CS 204: BGP

Jiasi Chen

Lectures: MWF 12:10-1pm in WCH 139

http://www.cs.ucr.edu/~jiasi/teaching/cs204_spring16/

Overview

- Inter-AS routing
  - BGP
  - Forwarding example
  - AS hierarchy
- Paper discussion

Autonomous Systems

- Autonomous system (AS)
  - Unit of routing policy
  - ~50k ASes in use
  - E.g., UCR has AS#6061, AT&T has AS#144, Princeton has AS#88

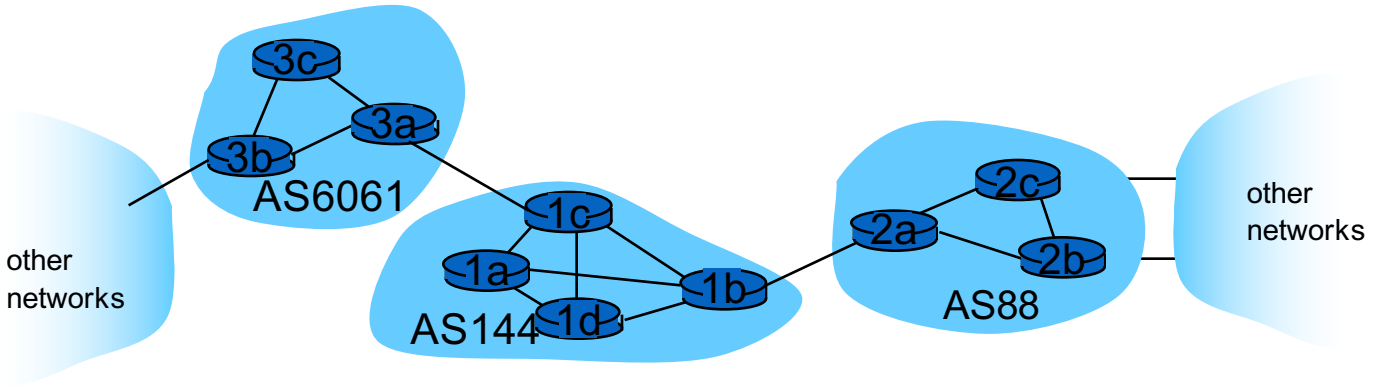

Review of Routing

- Inter-AS routing
  - BGP
- Intra-AS routing
  - RIP
  - OSPF

Link-state?

Distance vector?

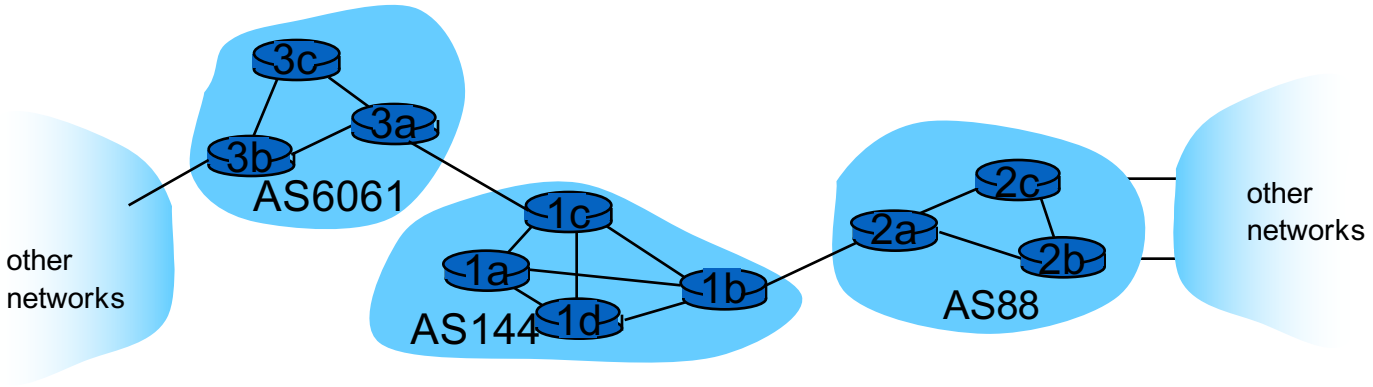

Why different Intra-, Inter-AS routing ?

policy:

- inter-AS: admin wants control over how its traffic routed, who routes through its net.
- intra-AS: single admin, so no policy decisions needed

scale:

- hierarchical routing saves table size, reduced update traffic

performance:

- intra-AS: can focus on performance
- inter-AS: policy may dominate over performance

