Master Course
Computer Networks
IN2097

Prof. Dr.-Ing. Georg Carle

Chair for Network Architectures and Services
Department of Computer Science
Technische Universität München
http://www.net.in.tum.de

Chair for Network Architectures and Services – Prof. Carle
Department of Computer Science
TU München

Internet Protocol

Network Prefix and Host Number

- Each IP network (often called subnetwork or subnet) has an IP address:

  IP address of a network = Host number is set to all zeros, e.g., 128.143.0.0

- IP routers are devices that forward IP datagrams between IP networks

- Delivery of an IP datagram proceeds in 2 steps:
  - Use network prefix to deliver datagram to the right network
  - Once the network is found, use the host number to deliver to the right interface

Interplay between routing and forwarding

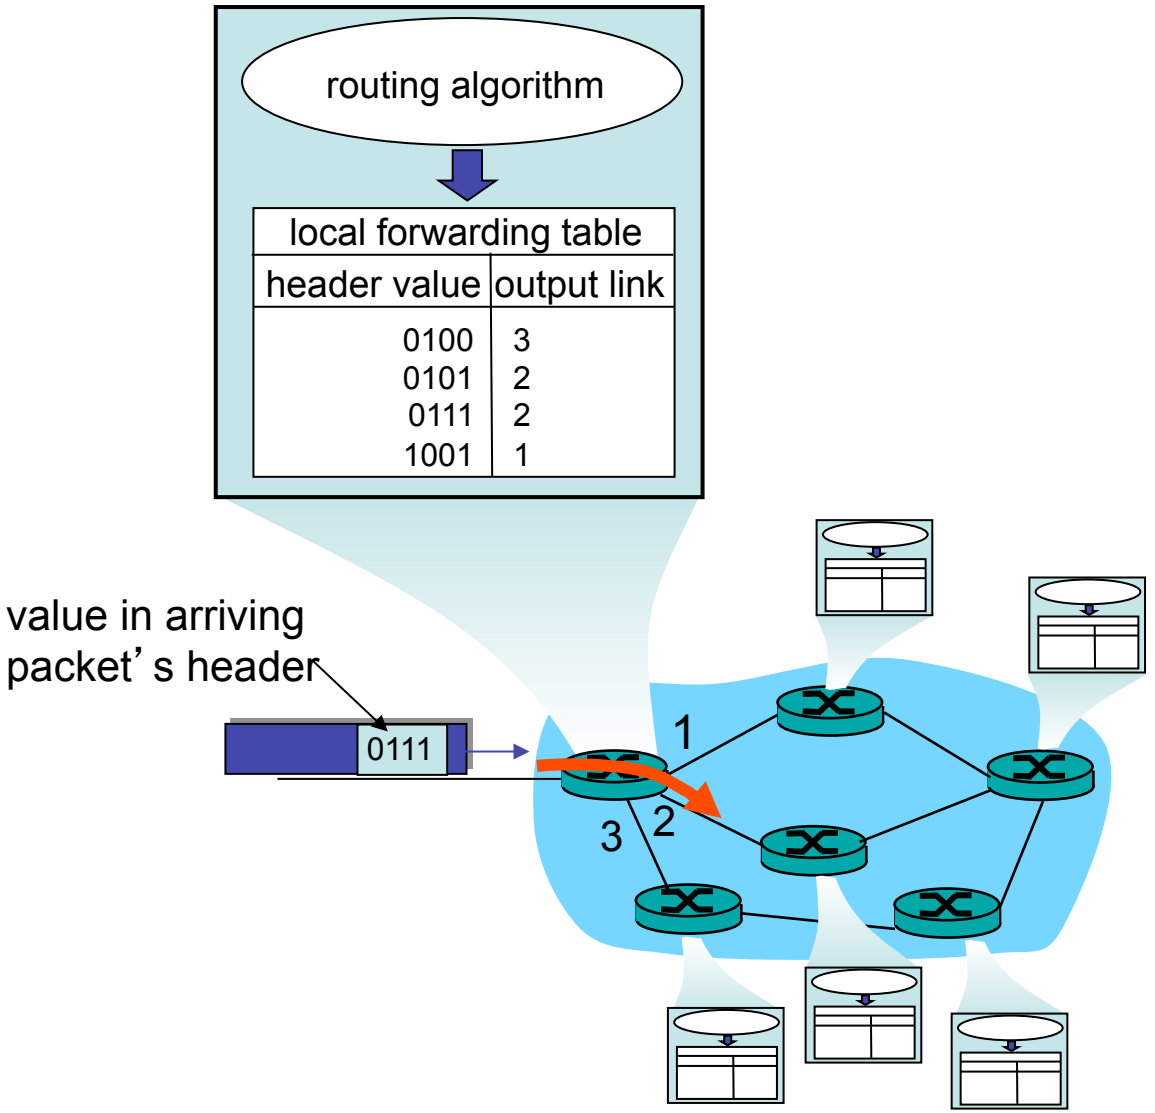

Routers: Forwarding and Routing

- Forwarding: data plane
  - Directing a data packet to an outgoing link
  - Individual router using a forwarding table
- Routing: control plane
  - Computing the paths the packets will follow
  - Routers talking amongst themselves
  - Individual router creating a forwarding table

Functional Components

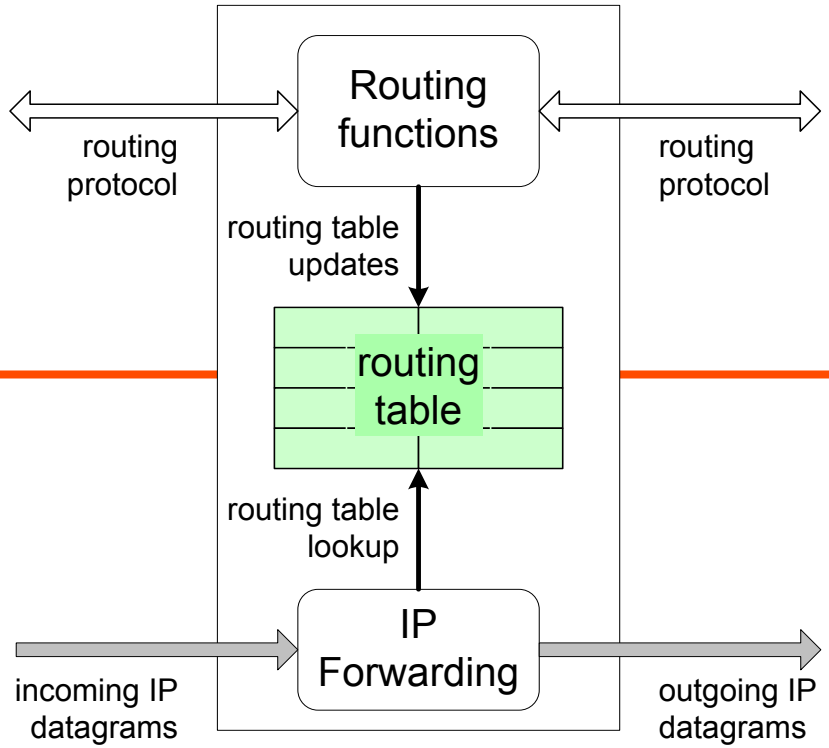

Control

Datapath: per-packet processing

Routing and Forwarding

Routing functions include:

- route calculation
- maintenance of the routing table
- execution of routing protocols

- On commercial routers handled by a single general purpose processor, called *route processor*

IP forwarding is per-packet processing

- On high-end commercial routers, IP forwarding is distributed (Most work is done on the interface cards)

The Internet Network layer

Host, router network layer functions:

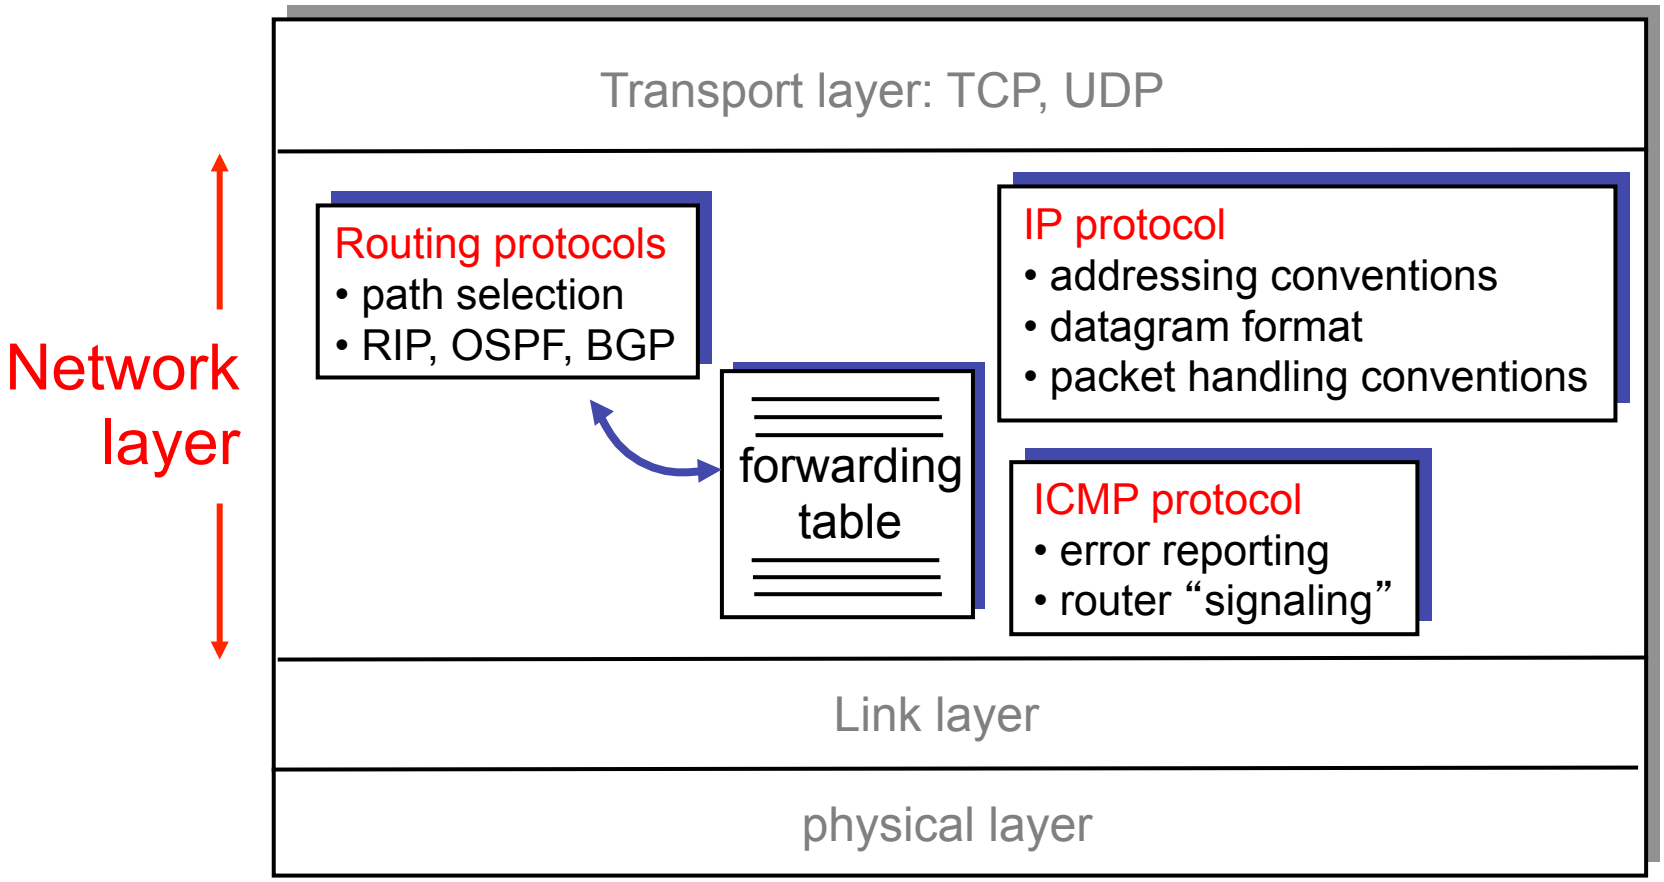

IP protocol model

- Many application specific protocols over IP
- IP (with best effort service model) over many medi^a specific (L2) protocols → „Hourglass“ model of IP
- Quality of Service added to IP as an ‚afterthought‘

