


CS 725/825 & IT 725

Lecture 6

Network Layer Routing

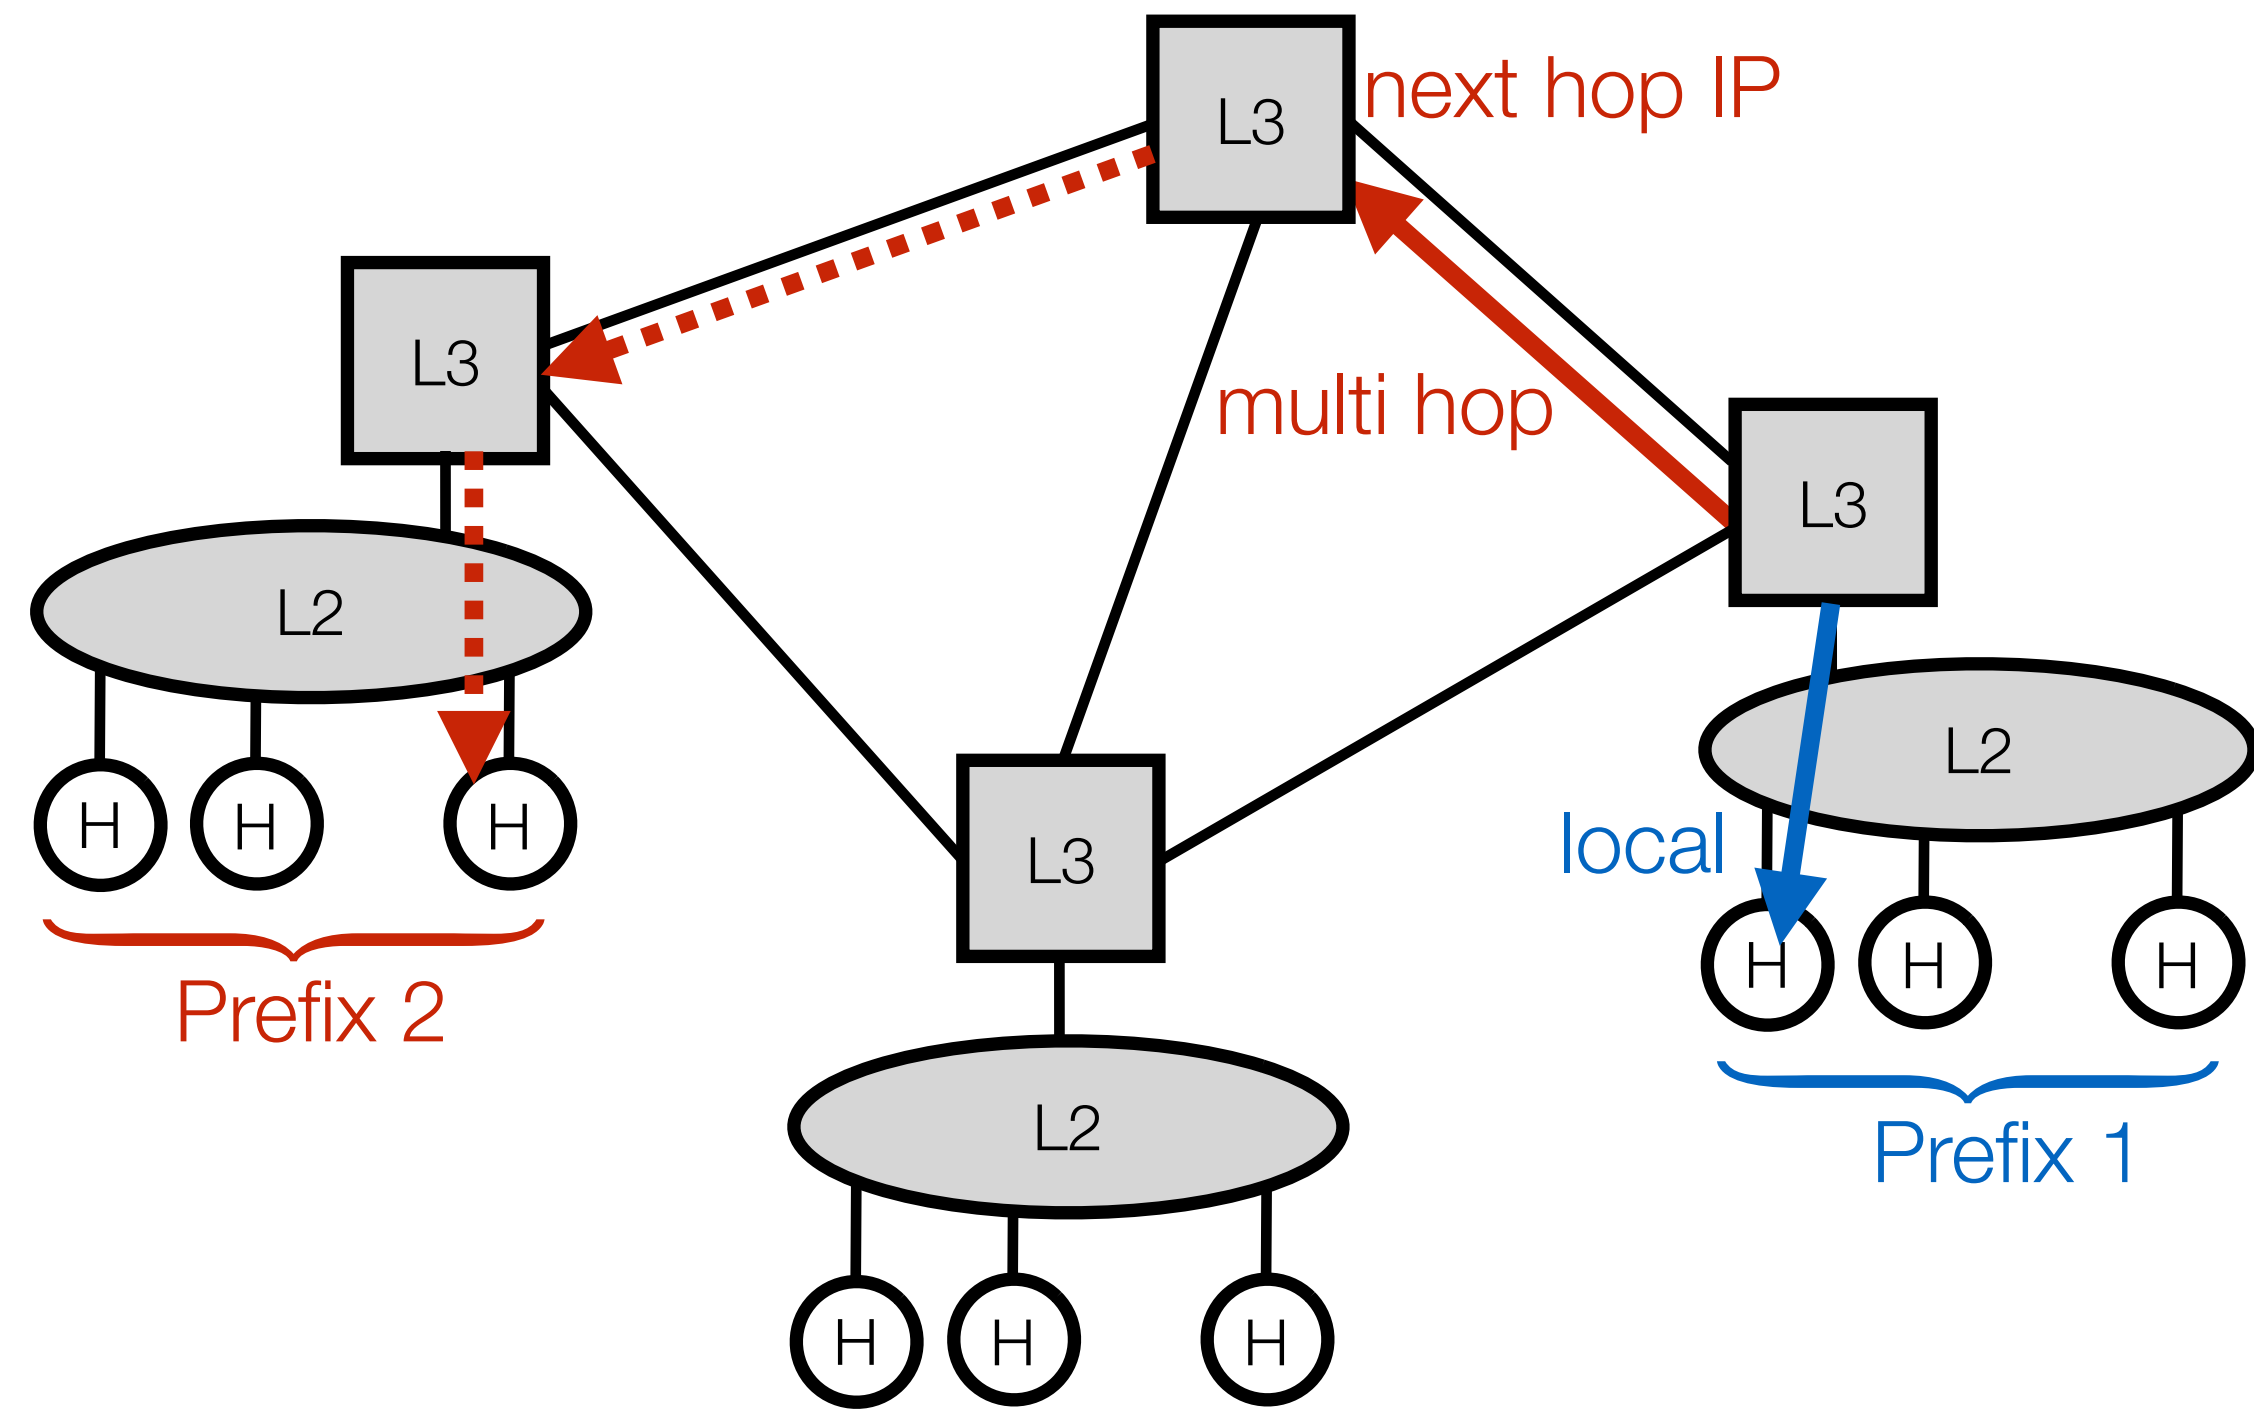
September 16, 2024

Network Layer Routing




Hop-by-hop routing: “your next hop is X, go there and ask for further directions”

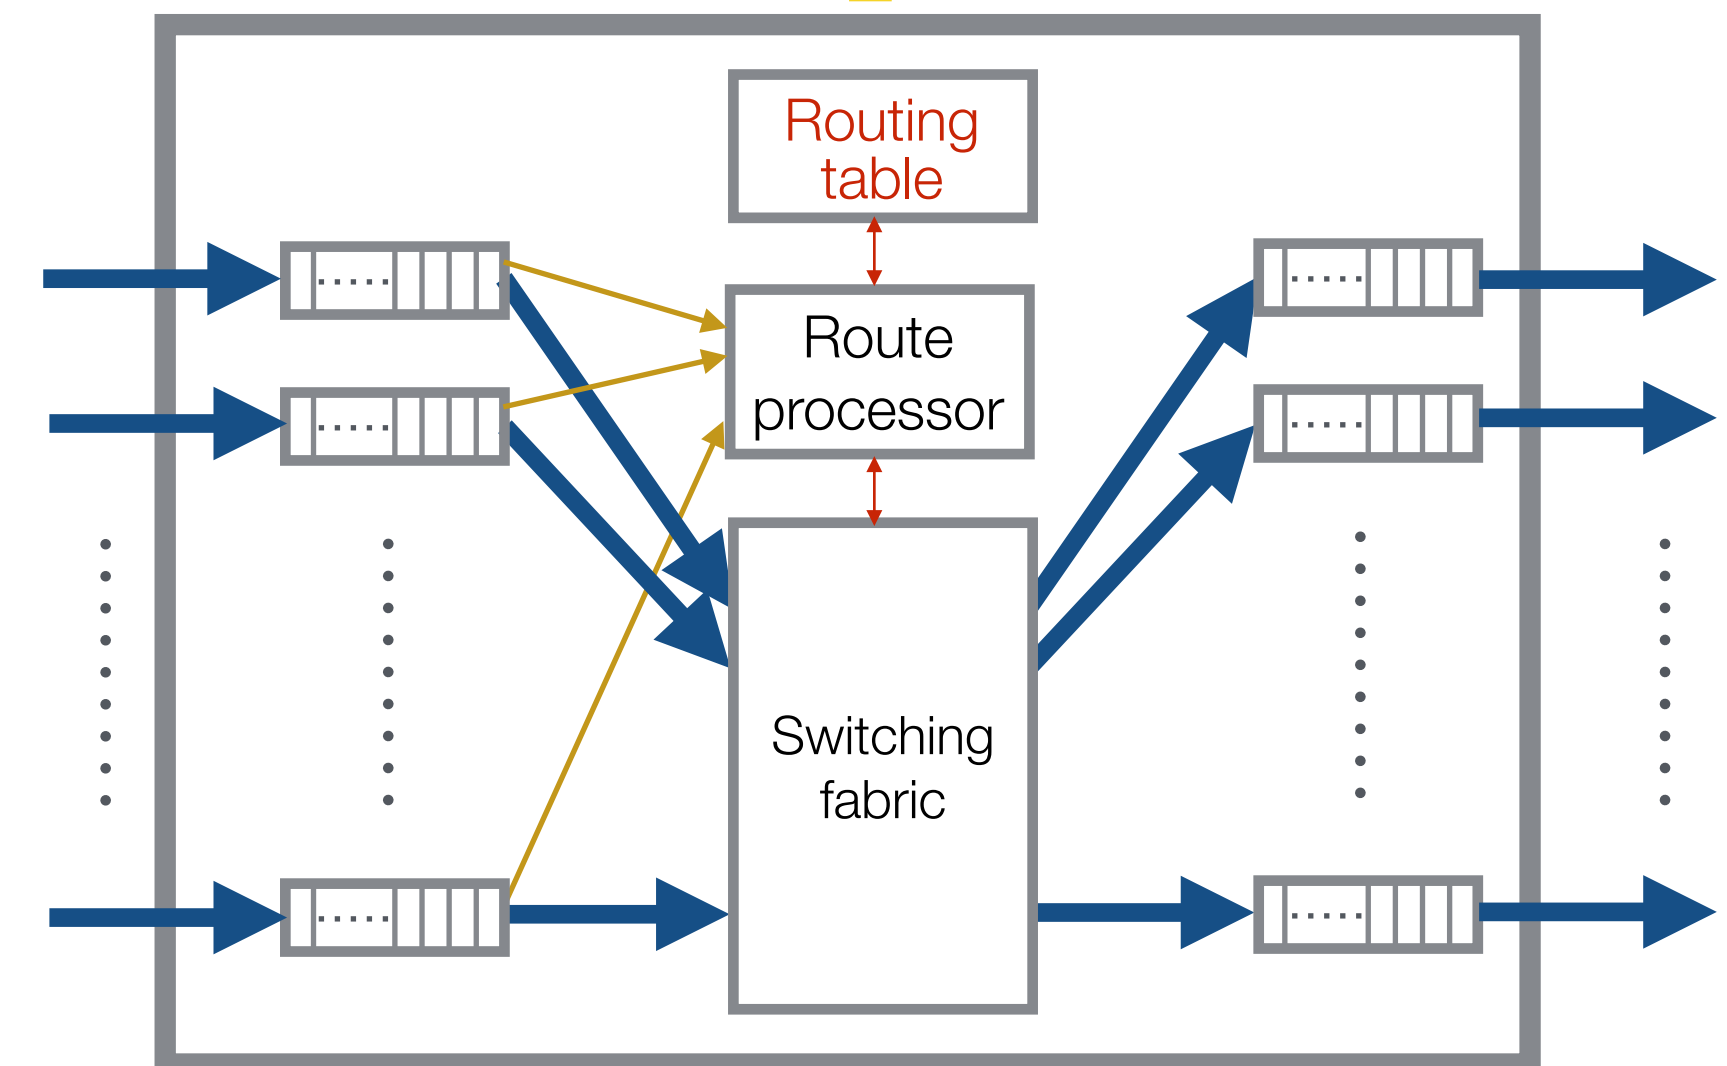
Network Layer Routing



Network Layer Routing



Prefix	Next hop
Prefix 1	L2 interface
Prefix 2	Next hop IP
...	...



Routing table

Prefix	Next hop
Prefix 1	L2 interface
Prefix 2	Next hop IP
...	...

use ARP to find the MAC address of the destination

search the routing table to find the L2 interface to reach the **next hop**, then use ARP to find out the next hop's MAC address

Multiple matching entries...

Destination IP: 132.177.5.123

Prefix	Next hop
132.177.4.0/22	
132.177.0.0/16	
...	...

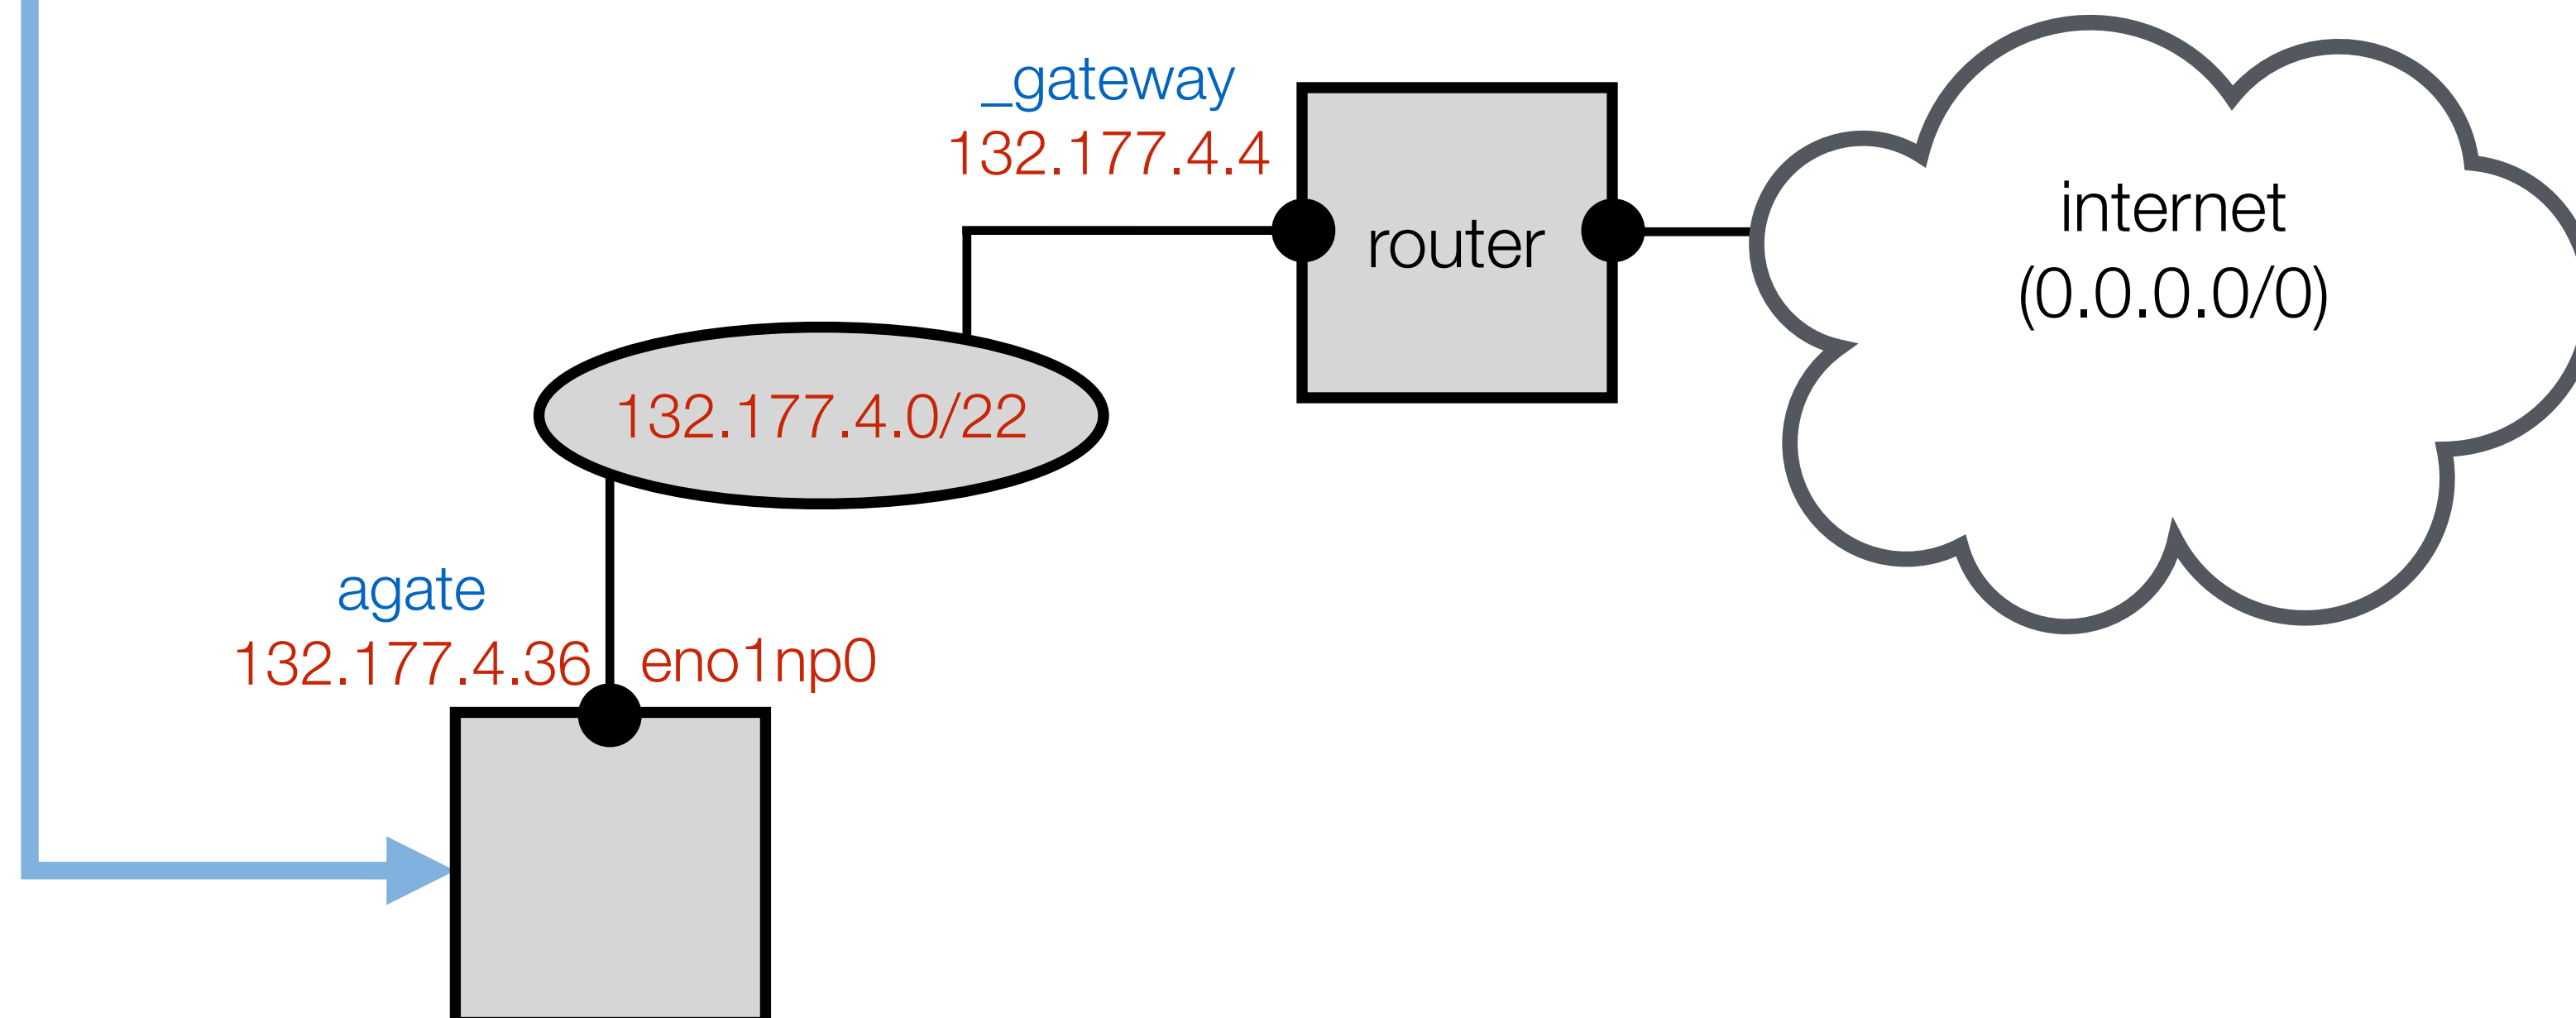
Longest prefix match

Host routing table

```
rbartos@agate ~$ route
Kernel IP routing table
Destination      Gateway          Genmask         Flags Metric Ref    Use Iface
default          _gateway        0.0.0.0         UG    100    0      0 eno1np0
unh-cs          0.0.0.0         255.255.252.0   U     100    0      0 eno1np0

rbartos@agate ~$ route -n
Kernel IP routing table
Destination      Gateway          Genmask         Flags Metric Ref    Use Iface
0.0.0.0          132.177.4.4     0.0.0.0         UG    100    0      0 eno1np0
132.177.4.0     0.0.0.0         255.255.252.0   U     100    0      0 eno1np0
```