Interconnected ASes

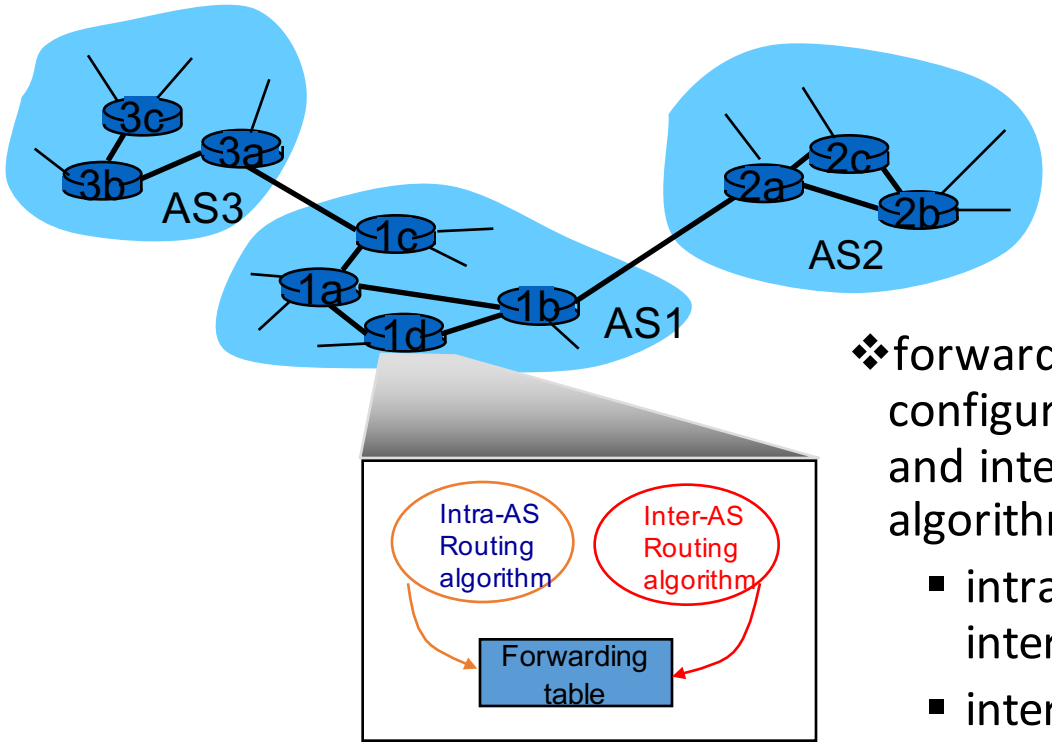

- forwarding table configured by both intra- and inter-AS routing algorithm
  - intra-AS sets entries for internal dests
  - inter-AS & intra-AS sets entries for external dests

Example: setting forwarding table in router 1d

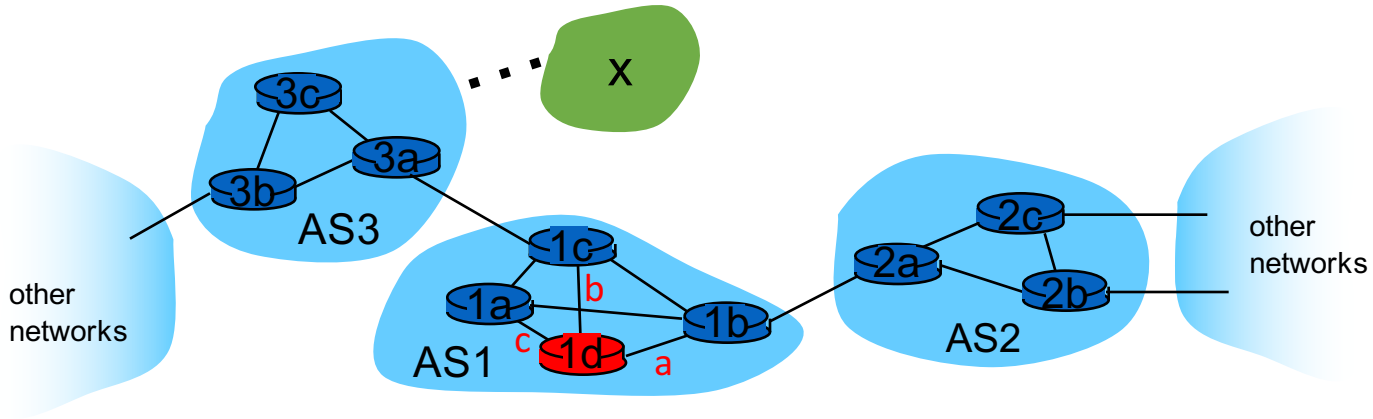

Send packet on interface a, b, or c?

Determine b is on least cost path, install forwarding table entry (x,b)

