Network Layer
Routing: RIP & OSPF

Prof. Anja Feldmann, Ph.D.

Internet Routing

Let's switch from the algorithms to the protocols that use the algorithms for routing in the Internet!

Interplay between routing & forwarding

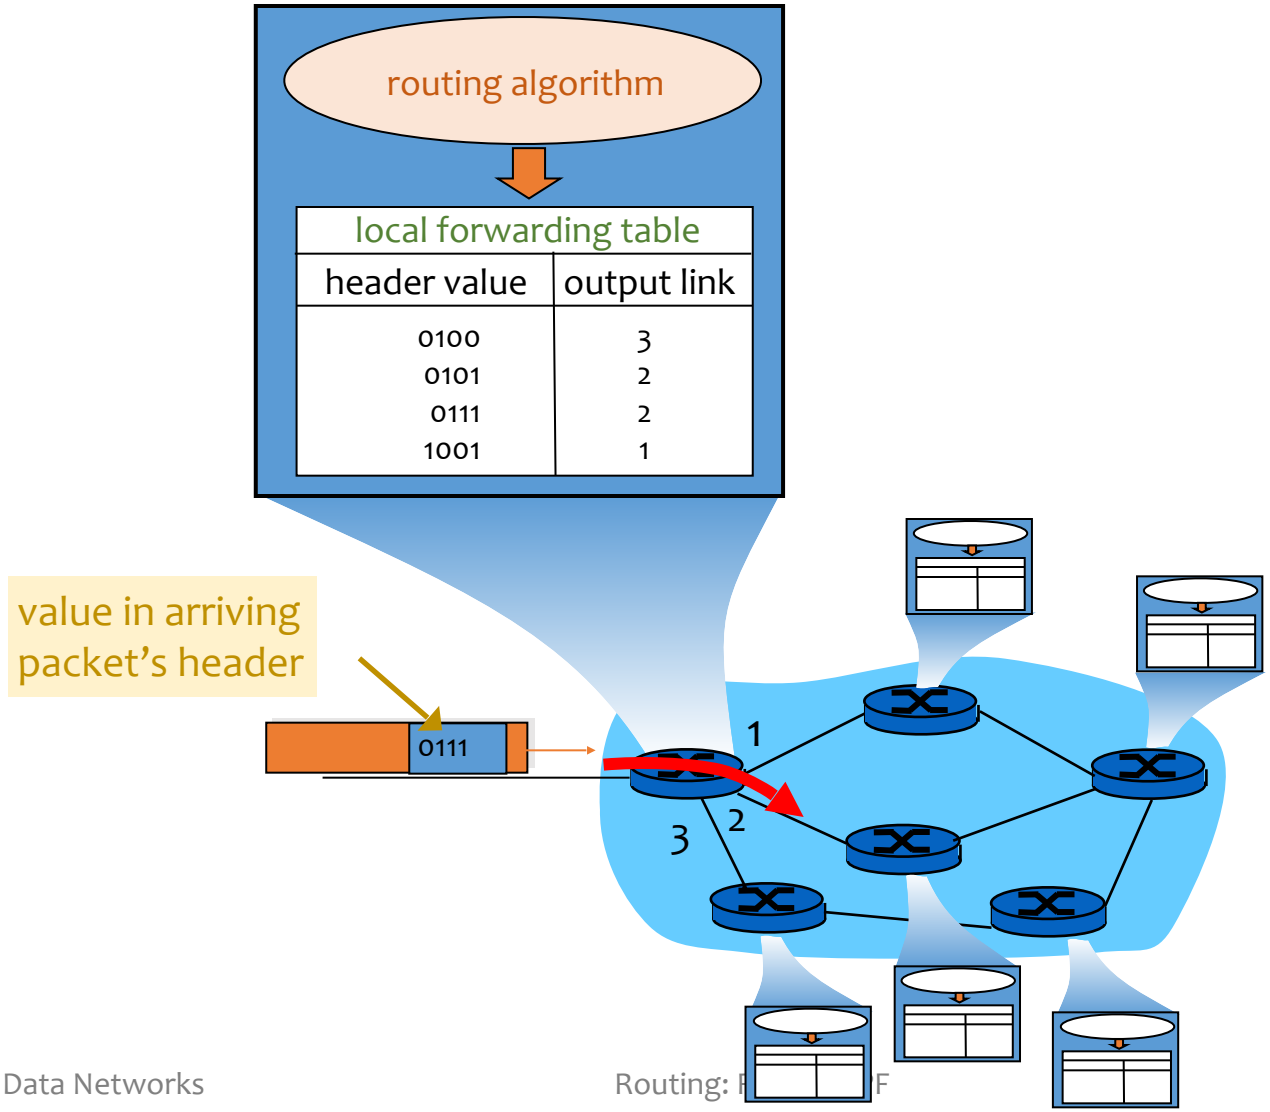

Internet Routing

So far – idealized picture

- All routers identical
- Network is "flat"

… not true in practice

Scale

With 850 million destinations:

- Can't store all destinations in routing tables!
- Routing table exchange would swamp links!

Internet Routing

Administrative autonomy

- Internet: network of networks
- Each network operator or administrator may want to control routing in its own network

Internet Routing

Administrative autonomy

- Internet: network of networks
- Each network operator or administrator may want to control routing in its own network

Autonomous Systems (AS)

- Aggregate routers into regions

- Routers in same AS run same "intra-AS" routing protocol
- Routers in different AS can run different intra-AS routing protocol

