BGP and BGP Security

16 March 2025
Lecture 1

Inter-AS Routing

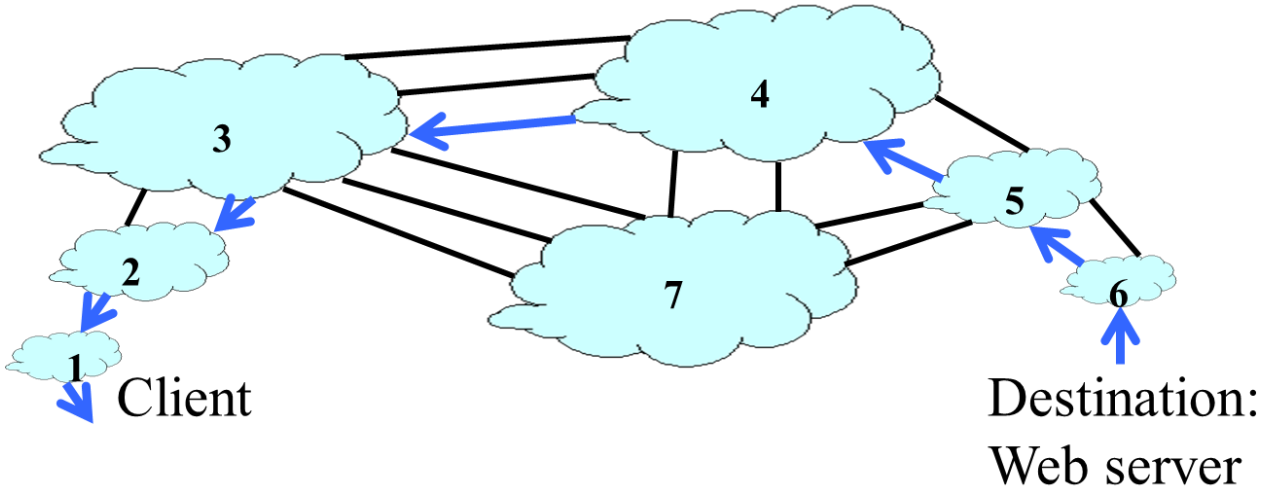

Idea: Provide an additional way to hierarchically aggregate routing information is a large internet.

We need to find routes to destinations

- What are destinations? IP Prefixes (12.X.X.X) (CIDR)
- What are nodes? AS (how many are there?)
- What are links? Connections and Business Relationships

Challenges for Inter-AS Routing

Scale (as of Apr 2024)

- Prefixes: 969,970 (no CIDR) or 542,956 (CIDR aggregated) and growing
- ASes: 75,852 visible ones, and growing
- Routers: at least in the millions…
- Border routers must know how to get anywhere in the world!

Source: http://www.cidr-report.org/as2.0/

Coordination with intra-AS protocols (OSPF)

- How to inject external routes to OSPF database

Challenges for Inter-AS Routing

Policy

- I want control over where I send traffic
- … and who send traffic through my AS *why?*
- AS don't want to expose internal topologies
- … or my business relations with neighbors

Trust:

- Provider A might be unwilling to believe advertisements from provider B
- See: http://www.cidr-report.org/as2.0/#Bogons

Example ASes (from cidr-report.org)

- AS11 HARVARD - Harvard University,US
- AS39 DNIC-AS-00039 - DoD Network Information Center,US
- AS5540 The Israel Electric Corporation Limited,IL
- AS5585 IIX-ASN Israel Internet Association,IL
- AS6810 BEZEK "Bezeq"- THE ISRAEL TELECOMMUNICATION CORP. LTD.,IL
- AS8738 VISA-ISRAEL-AS Israel Credit Cards Ltd,IL
- AS8867 TEHILA-AS Government of Israel, IL
- AS12736 IAA-AS Israel Airports Authority, IL
- AS21486 SYNAMEDIA-AS Synamedia Israel Technologies Ltd, IL
- AS34380 AMDOCS AMDOCS (ISRAEL) LTD,IL
- AS43423 ISRAEL-POST-LTD Israel Postal Company Ltd, IL
- AS1680 NV-ASN 013 NetVision Ltd., IL
- AS8584 BARAK Netvision 013 Barak - Barak Network, IL
- AS9117 CELLCOM-AS 013 NetVision Ltd, IL
- AS7432 EGENIUS - Evil Geniuses for a Better Tomorrow, US
- AS8551 BEZEQ-INTERNATIONAL-AS Bezeqint Internet Backbone, IL

Routing Requirements

Divide the routing problem in two parts:

- Routing within a single autonomous system
- Routing between autonomous systems

Two-level route propagation hierarchy

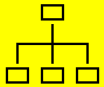

- Inter-domain routing protocol (Internet-wide standard)
- Intra-domain routing protocol (each AS selects its own)

Principle: Information hiding

Inter-AS Routing History: EGP

- Exterior Gateway Protocol (EGP) (RFC 904, 1984)

Forced tree-like topology

- Single backbone
- AS's connected only as parents and children (not peers)

