




Routing: BGP

Overview continued

Prof. Anja Feldmann, Ph.D.




BGP: A path-vector protocol

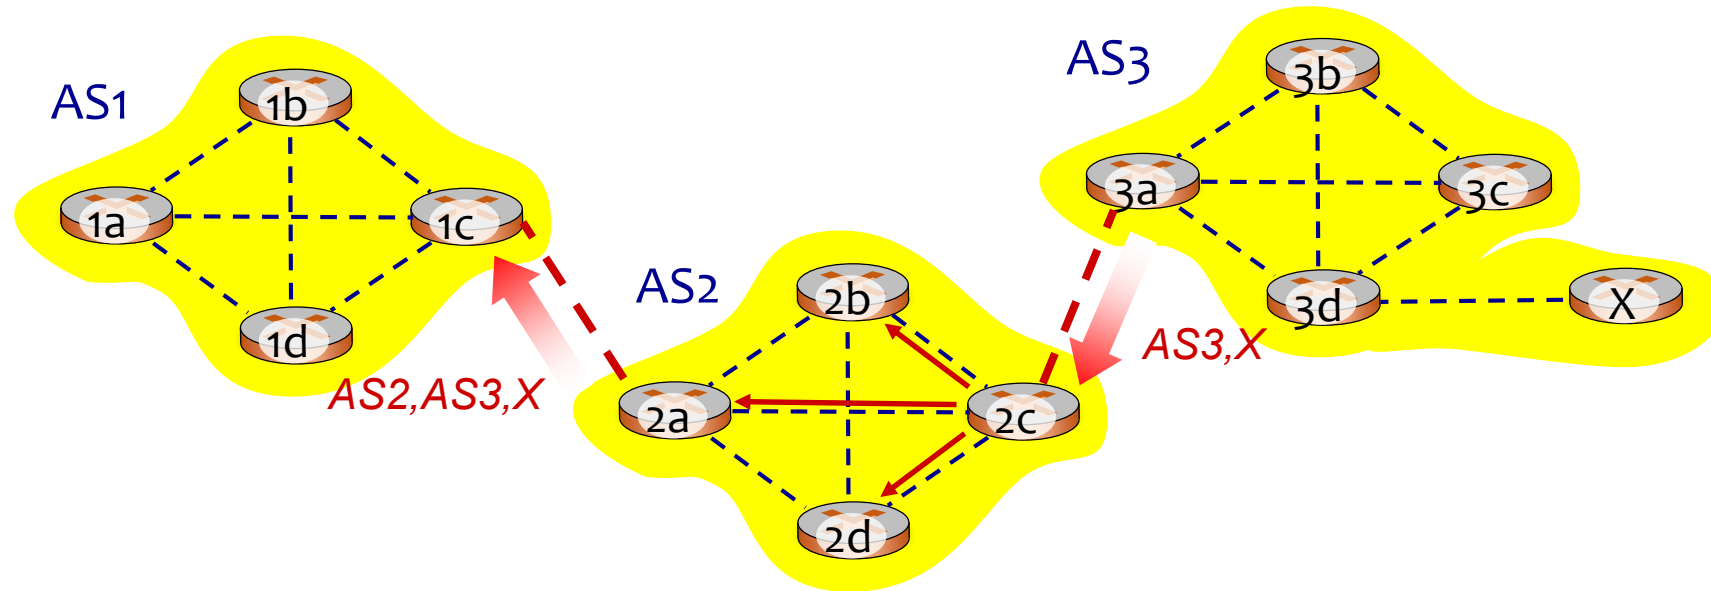


Distance vector algorithm with extra information


- Two important attributes:
 - **AS-PATH**: Contains all ASs along the way: AS 67 AS 17
 - **NEXT-HOP**: Indicates the specific internal-AS router to next-hop AS.
- Path can be used to make routing decisions, e.g., to avoid loops
- Pure distance vector does not enable policies
- Link state does not scale and exposes policies