Interconnected ASes

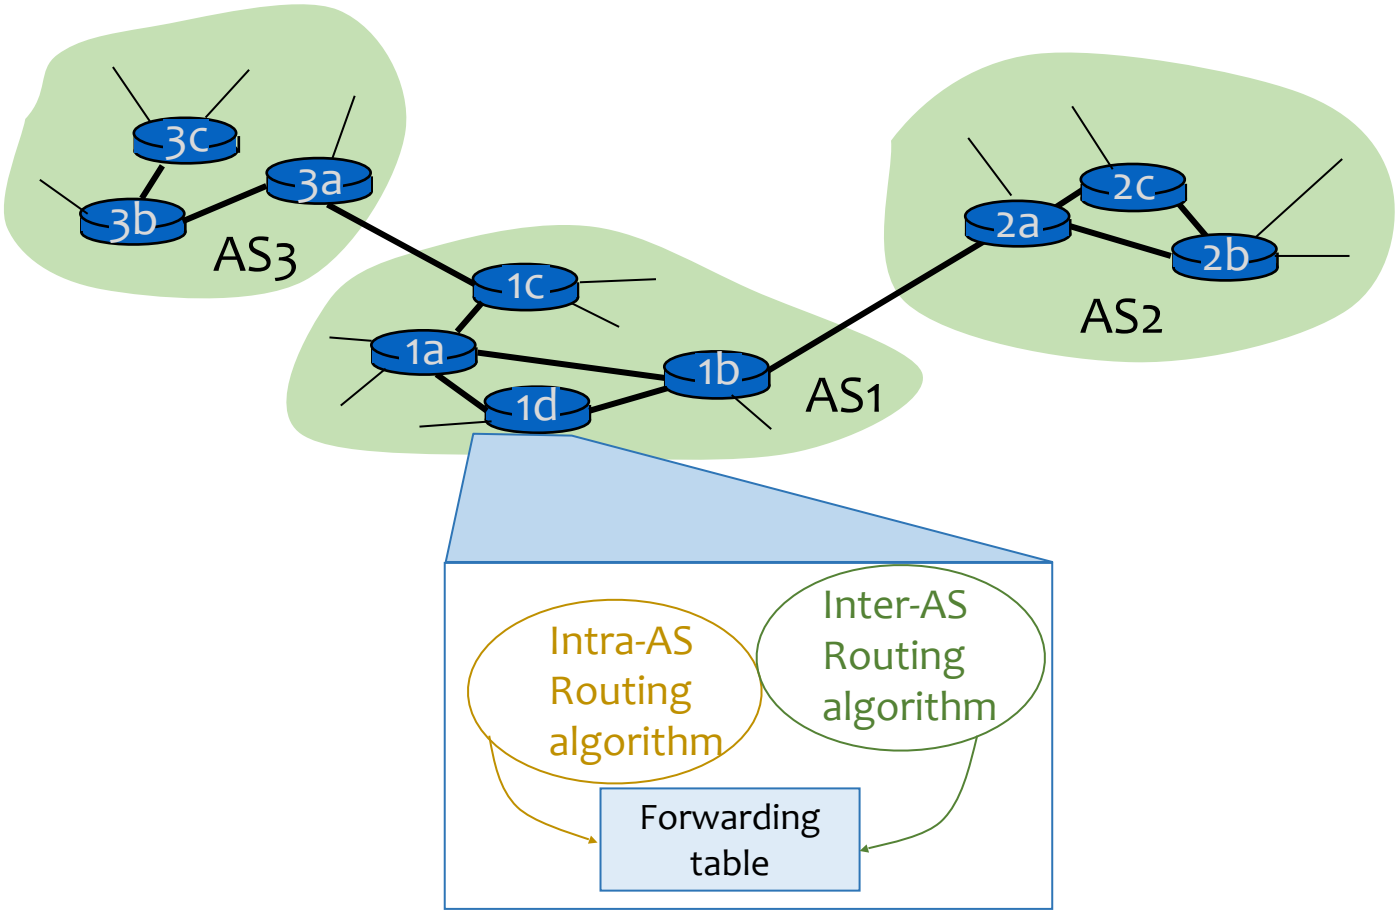

Interconnected ASes

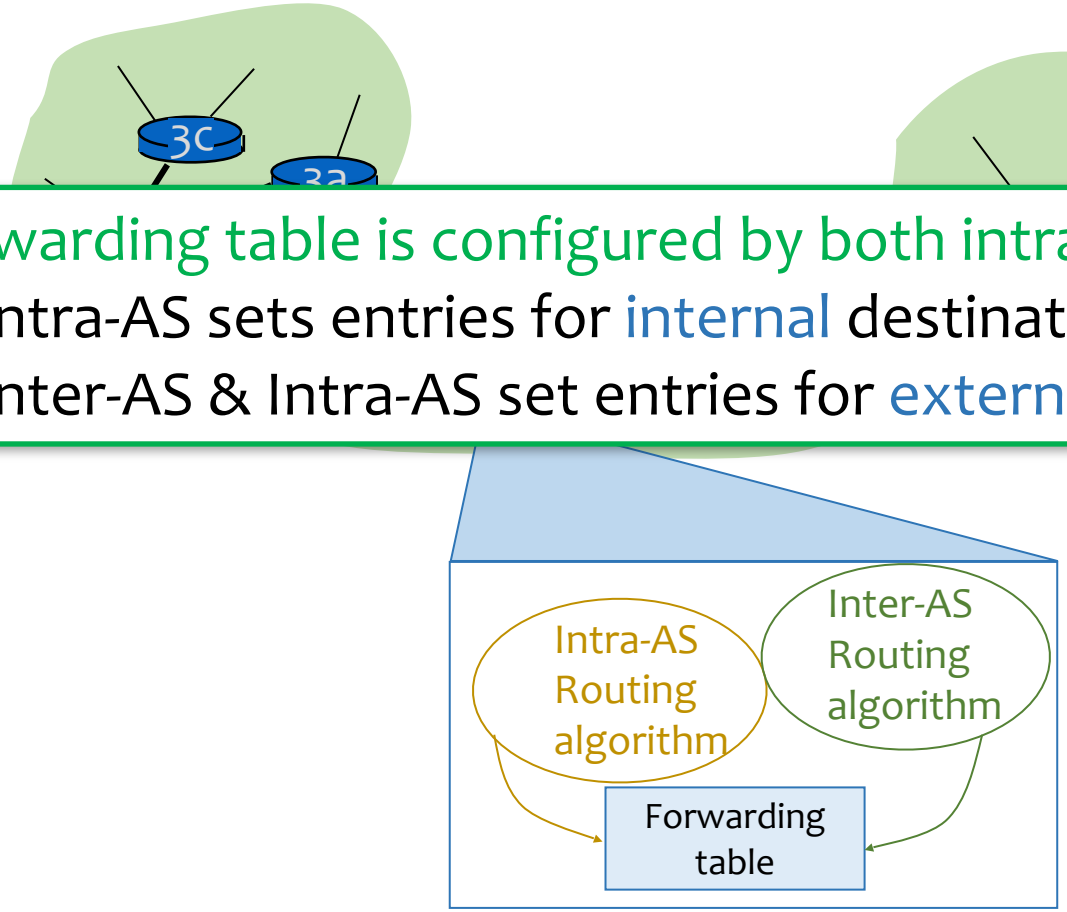

Forwarding table is configured by both intra-AS and inter-AS routing algorithm.

- Intra-AS sets entries for internal destinations
- Inter-AS & Intra-AS set entries for external destinations