Example: choosing among multiple ASes

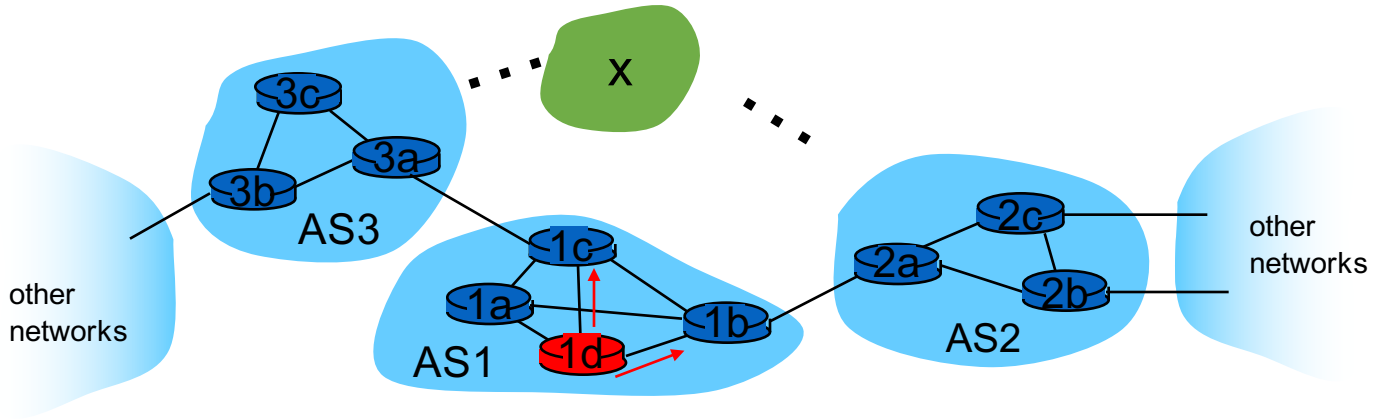

Which path to choose?

This is also the job of the inter-AS protocol

Example: choosing among multiple ASes

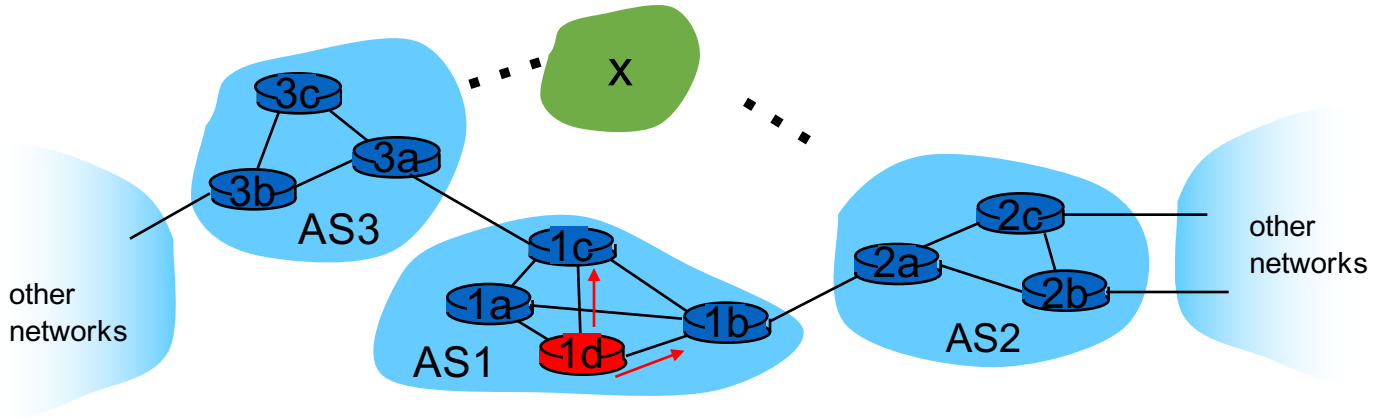

learn from inter-AS protocol that subnet *x* is reachable via multiple gateways

use routing info from intra-AS protocol to determine costs of least-cost paths to each of the gateways

hot potato routing: choose the gateway that has the smallest least cost

determine from forwarding table the interface that leads to least-cost gateway. Enter that interface in forwarding table

Overview

- Inter-AS routing
  - BGP
  - Forwarding example
  - AS hierarchy
- Paper discussion

Internet inter-AS routing: BGP

- BGP (Border Gateway Protocol): *the* de facto inter-domain routing protocol
  - "glue that holds the Internet together"
- BGP provides each AS a means to:
  - eBGP: obtain subnet reachability information from neighboring ASs.
  - iBGP: propagate reachability information to all AS-internal routers.
  - determine "good" routes to other networks based on reachability information and policy.
- allows subnet to advertise its existence to rest of Internet: *"I am here"*

BGP basics

- BGP session: two BGP routers ("peers") exchange BGP messages:
  - advertising *paths* to different destination network prefixes ("path vector" protocol)
  - exchanged over semi-permanent TCP connections

- when AS3 advertises a prefix to AS1:
  - AS3 *promises* it will forward datagrams towards that prefix
  - AS3 can aggregate prefixes in its advertisement