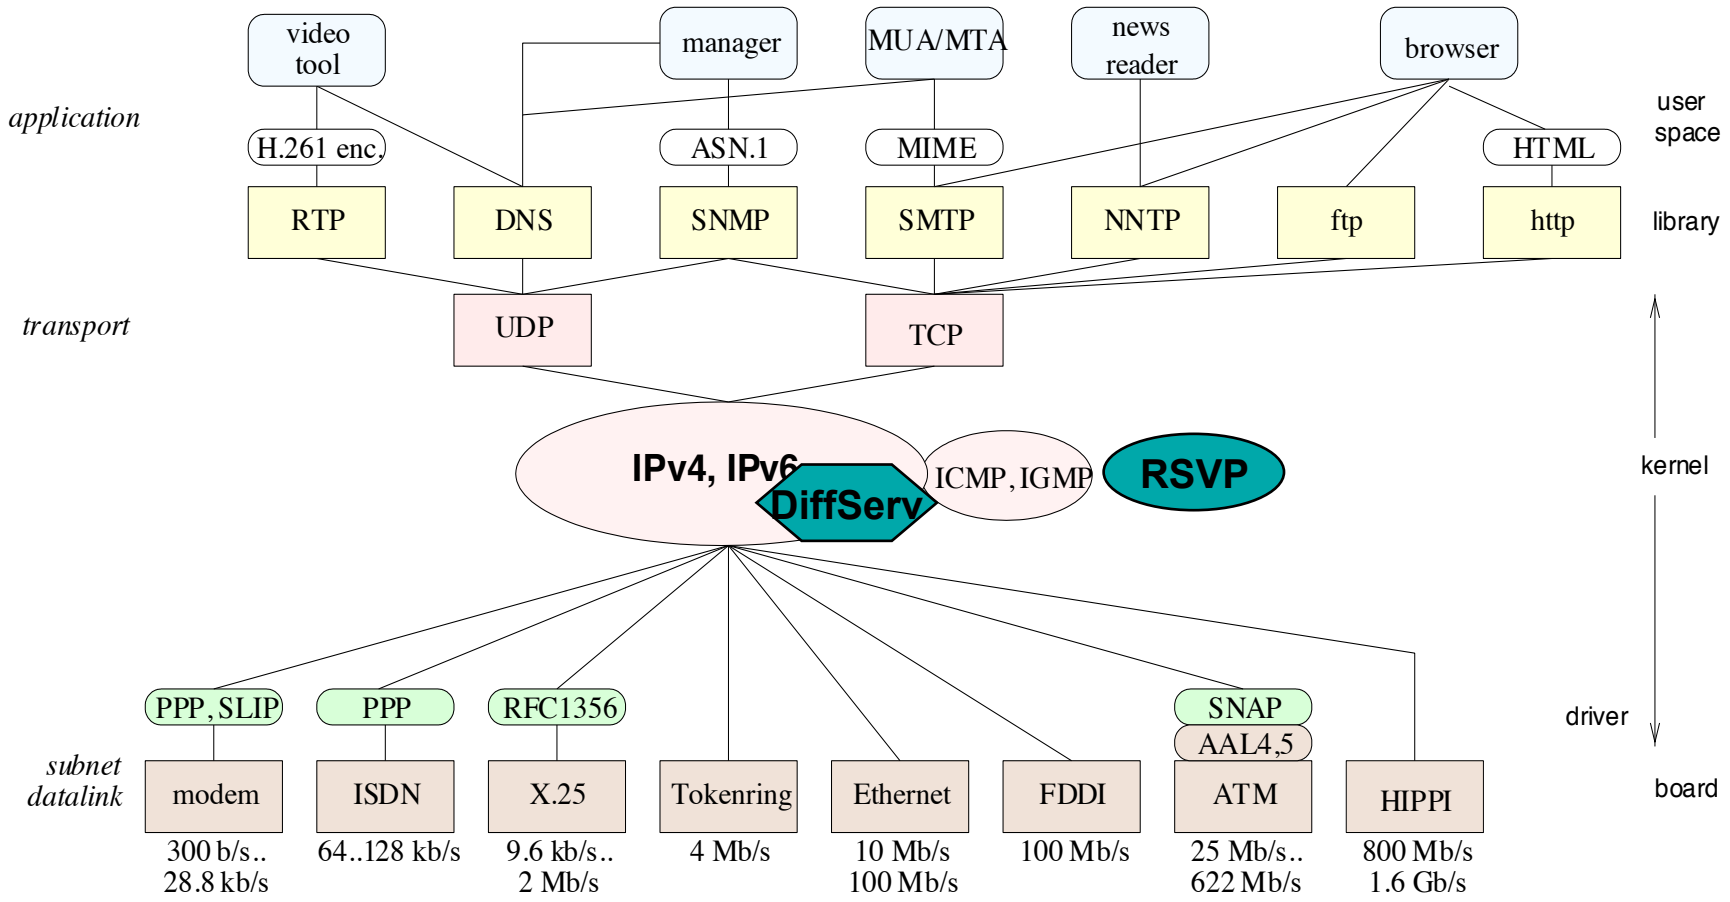

Layer 3 Addresses

2^{32} (~4 billion) possible entries

Destination Address Range	Link Interface
11001000 00010111 00010000 00000000 through 11001000 00010111 00010111 11111111	0
11001000 00010111 00011000 00000000 through 11001000 00010111 00011000 11111111	1
11001000 00010111 00011001 00000000 through 11001000 00010111 00011111 11111111	2
otherwise	3