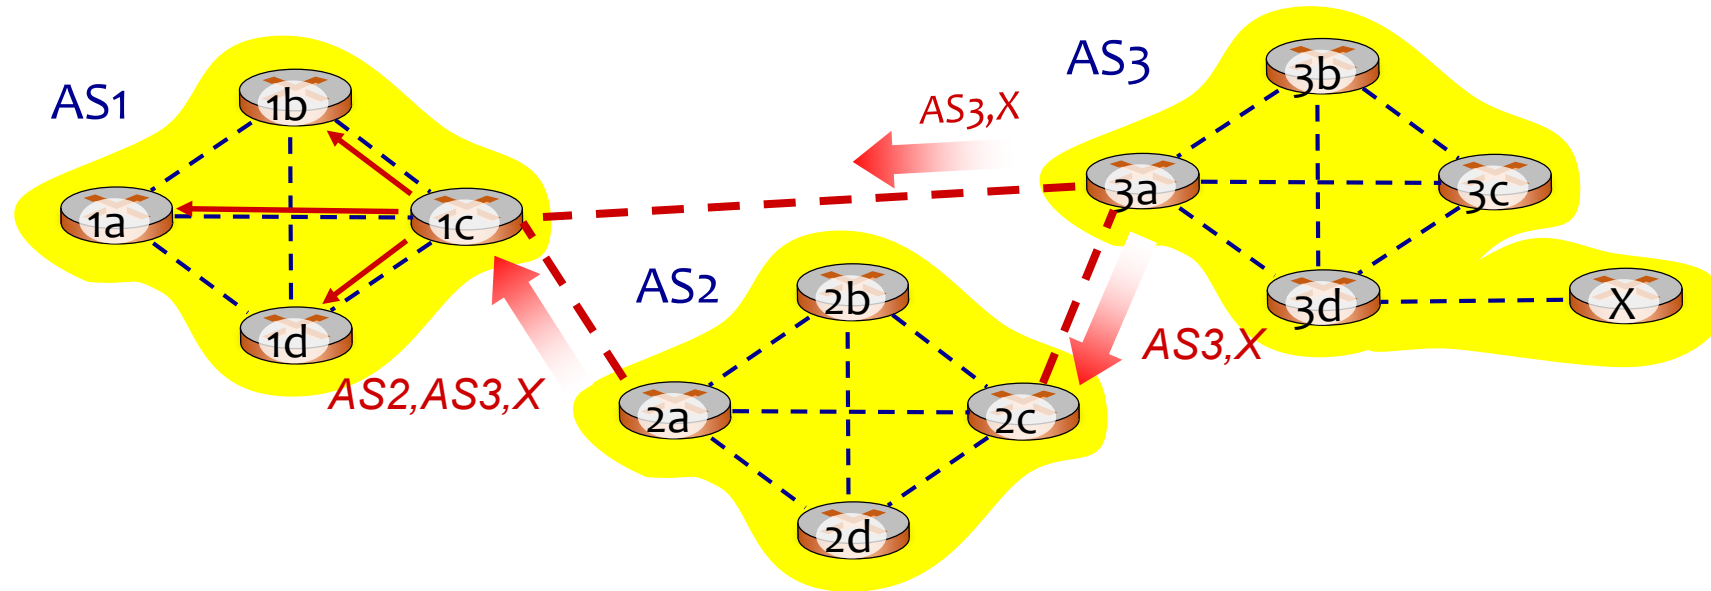
BGP path advertisement



- AS2 router 2c receives path advertisement AS_3, X (via eBGP) from AS3 router 3a
- Based on AS2 policy, AS2 router 2c accepts path AS_3, X , propagates (via iBGP) to all AS2 routers
- Based on AS2 policy, AS2 router 2a advertises (via eBGP) path AS_2, AS_3, X to AS1 router 1c