“-n” = do not resolve IP addresses



Host routing table

```
rbartos@agate ~$ route
Kernel IP routing table
Destination      Gateway          Genmask          Flags Metric Ref    Use Iface
default          _gateway        0.0.0.0          UG    100    0      0 eno1np0
unh-cs          0.0.0.0        255.255.252.0   U     100    0      0 eno1np0

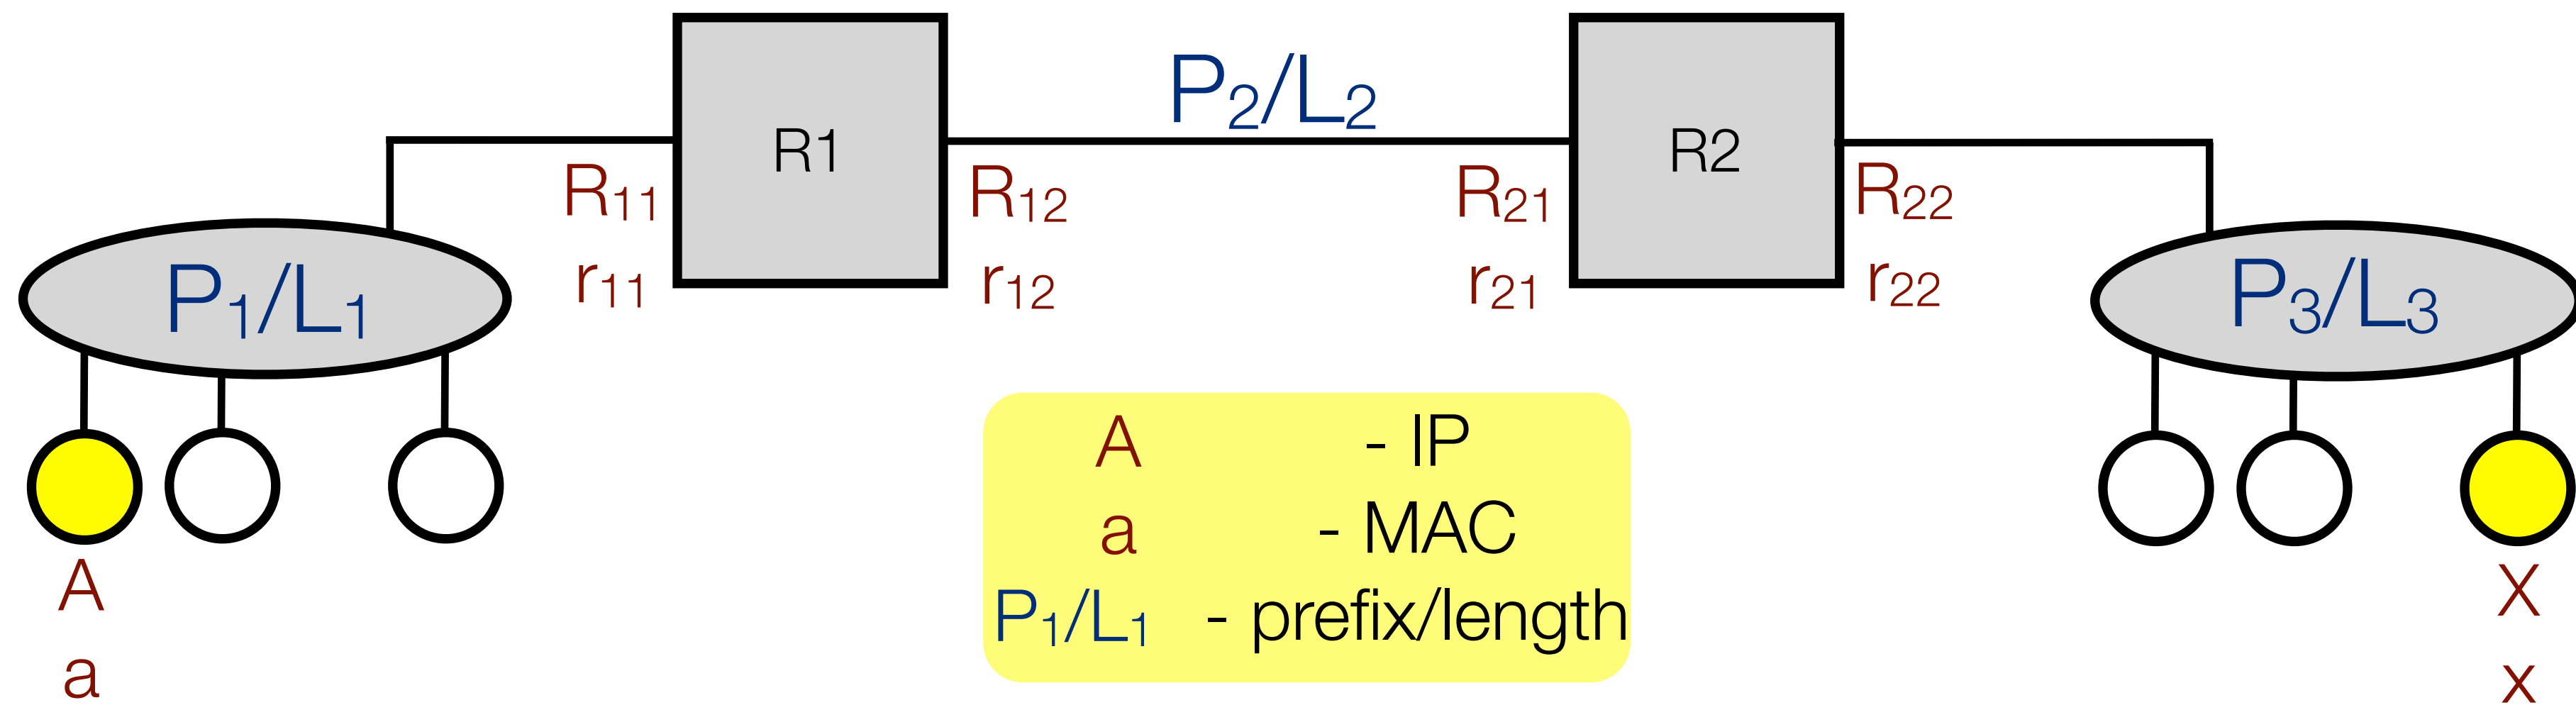
rbartos@agate ~$ route -n
Kernel IP routing table
Destination      Gateway          Genmask          Flags Metric Ref    Use Iface
0.0.0.0         132.177.4.4     0.0.0.0          UG    100    0      0 eno1np0
132.177.4.0     0.0.0.0        255.255.252.0   U     100    0      0 eno1np0
```

“-n” = do not resolve IP addresses

132.177.4.0/22 → local delivery via interface **eno1np0**
(do ARP for the destination IP)

0.0.0.0/0 (everything else) → send to **_gateway**
132.177.4.4 via **eno1np0**
(do ARP for 132.177.4.4)

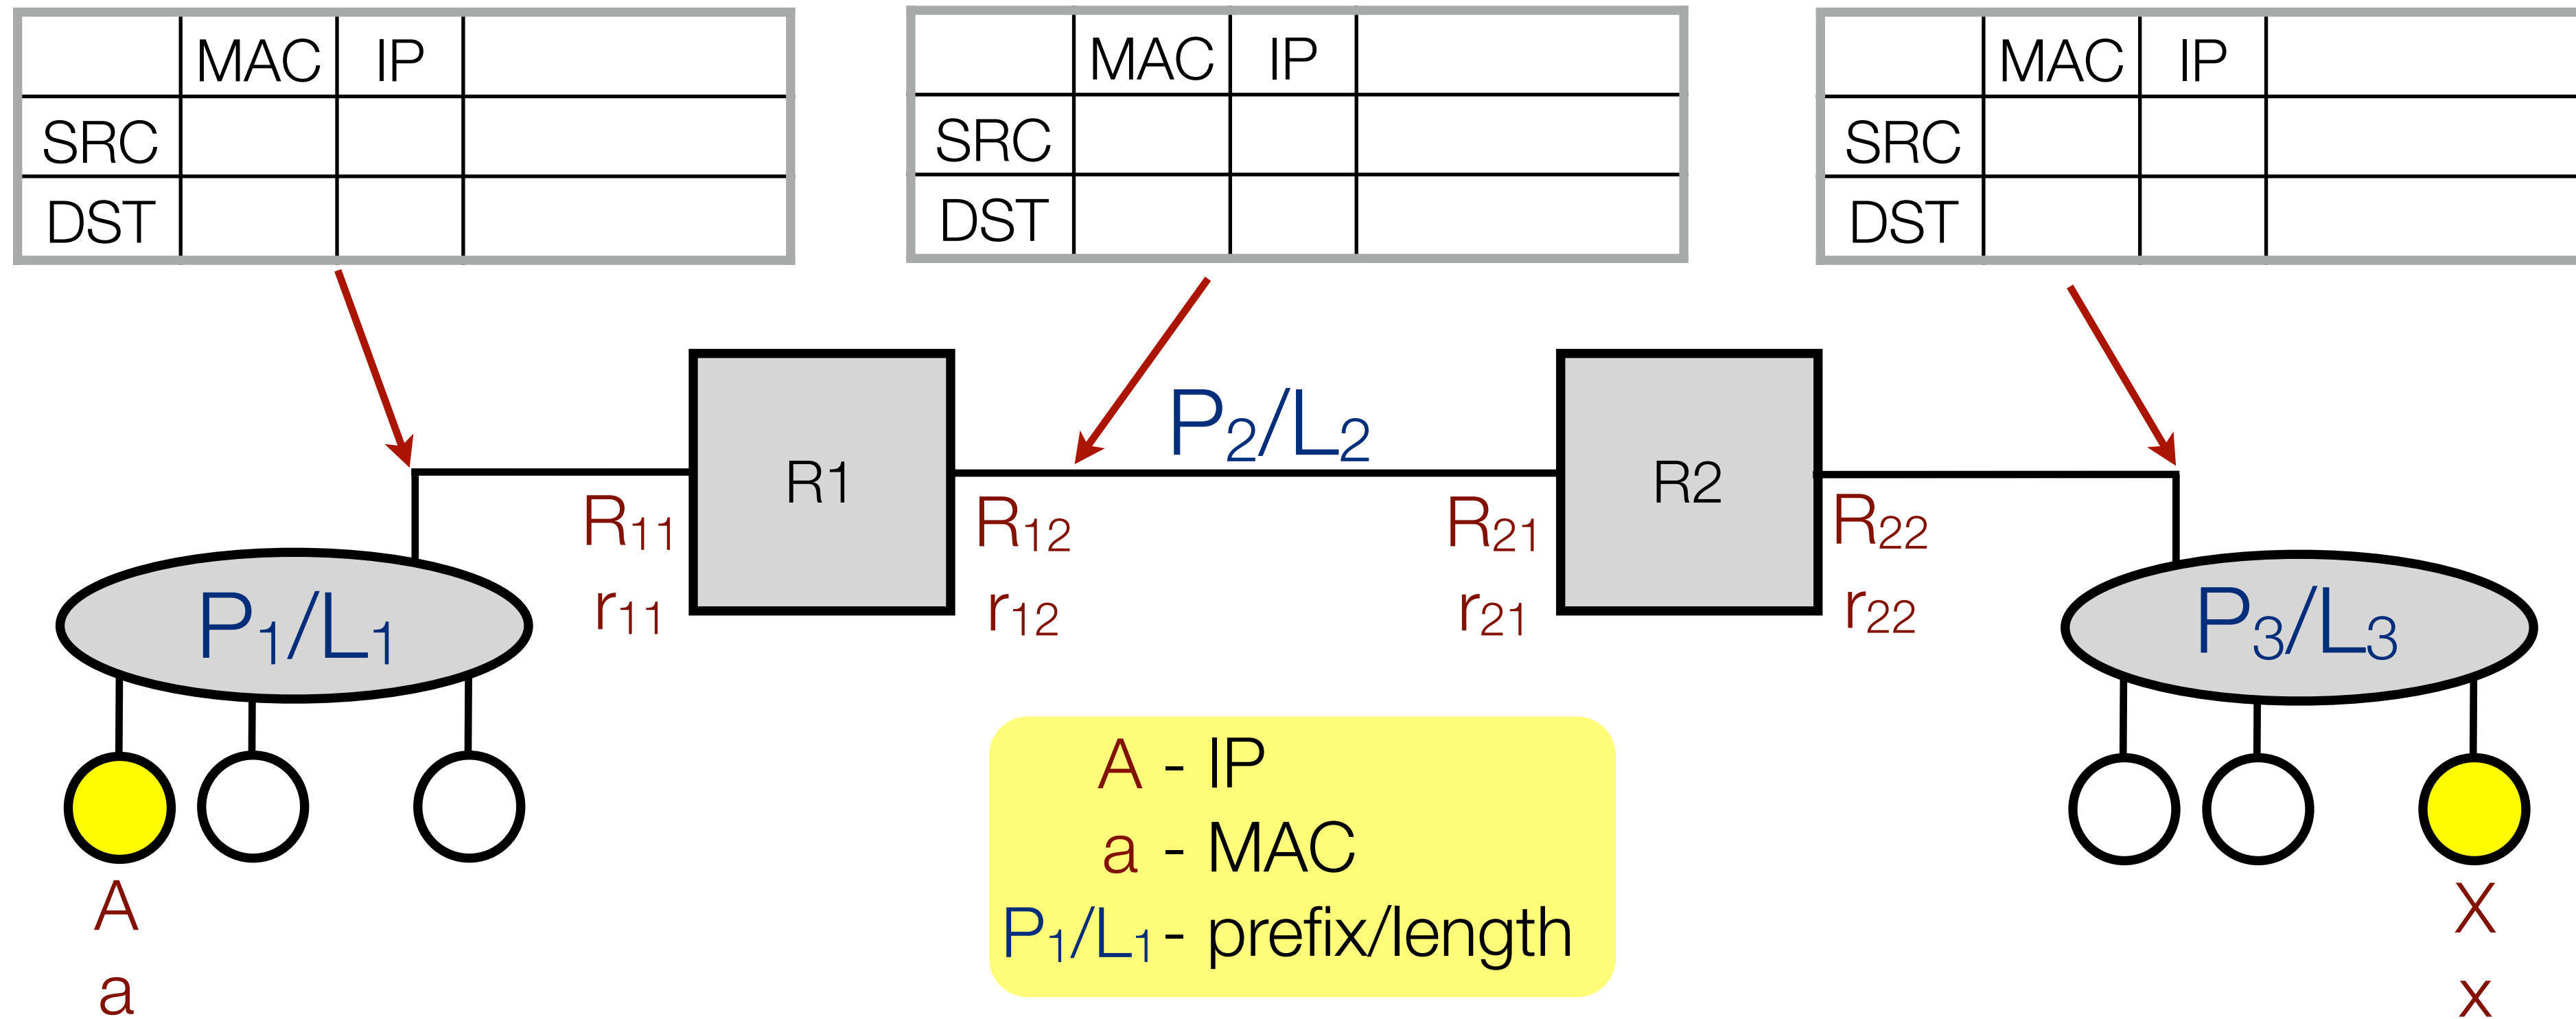
Static IP Routing



For example:

	P_1/L_1	10.0.1.0/24
A	10.0.1.101	R_{11} 10.0.1.1
a	f8:ff:c2:24:6a:0f	r_{11} aa:c2:74:08:cf:12