IP Datagram

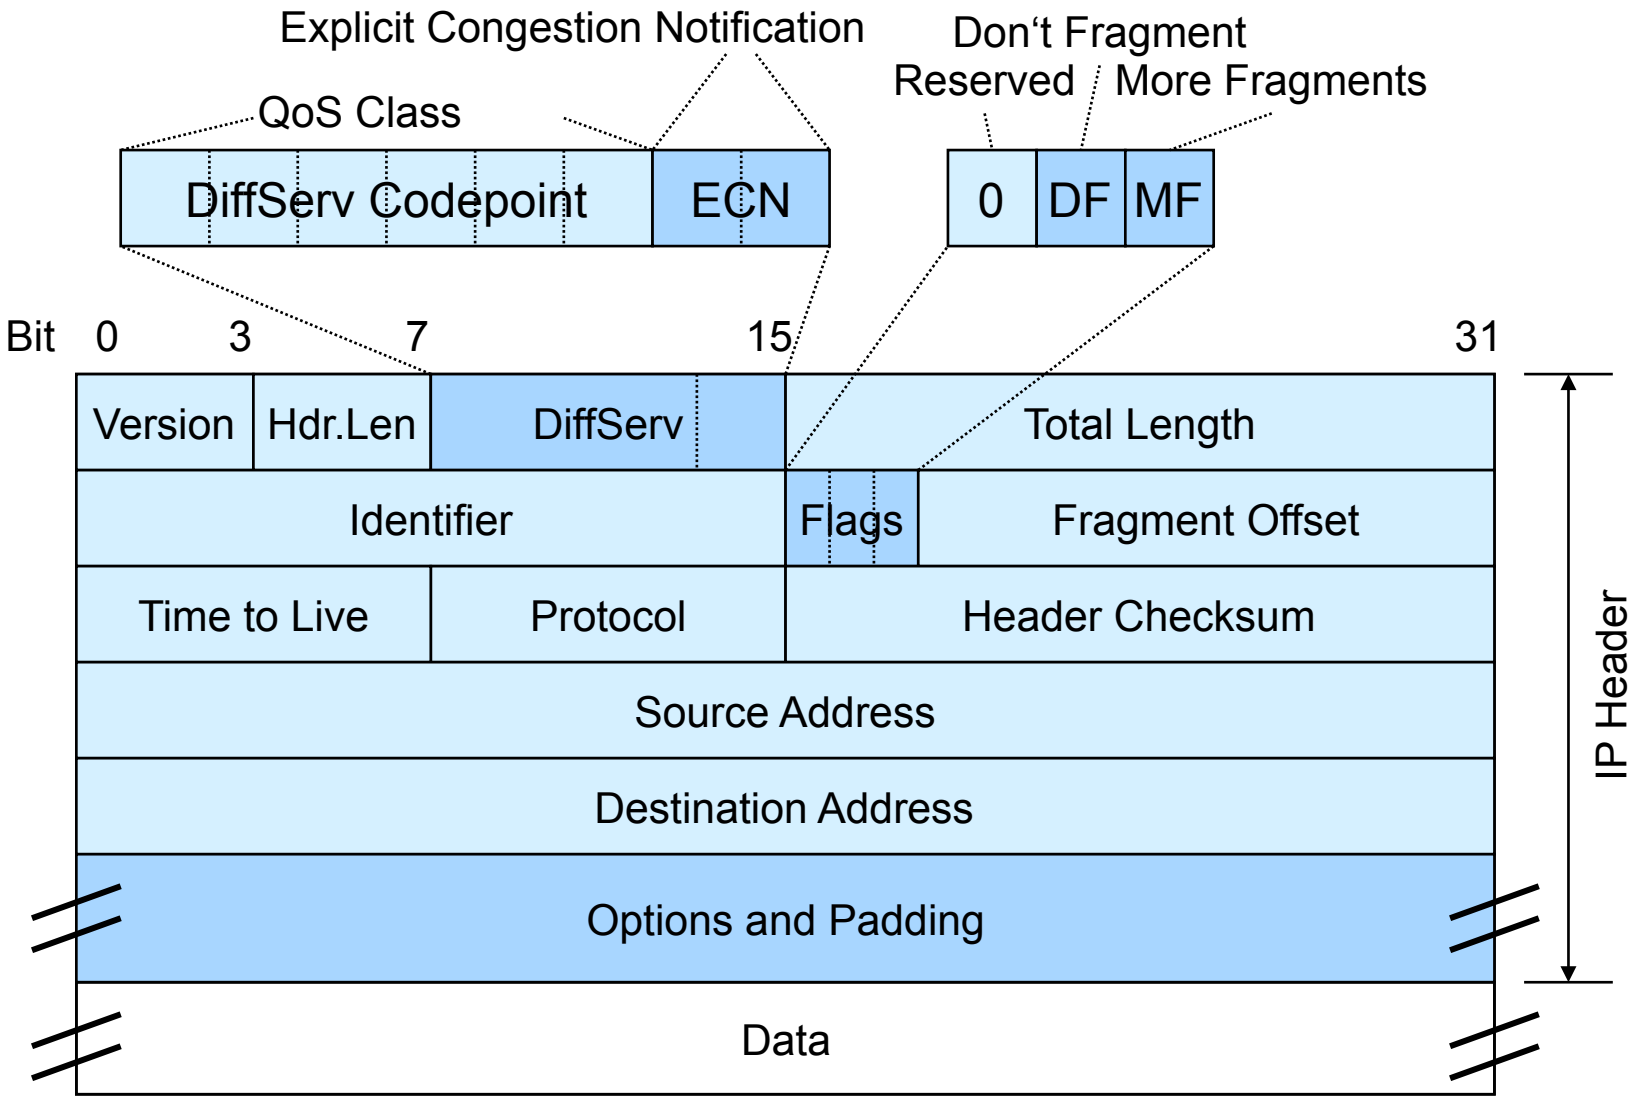
Explicit Congestion Notification
Don‘t Fragment
Reserved
More Fragments
QoS Class
DiffServ Codepoint
ECN
0
DF
MF
Bit 0 3 7 15 31
Version
Hdr.Len
DiffServ
Total Length
Identifier
Flags
Fragment Offset
Time to Live
Protocol
Header Checksum
Source Address
Destination Address
Options and Padding
Data
IP Header

Longest prefix matching

Prefix Match	Link Interface
11001000 00010111 00010	0
11001000 00010111 00011000	1
11001000 00010111 00011	2
otherwise	3

Examples

DA: 11001000 00010111 00010110 10100001 Which interface?

DA: 11001000 00010111 00011000 10101010 Which interface?

IP Addressing

- IP address: 32-bit identifier for host, router *interface*
- *Interface:* connection between host/router and physical link
  - IP addresses associated with each interface

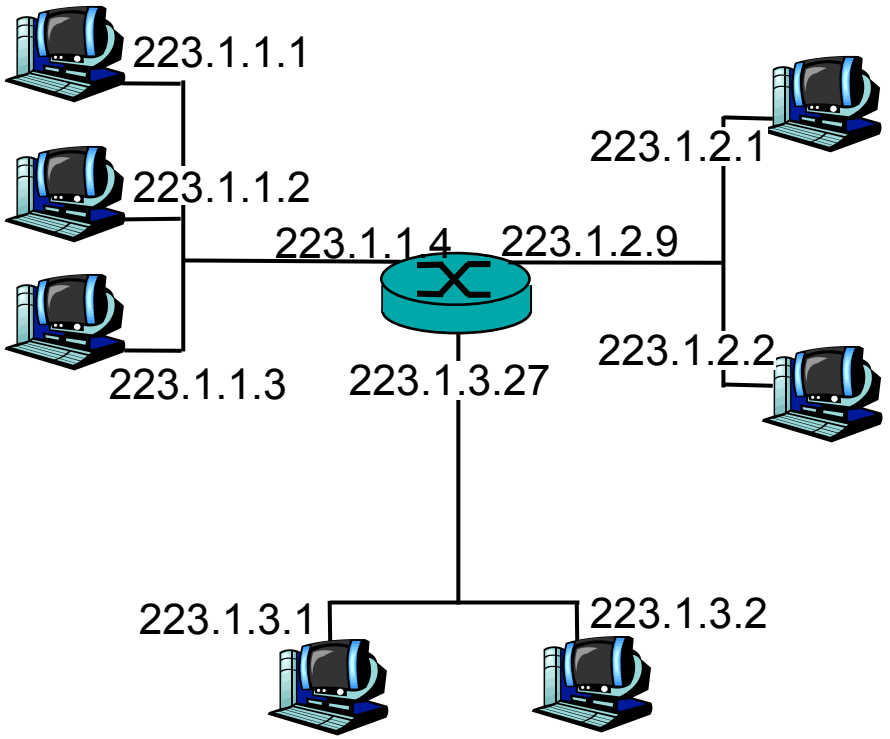

223.1.1.1 = 11011111 00000001 00000001 00000001

223 1 1 1

Subnets

- IP address:
  - subnet part (high order bits)
  - host part (low order bits)
- *What's a subnet ?*
  - device interfaces with same subnet part of IP address
  - can physically reach each other without intervening router