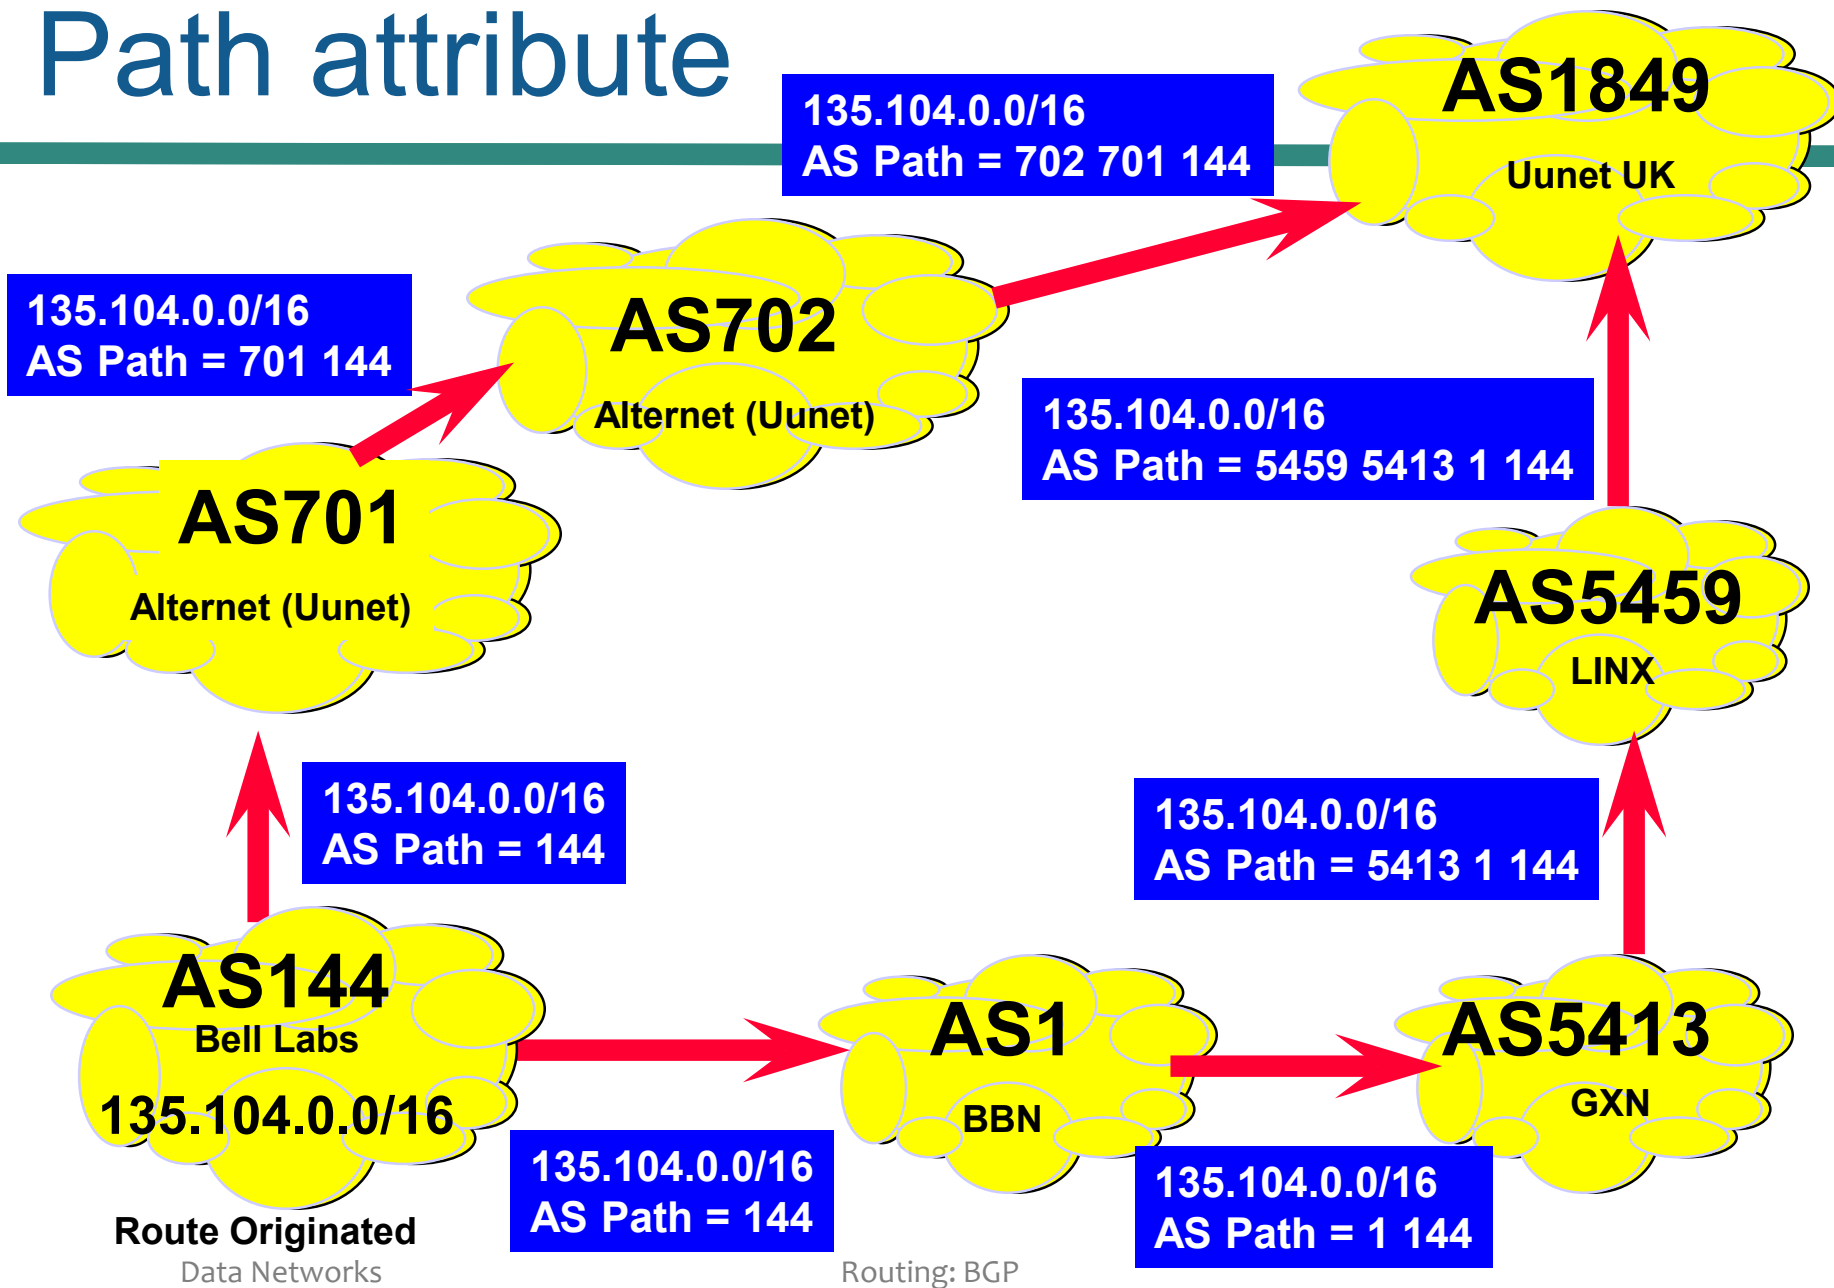
BGP path advertisement