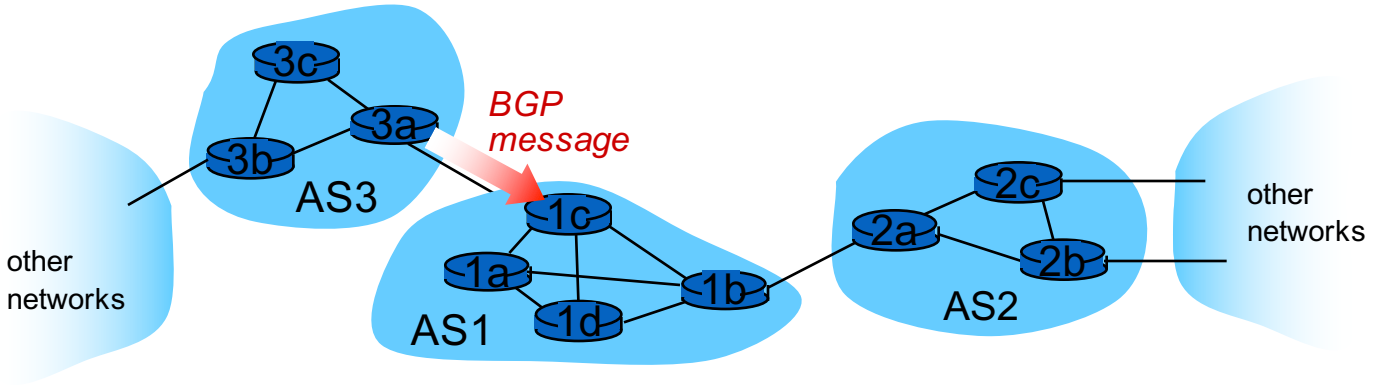

Path attributes and BGP routes

- advertised prefix includes BGP attributes
  - prefix + attributes = "route"
- two important attributes:
  - AS-PATH: contains ASs through which prefix advertisement has passed
  - NEXT-HOP: indicates specific internal-AS router to next-hop AS. (may be multiple links from current AS to next-hop-AS)

BGP basics: distributing path information

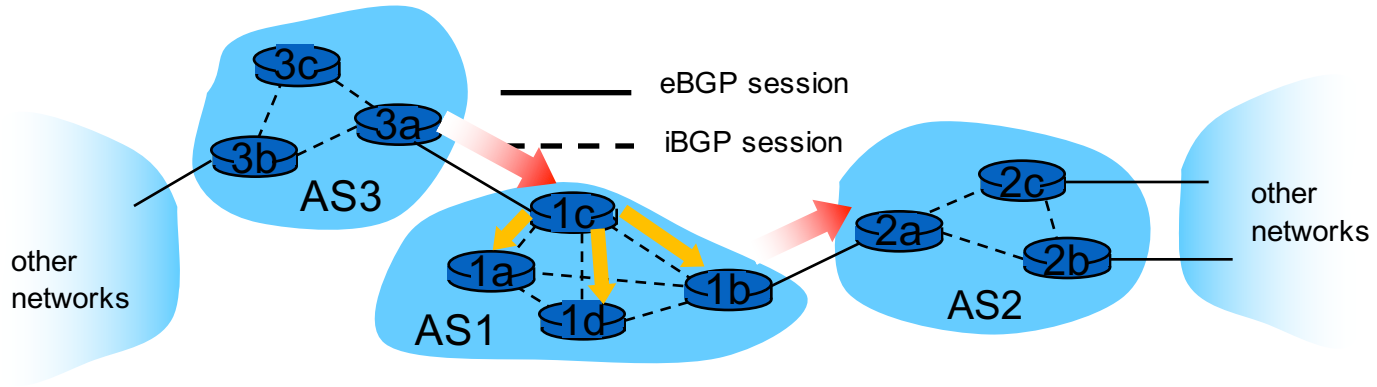

BGP route selection

- router may learn about more than 1 route to destination AS, selects route based on:
  1. local preference value attribute: policy decision
  2. shortest AS-PATH
  3. closest NEXT-HOP router: hot potato routing
  4. additional criteria

- gateway router receiving route advertisement uses import policy to accept/decline
  - e.g., never route through AS x
  - *policy-based* routing

BGP messages

- BGP messages exchanged between peers over TCP connection
- BGP messages:
  - OPEN: opens TCP connection to peer and authenticates sender
  - UPDATE: advertises new path (or withdraws old)
  - KEEPALIVE: keeps connection alive in absence of UPDATES; also ACKs OPEN request
  - NOTIFICATION: reports errors in previous msg; also used to close connection

Overview

- Inter-AS routing
  - BGP
  - Forwarding example
  - AS hierarchy
- Paper discussion

How does entry get in forwarding table?