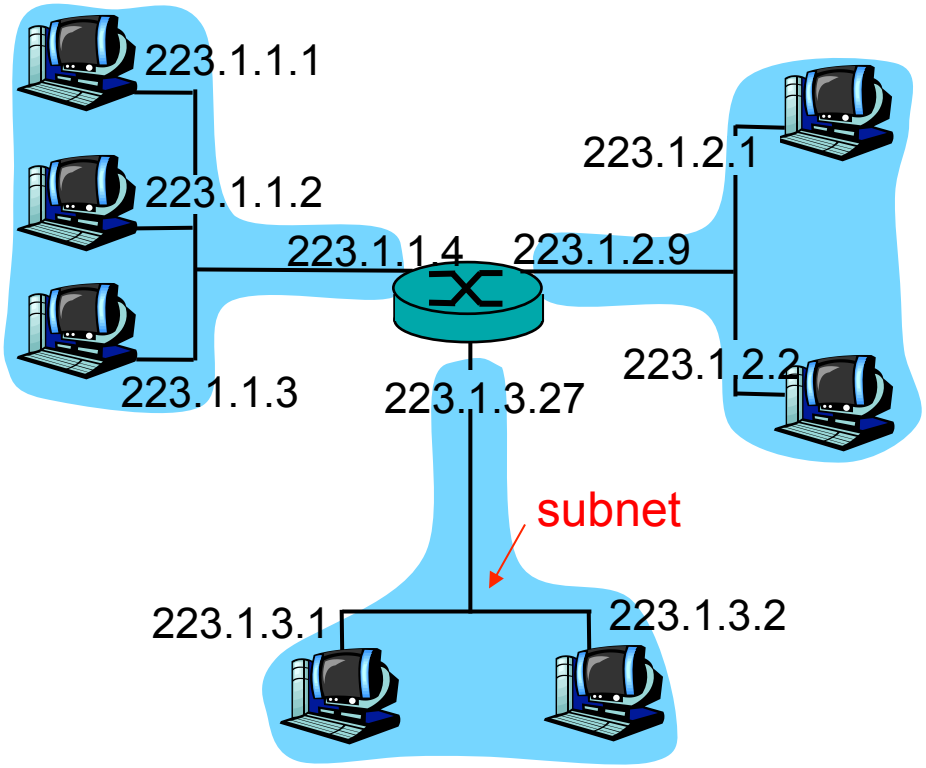

network with 3 subnets

Subnets

- To determine subnets, detach interfaces from host or router

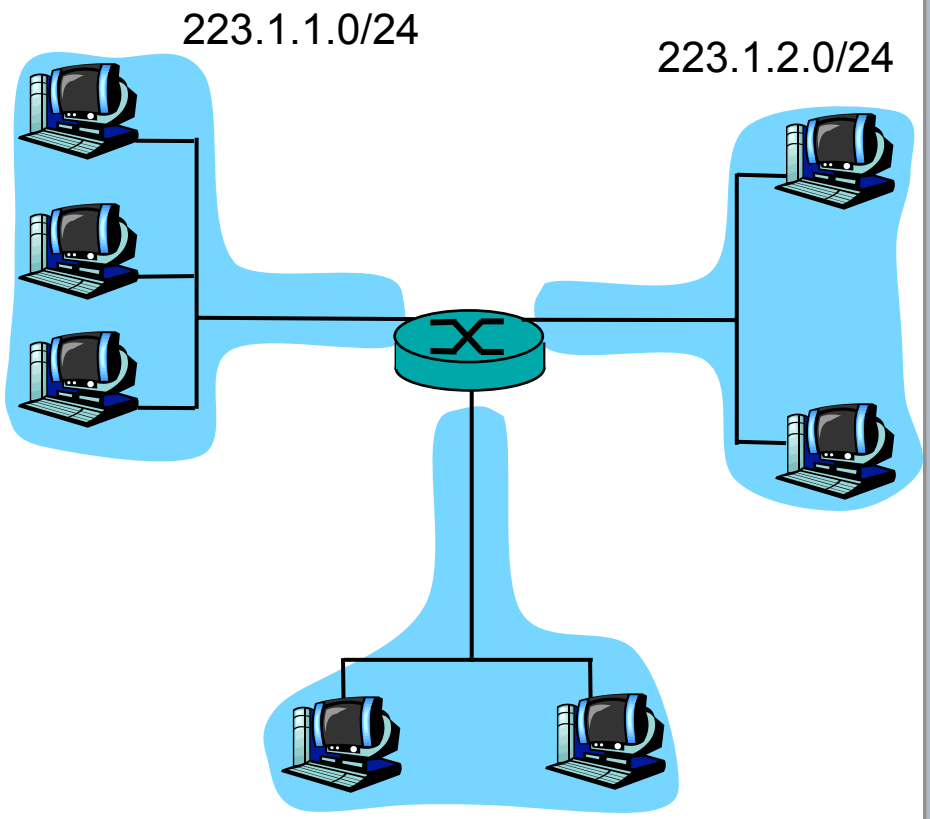

Subnet mask: /24

Subnets

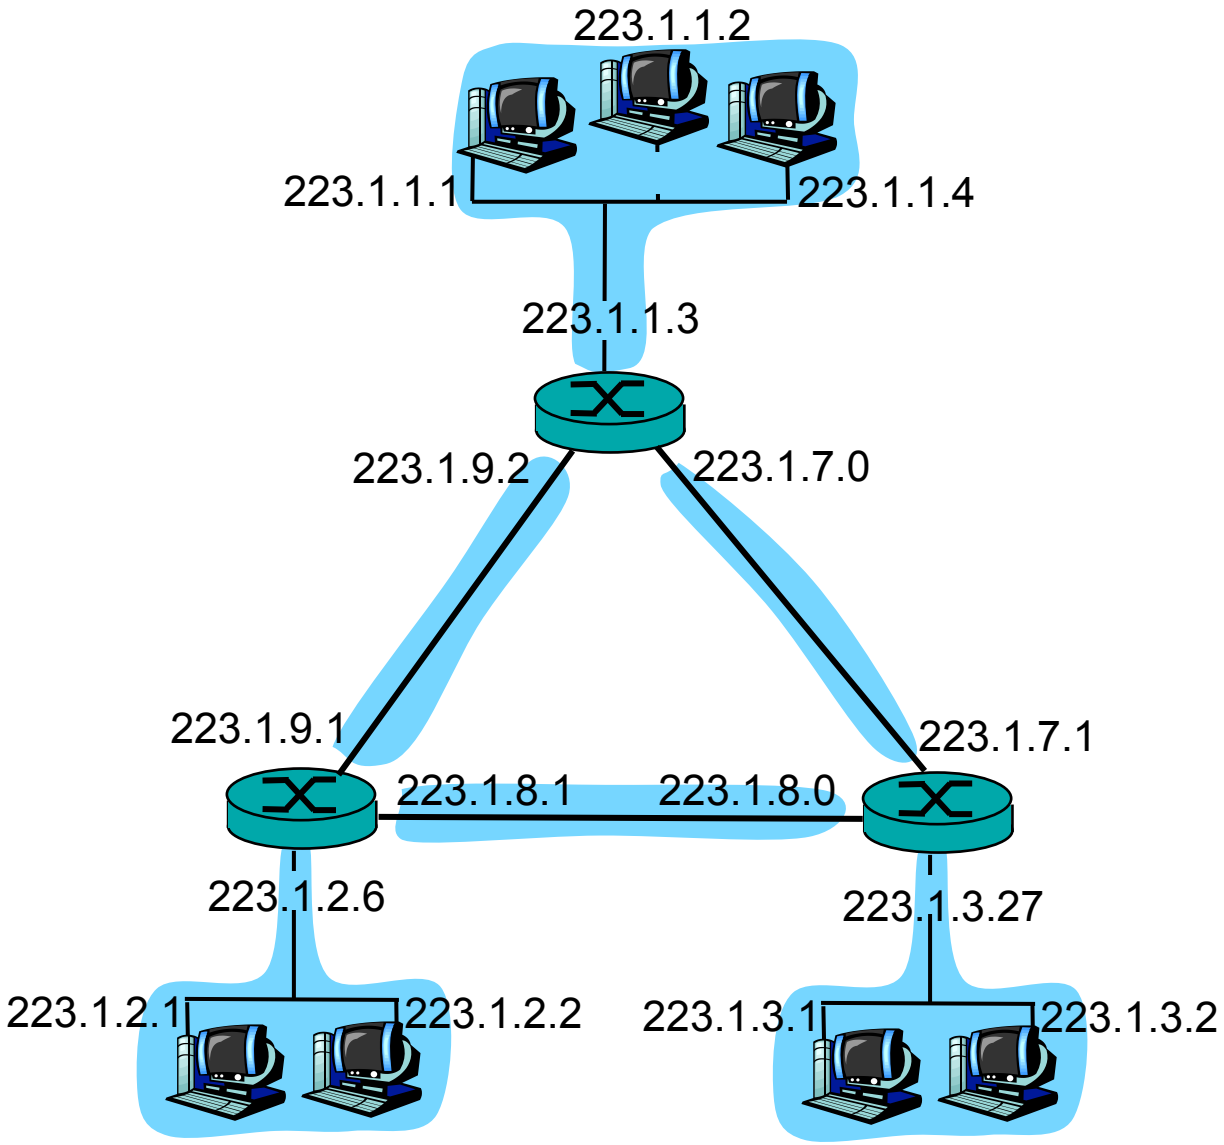
223.1.1.2
223.1.1.1
223.1.1.4
223.1.1.3
223.1.9.2
223.1.7.0
223.1.9.1
223.1.8.1
223.1.8.0