Did not allow for the topology to become general

No aggregation

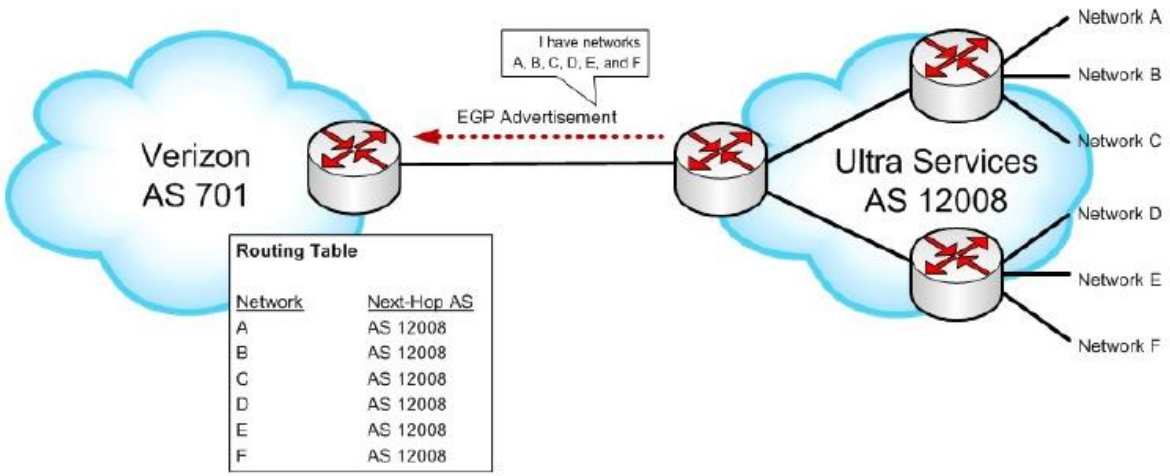

Image source https://www.networxsecurity.org/fileadmin/user_upload/images/2015-08/Exterior_gateway_protocol.jpg

Border Gateway Protocol (BGP)

Assumes the Internet is an arbitrarily interconnected set of ASs.

Today, the Internet consists of an interconnection of multiple backbone networks

Usually called *service provider networks*

Operated by private companies, not governments

Sites are connected to each other in arbitrary ways

Current version BGP-4 (RFC 4271)

Published 2006, some updates

Inter-AS Routing Options

Distance Vector Routing

Example: RIP

Problems:

- Distance (?)
- Loops
- Slow convergence (bad news travels slowly)
- Scalable?

Advantages:

- Hides total topology
- Nodes only know the “next hop”

Link State Routing

- Example: OSPF

Problems:

- Link costs
  - What’s cost? Distance? Business relationships?
- Shortest path (?)
  - Every node must agree on link cost algorithm
- Flooding
- High bandwidth and storage overhead
- Nodes must tell a lot about themselves
- Each node computes the whole network graph to make a spanning tree

Path Vector Routing

- Extension of distance-vector routing
  - Support flexible routing policies
  - Faster convergence (avoid count-to-infinity)
- Key idea: advertise the entire path
  - Distance vector: send *distance metric* per dest d
  - Path vector: send the *entire path* for each dest d

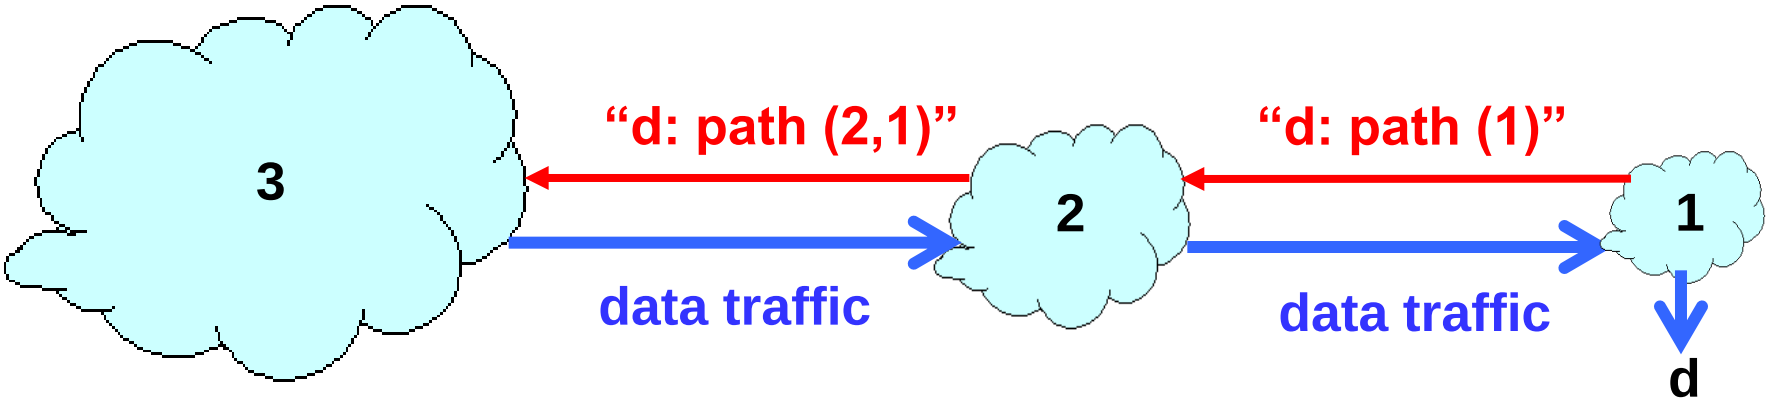

Faster Loop Detection

- Node can easily detect a loop
  - Look for its own node identifier in the path
  - E.g., node 1 sees itself in the path “3, 2, 1”
- Node can simply discard paths with loops
  - E.g., node 1 simply discards the advertisement