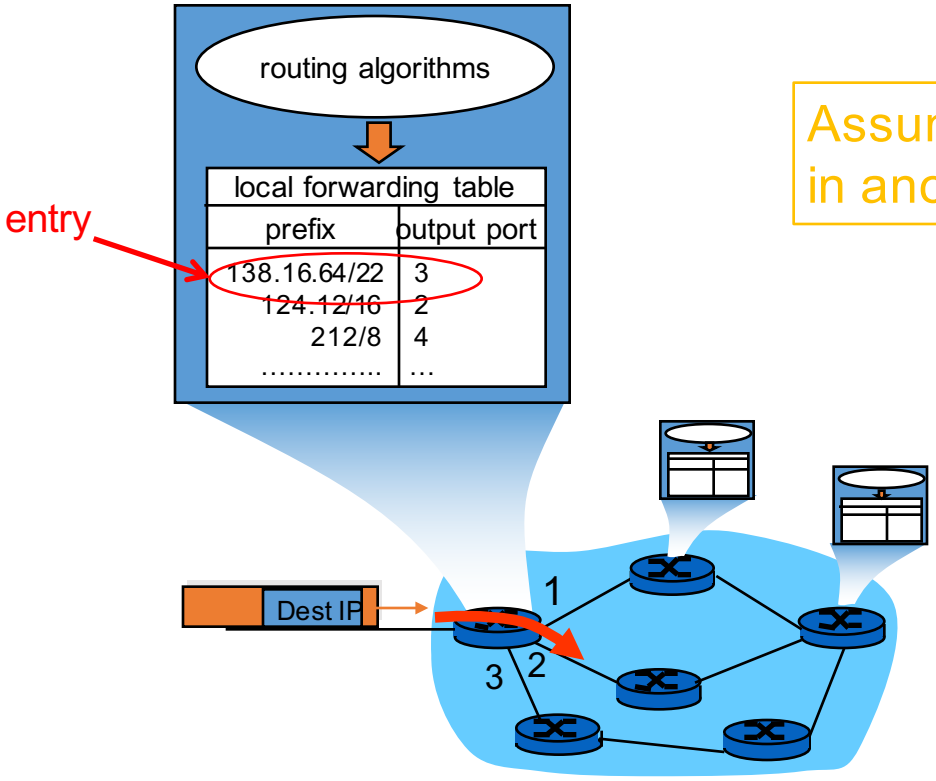

Assume prefix is in another AS.

How does entry get in forwarding table?

High-level overview

1. Router becomes aware of prefix
2. Router determines output port for prefix
3. Router enters prefix-port in forwarding table

Router becomes aware of prefix

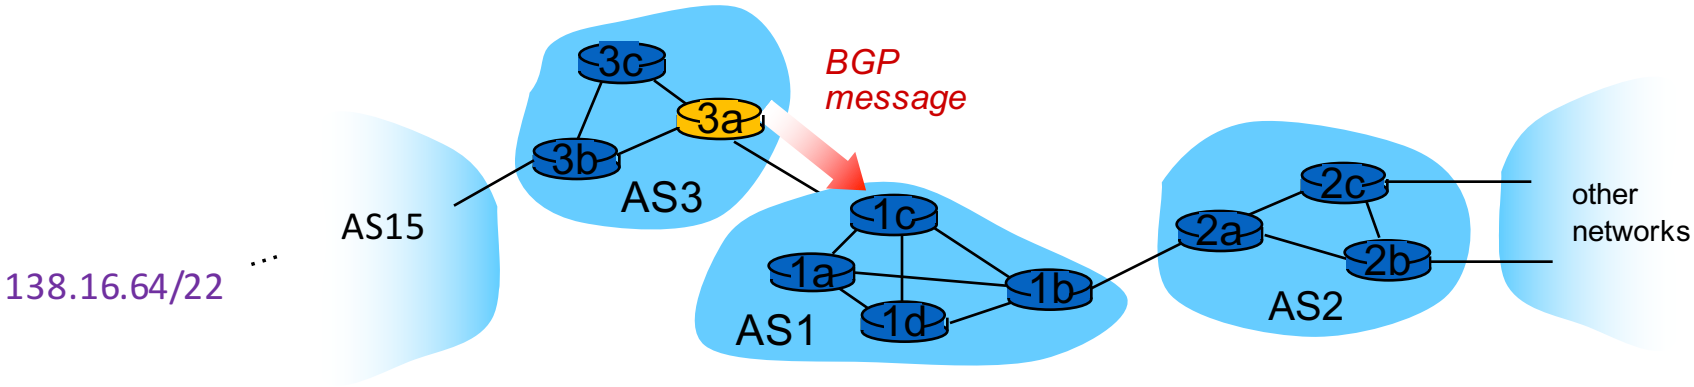

- BGP message contains "routes"
- "route" is a prefix and attributes: AS-PATH, NEXT-HOP,…
- Example
  - Prefix: 138.16.64/22
  - AS-PATH: AS3 AS15 …
  - NEXT-HOP: 201.44.13.125