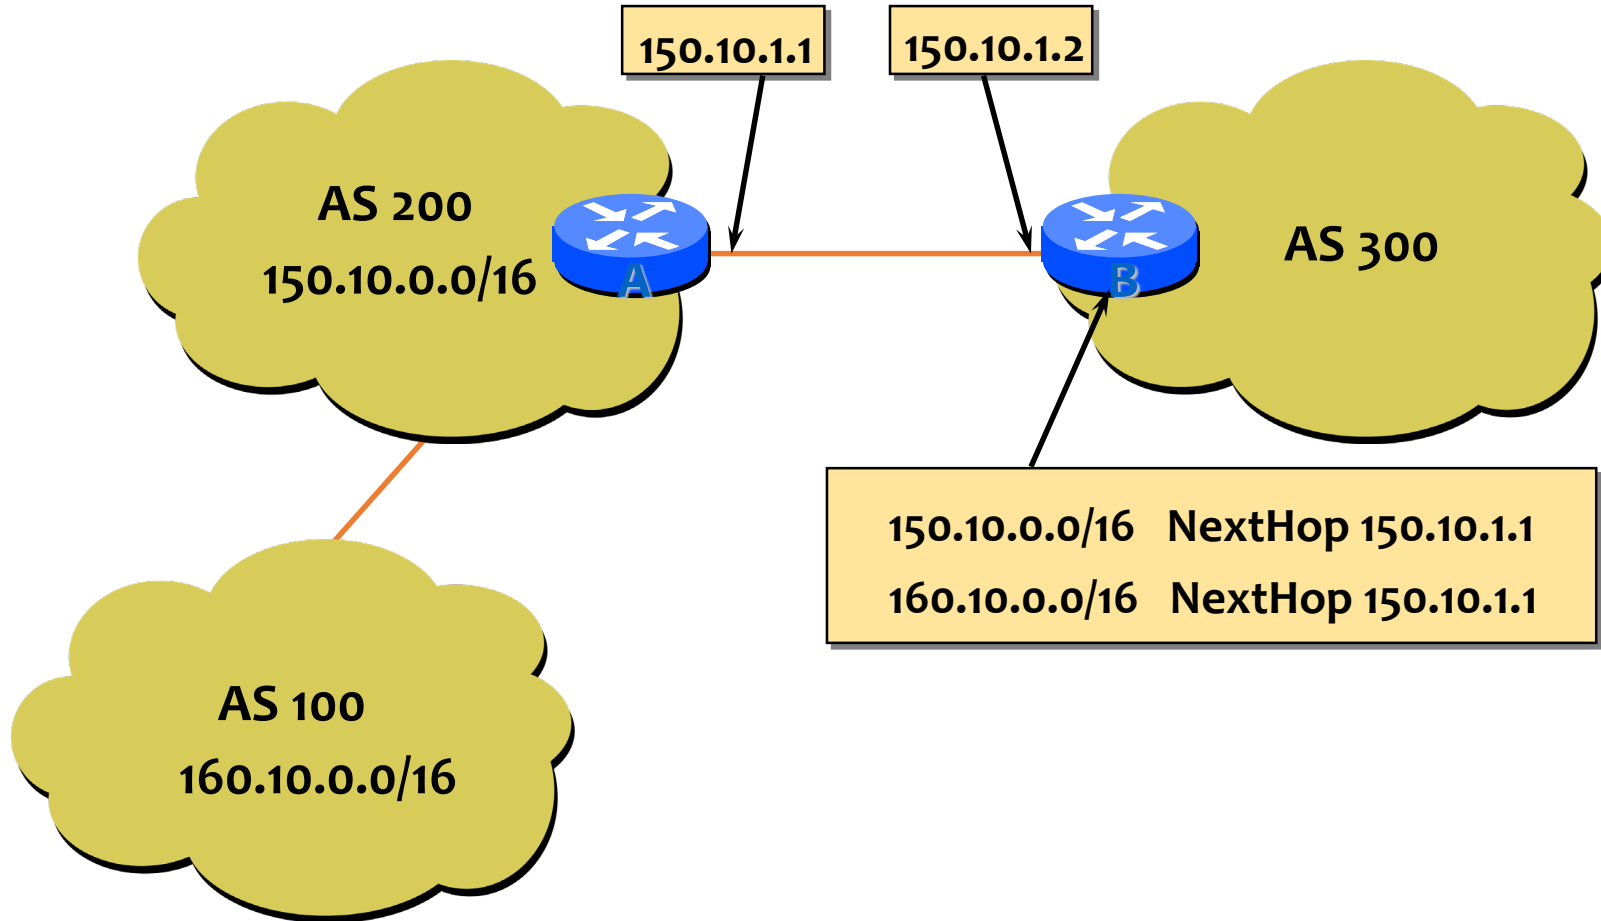
Gateway router may learn about **multiple** paths to destination:

- AS1 gateway router 1c learns path **AS2,AS3,X** from 2a
- AS1 gateway router 1c learns path **AS3,X** from 3a

AS Path attribute



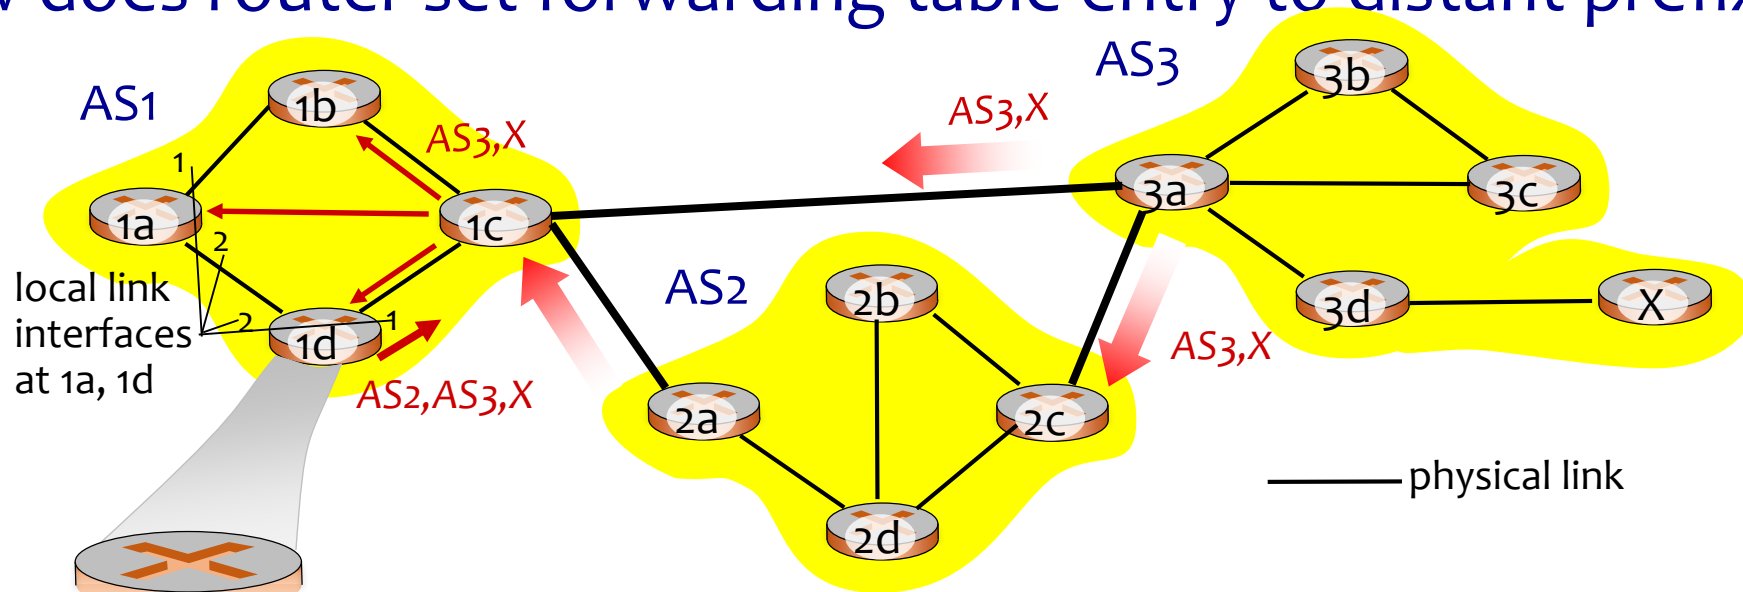
Next Hop attribute



BGP, OSPF, forwarding table entries




Q: How does router set forwarding table entry to distant prefix?