223.1.7.1
223.1.2.6
223.1.3.27
223.1.2.1
223.1.2.2
223.1.3.1
223.1.3.2

Classful IP Addresses (Historic)

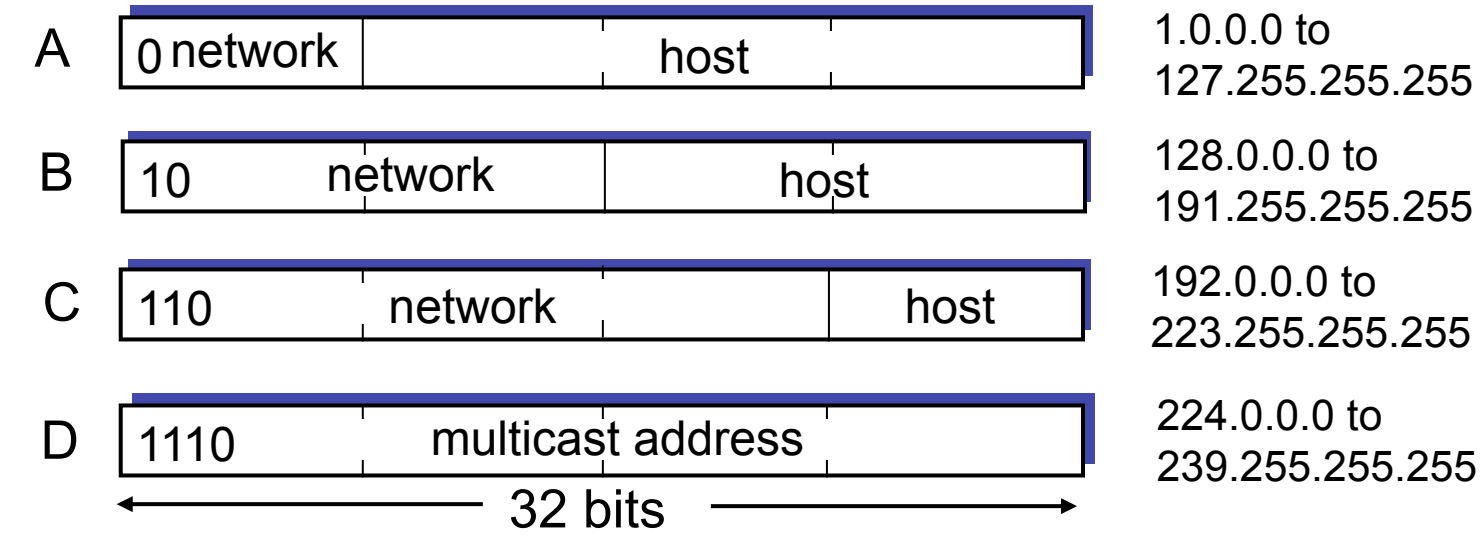

Class	Range (as "dotted quad")		
A	0.0.0.0	to	127.255.255.255
B	128.0.0.0	to	191.255.255.255
C	192.0.0.0	to	223.255.255.255
D	224.0.0.0	to	239.255.255.255
E	240.0.0.0	to	255.255.255.255