Router may receive multiple routes

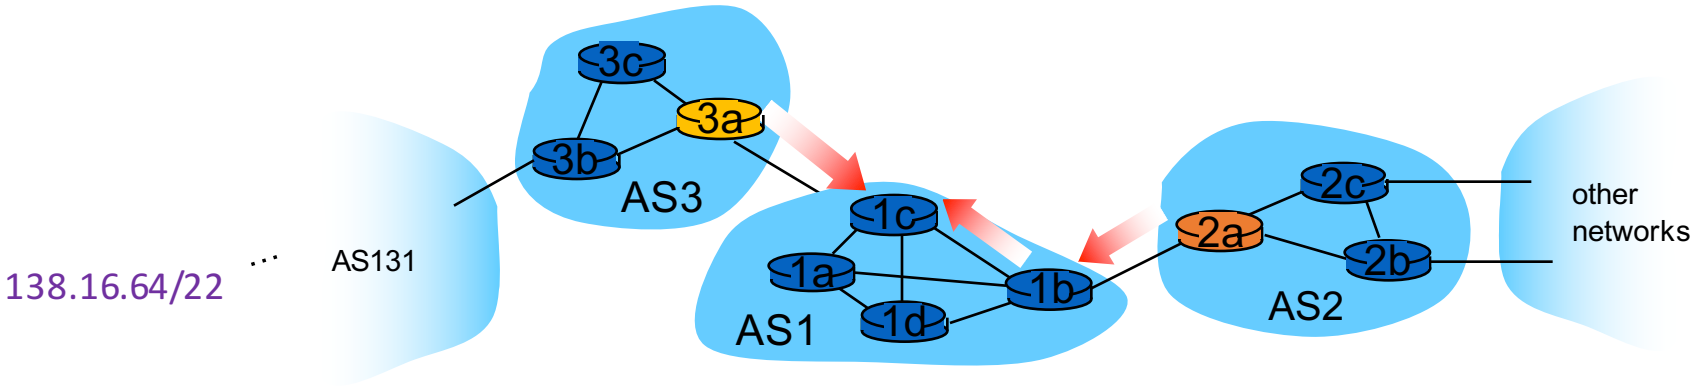

- Router may receive multiple routes for same prefix
- Has to select one route

Select best BGP route to prefix

- Router selects route based on shortest AS-PATH

- Example:
  - AS2 AS17 to 138.16.64/22
  - AS3 AS131 AS201 to 138.16.64/22

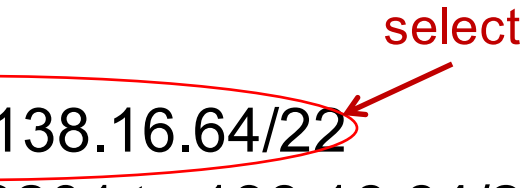

- What if there is a tie? We'll come back to that!

Find best intra-route to BGP route

- Use selected route's NEXT-HOP attribute
  - Route's NEXT-HOP attribute is the IP address of the router interface that begins the AS PATH.
- Example:
  - AS-PATH: AS2 AS17 …; NEXT-HOP: 111.99.86.55
- Router uses OSPF to find shortest path from 1c to 111.99.86.55

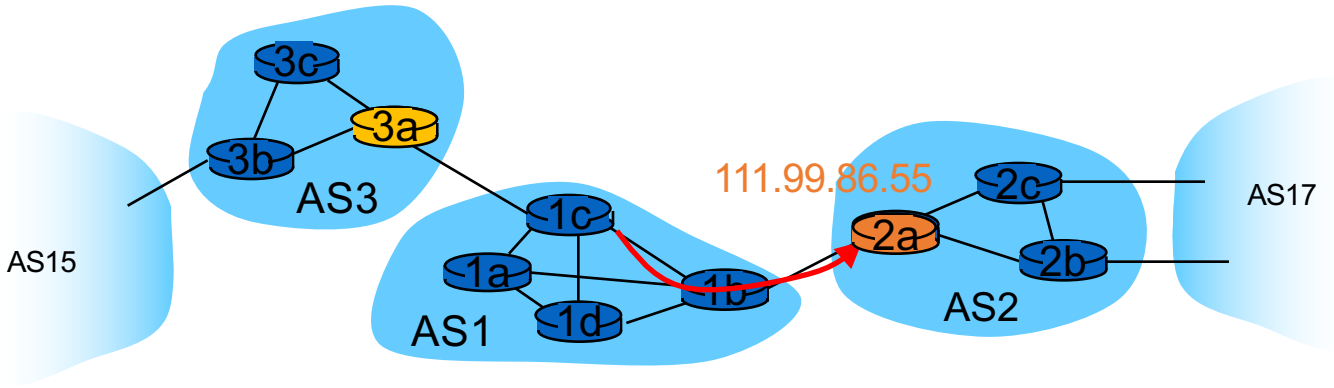

Router identifies port for route

- Identifies port along the OSPF shortest path
- Adds prefix-port entry to its forwarding table:
  - (dest IP , port 4)

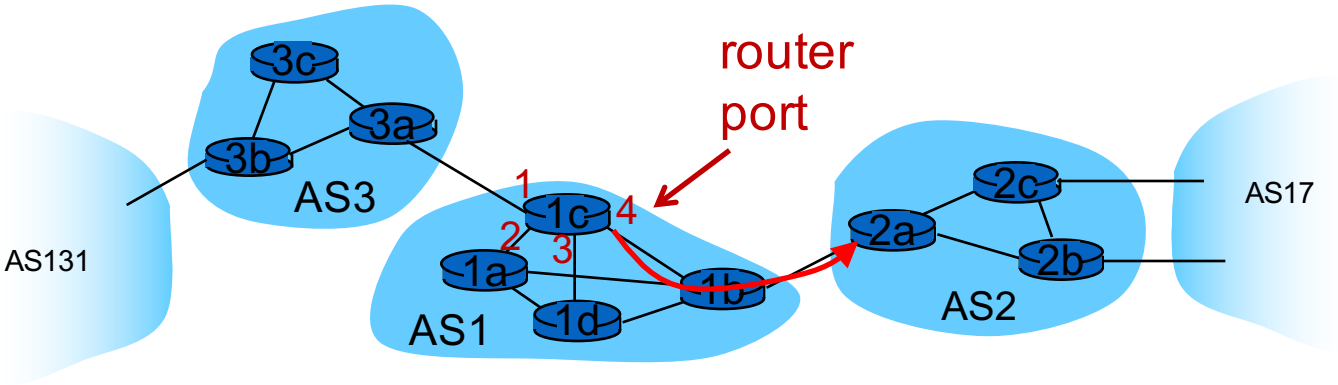

Hot Potato Routing

- Suppose there two or more best inter-routes.
- Then choose route with closest NEXT-HOP
  - Use OSPF to determine which gateway is closest
  - Q: From 1c, chose AS3 AS131 or AS2 AS17?
  - A: route AS3 AS131 since it is closer