Static IP Routing

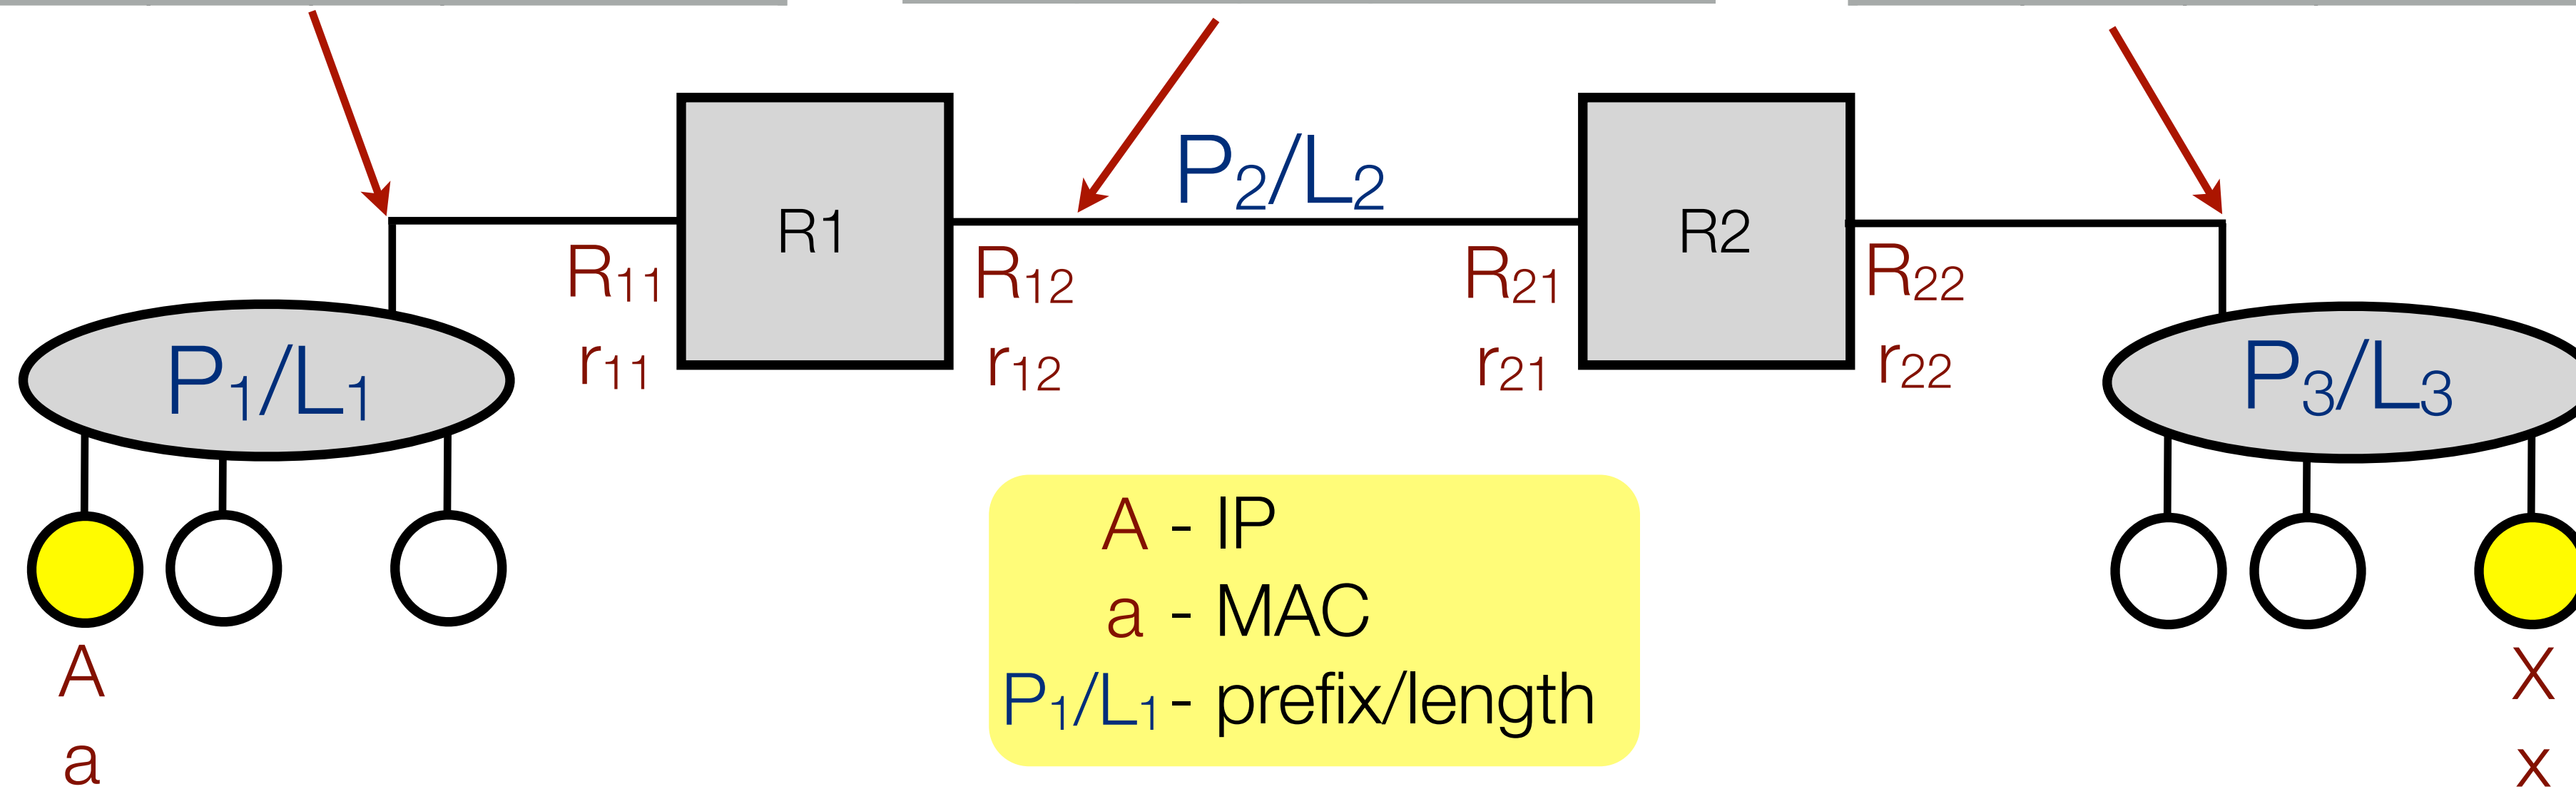


Static IP Routing

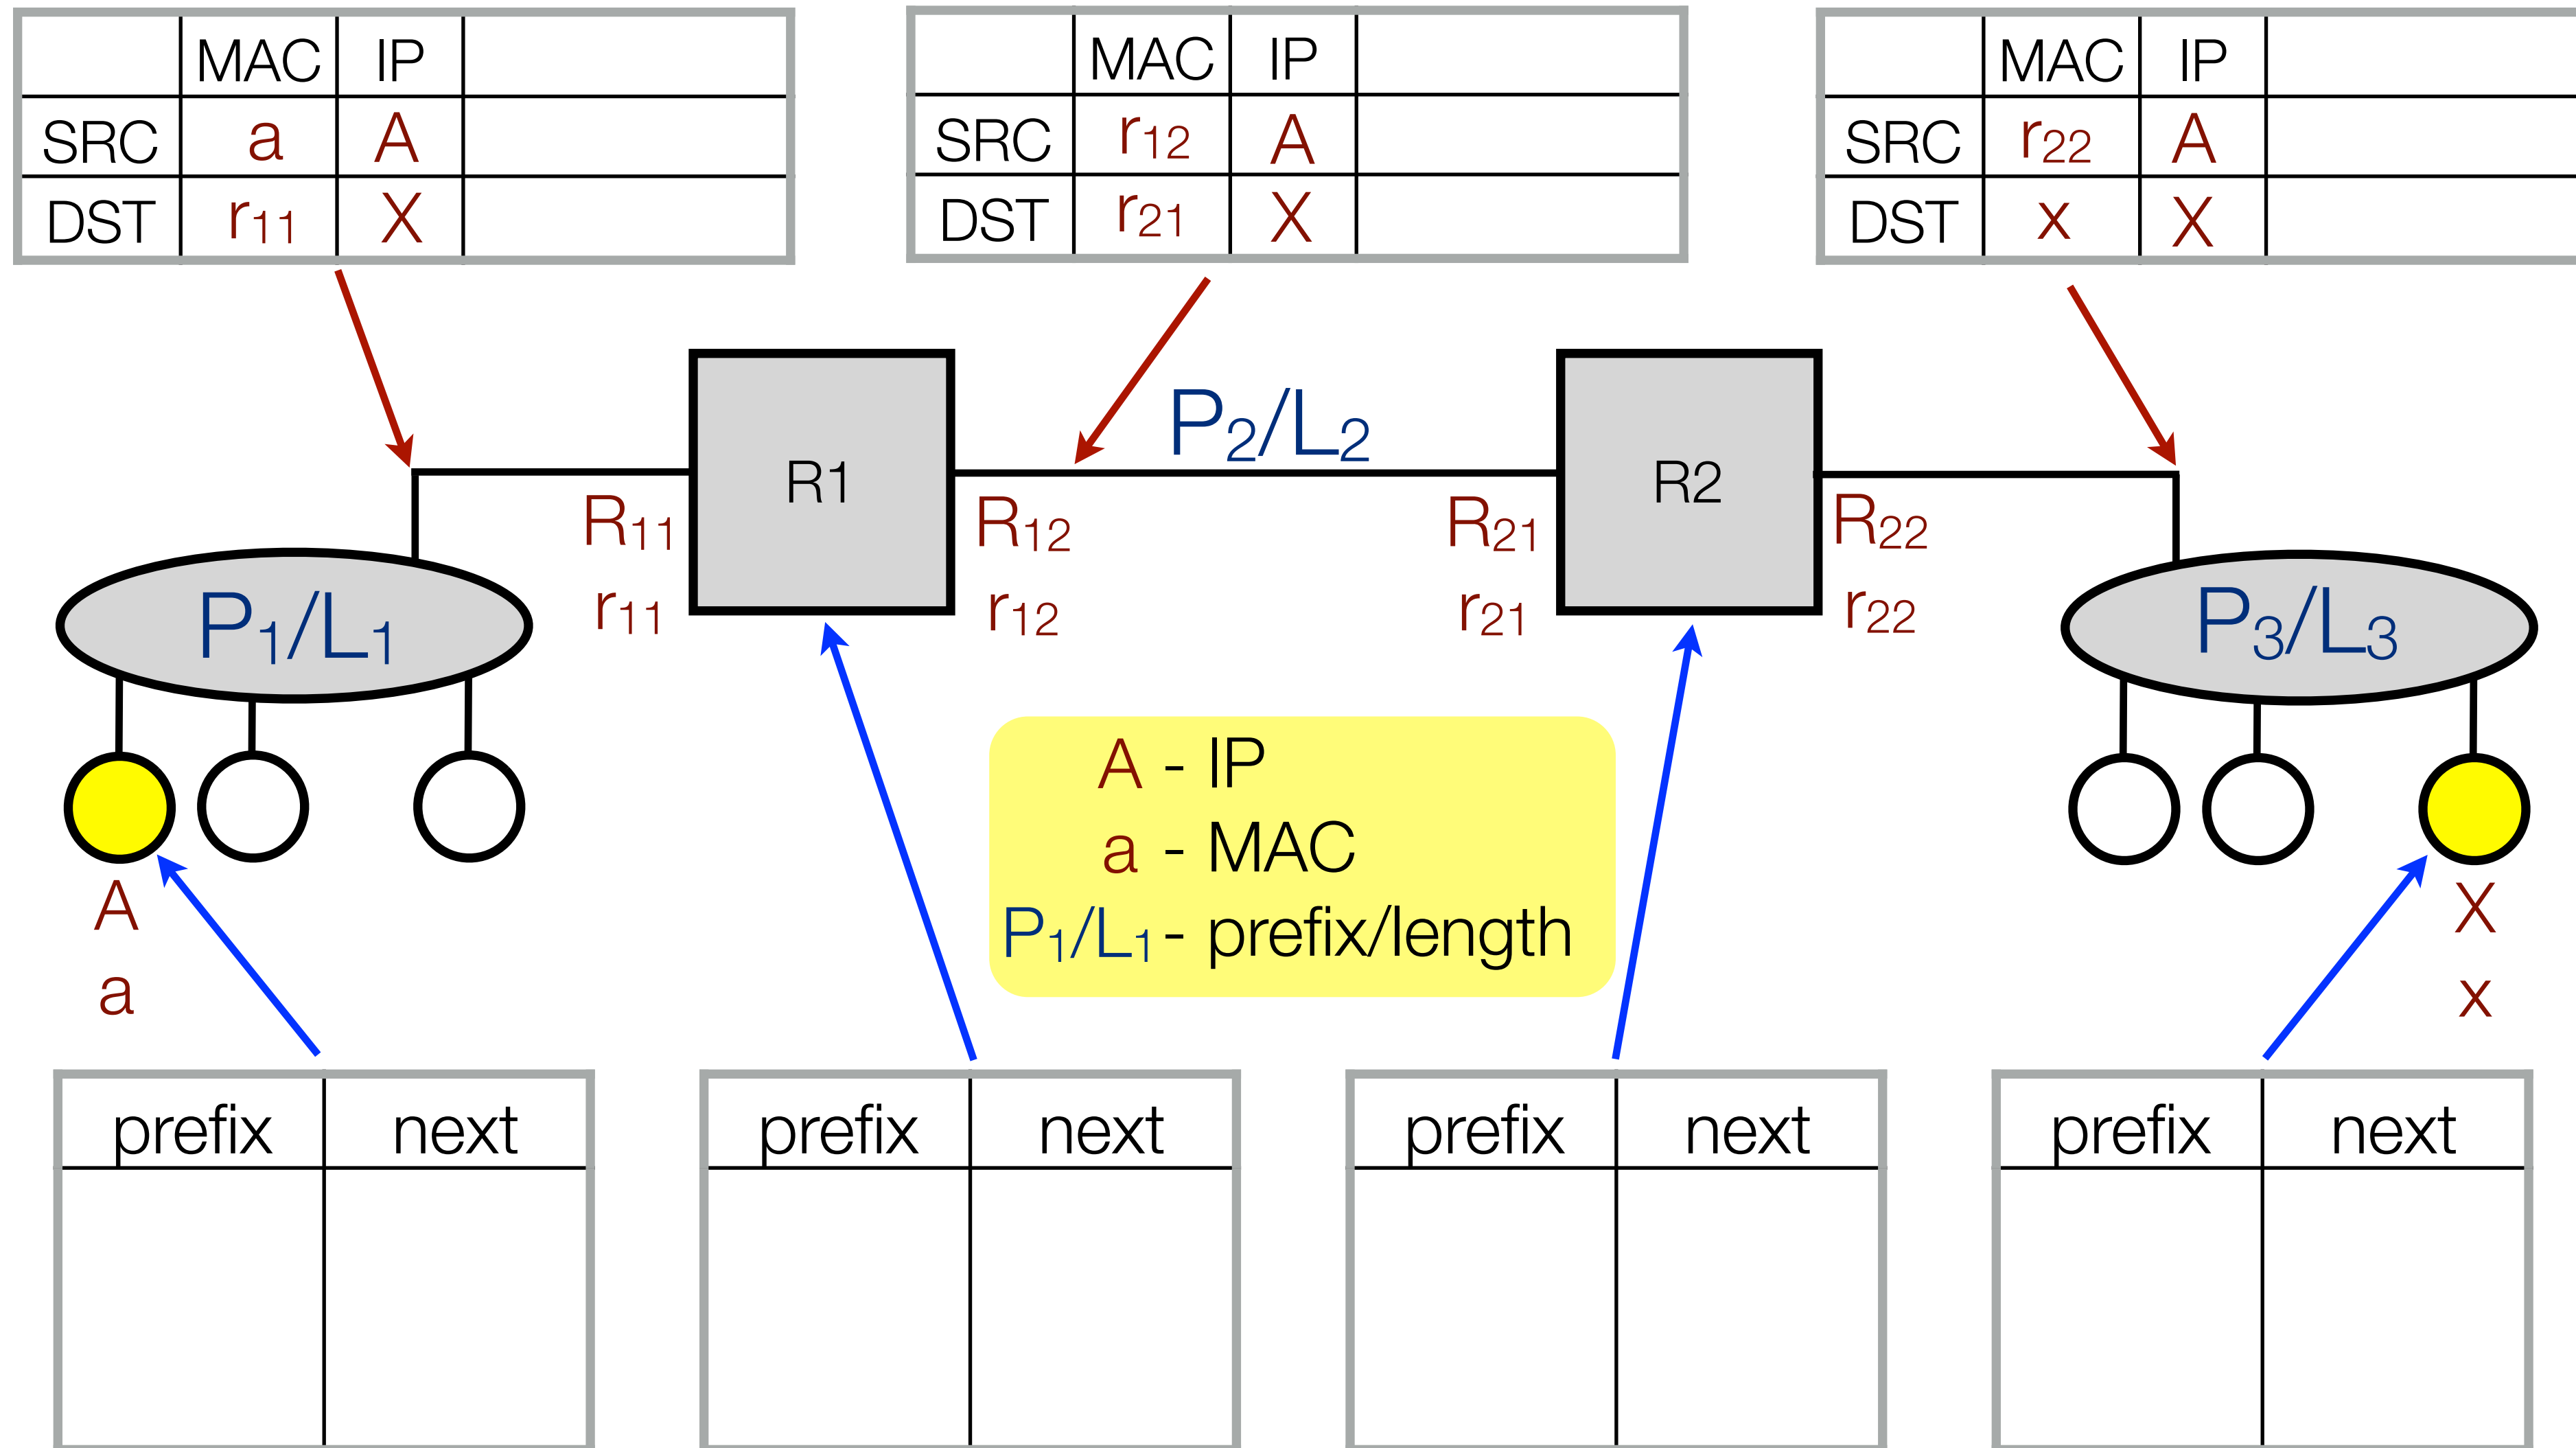
	MAC	IP	
SRC	a	A	
DST	r ₁₁	X	

	MAC	IP	
SRC	r ₁₂	A	
DST	r ₂₁	X	

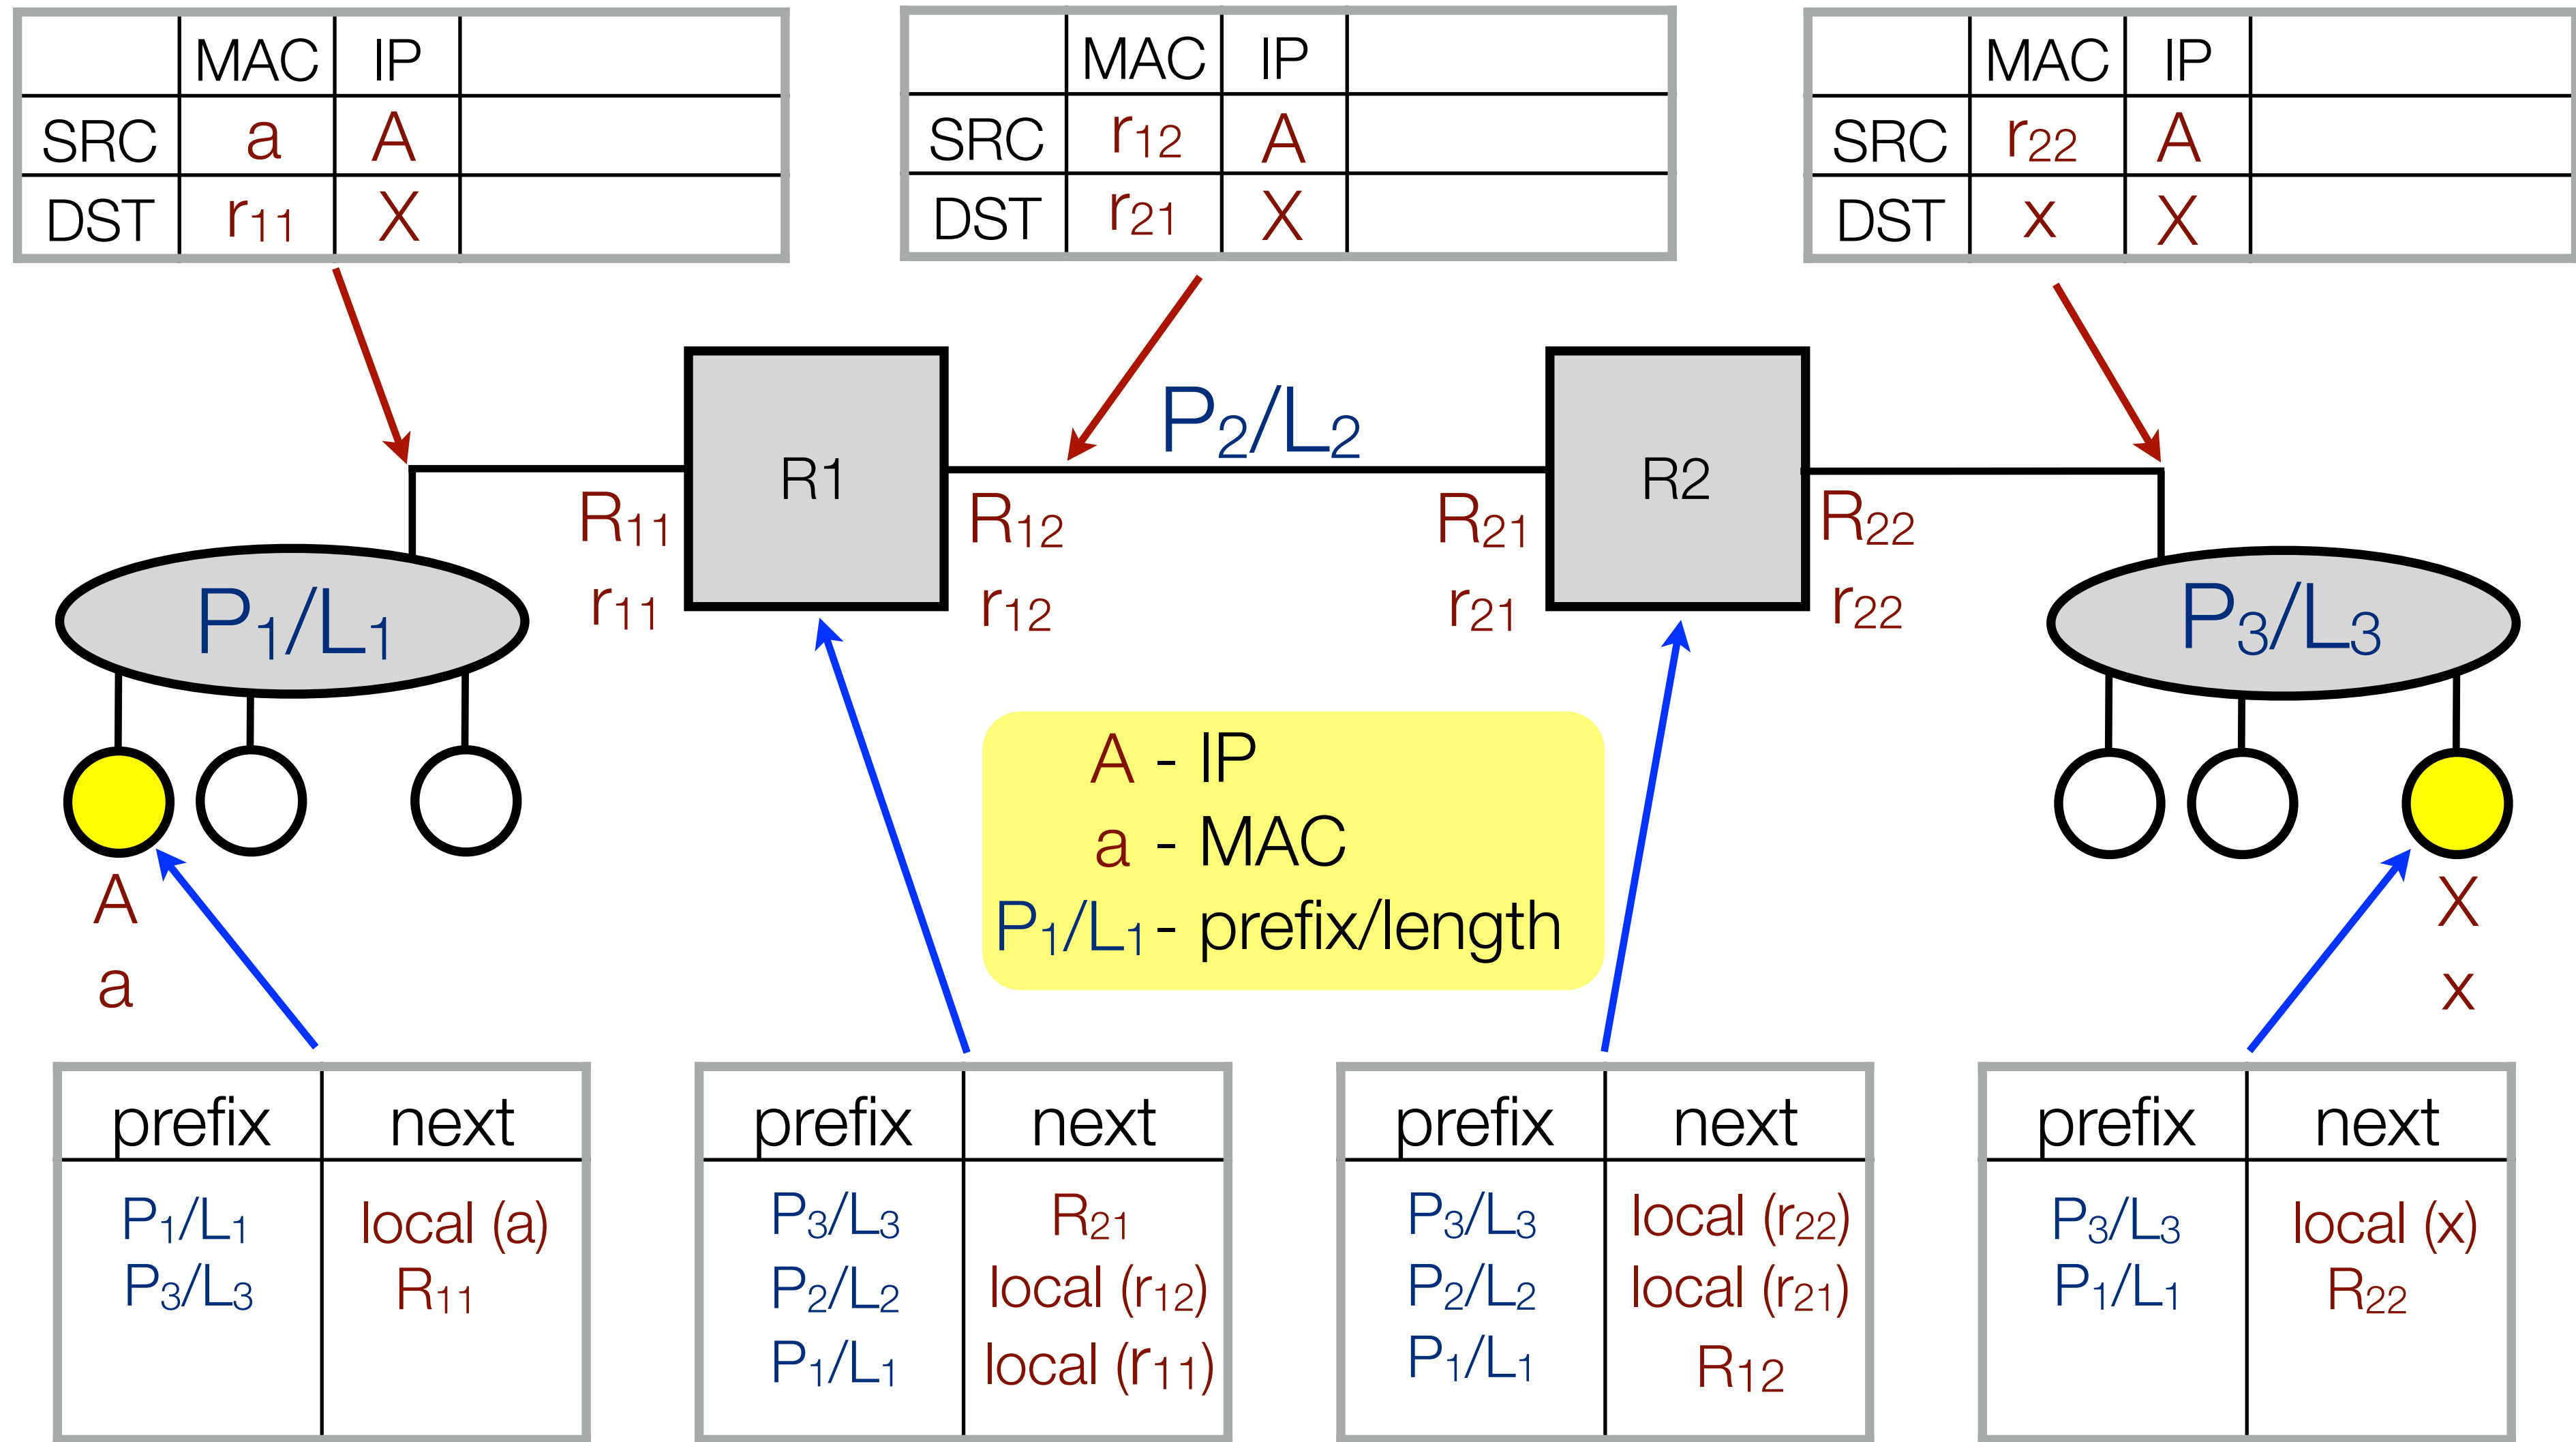
	MAC	IP	
SRC	r ₂₂	A	
DST	x	X	



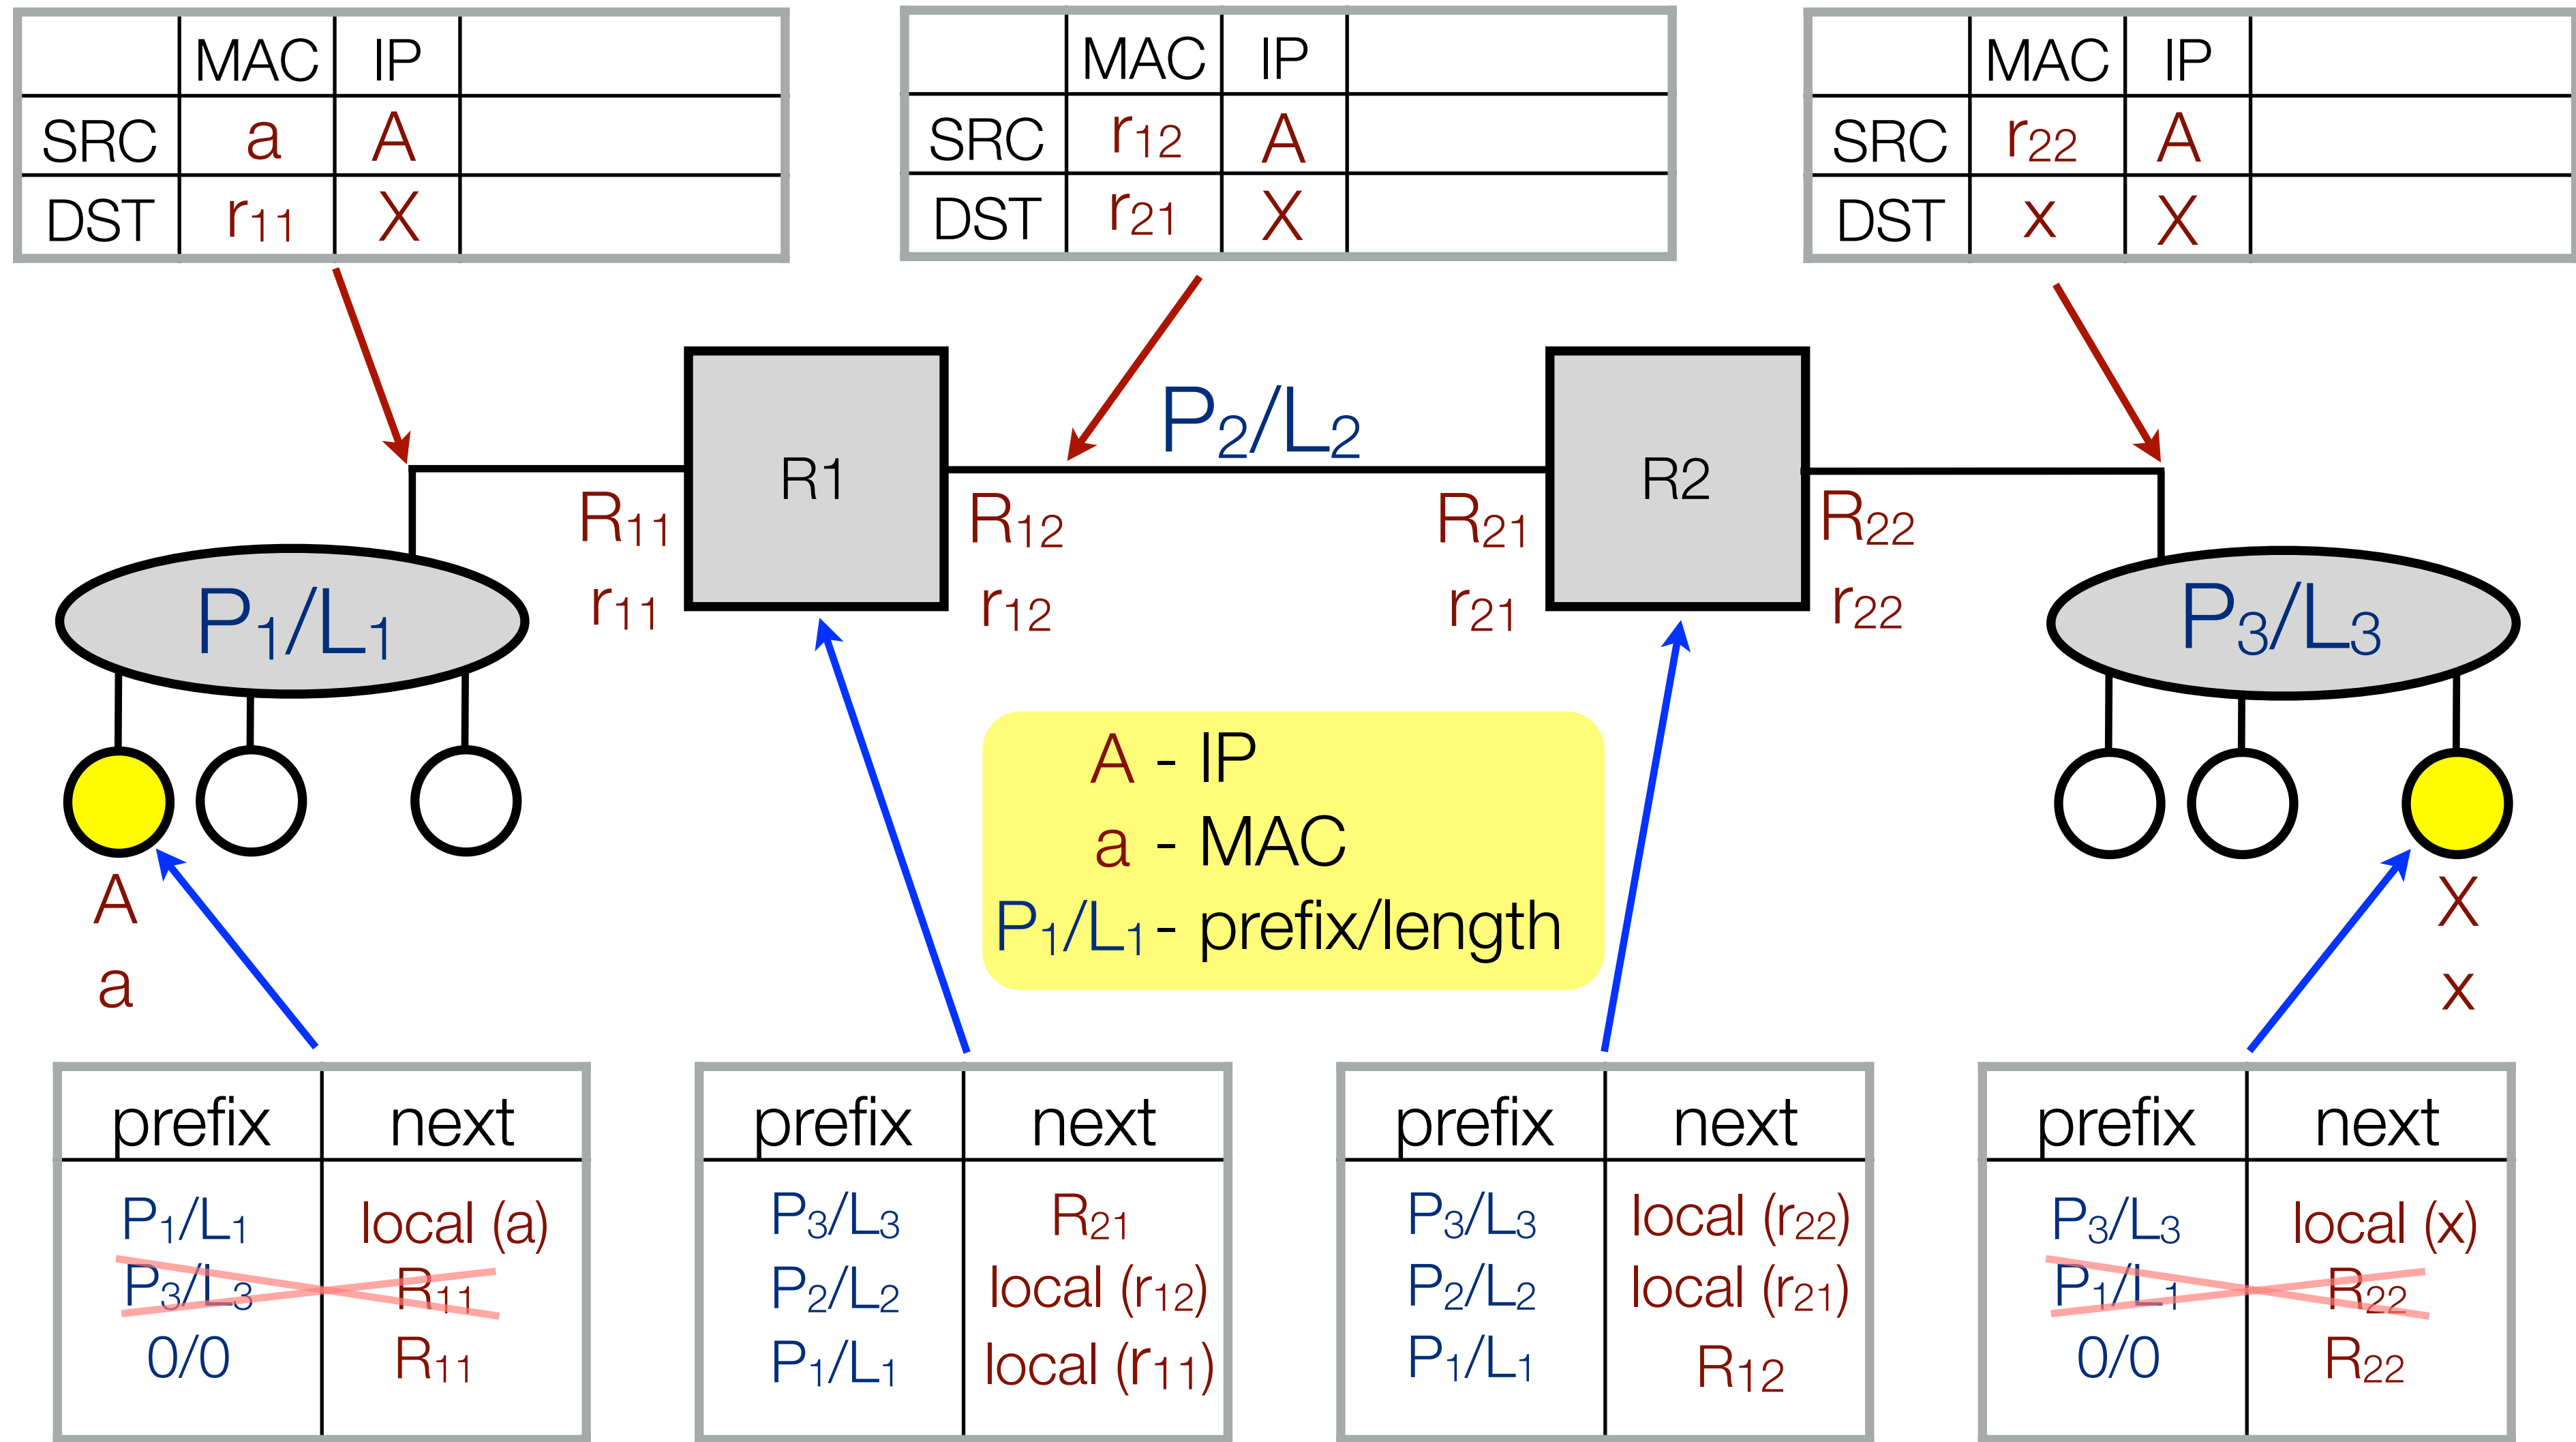
Static IP Routing



Static IP Routing



Static IP Routing



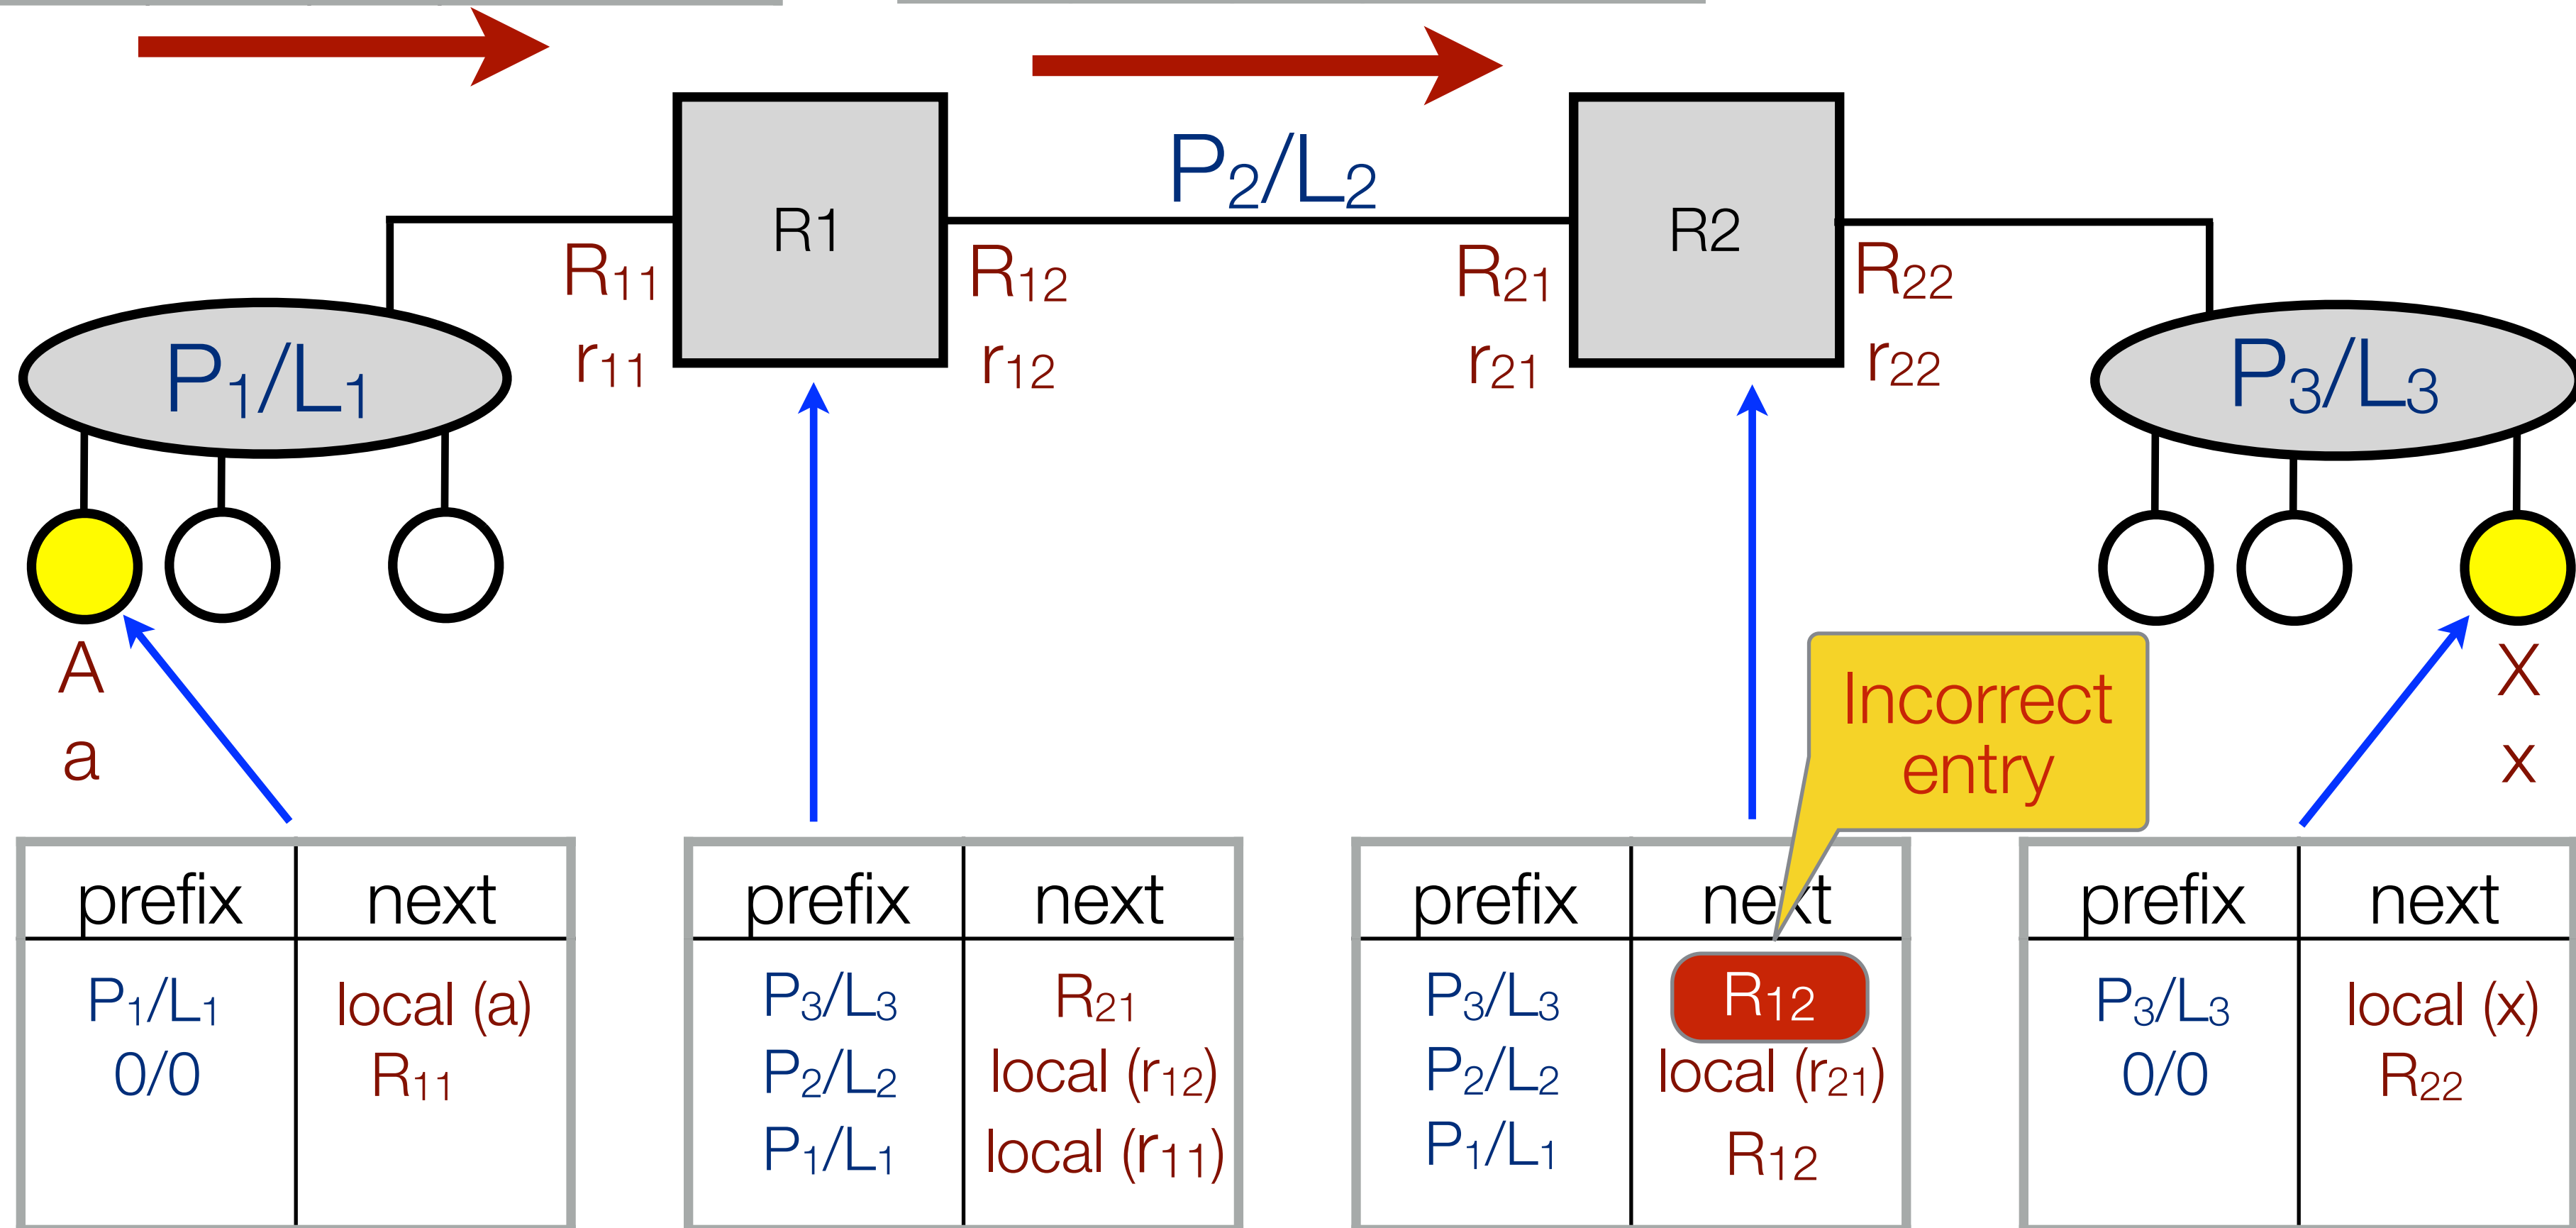
Routing loop



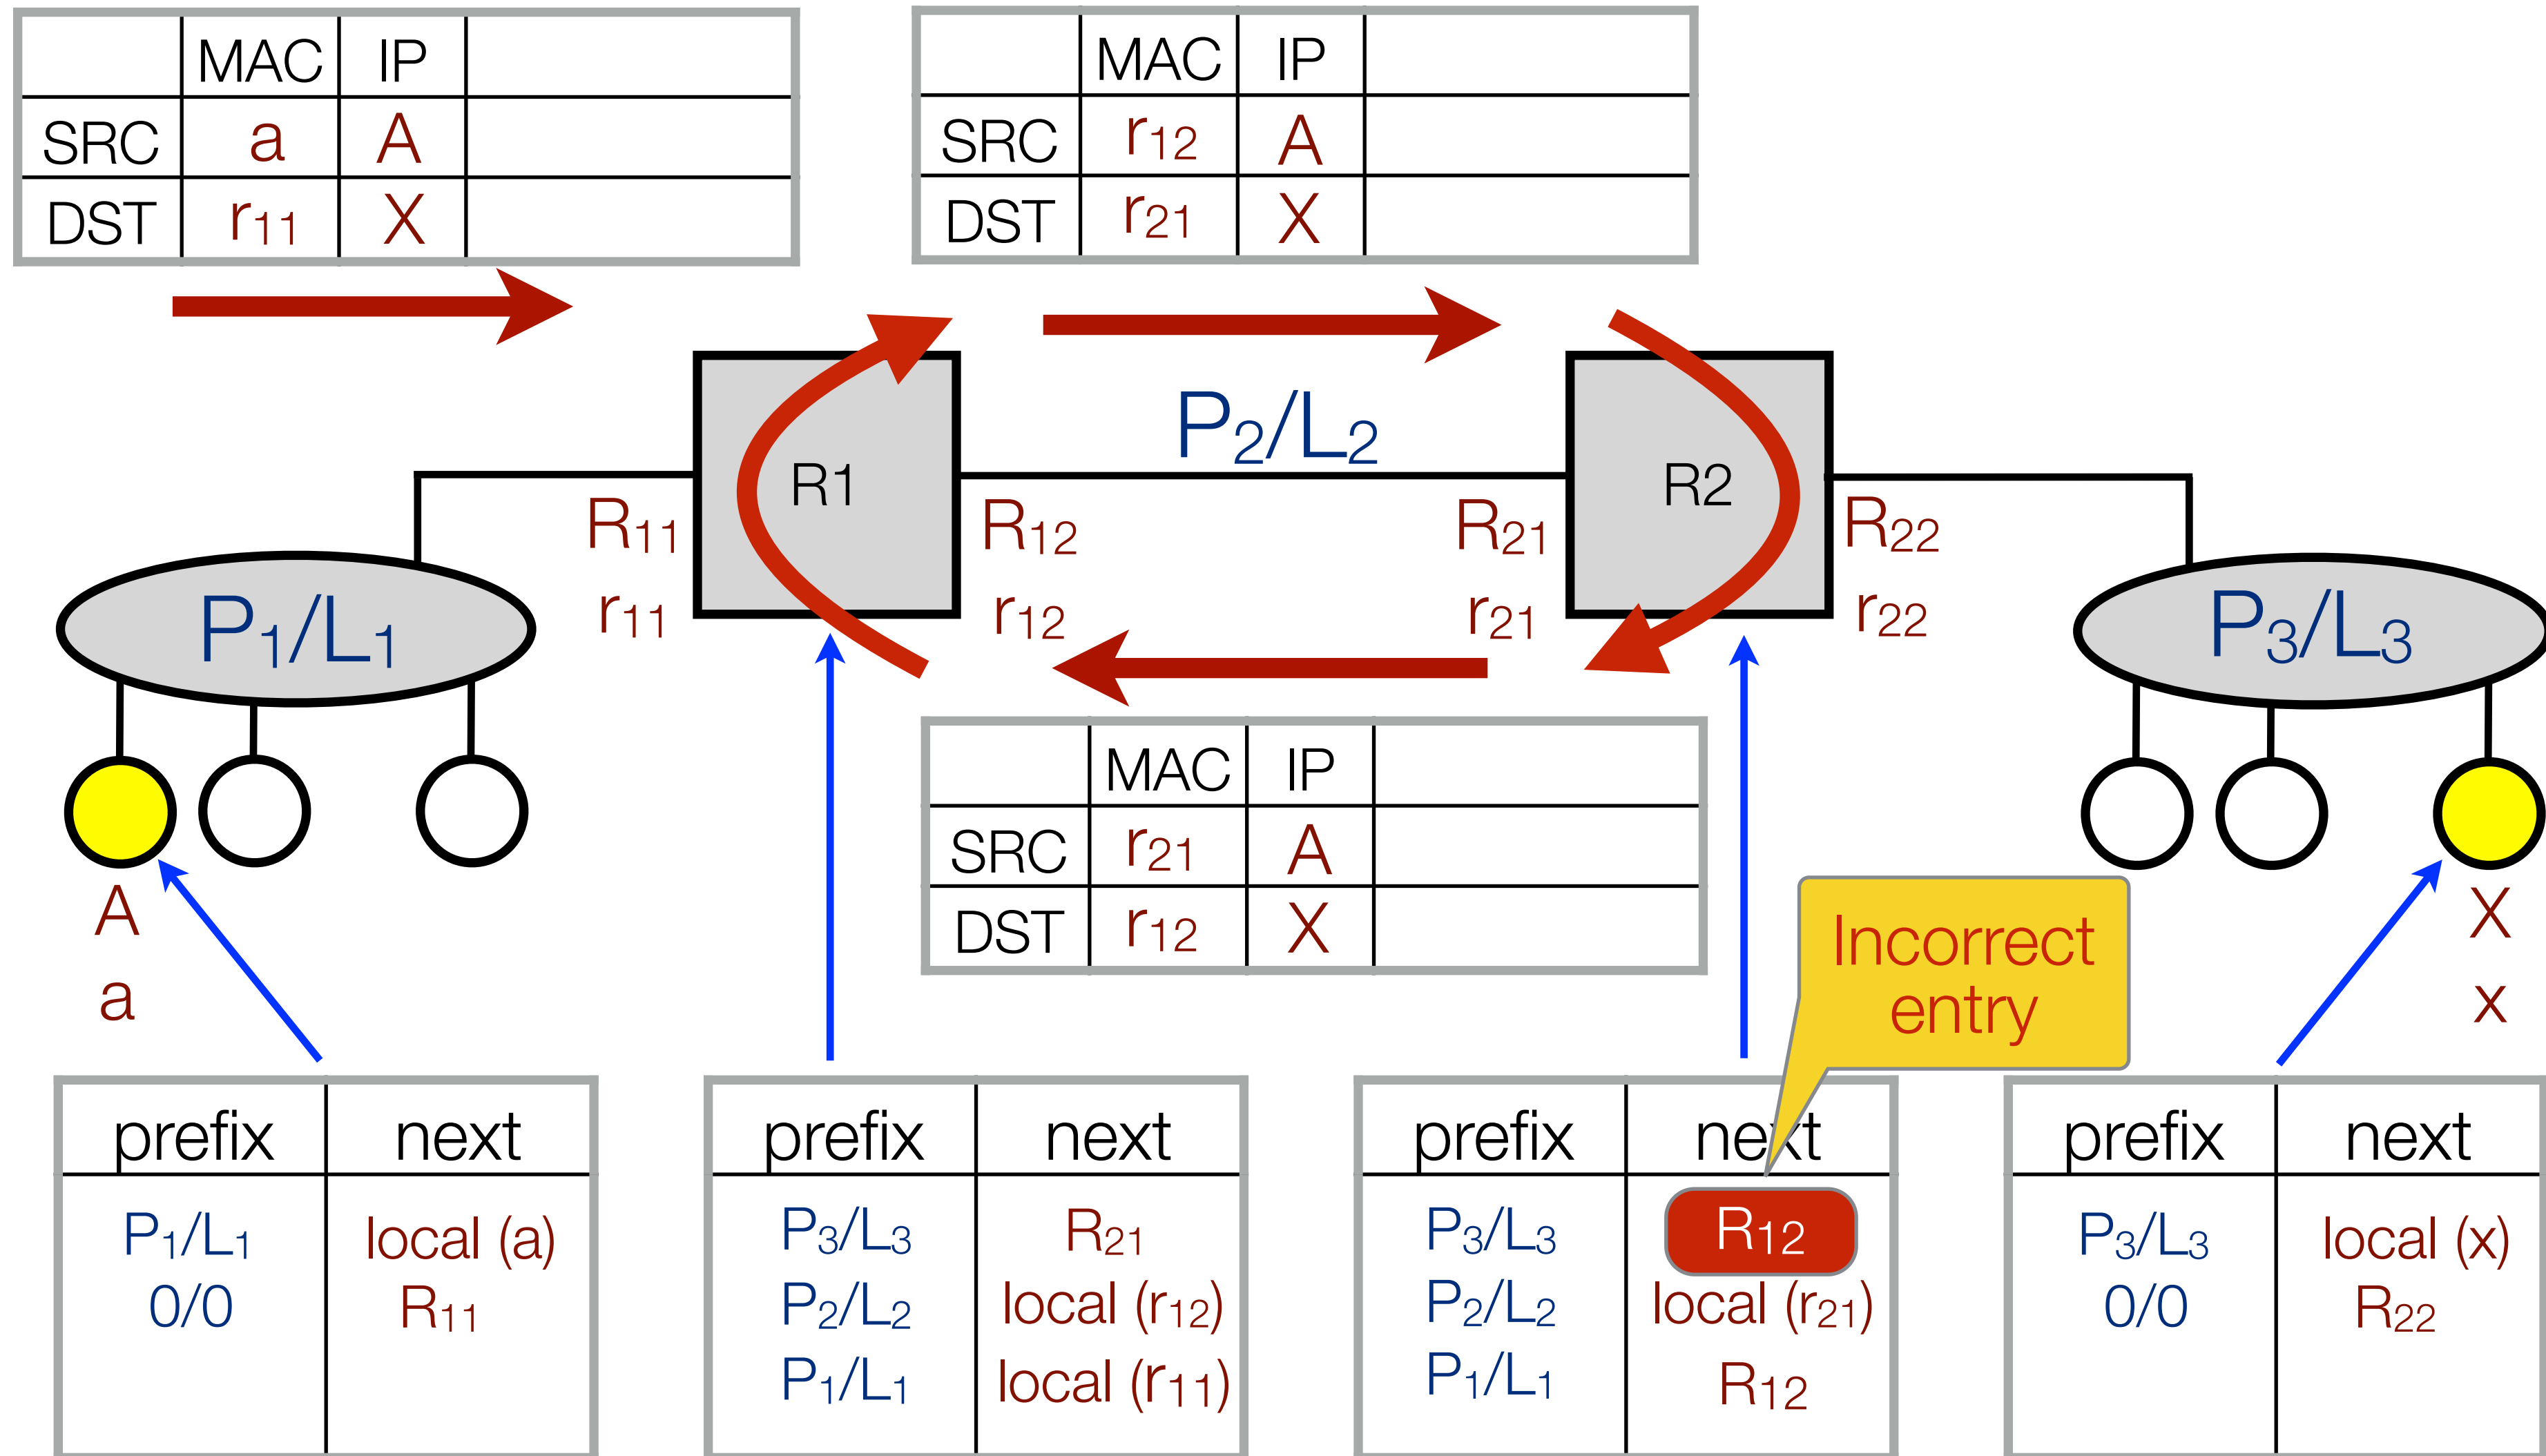
Routing loop

	MAC	IP	
SRC	a	A	
DST	r ₁₁	X	

	MAC	IP	