Inter-AS tasks

Suppose router in AS_1 receives datagram for dst. outside of AS_1

- Router should forward packet towards an AS border router, but which one?

Job of Inter-AS routing

AS_1 needs …

- … to learn which dsts. are reachable through AS_2 and which through AS_3
- … to propagate this reachability information to all routers in AS_1

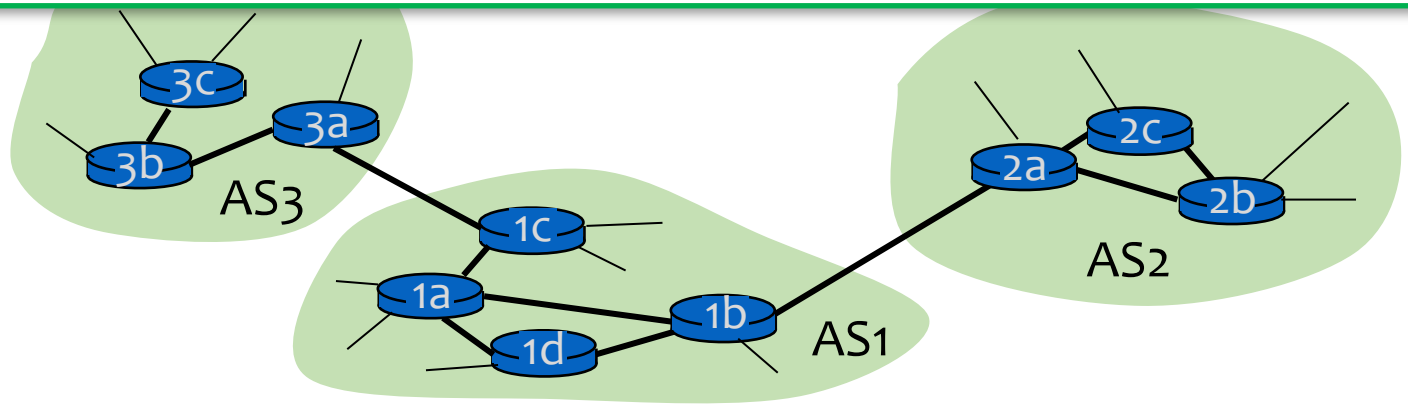

Intra-AS routing

Also known as Interior Gateway Protocols (IGP)

Most common Intra-AS routing protocols:

- Routing Information Protocol (RIP)
- Open Shortest Path First (OSPF)