CIDR: Classless InterDomain Routing

- subnet portion of address of arbitrary length
- address format: a.b.c.d/x, where x is # bits in subnet portion of address

200.23.16.0/23

IP addressing: CIDR

Class-based:

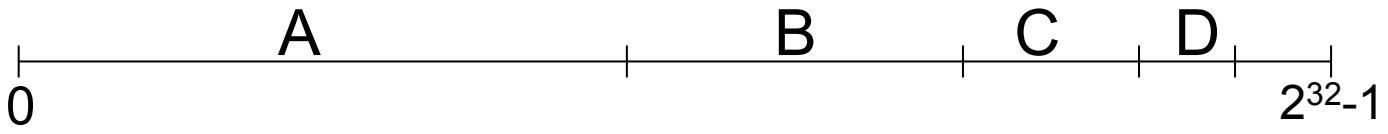

Classless:

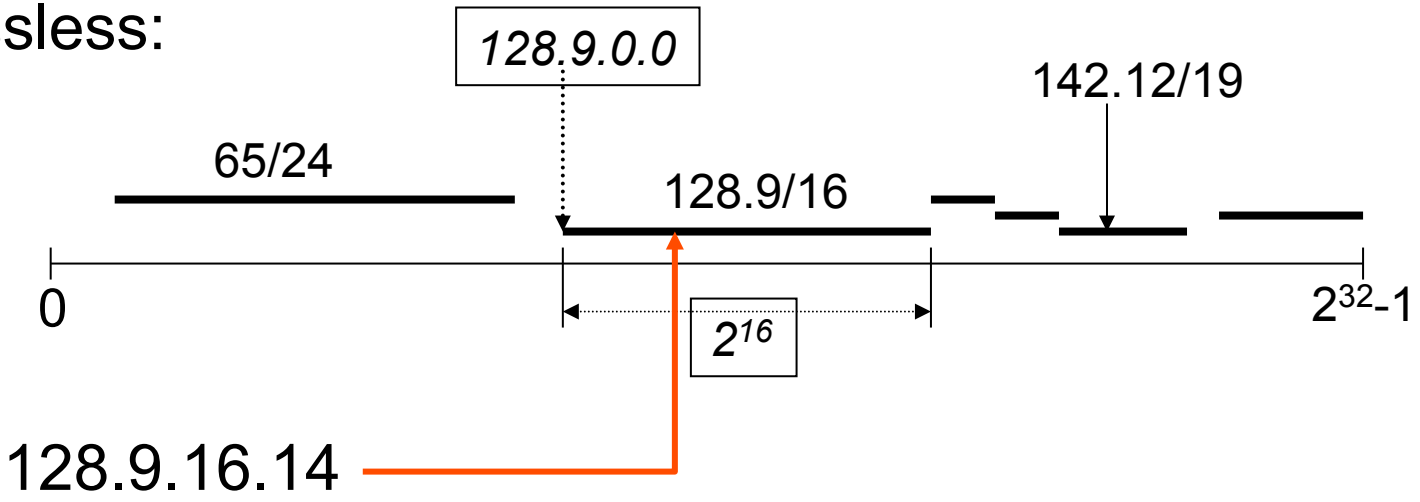

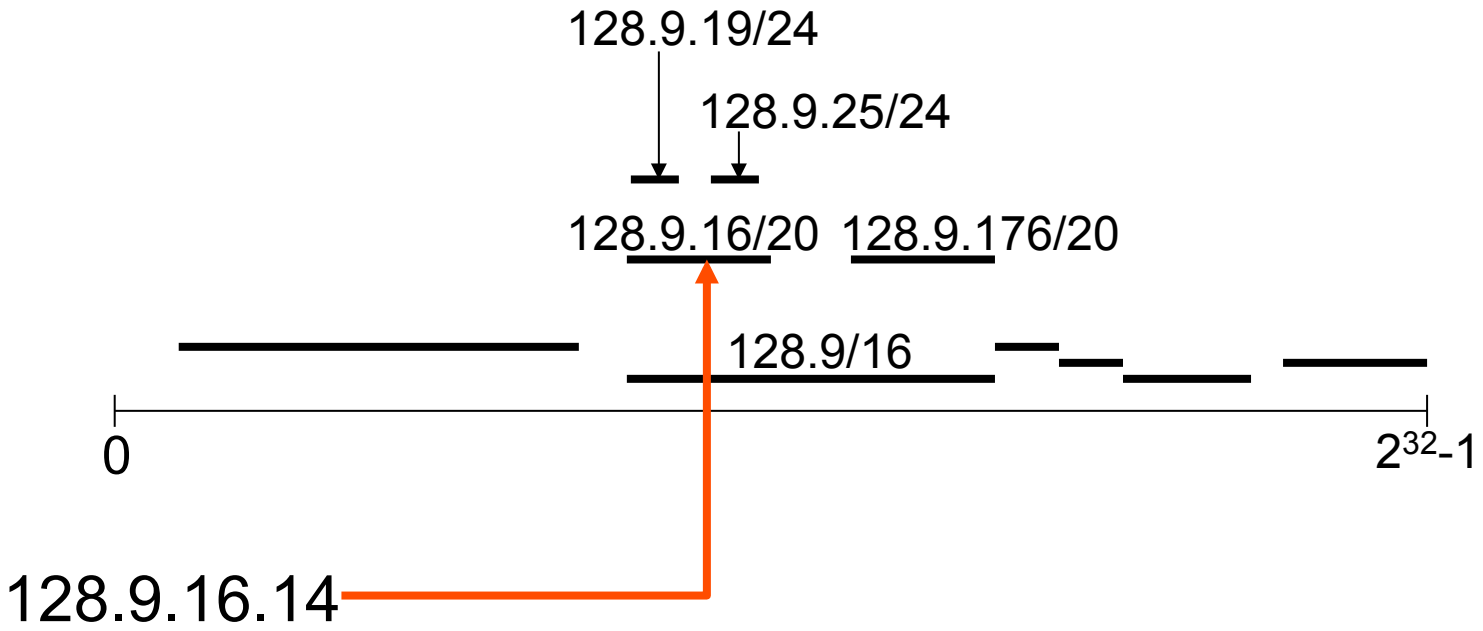

Most specific route = “longest matching prefix”