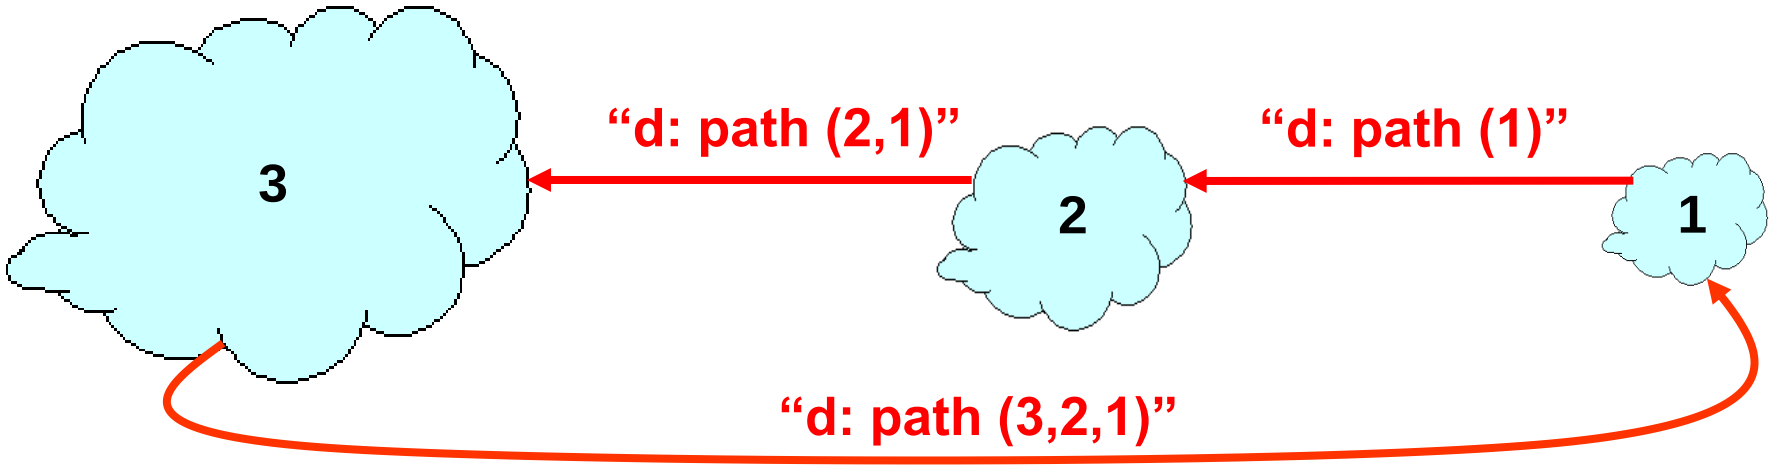

Flexible Policies

- Each node can apply local policies
  - Path selection: Which path to use?
  - Path export: Whether to advertise the path?
- Examples
  - Node 2 may prefer the path “2, 3, 1” over “2, 1”
  - Node 1 may not let node 3 hear the path “1, 2”