Intra-AS routing: RIP

Routing Information Protocol (RIP)

- Distance vector protocol (based on Bellman-Ford)
- Routers periodically exchange reachability info with their neighbors
- Included in BSD-UNIX Distribution in 1982
- Distance metric: Hop count or number of hops (max = 15 hops)

RIP: Hop count

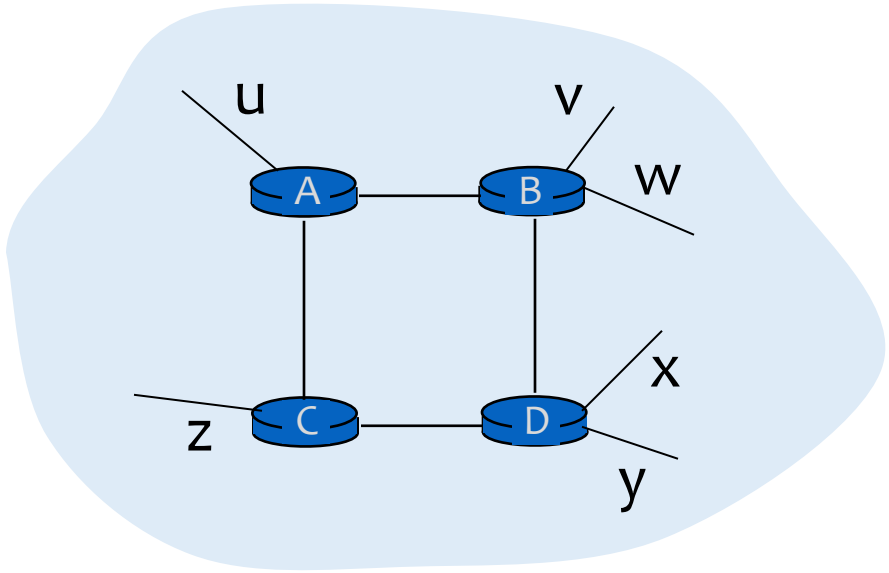

Destination subnet (from A)	Hop count
u	1
v	2
w	2
x	3
y	3
z	2

RIP: Advertisements

- Distance vectors exchanged among neighbors every 30 s via Response Message (also called advertisement)

