  - As early as possible
- Asymmetric routing
  - Traffic does not flow on same path in both directions

Joining BGP and IGP Information

- Border Gateway Protocol (BGP)
  - Announces reachability to external destinations
  - Maps a destination prefix to an egress point
    - 128.112.0.0/16 reached via 192.0.2.1
- Interior Gateway Protocol (IGP)
  - Used to compute paths within the AS
  - Maps an egress point to an outgoing link
    - 192.0.2.1 reached via 10.1.1.1

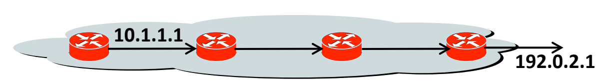

Joining BGP with IGP Information

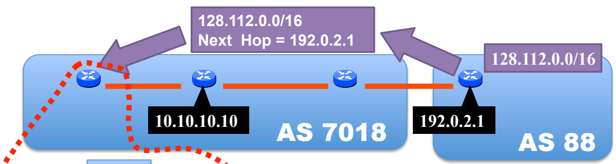

IGP

destination	next hop
192.0.2.0/30	10.10.10.10

\+

BGP

destination	next hop
128.112.0.0/16	192.0.2.1

(A) True (B) False

- The FIB of internal routers are of size O(all dest prefixes known to ISP)
- The FIB of internal routers point to border router to neighbor ISP

17

Joining BGP with IGP Information

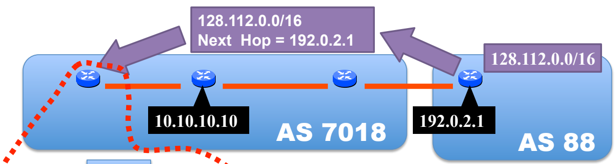

IGP

destination	next hop
192.0.2.0/30	10.10.10.10

\+

BGP

destination	next hop
128.112.0.0/16	192.0.2.1

Forwarding Table

destination	next hop
128.112.0.0/16	10.10.10.10
192.0.2.0/30	10.10.10.10

18

Interdomain Routing Policy

19

Selecting a Best Path

- Routing Information Base
  - Store all BGP routes for each destination prefix
  - Withdrawal: remove the route entry
  - Announcement: update the route entry
- BGP decision process
  - Highest local preference
  - Shortest AS path
  - Closest egress point
  - Arbitrary tie break

20

Import Policy: Local Preference

- Favor one path over another
  - Override the influence of AS path length
- Example: prefer customer over peer

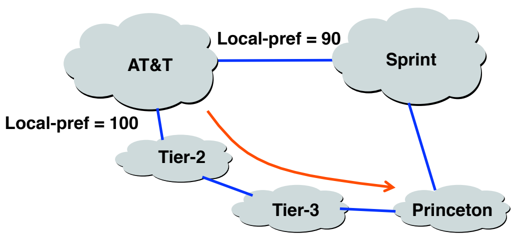

21

Import Policy: Filtering

- Discard some route announcements
  - Detect configuration mistakes and attacks
- Examples on session to a customer
  - Discard route if prefix not owned by the customer
  - Discard route with other large ISP in the AS path

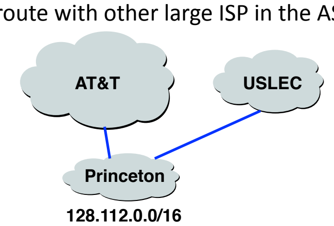

22

Export Policy: Filtering

- Discard some route announcements
  - Limit propagation of routing information
- Examples
  - Don't announce routes from one peer to another
  - Don't announce routes for management hosts

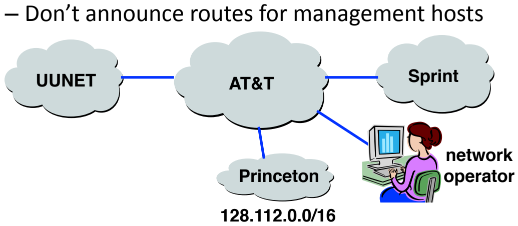

23

Export Policy: Attribute Manipulation

- Modify attributes of the active route
  - To influence the way other ASes behave
- Example: AS prepending
  - Artificially inflate AS path length seen by others
  - Convince some ASes to send traffic another way

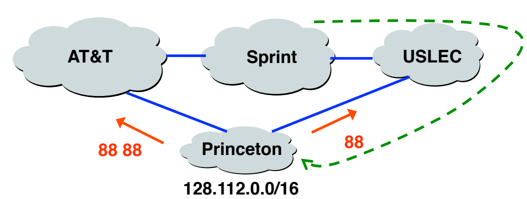

24

Business Relationships

- Common relationships
  - Customer-provider
  - Peer-peer
  - Backup, sibling, …
- ISP terminology:
  - Tier-1 (~15 worldwide): No settlement or transit
  - Tier-2 ISPs: Widespread peering, still buy transit
- Policies implementing in BGP, e.g.,
  - Import: Ranking customer routes over peer routes
  - Export: Export only customer routes to peers and providers

BGP Policy

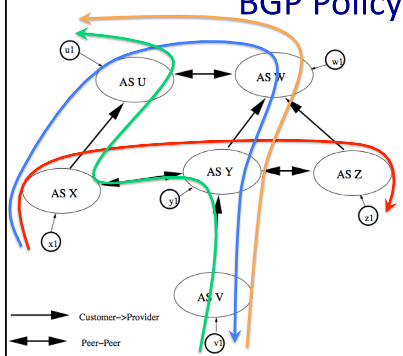

Tier 1 ISPs?