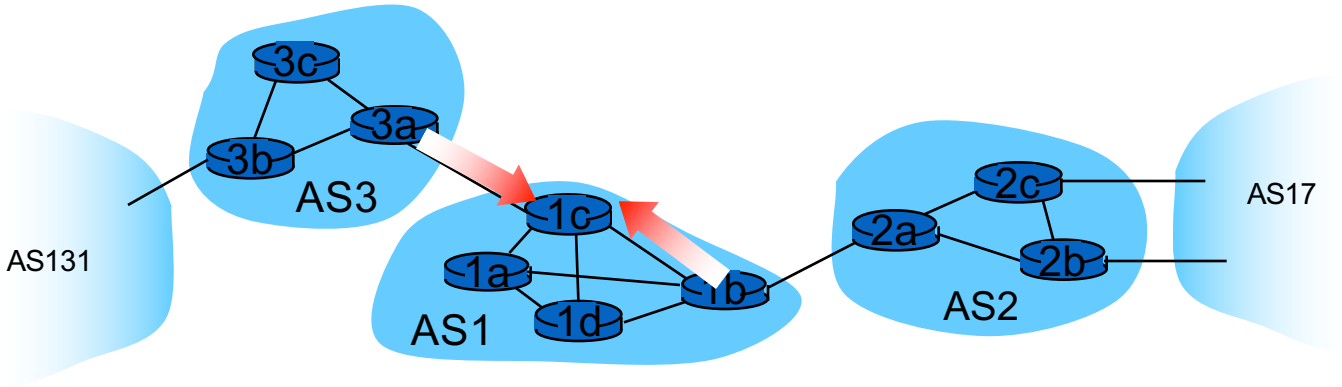

How does entry get in forwarding table?

Summary

1. Router becomes aware of prefix
   - via BGP route advertisements from other routers
2. Determine router output port for prefix
   - Use BGP route selection to find best inter-AS route
   - Use OSPF to find best intra-AS route leading to best inter-AS route
   - Router identifies router port for that best route
3. Enter prefix-port entry in forwarding table

Overview

- Inter-AS routing
  - BGP
  - Forwarding example
  - AS hierarchy
- Paper discussion

Overview

- Inter-AS routing
  - BGP
  - Forwarding example
  - AS hierarchy
- Paper discussion

BGP routing policy

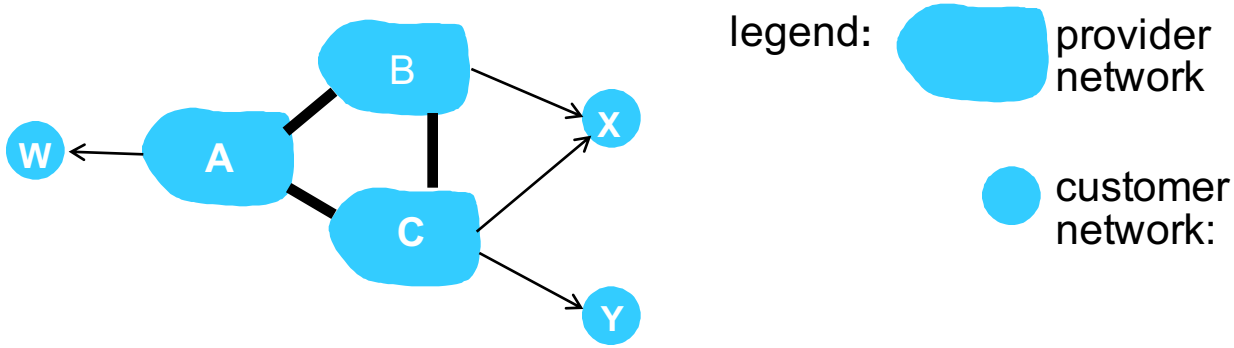

- A,B,C are *provider networks*
- X,W,Y are customer (of provider networks)
- X is *dual-homed:* attached to two networks
  - X does not want to route from B via X to C
  - .. so X will not advertise to B a route to C

BGP routing policy (2)

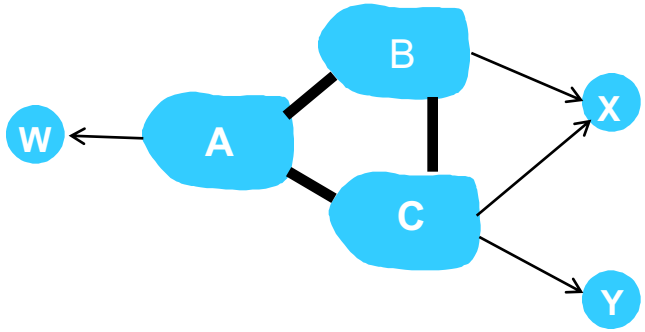

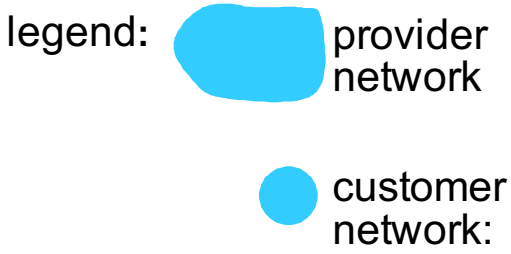

- A advertises path AW to B
- B advertises path BAW to X
- Should B advertise path BAW to C?
  - No way! B gets no "revenue" for routing CBAW since neither W nor C are B's customers
  - B wants to force C to route to w via A
  - B wants to route *only* to/from its customers!