SRC	r ₁₂	A	
DST	r ₂₁	X	



Routing loop

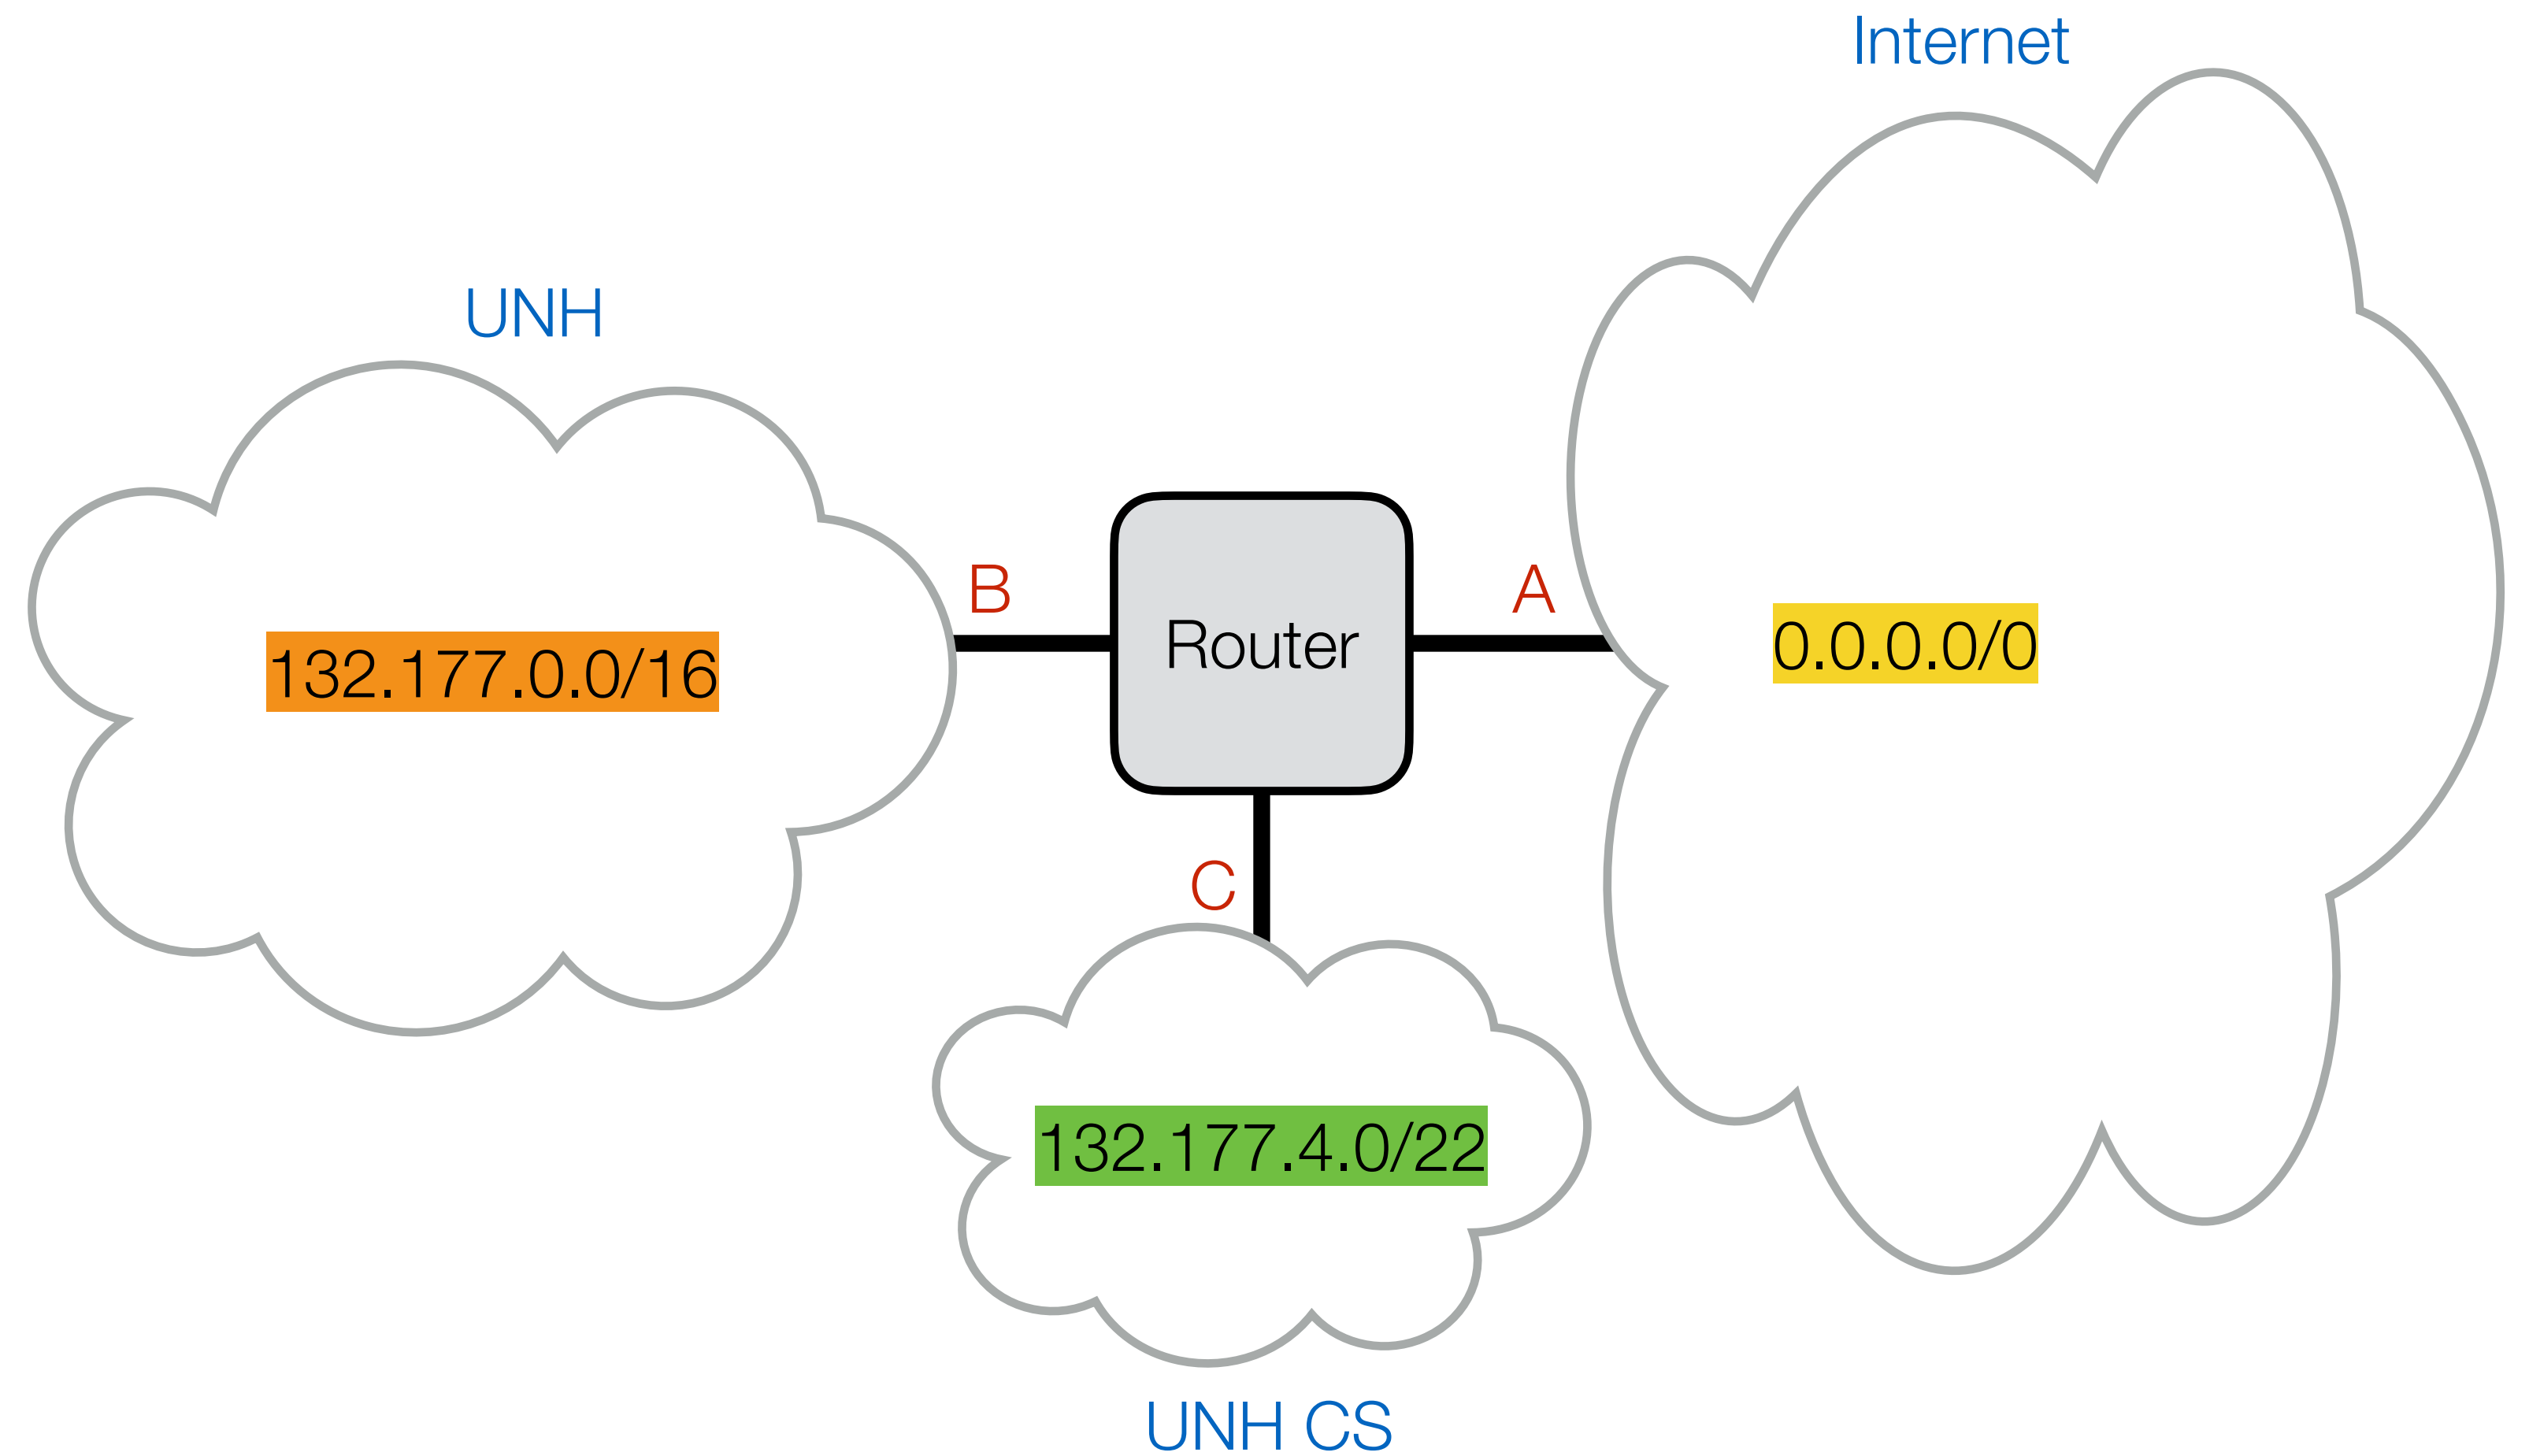


No route to host

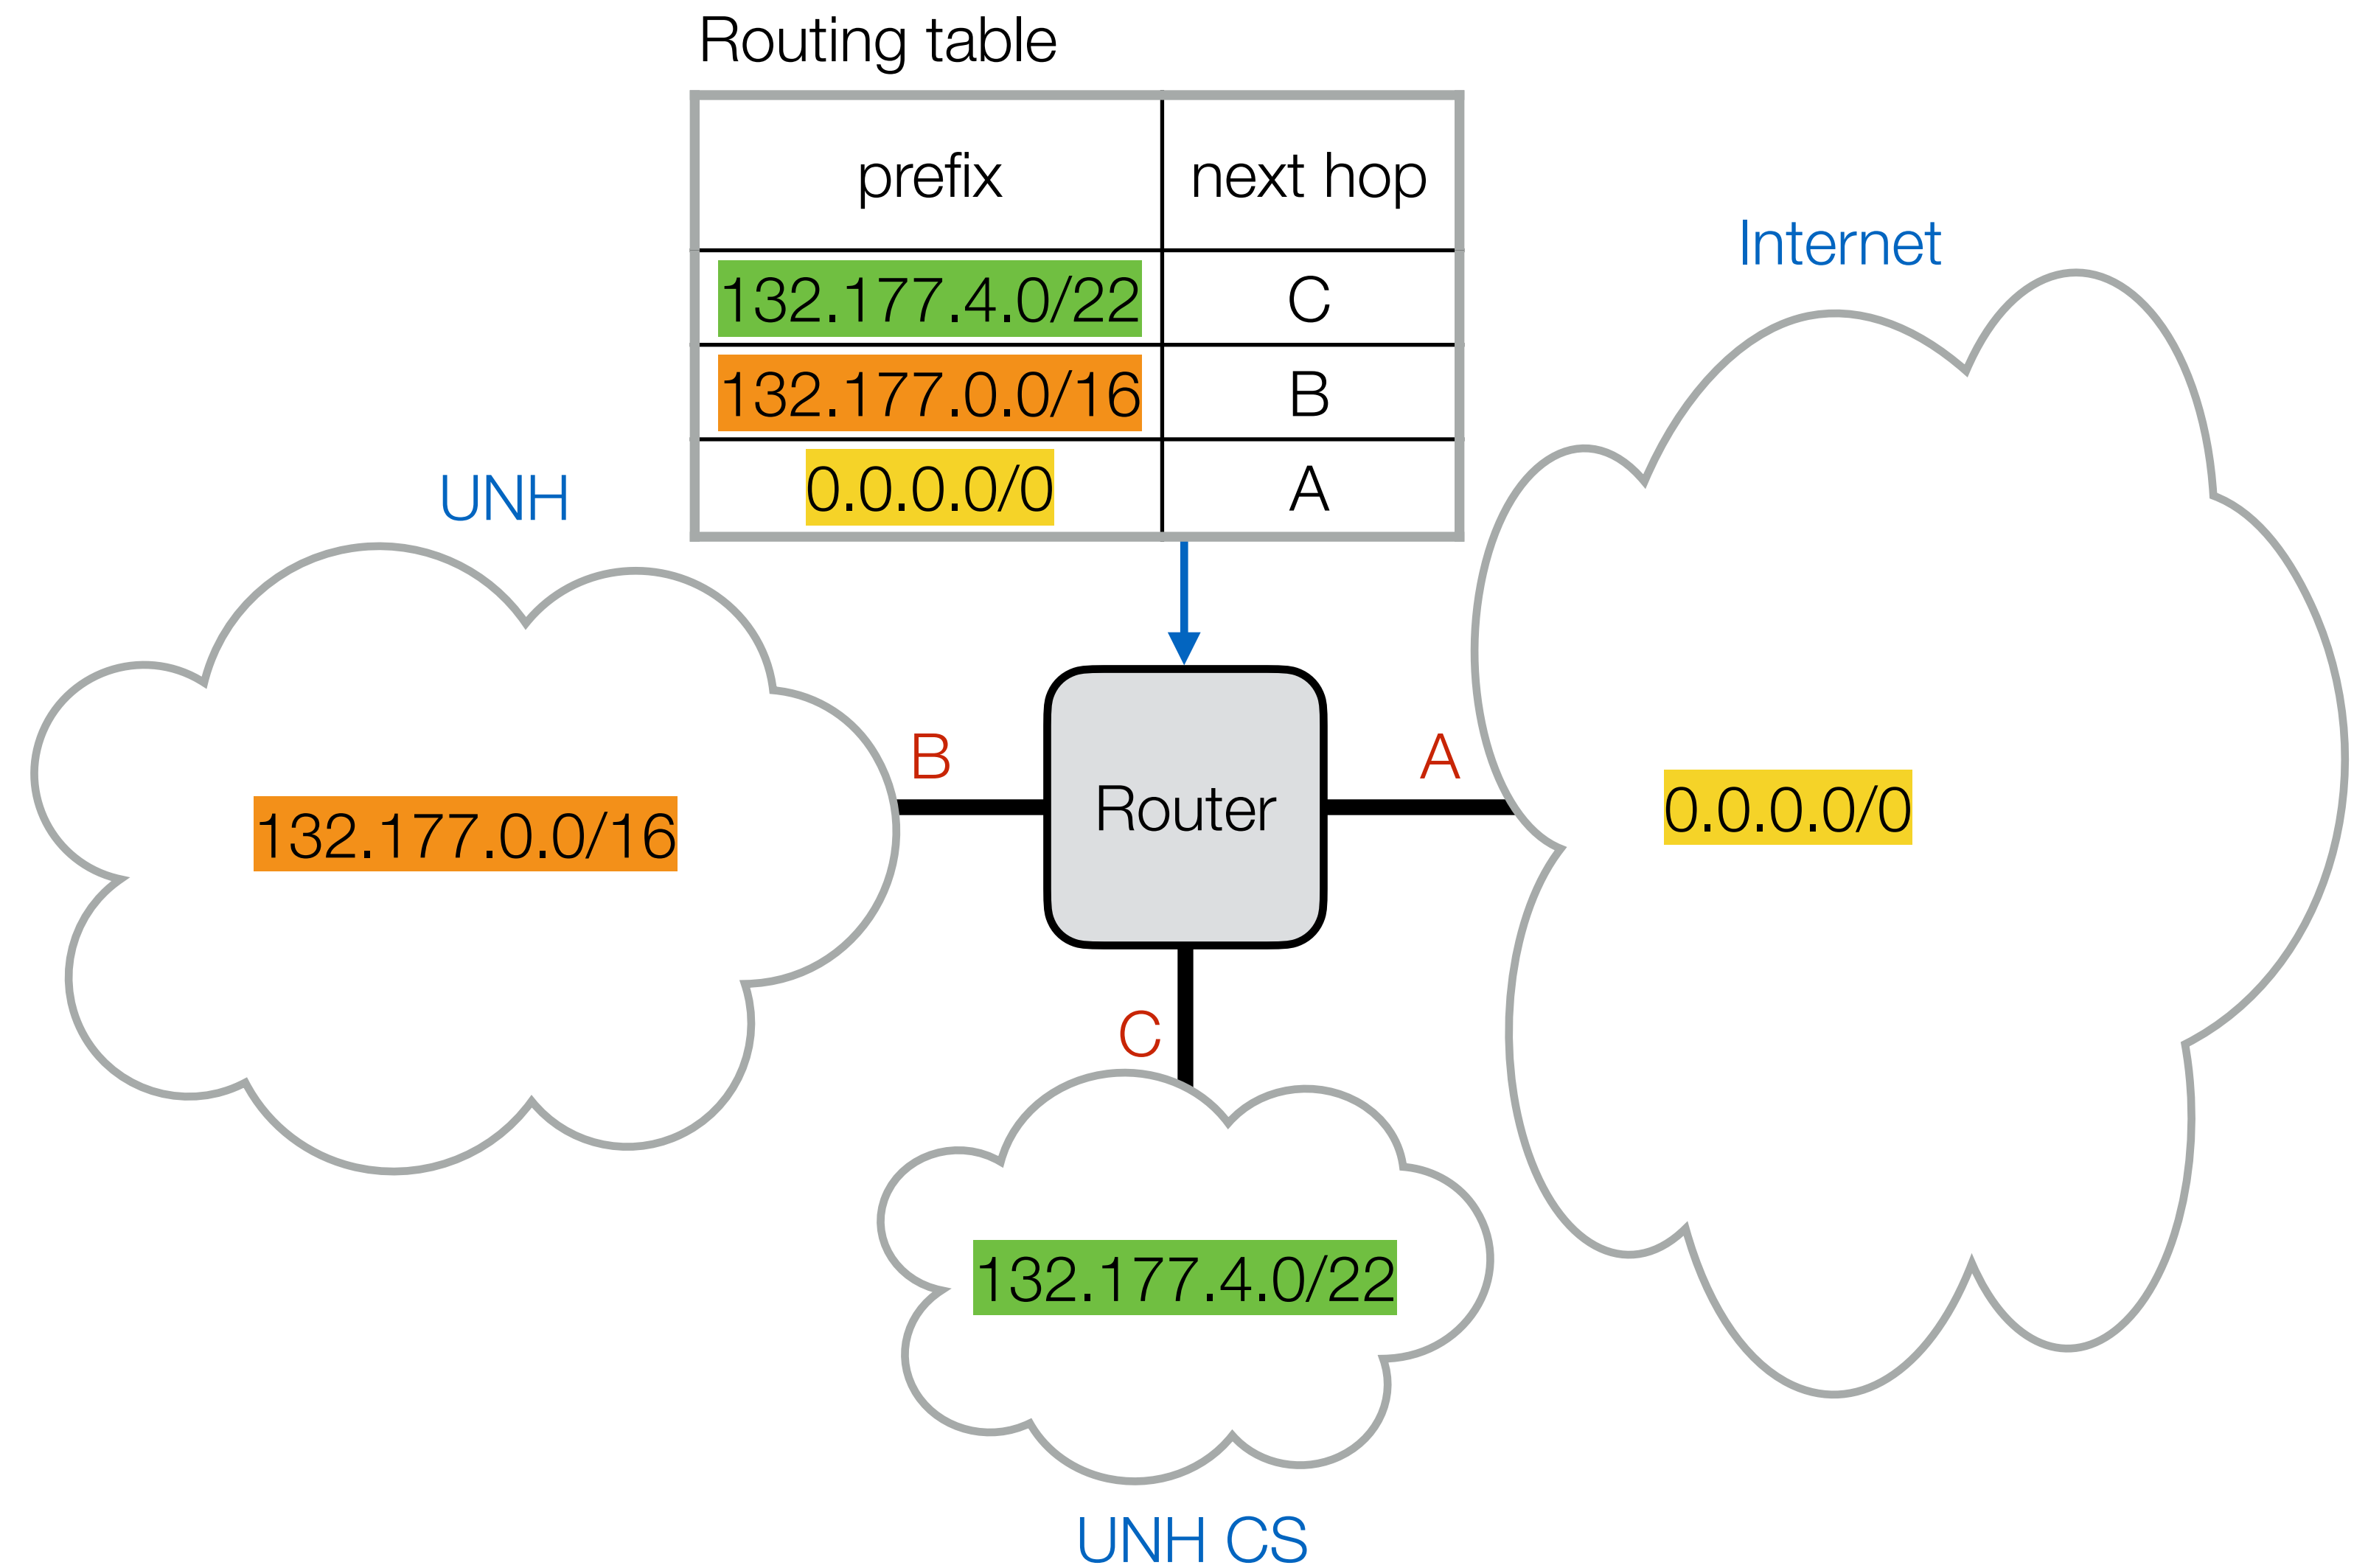


Photo courtesy of Chip McNaughton, UNH-CS

Routing example

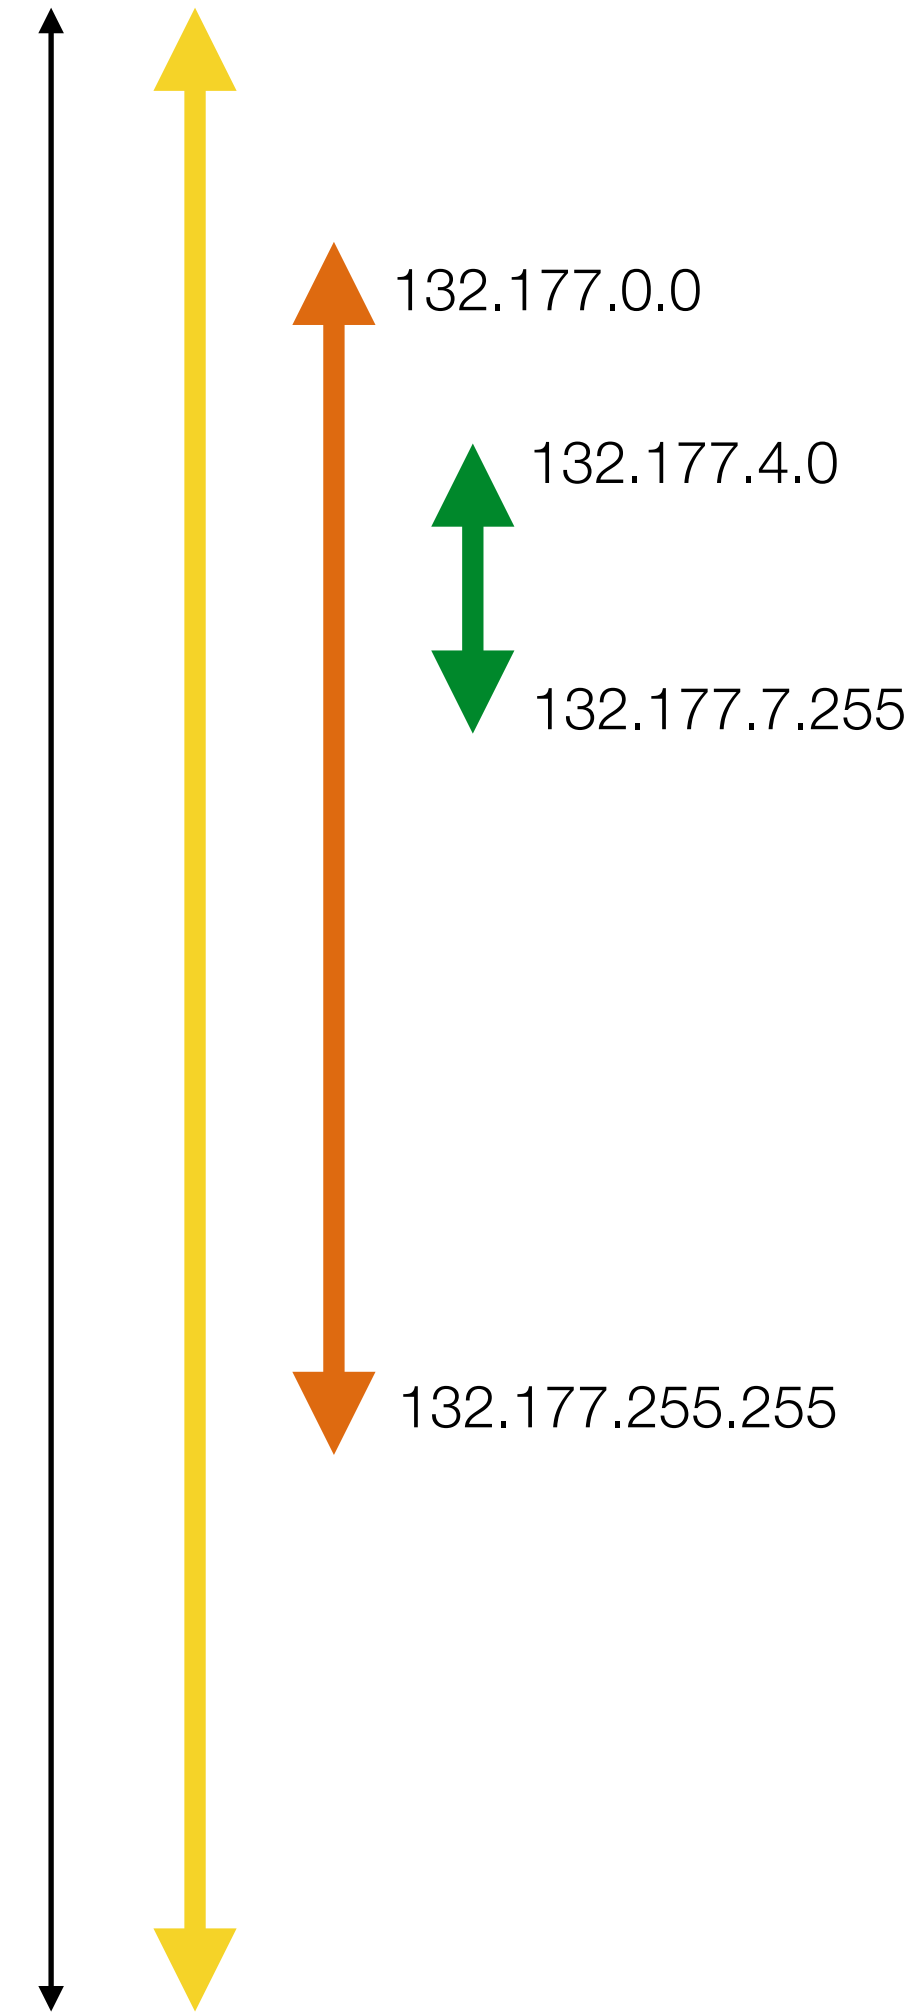


Routing example



Routing example

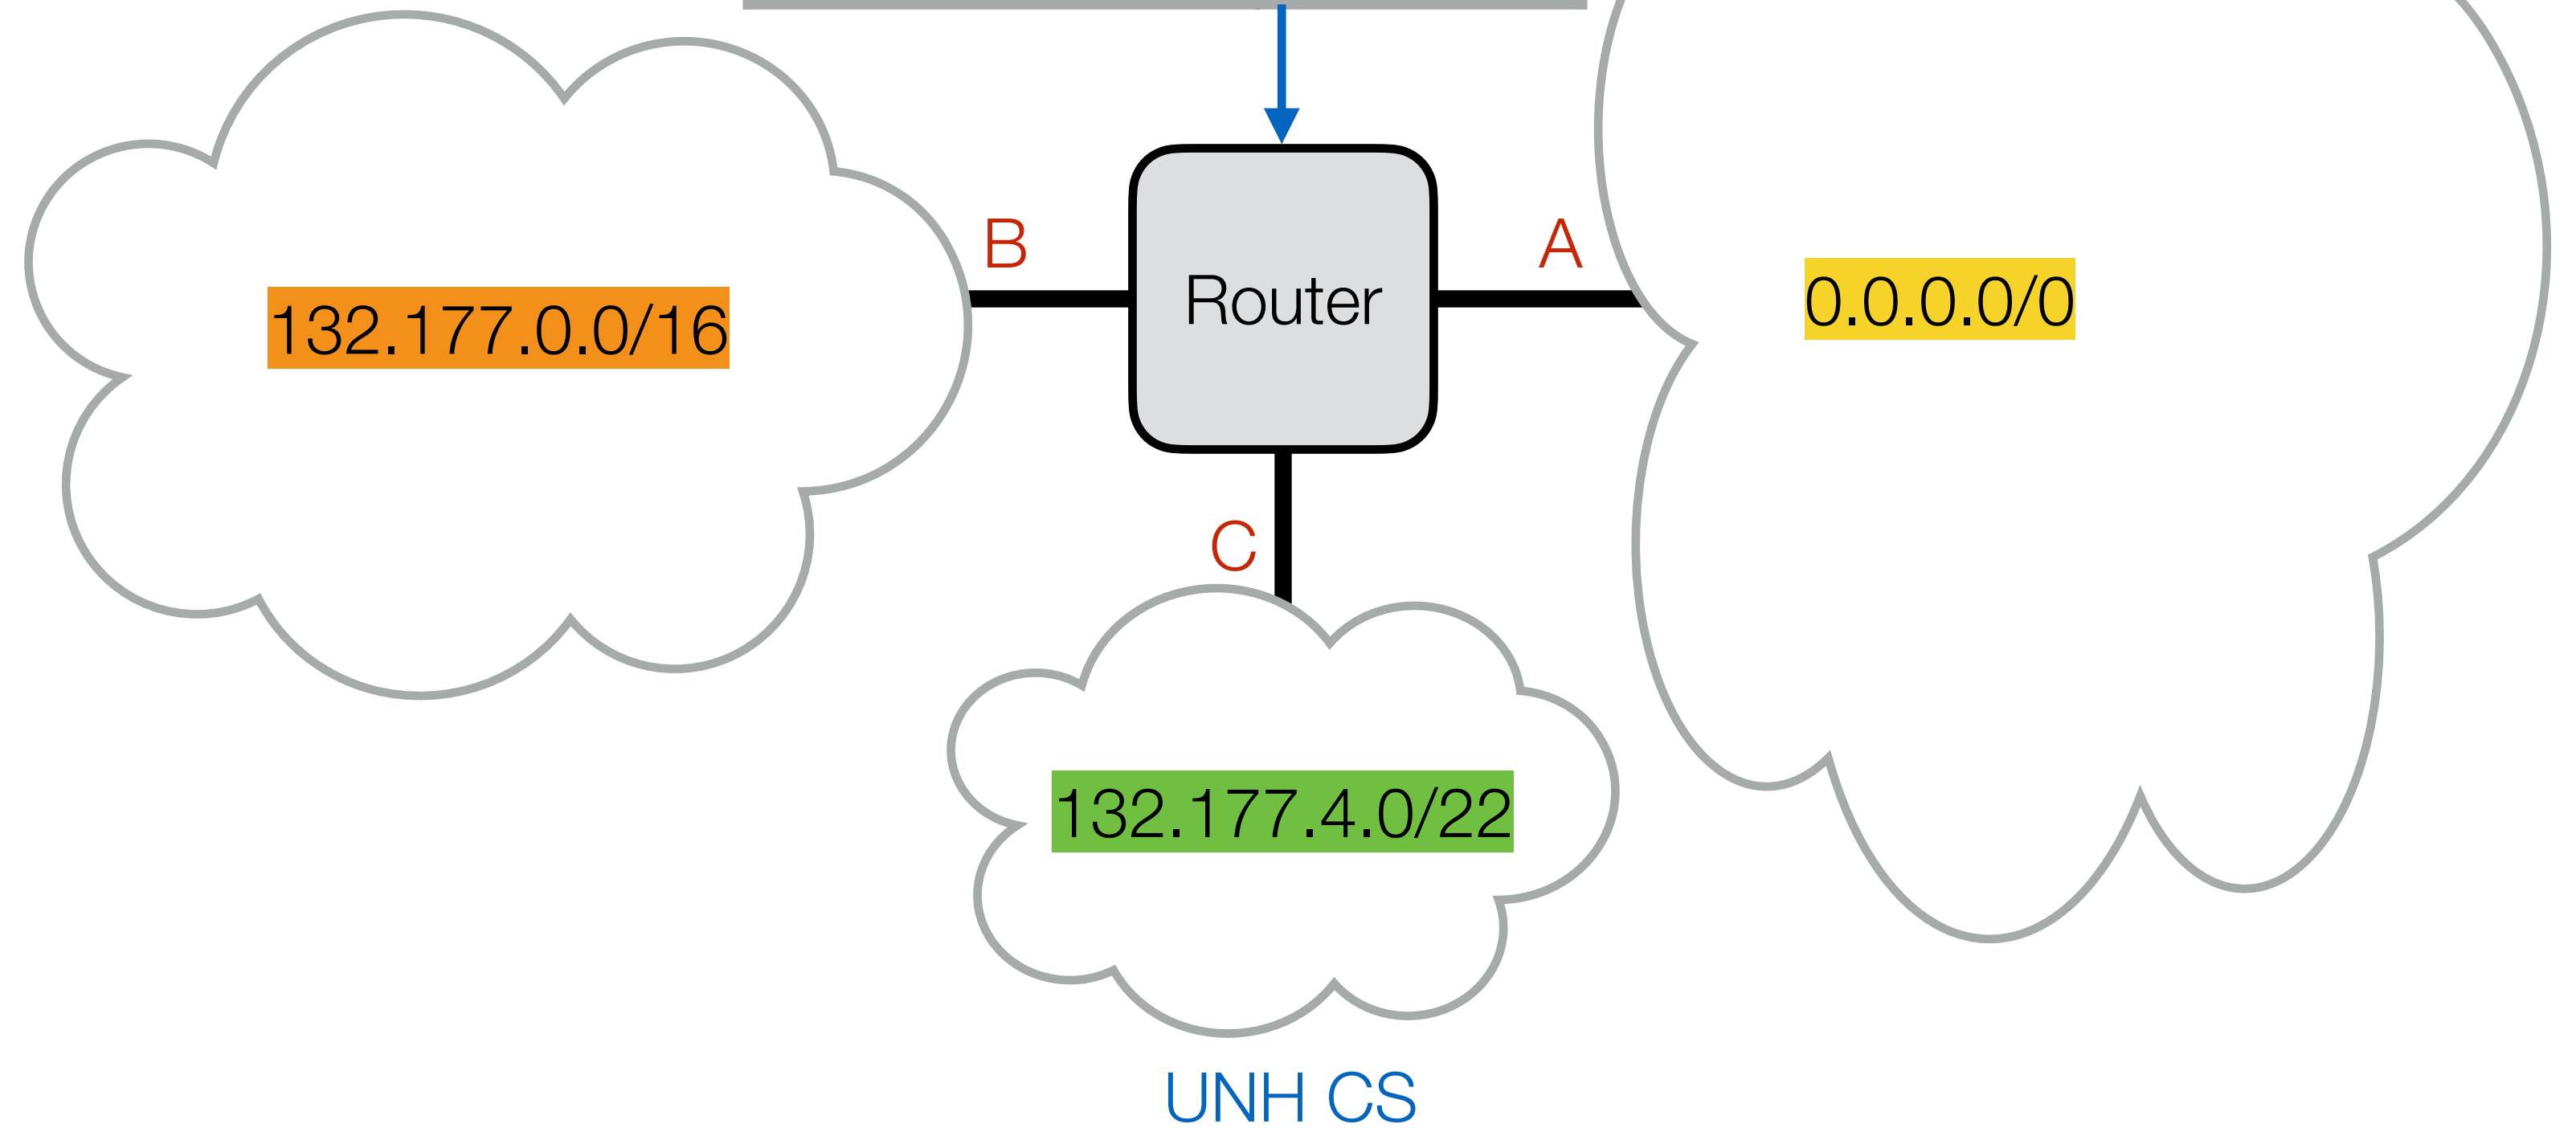
IP address space



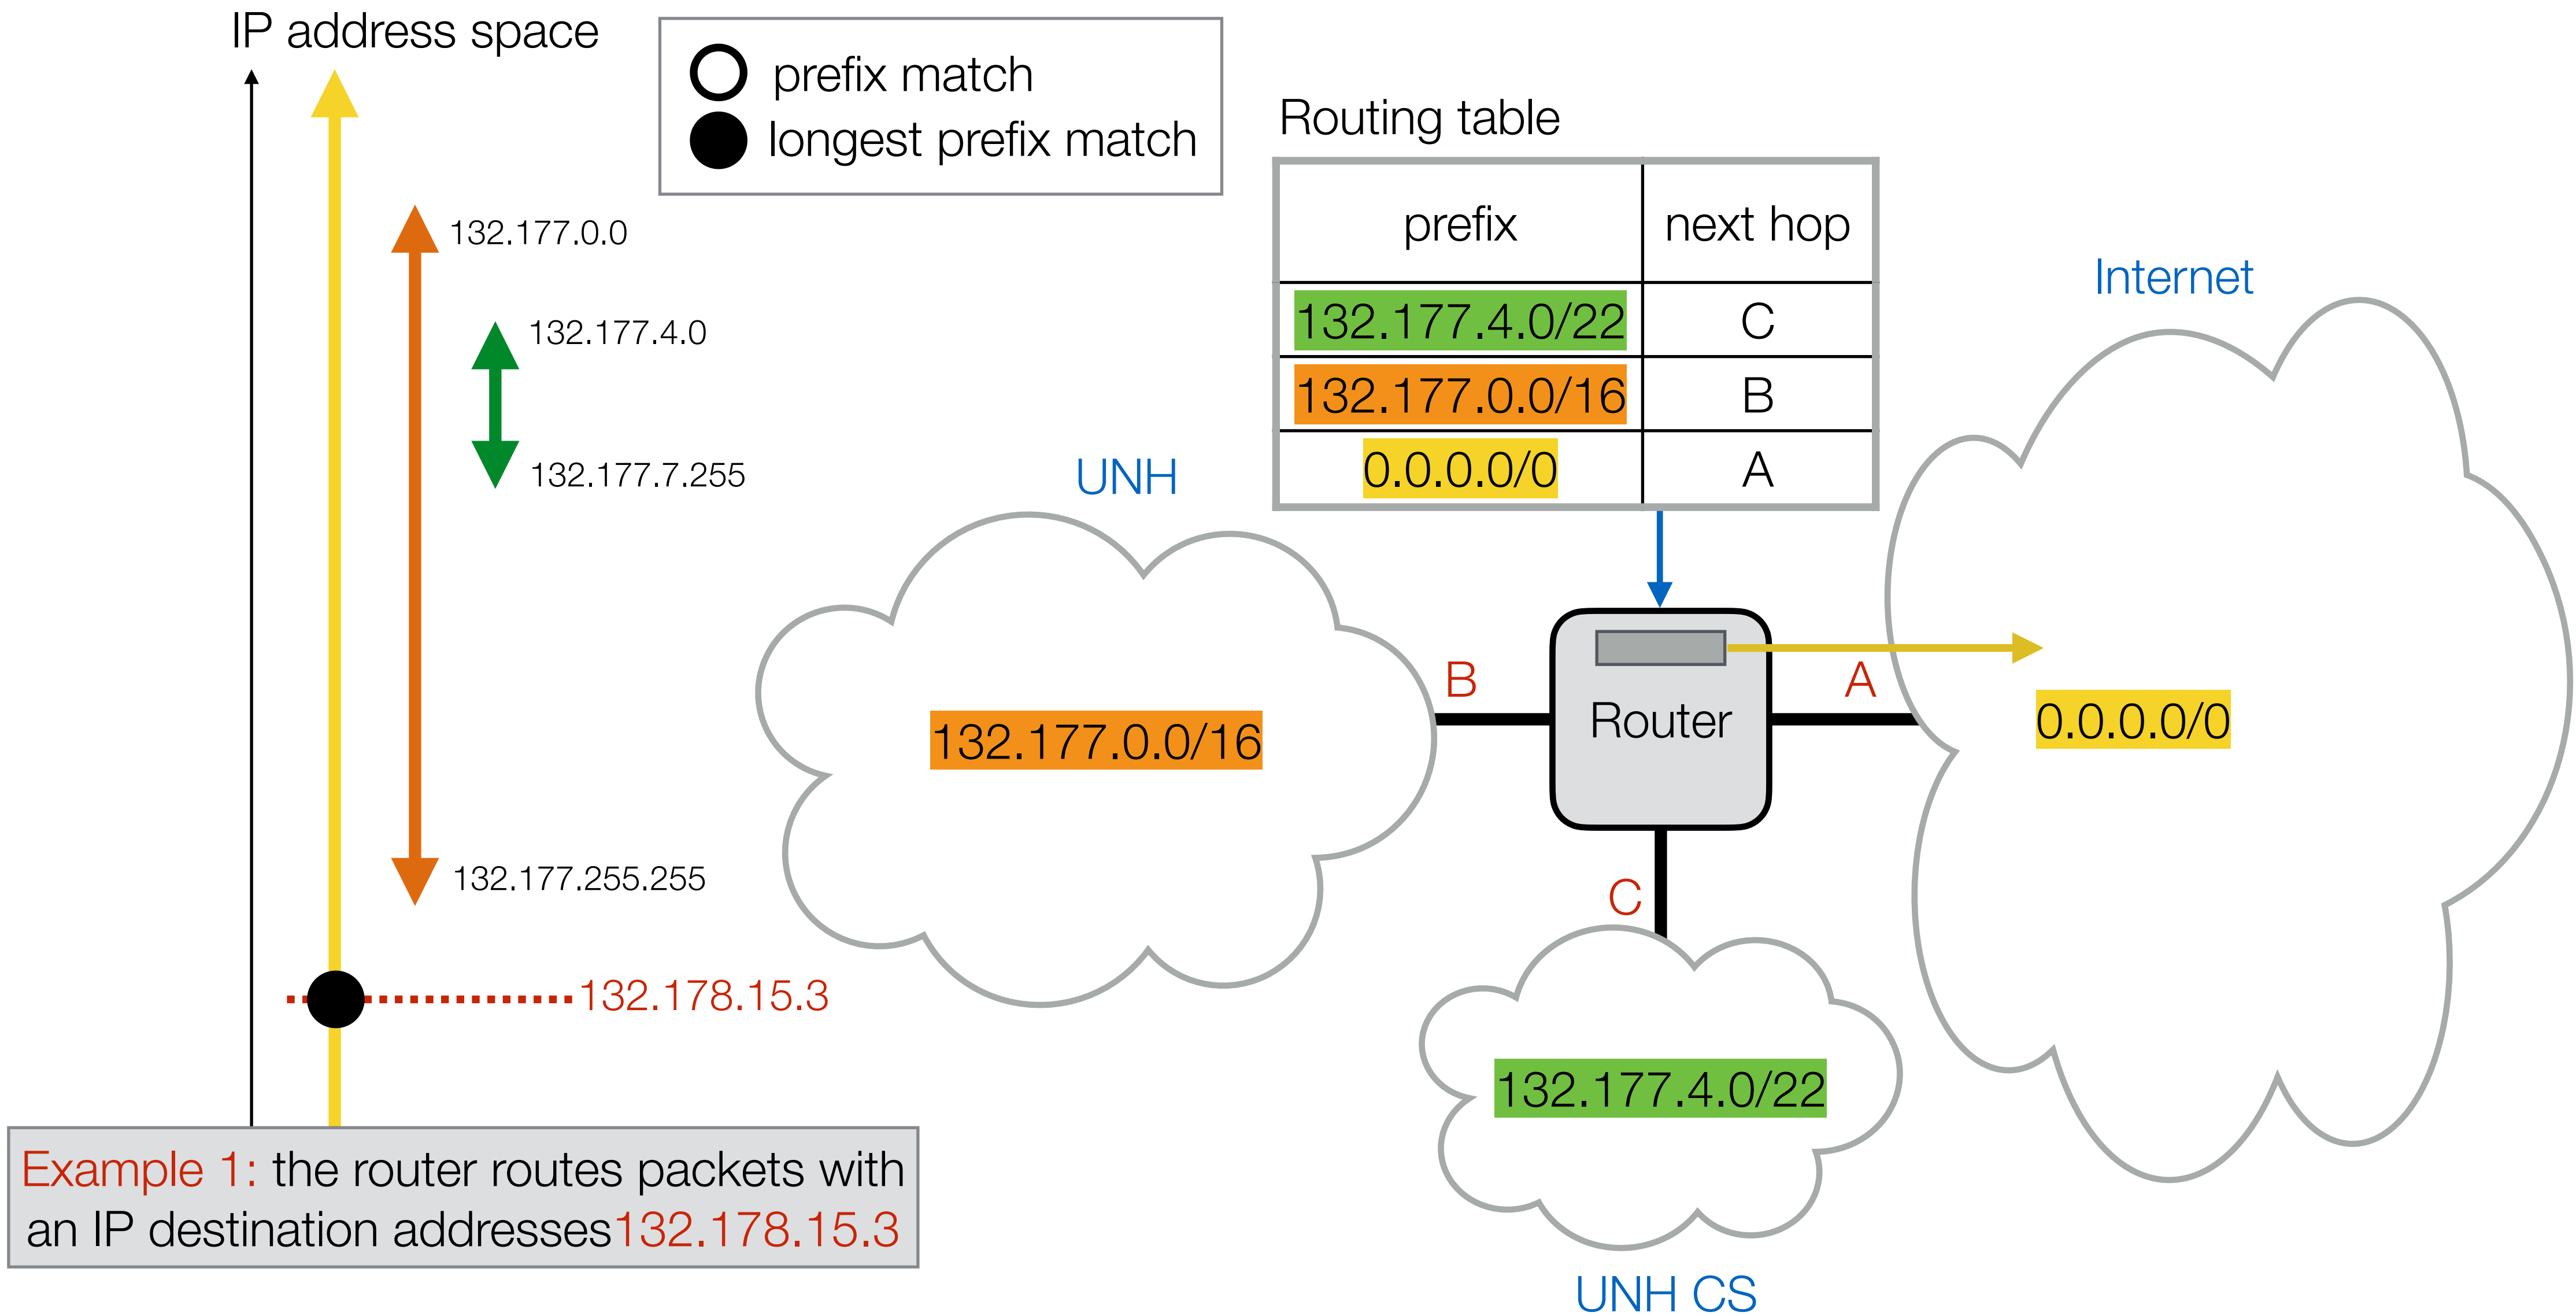
Routing table

prefix	next hop
132.177.4.0/22	C
132.177.0.0/16	B
0.0.0.0/0	A

UNH

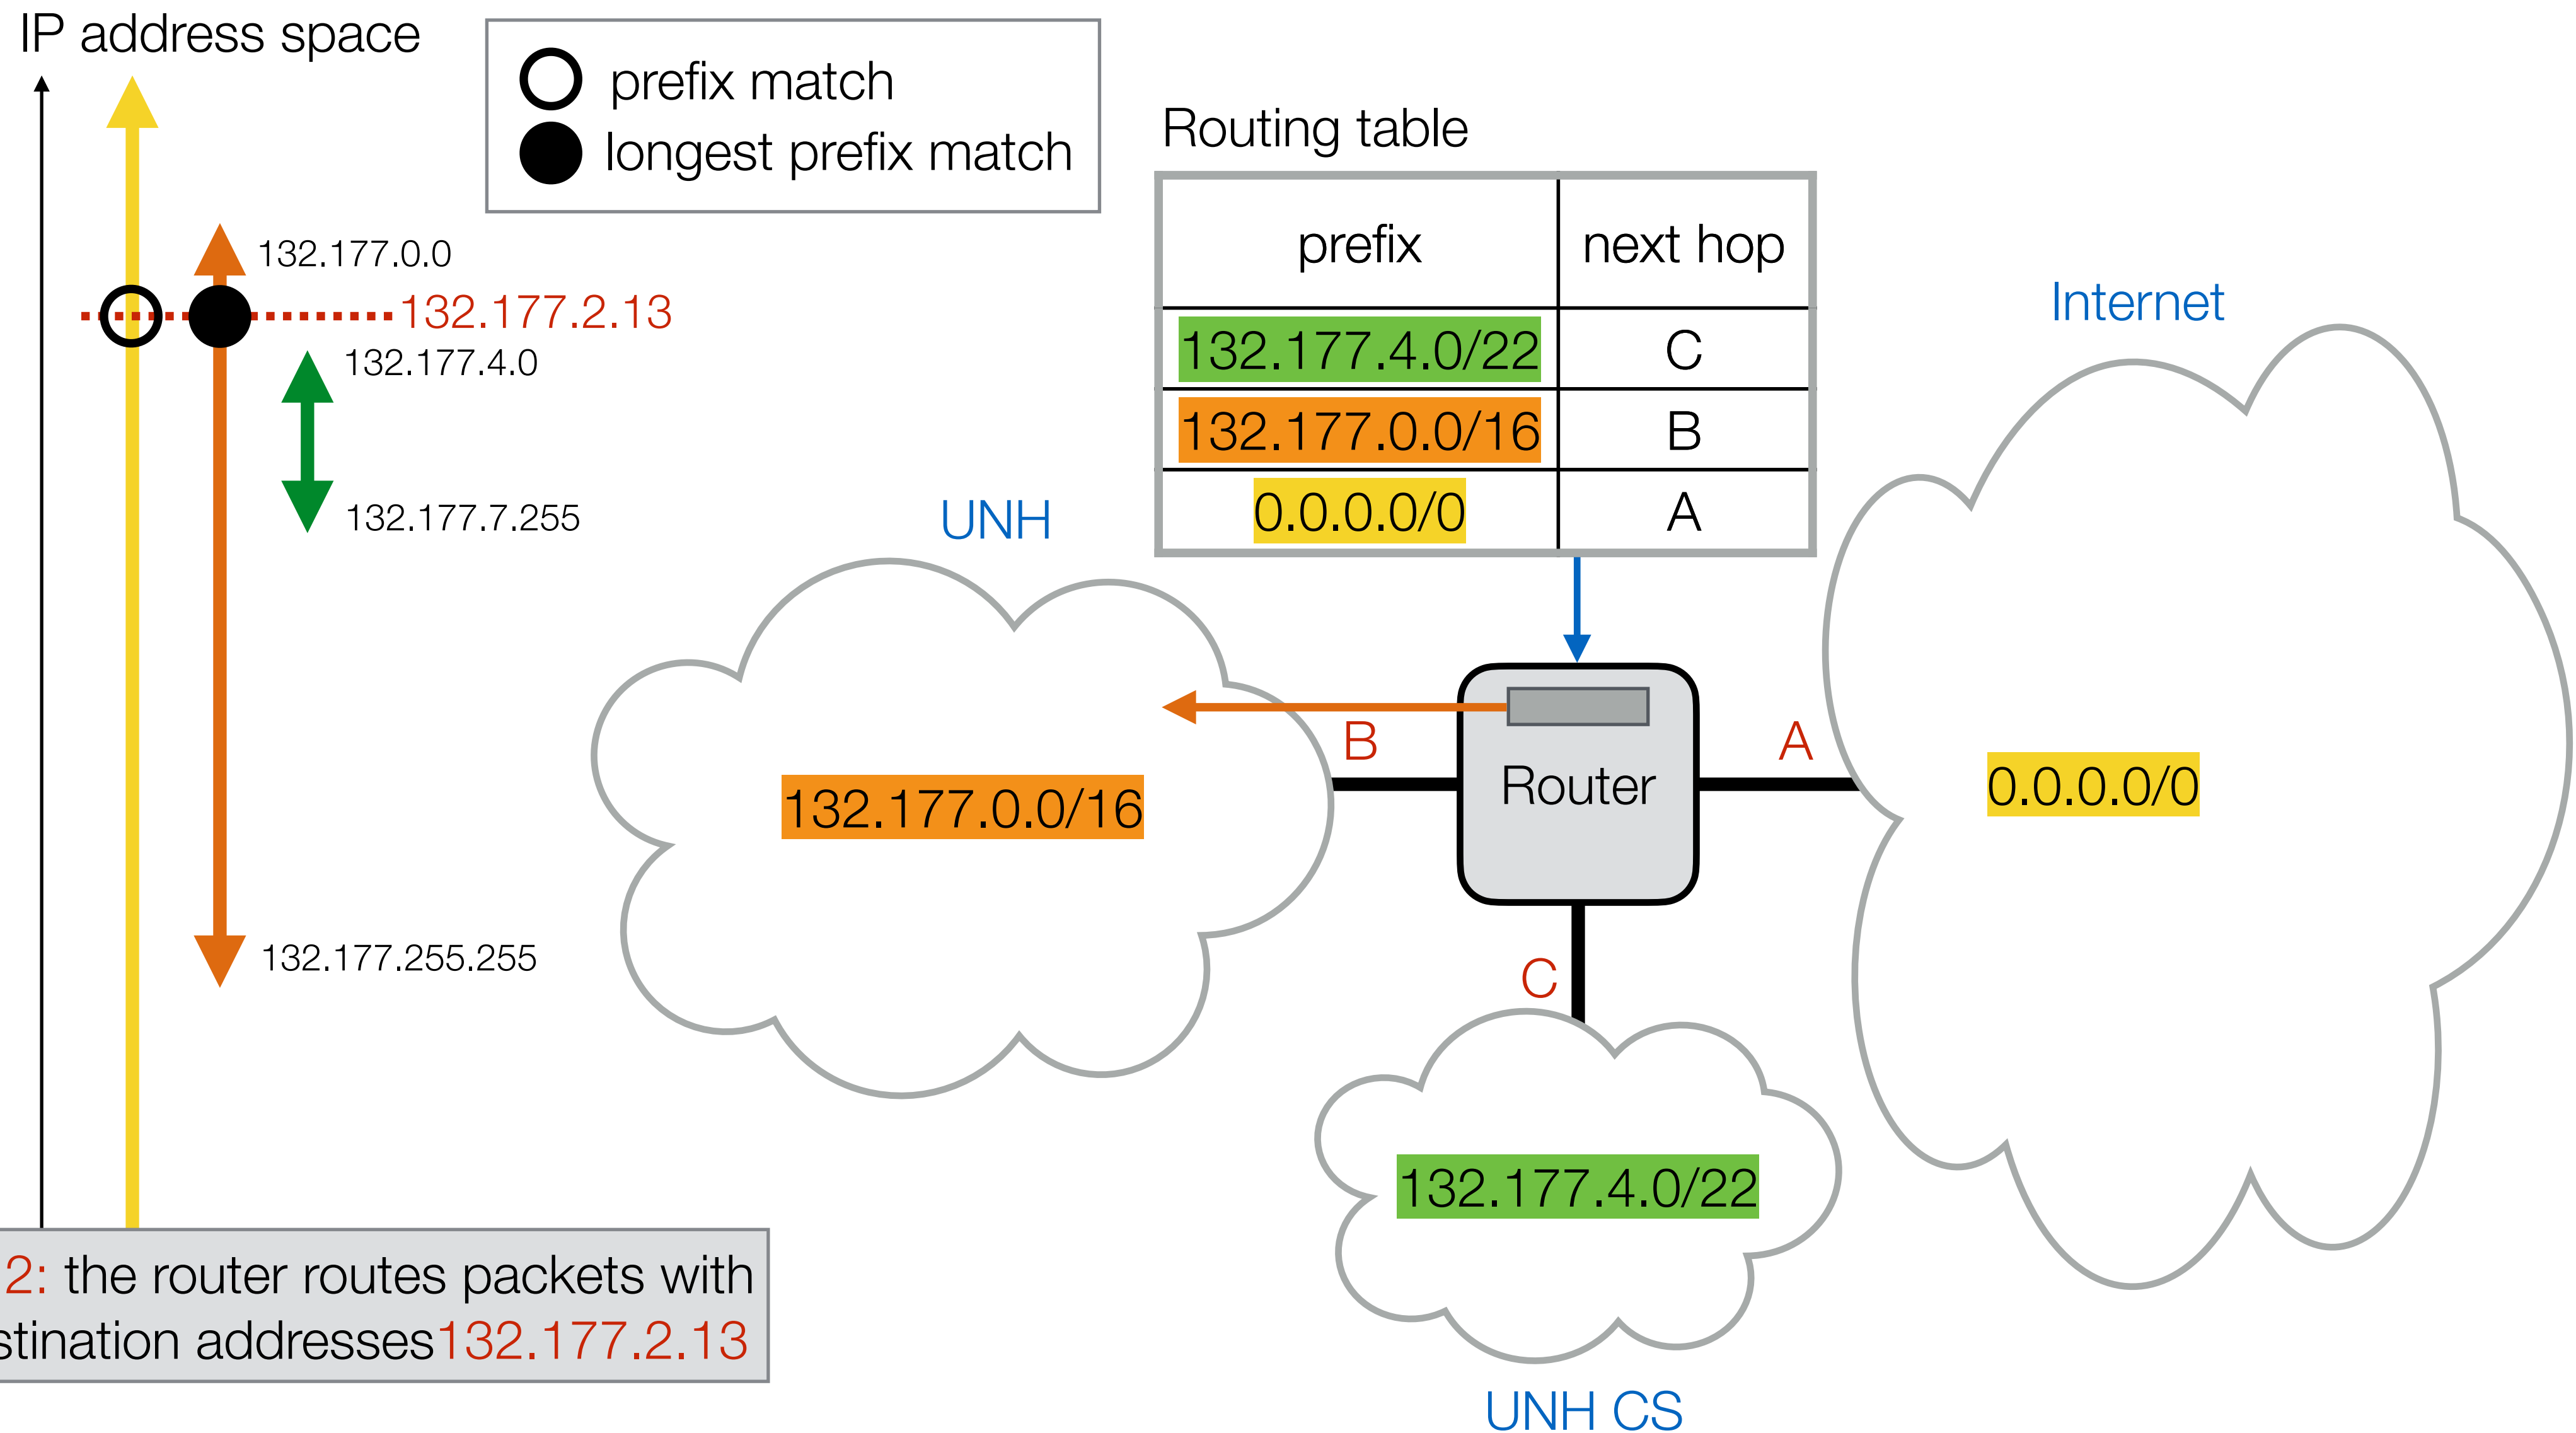