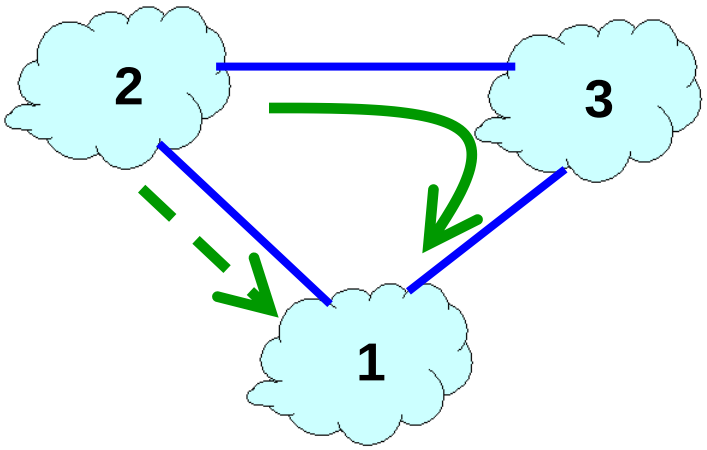

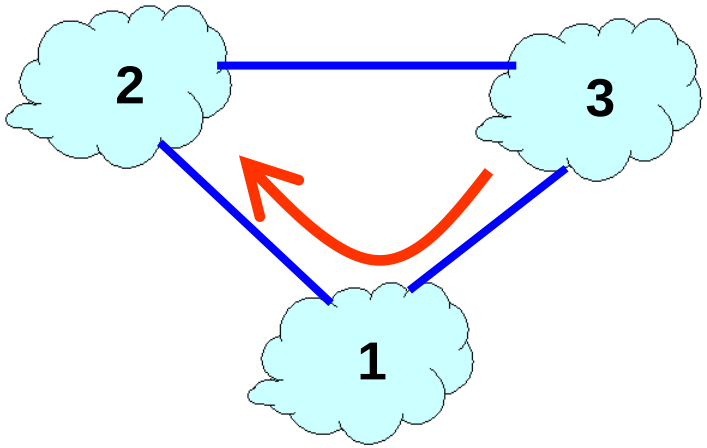

Path Vector Routing

- Extension of distance-vector routing
  - Support flexible routing policies
  - Faster convergence (avoid count-to-infinity)
- Key idea: advertise the entire path
  - Distance vector: send *distance metric* per dest d
  - Path vector: send the *entire path* for each dest d

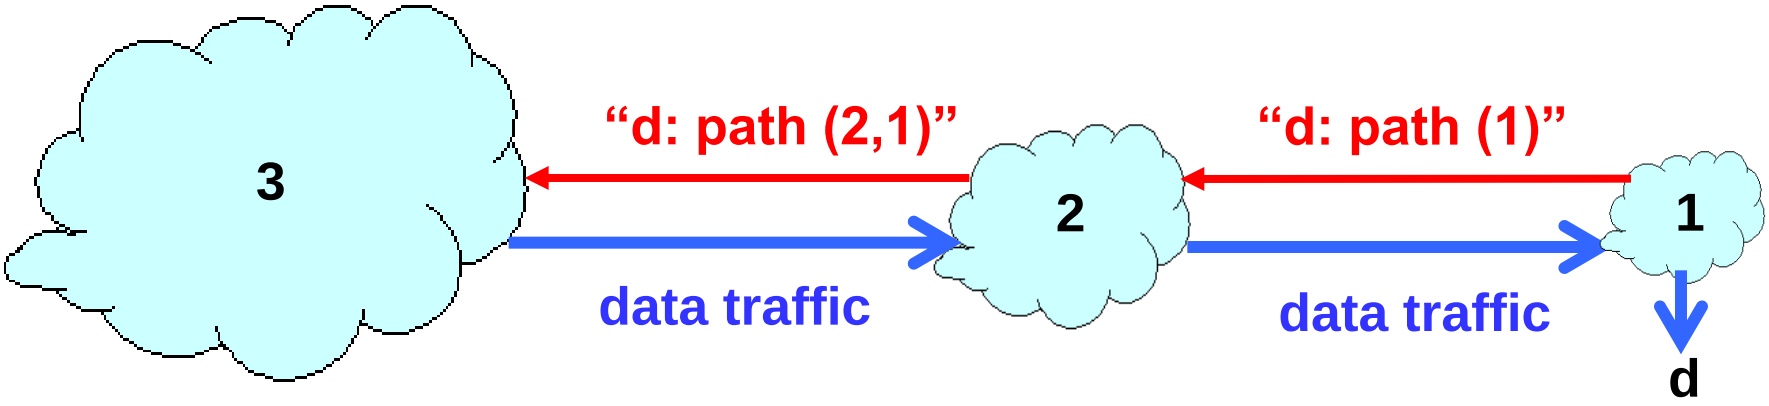

Faster Loop Detection

- Node can easily detect a loop
  - Look for its own node identifier in the path
  - E.g., node 1 sees itself in the path “3, 2, 1”
- Node can simply discard paths with loops
  - E.g., node 1 simply discards the advertisement