- Each advertisement:
  - List of up to 25 destination networks within AS

A-to-D advertisement

RIP: Example

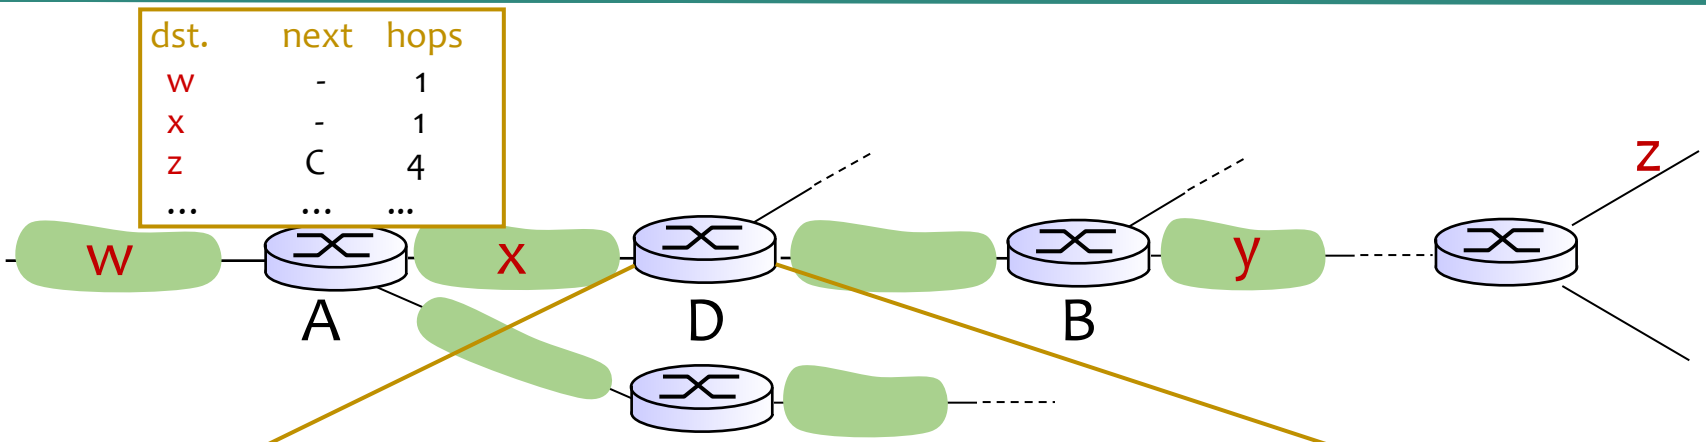

Routing table at D

destination subnet	next router	#hops to dst.
w	A	2
y	B	2
z	B	7
x	--	1
...	...	...

A-to-D advertisement

RIP: Example

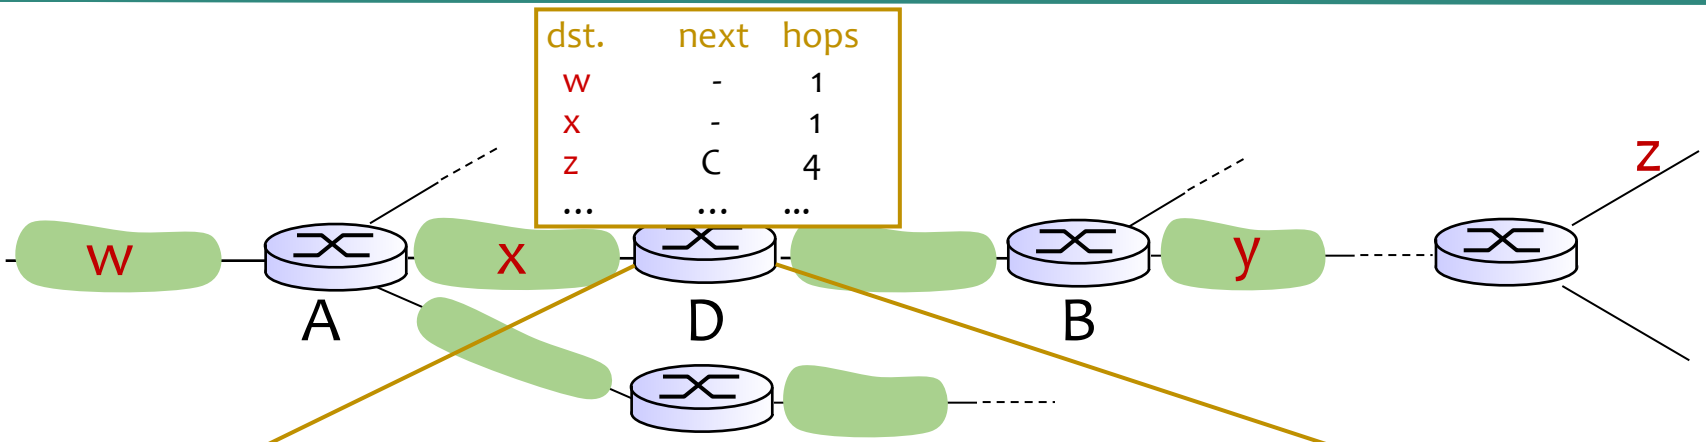

Routing table at D

destination subnet	next router	#hops to dst.
w	A	2
y	B	2
z	~~B~~A	~~7~~5
x	--	1
…	…	…