Routing example 1



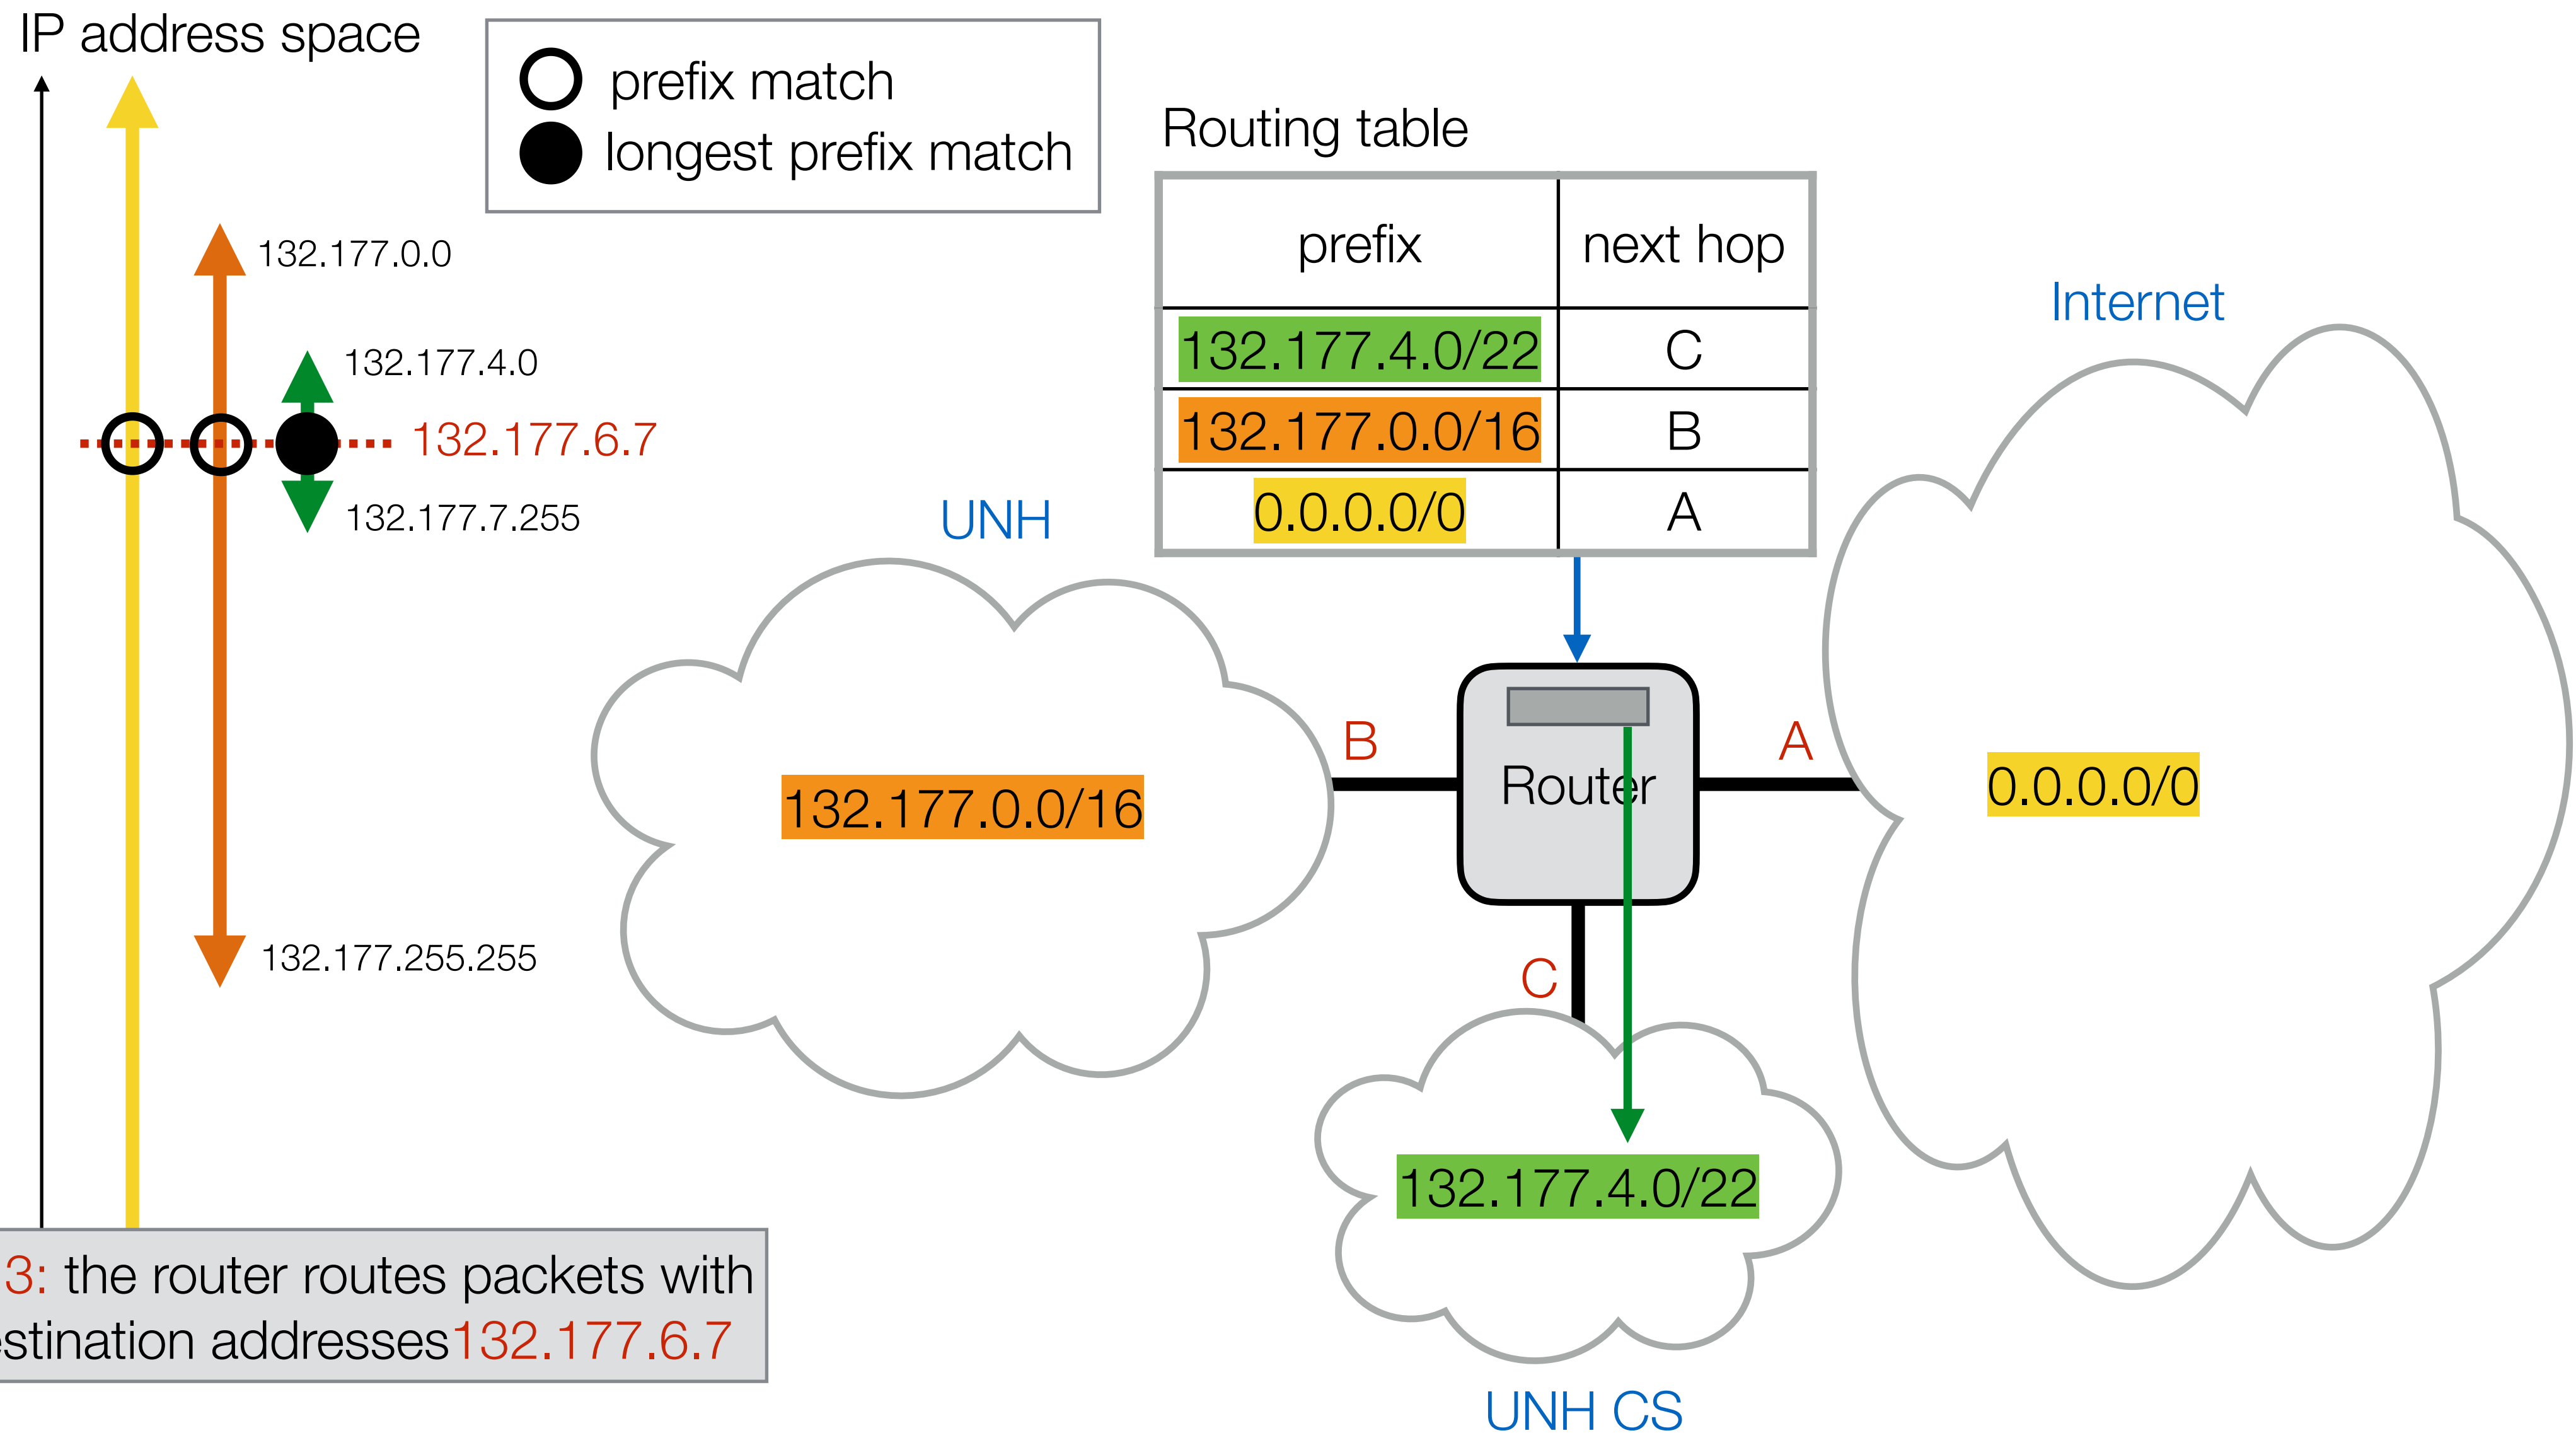
Example 1: the router routes packets with an IP destination addresses 132.178.15.3

Routing example 2



Example 2: the router routes packets with an IP destination addresses **132.177.2.13**

Routing example 3

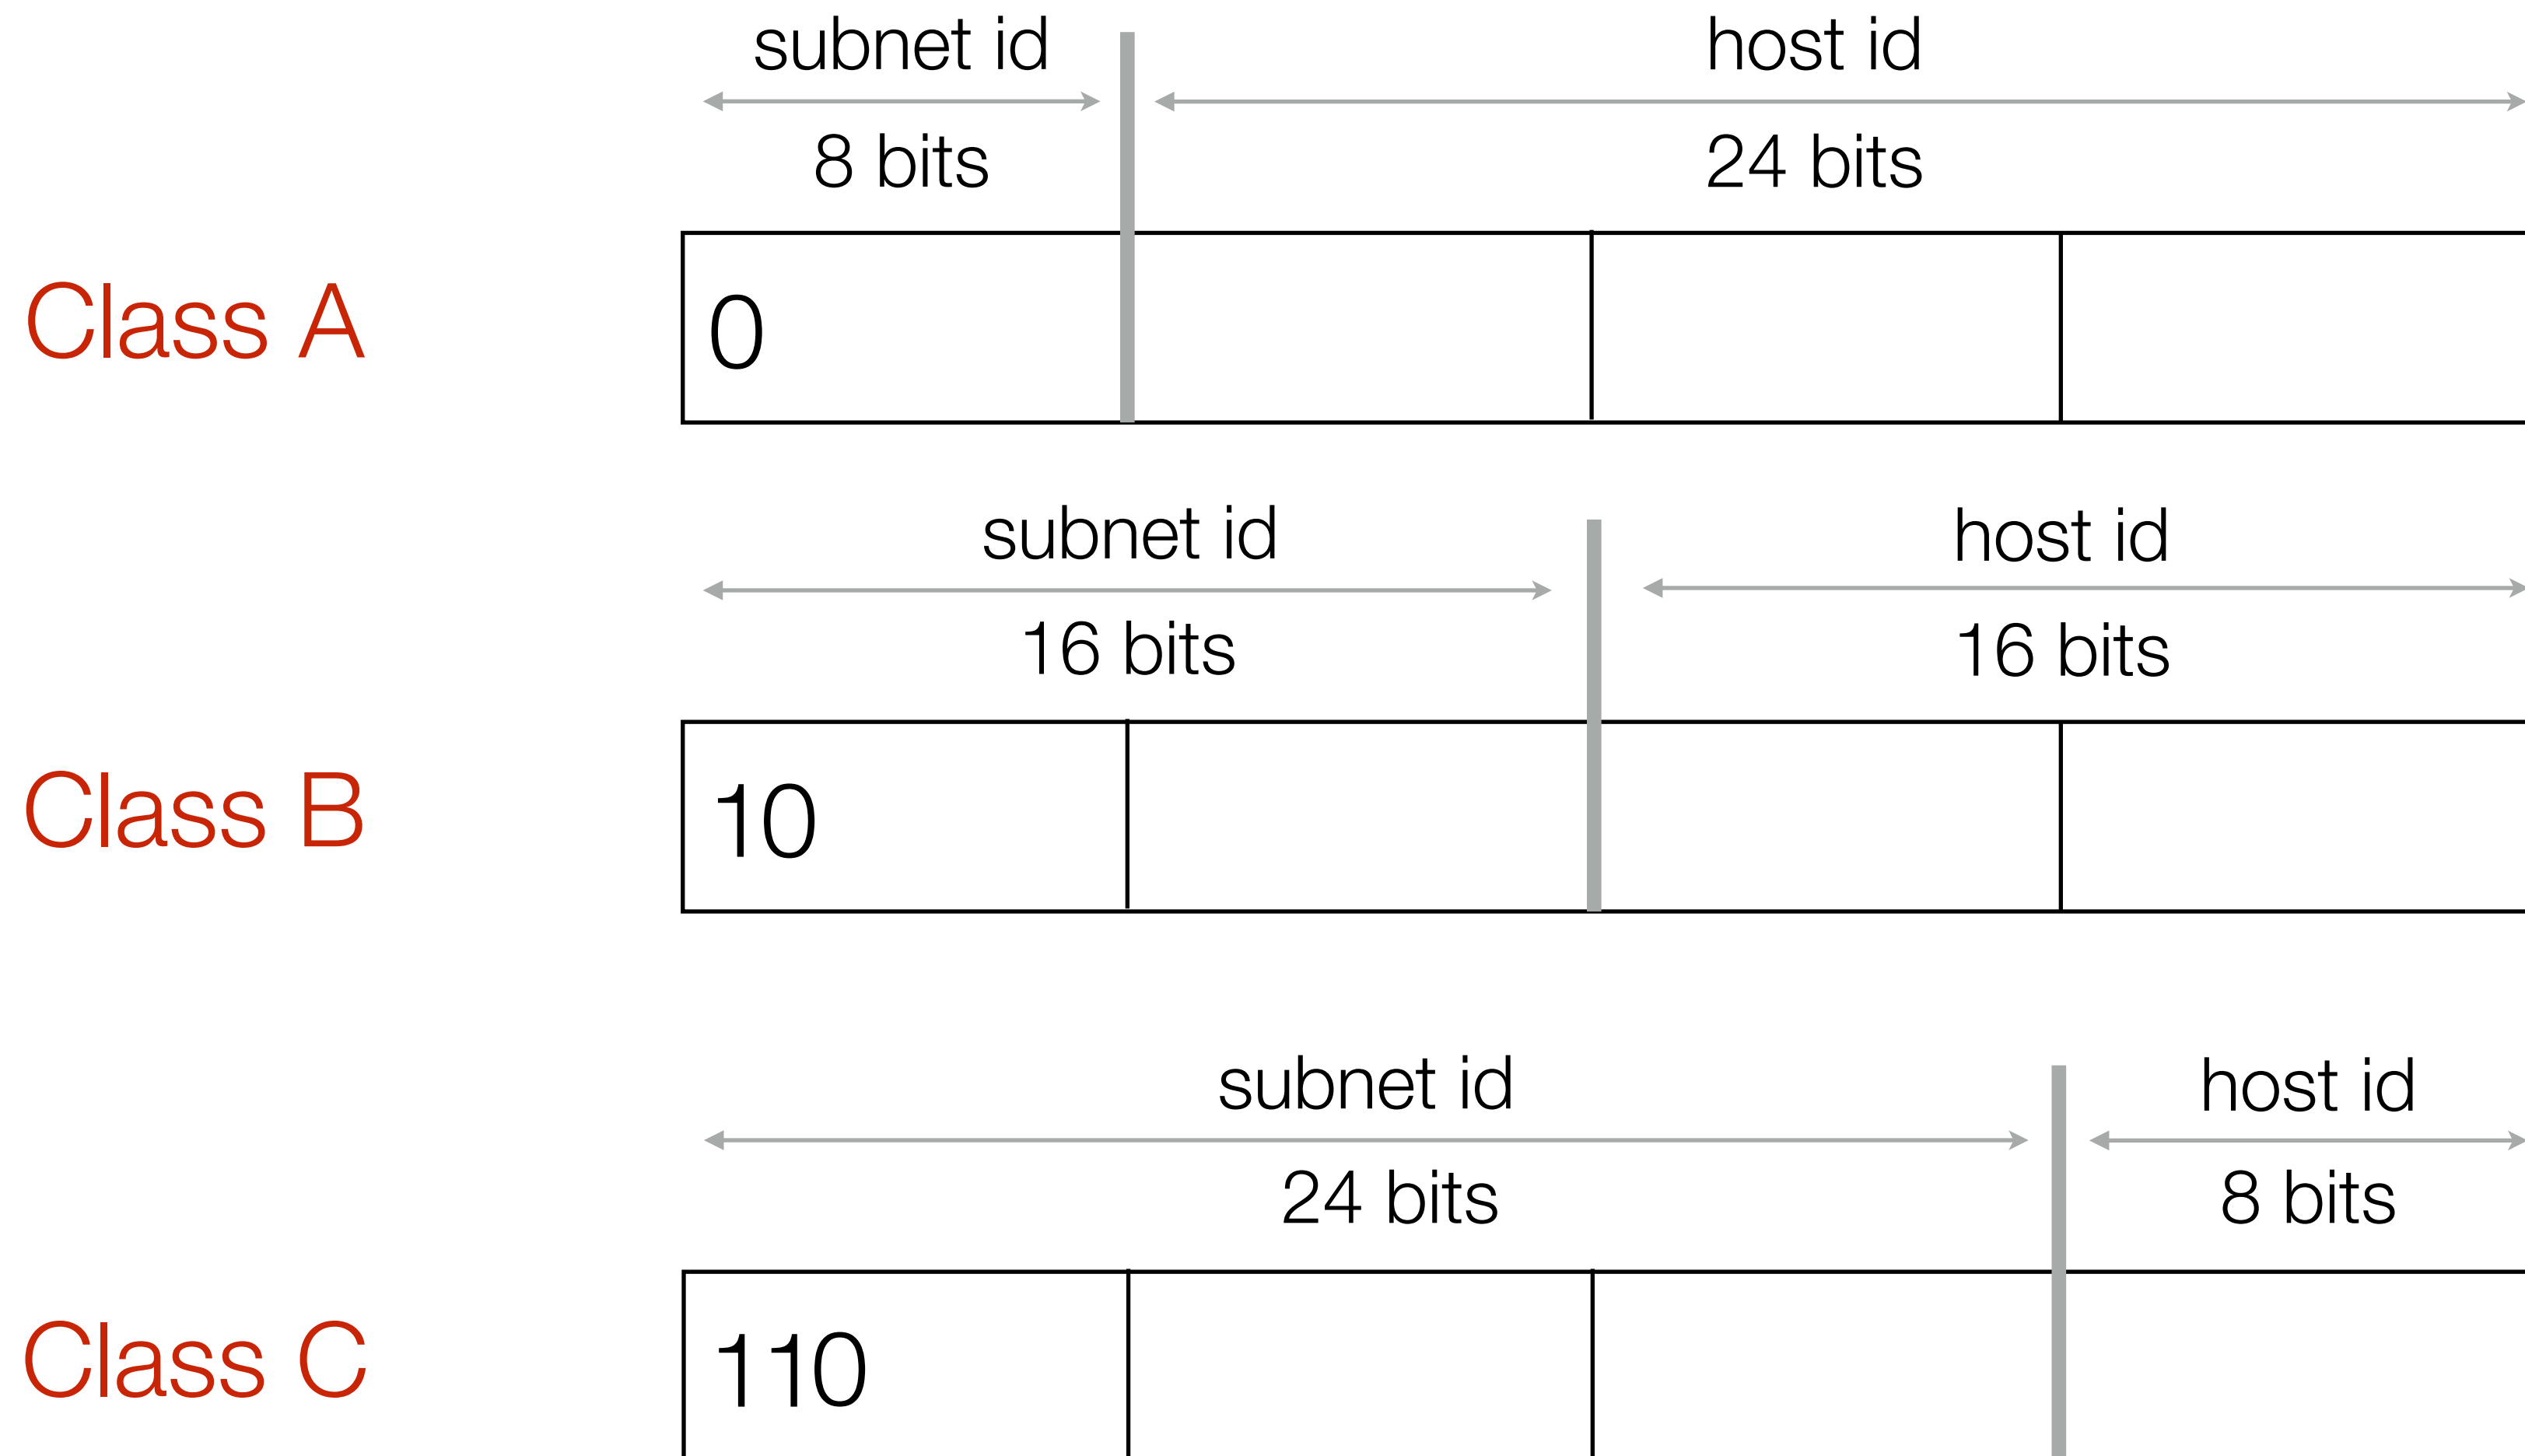


Example 3: the router routes packets with an IP destination addresses **132.177.6.7**

Special Use IPv4 Addresses

- ▶ **0.0.0.0/8** - "This" network
- ▶ **127.0.0.0/8** - Loopback
- ▶ **10.0.0.0/8,**
172.16.0.0/12,