dest	interface
...	...
X	1
...	...

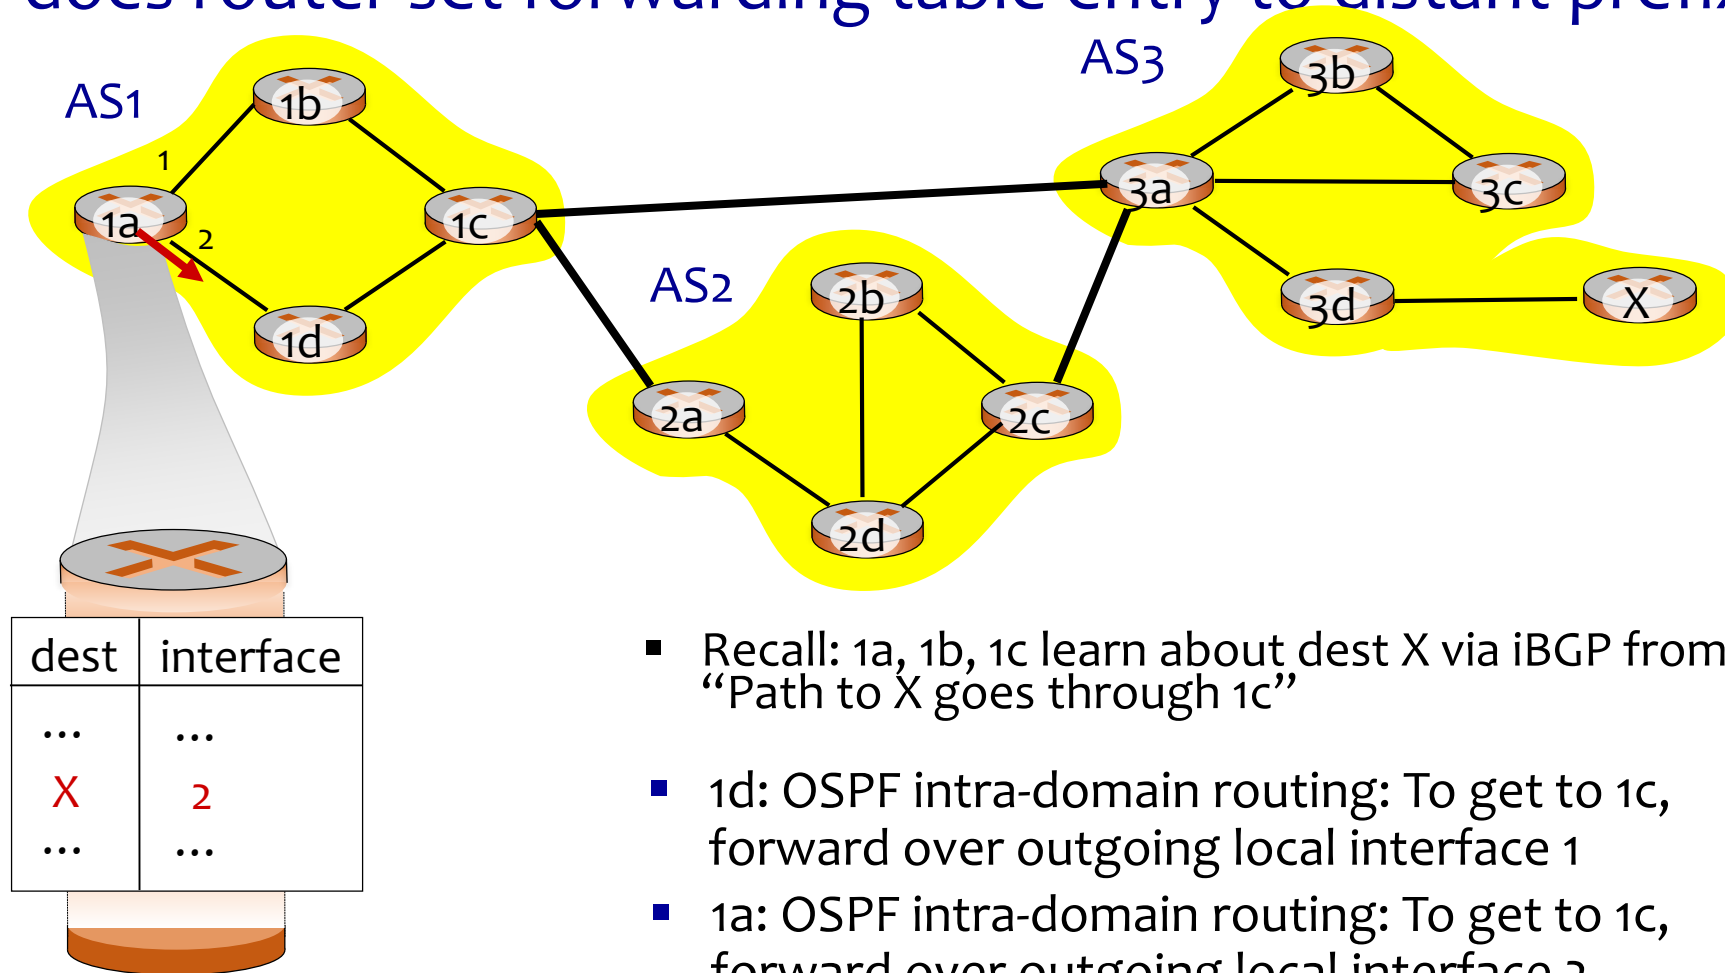
- Recall: 1a, 1b, 1c learn about dest X via iBGP from 1c: “Path to X goes through 1c”
- 1d: OSPF intra-domain routing: To get to 1c, forward over outgoing local interface 1