RIP: Failure and recovery

If no advertisement heard after 180 s, neighbor or link is declared dead

- Routes via the neighbor are invalidated
- New advertisements sent to neighbors
- Neighbors in turn send out new advertisements (if their tables changed)
- Link failure information quickly propagates to entire network

- Poison reverse used to prevent ping-pong loops (infinite distance = 16 hops)
  - Advertise an infinite metric for a route on the interface over which it was received

RIPv2: Message format

- Transported over UDP; port 520

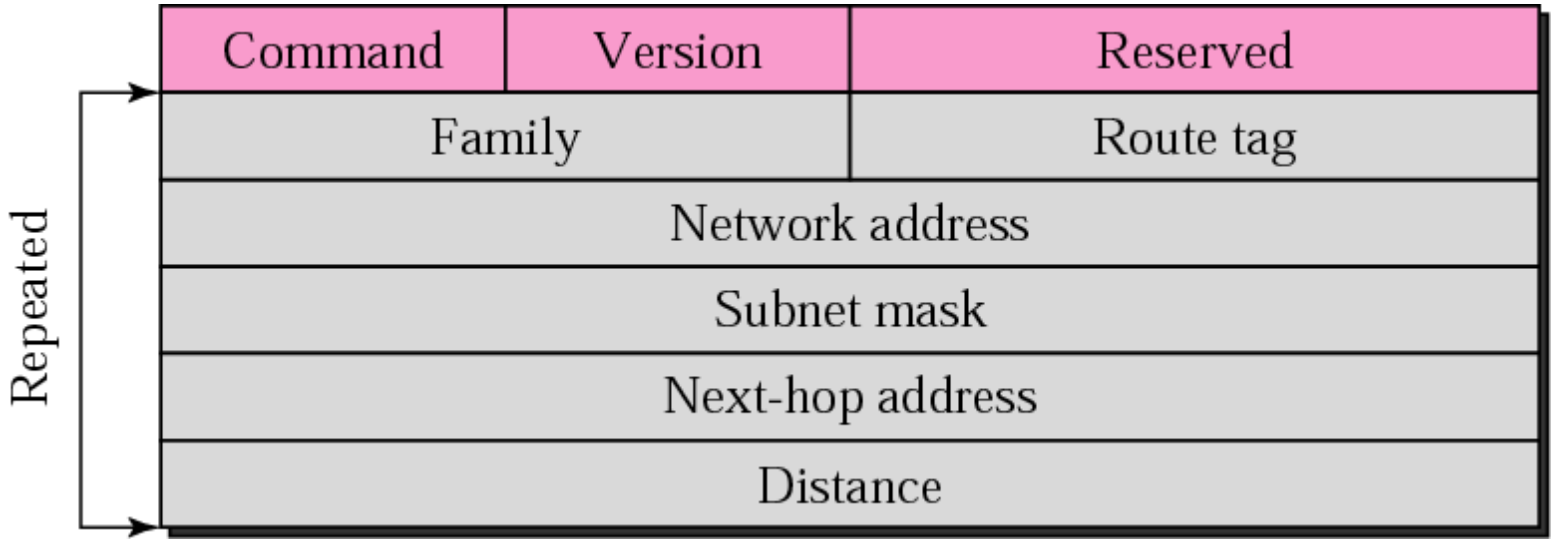

RIP: Table processing

- RIP routing tables managed by an application process (route-d daemon in Unix)
- Advertisements are
  - encapsulated in UDP packets (no reliable delivery required)
  - periodically repeated