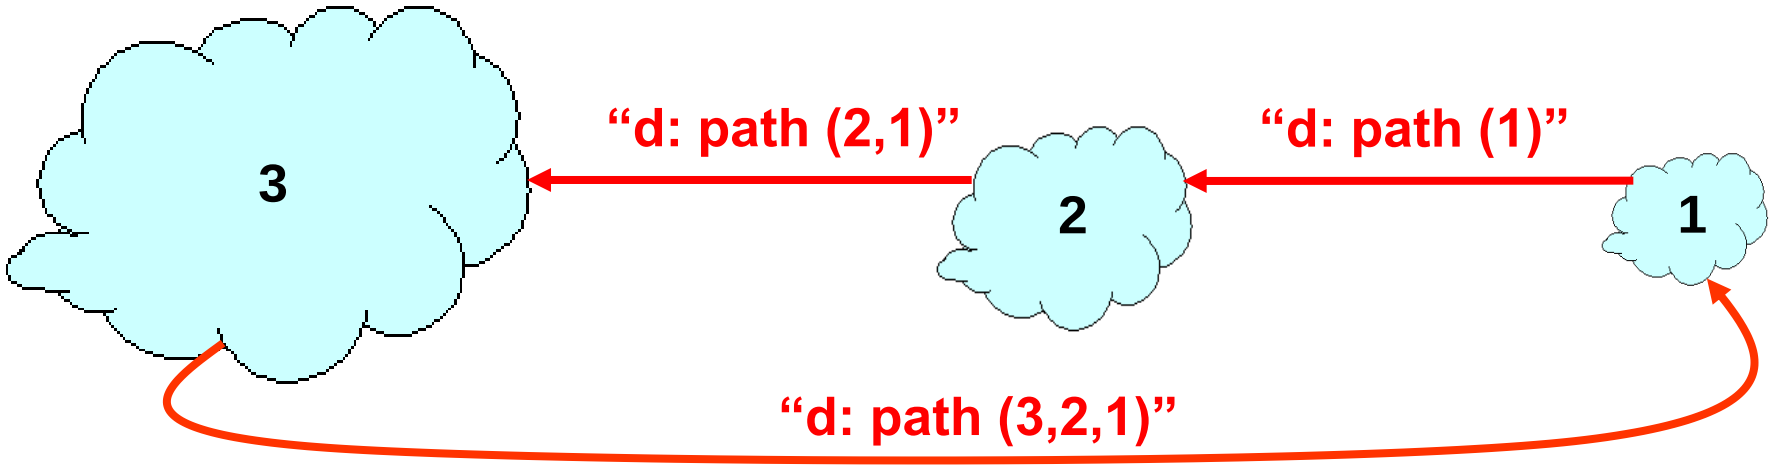

Flexible Policies

- Each node can apply local policies
  - Path selection: Which path to use?
  - Path export: Whether to advertise the path?
- Examples
  - Node 2 may prefer the path “2, 3, 1” over “2, 1”
  - Node 1 may not let node 3 hear the path “1, 2”

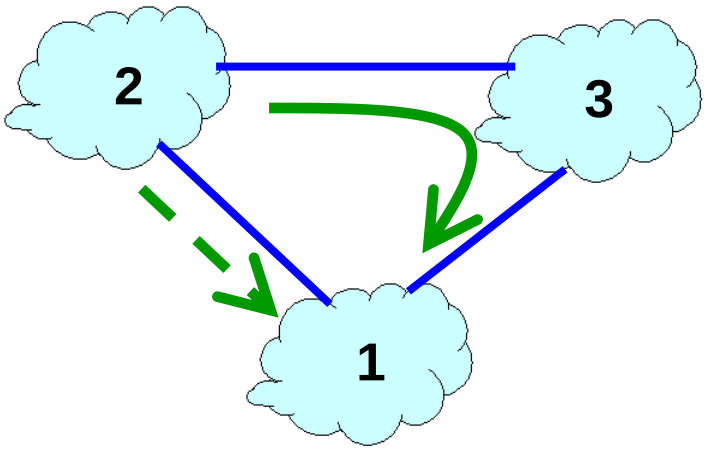

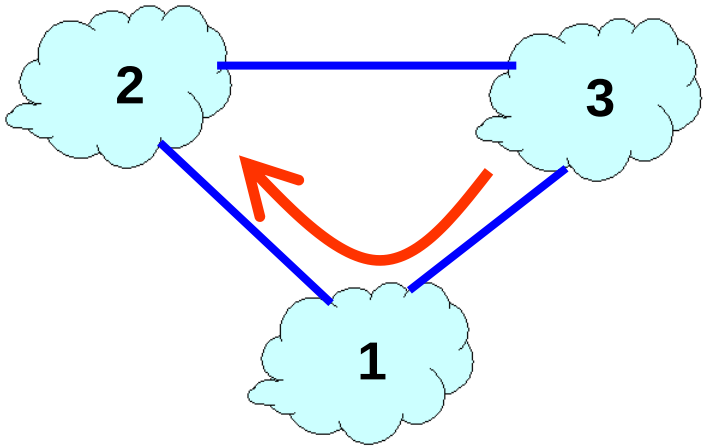

Path Vectors in BGP

Each AS has:

1+ BGP *speaker* that gives *path* information and advertises:

- local networks
- other reachable networks (transit AS only)

1+ border “gateways” which need not be the same as the speakers

- Border gateways are routers through which packets enter and leave the AS

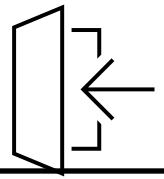

Distributing BGP Data

Long lived TCP Sessions between BGP Speakers (port 179)

- Exchange all active routes
- Exchange incremental updates (ALIVE messages, UPDATE)
  - Announce new routes (add IDs to new or existing paths)
  - Withdraw routes