BGP, OSPF, forwarding table entries




Q: How does router set forwarding table entry to distant prefix?



dest	interface
...	...
X	2
...	...

- Recall: 1a, 1b, 1c learn about dest X via iBGP from 1c: “Path to X goes through 1c”
- 1d: OSPF intra-domain routing: To get to 1c, forward over outgoing local interface 1
- 1a: OSPF intra-domain routing: To get to 1c, forward over outgoing local interface 2



Routing policy




- Reflects goals of network provider
 - Which routes to accept from other ASes
 - How to manipulate the accepted routes
 - How to propagate routes through network
 - How to manipulate routes before they leave the AS
 - Which routes to send to another AS



Policies with BGP




- BGP provides capabilities for enforcing various policies
- Policies are **not** part of BGP!
- Policies are used to configure BGP
- BGP enforces policies by **choosing paths from multiple alternatives** and **controlling advertisements to other AS's**




BGP attributes




- AS path (well-known, mandatory)
- Next Hop (well-known, mandatory)
- Origin (well-known, mandatory)
- Multiple Exit Discriminator (MED)
(Optional, nontrans, eBGP)
- Local Preference (LocPref)
(well-known, discretionary, iBGP)
- Community (Optional, transitive)
- Atomic Aggregate (well-known, discretionary)
- Aggregator (Optional, transitive)
- Originator ID (Optional, nontransitive, Cisco)
- Other vendor-specific optional attributes ...




BGP: Local preference



- Path with *highest* local preference wins
- Allows providers to *prefer* routes