192.168.0.0/16 - Private networks
- ▶ **255.255.255.255/32** - Limited broadcast
- ▶ **169.254.0.0/16** - Link local
- ▶ **(first and last IP address of a prefix: subnet broadcast)**
 - i.e., host id consisting of all zeros or all ones


Class-based Routing



only for historical reference...

Class-based Routing

(continued...)



CIDR Addresses

notation	addrs/block	# blocks	
n.n.n.n/32	1	4294967296	"host route"
n.n.n.x/31	2	2147483648	"p2p link"
n.n.n.x/30	4	1073741824	
n.n.n.x/29	8	536870912	
n.n.n.x/28	16	268435456	
n.n.n.x/27	32	134217728	
n.n.n.x/26	64	67108864	
n.n.n.x/25	128	33554432	
n.n.n.0/24	256	16777216	legacy "Class C"
n.n.x.0/23	512	8388608	
n.n.x.0/22	1024	4194304	
n.n.x.0/21	2048	2097152	
n.n.x.0/20	4096	1048576	
n.n.x.0/19	8192	524288	
n.n.x.0/18	16384	262144	
n.n.x.0/17	32768	131072	
n.n.0.0/16	65536	65536	legacy "Class B"
n.x.0.0/15	131072	32768	
n.x.0.0/14	262144	16384	
n.x.0.0/13	524288	8192	
n.x.0.0/12	1048576	4096	
n.x.0.0/11	2097152	2048	
n.x.0.0/10	4194304	1024	
n.x.0.0/9	8388608	512	
n.0.0.0/8	16777216	256	legacy "Class A"
x.0.0.0/7	33554432	128	
x.0.0.0/6	67108864	64	
x.0.0.0/5	134217728	32	
x.0.0.0/4	268435456	16	
x.0.0.0/3	536870912	8	
x.0.0.0/2	1073741824	4	
x.0.0.0/1	2147483648	2	
0.0.0.0/0	4294967296	1	"default route"

Classless Inter-Domain Routing

From **RFC 4632**:

n is an 8-bit decimal octet value.

x is a 1- to 7-bit value, based on the prefix length, shifted into the most significant bits of the octet and converted into decimal form; the least significant bits of the octet are zero.