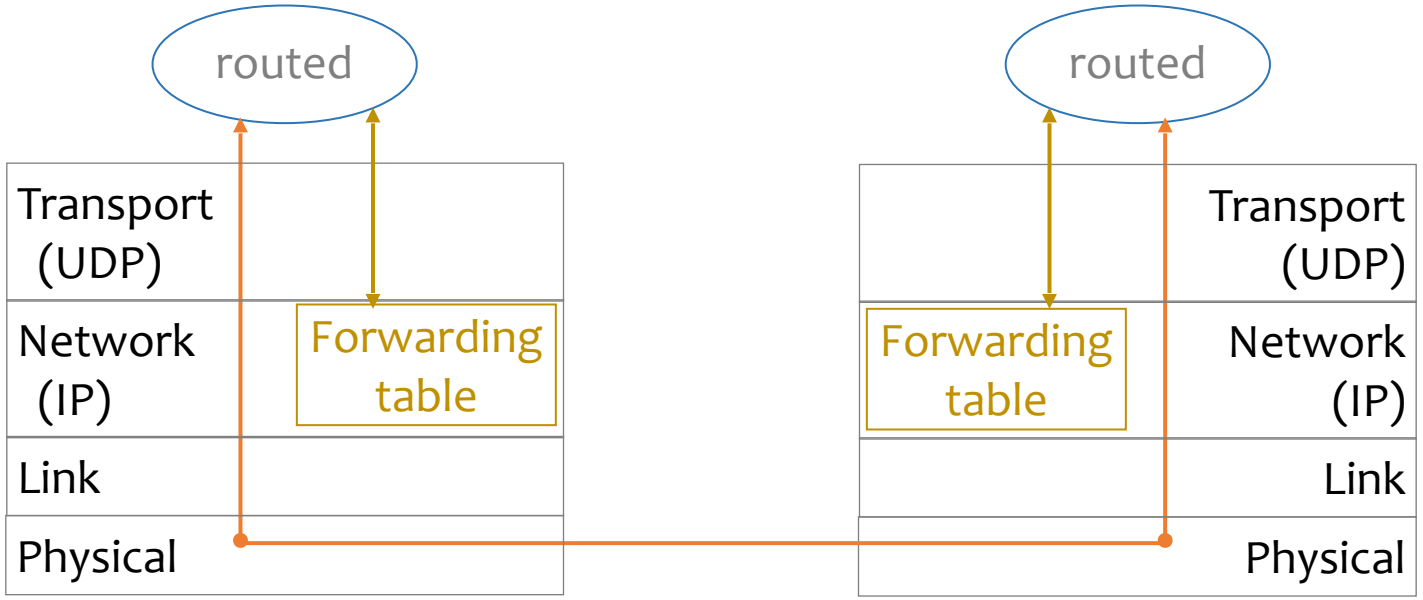

RIP: Table example

- Three attached Class C networks (LANs)
- Router only knows routes to attached LANs
- Default router used to “go up”
- Route multicast address: 224.0.0.0
- Loopback interface

Router: giroflee.eurocom.fr

Destination	Gateway	Flags	Ref	Use	Interface
127.0.0.1	127.0.0.1	UH	0	26492	lo0
192.168.2.	192.168.2.5	U	2	13	fa0
193.55.114.	193.55.114.6	U	3	58503	le0
192.168.3.	192.168.3.5	U	2	25	qaa0
224.0.0.0	193.55.114.6	U	3	0	le0
default	193.55.114.129	UG	0	143454	

RIP: Avoid “count-to-infinity” problem

Split Horizon

- Don‘t advertise a route for an interface on which it was received

Split Horizon with Poisoned Reverse

- Put an infinite metric on routes out the interface on which is was received

Triggered Updates

- After metric change: Send update immediately

Hold-down Timer (Cisco)

- After invalidation of route:
  - For some seconds ignore all updates for route

RIP: Routing tasks

Neighbor?

- Discovery
- Maintenance

Database?

- Granularity
- Maintenance – updates
- Synchronization

Routing table?

- Metric
- Calculation
- Update

RIP: Summary

Routing Information Protocol (RIP)

- Distance vector protocol (based on Bellman-Ford)
- Routers periodically exchange reachability info with their neighbors
- Distance metric: Hop count or number of hops (max = 15 hops)

- Advantage: Simple, minimal communication overhead

- Disadvantage: Long convergence times, loop detection