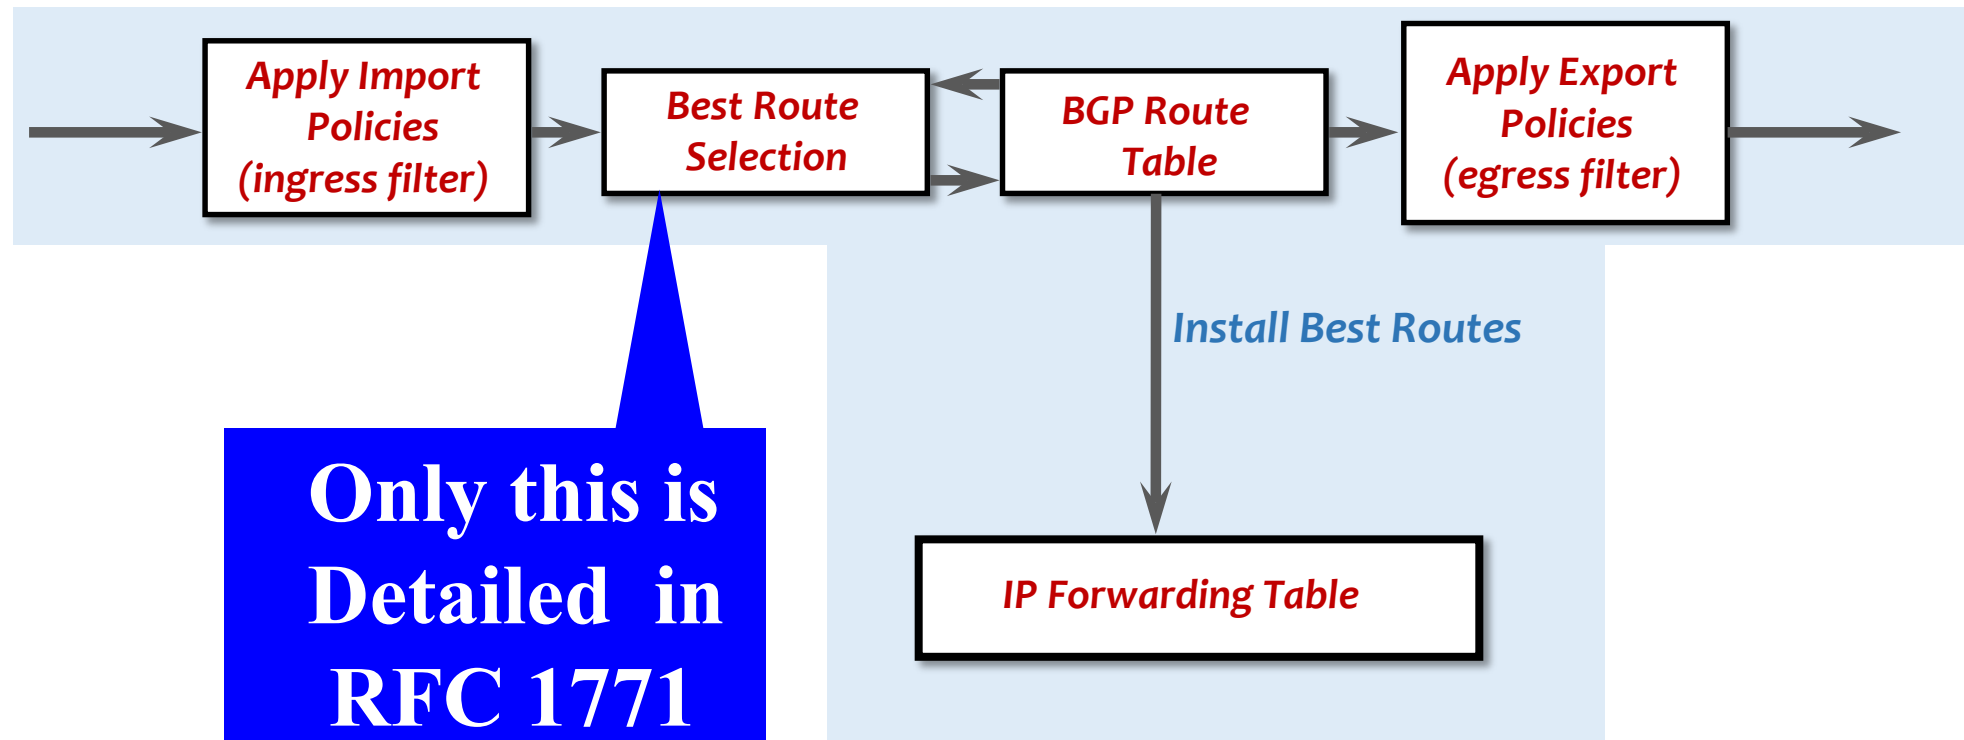
160.10.0.0/16	500
> 160.10.0.0/16	800



BGP route processing




Receive BGP Updates Apply Policy = filter routes & tweak attributes Based on Attribute Values Best and Alternate Routes Apply policies to Best Routes! Transmit BGP Updates



BGP route selection



- Router learn *more than 1 route* to some prefix
- Router **must** select *best route*
- **Elimination rules:**
 - Local preference value attribute: Policy decision
 - Shortest AS-PATH
 - Best MED (multi-exit-discriminator)
 - Closest NEXT-HOP router: Hot potato routing
 - Additional criteria
 - IP address of peer



BGP: A path-vector protocol



Distance vector algorithm with extra information

- When advertising a prefix, advert includes BGP attributes
 - *Prefix + other attributes = “route”*
- When gateway router receives route advertisement, uses *ingress filters* to accept/decline
 - Can make decision based on ASes on path, e.g., to avoid loops