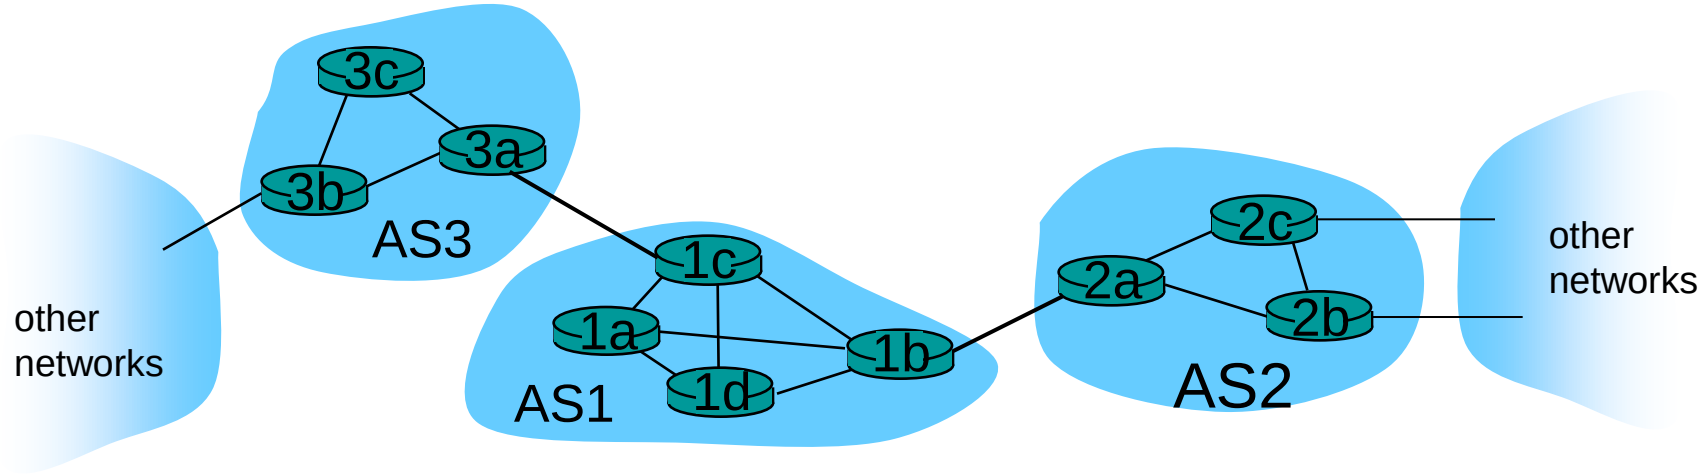

Distributing BGP Data

Internal BGP (iBGP): Sessions between routers in a single AS

- Ex. 1c talks to 1b to coordinate

External BGP (eBGP): Sessions between routers in different AS

- Ex. 1c talks to 3a

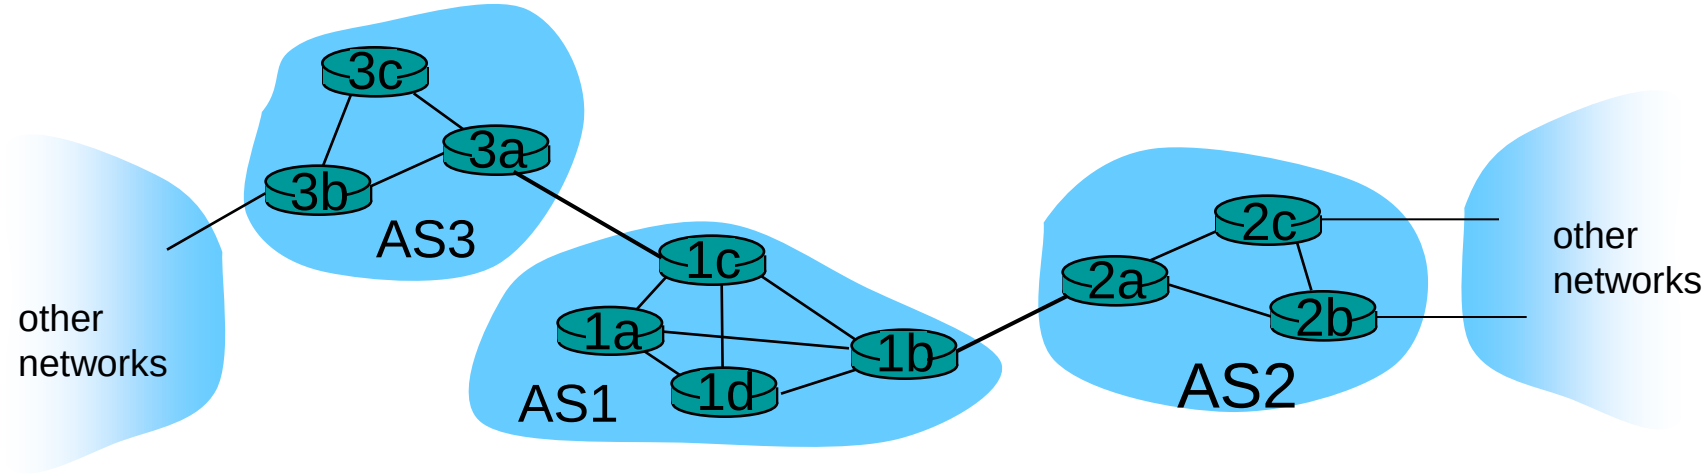

BGP Path Selection

- Simplest case: Shortest AS path
  - Example: 1129 to 88 on previous slides
  - Break ties by flipping a coin
- Hot potato routing: Leave via closest internal router
  - iBGP at work!