Overview

- Inter-AS routing
  - BGP
  - Forwarding example
  - AS hierarchy
- Paper discussion

Stable BGP Routing

- Trying to get to destination A
- Routes listed in order of preference

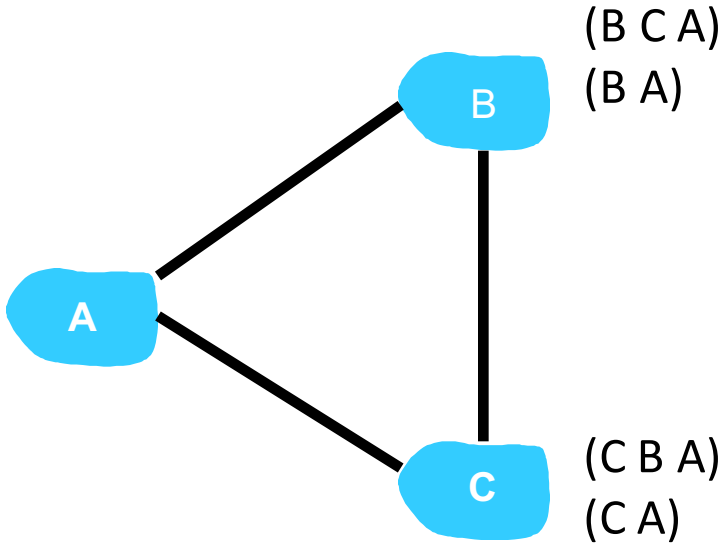

Stable BGP routing

- Suppose we start off with the second choice options…

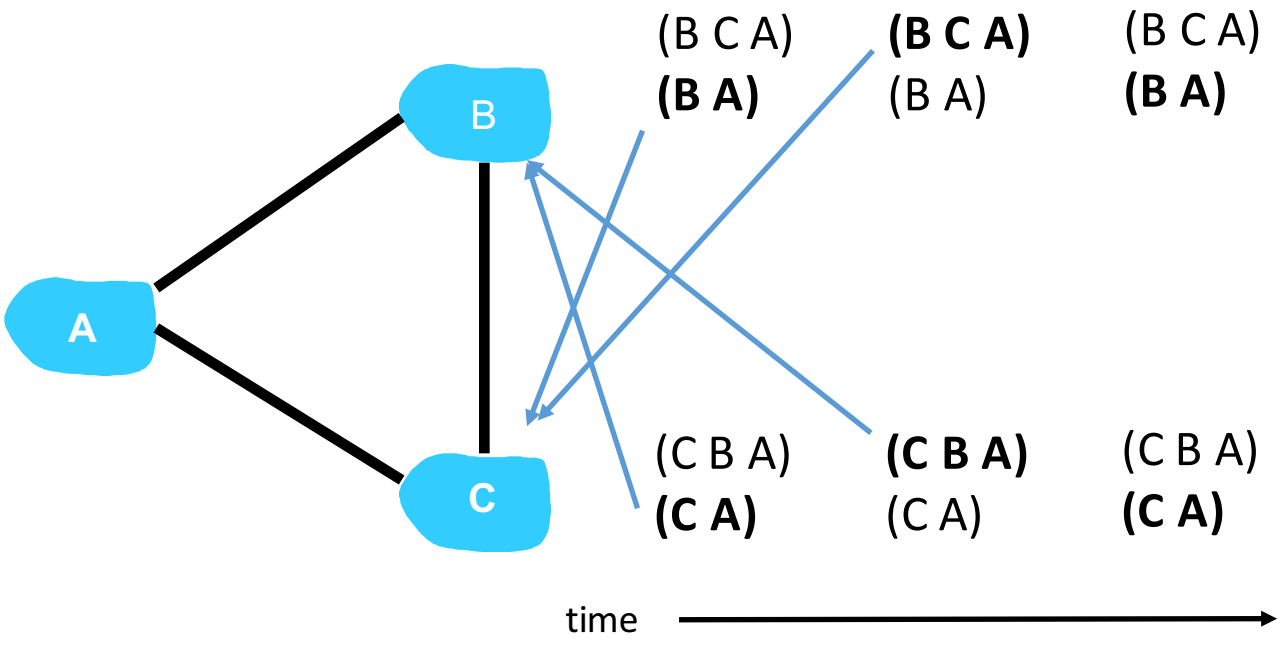

Paper Discussion

- What are implicit and explicit policies?
- What are the underlying assumptions?

Sources

- *Computer Networking: A Top-Down Approach*, Kurose & Ross
- Lixin Gao and Jennifer Rexford, "Stable Internet Routing Without Global Coordination," *IEEE Trans. Networking*, 2001.