A. U, W
B. U, X
C. X, Y, Z

Which path may packets take (given commercial policies)?

A. Red
B. Blue
C. Green
D. Orange

BGP Policy Configuration

- Routing policy languages are vendor-specific
  - Not part of the BGP protocol specification
  - Different languages for Cisco, Juniper, etc.
- Still, all languages have some key features
  - List of clauses matching on route attributes
  - … and discarding or modifying the matching routes
- Configuration done by human operators
  - Implementing the policies of their AS
  - Business relationships, traffic engineering, security

Backbone Traffic Engineering

Routing With "Static" Link Weights

- Routers flood information to learn topology
  - Determine "next hop" to reach other routers…
  - Compute shortest paths based on link weights
- Link weights configured by network operator

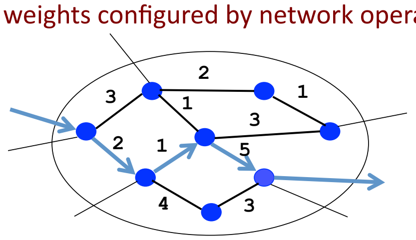

29

Setting the Link Weights

- How to set the weights
  - Inversely proportional to link capacity?
  - Proportional to propagation delay?
  - Network-wide optimization based on traffic?

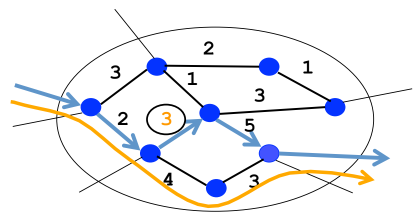

30

Measure, Model, and Control

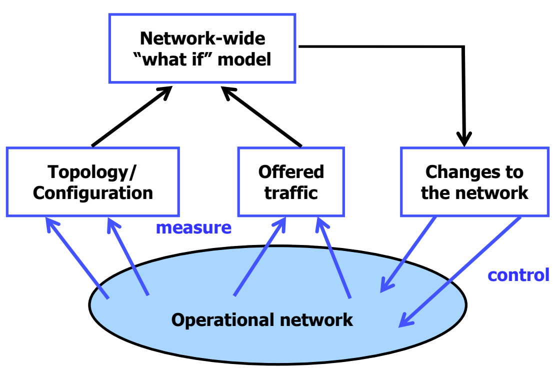

31

Limitations of Shortest-Path Routing

- Sub-optimal traffic engineering
  - Restricted to paths expressible as link weights
- Limited use of multiple paths
  - Only equal-cost multi-path, with even splitting
- Disruptions when changing the link weights
  - Transient packet loss and delay, and out-of-order
- Slow adaptation to congestion
  - Network-wide re-optimization and configuration
- Overhead of the management system

32

Constrained Shortest Path First

- Run a link-state routing protocol
  - Configurable link weights
  - Plus other metrics like available bandwidth
- Constrained shortest-path computation
  - Prune unwanted links (e.g., not enough bw)
  - Compute shortest path on the remaining graph

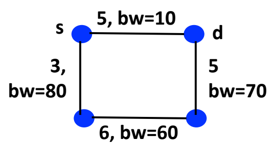

Constrained Shortest Path First

- Signal along the path
  - Source router sends msg to pin path to dest
  - Revisit decisions periodically, in case better options exist

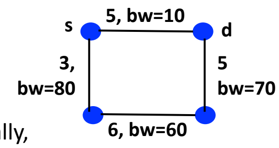

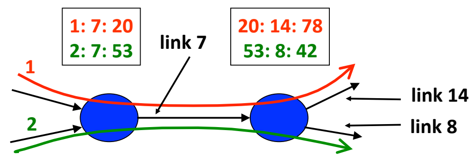

Challenges for Backbone Networks

Challenges

- Routing protocol scalability
  - Thousands of routers
  - Hundreds of thousands of address blocks
- Fast failover
  - Slow convergence disrupts user performance
  - Backup paths for faster recovery
  - E.g., backup path around a failed link

Challenges

- Router configuration
  - Adding customers, planned maintenance, traffic engineering, access control, …
  - Manual configuration is very error prone
- Measurement
  - Measuring traffic, performance, routing, etc.
  - To detect attacks, outages, and anomalies
  - To drive traffic-engineering decisions

37

Challenges

- Diagnosing performance problems
  - Incomplete control and visibility
  - Combining measurement data
- Security
  - Defensive packet and route filtering
  - Detecting and blocking denial-of-service attacks
  - DNS security, detecting and blocking spam, etc.
- New services
  - IPv6, IPTV, …

38

Conclusions

- Backbone networks
  - Transit service for customers
  - Glue that holds the Internet together
- Routing challenges
  - Interdomain routing policy
  - Intradomain traffic engineering

39