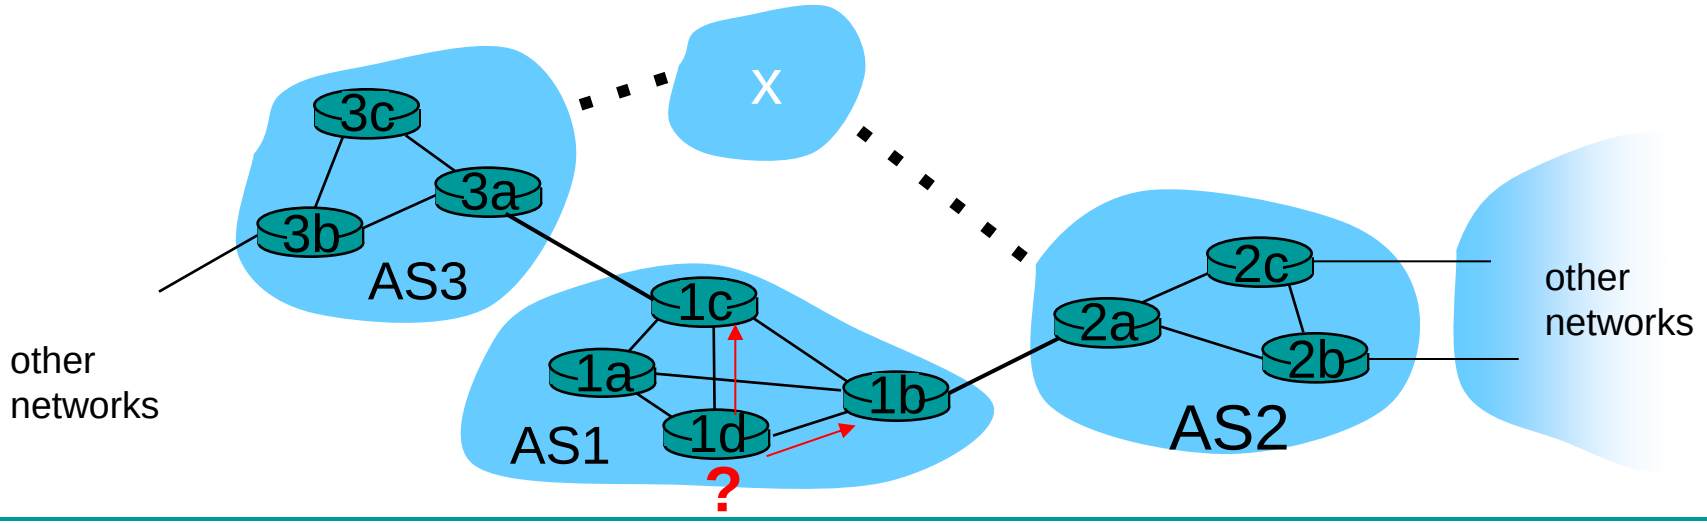

Policy Based Path Selection

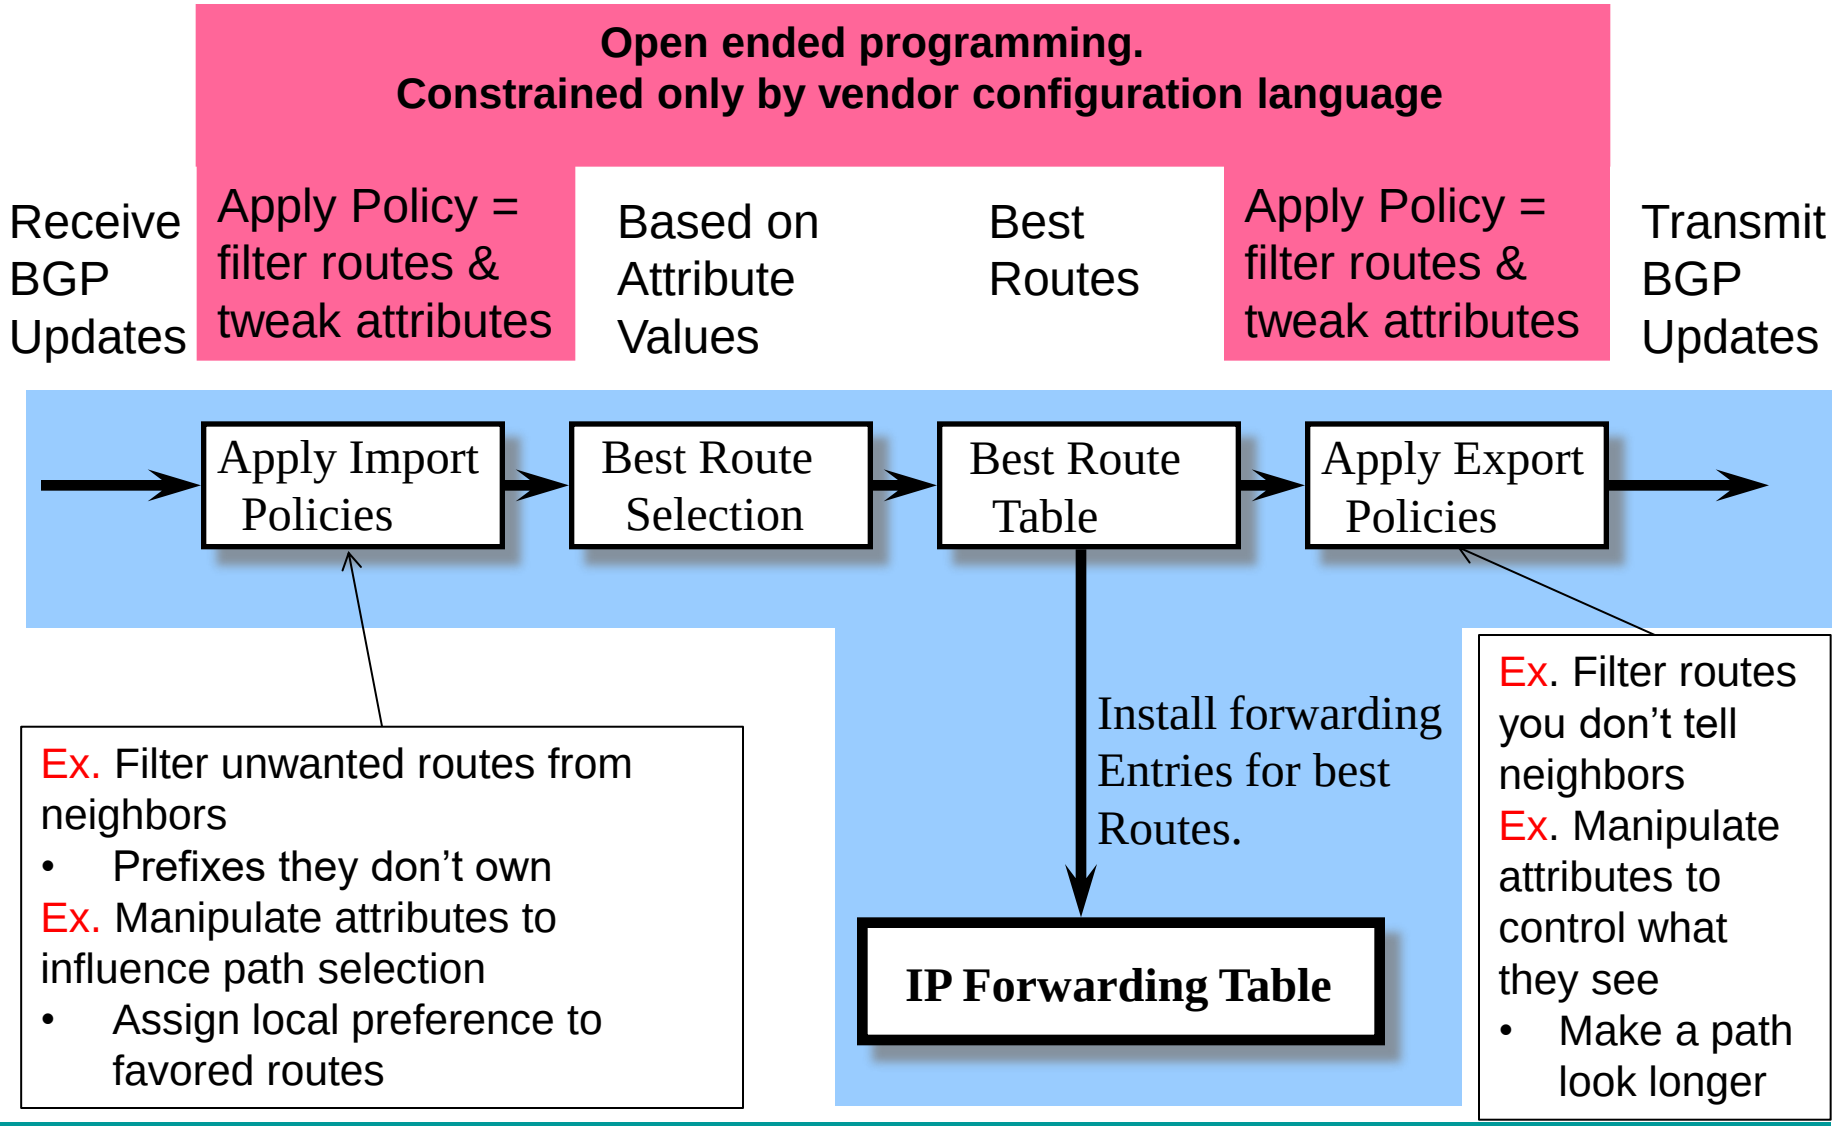

A Paper on Attacking BGP

Birge-Lee, H., Wang, L., Rexford, J., & Mittal, P. (2019). SICO: Surgical interception attacks by manipulating BGP communities. *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (CCS '19)*, 431–448. Association for Computing Machinery. https://doi.org/10.1145/3319535.3363197