Multi-domain Structure of the Internet

- Hierarchical network structure

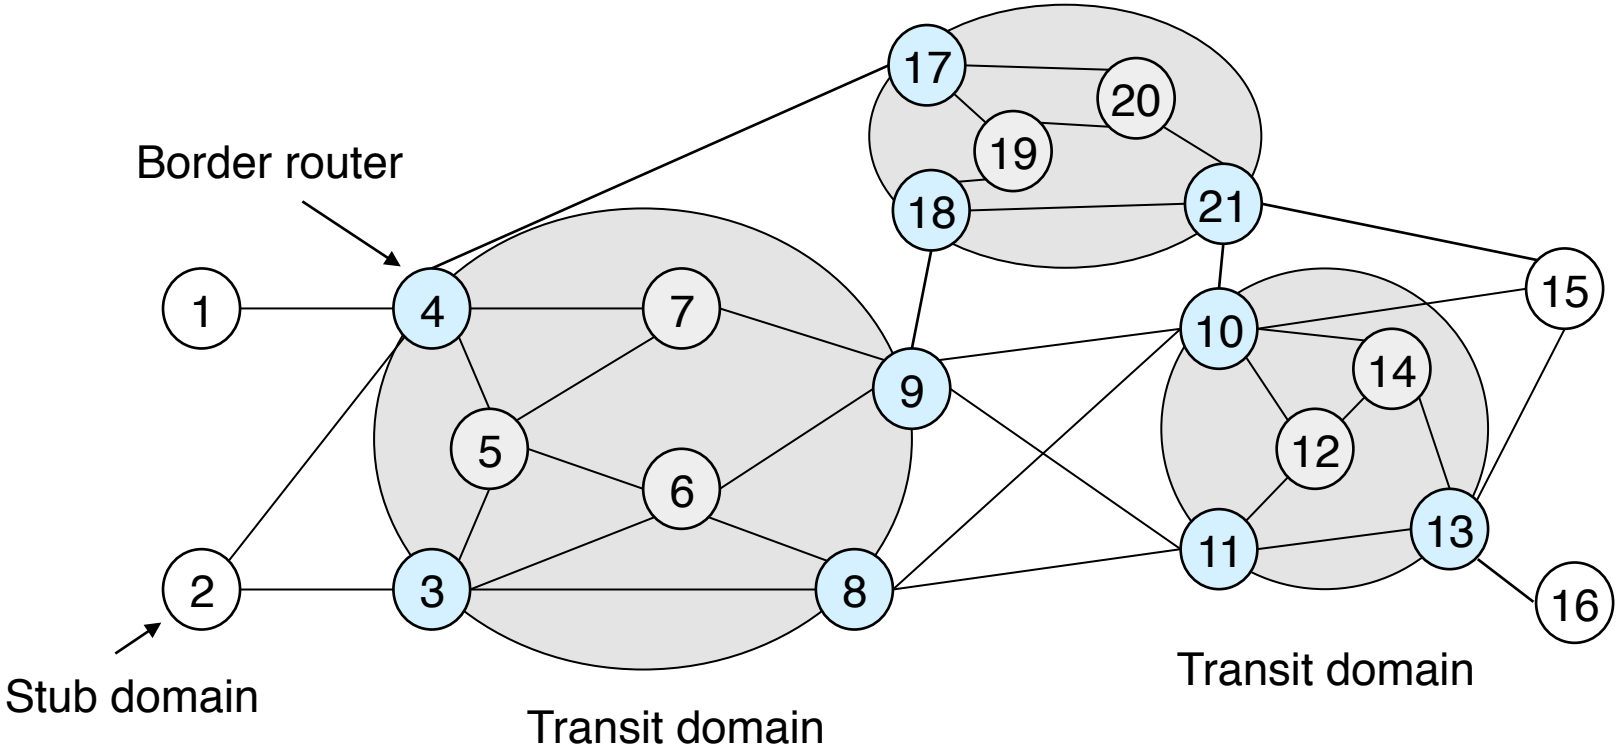

Border router: routers with direct connectivity to other domains
Stub domain: domain that originates or sinks traffic
Transit domain: domain that forwards transit traffic

Visualisation of IP Addresses

- **Problem**: how to visualize 4 billion IP adresses?
  - Number line: length 2^{32} pixels not feasible (> 300 km with 300 DPI)
  - Bitmap: 2^{16} x 2^{16} pixels (25 m^2 with 300 DPI)
  - Visualisation of /24 networks (2^8 IP adresses per pixel)
    ⇨ bitmap with 2^{12} x 2^{12} Pixel (16 MPixel, A4 with 300 DPI)

- Requirement: meaningful neighbourhood properties of addresses in bitmap
  - Number line: neighbourhood properties correct
  - Bitmap: neighbourhood properties depend on 2D mapping
  - Approach: room-filling curves

Room-Filling Curves

- **Approach**
  - Map curve to n-dimensional space
  - Requirement: complete filling of space with steady function
  - Recursion

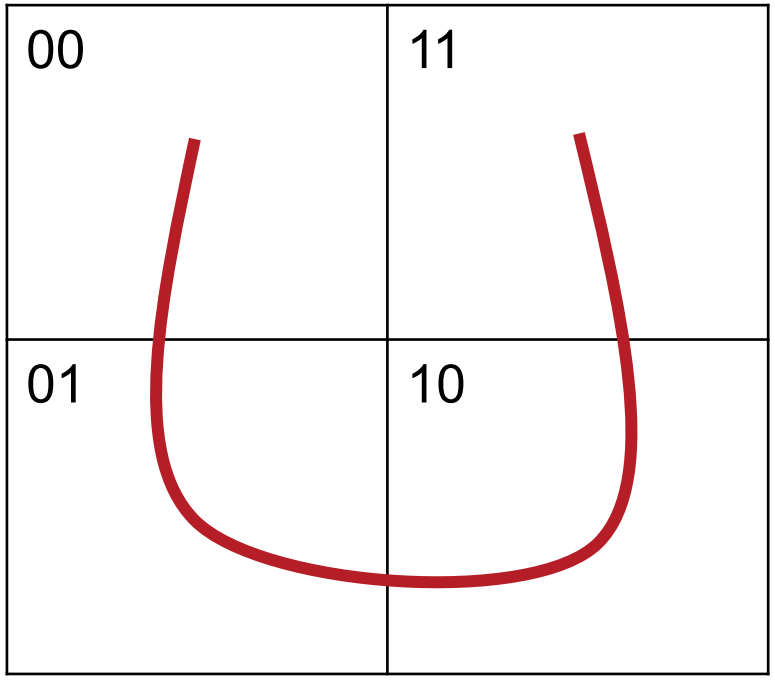

Room-Filling Curves

- **Approach**
  - Map curve to n-dimensional space
  - Requirement: complete filling of space with steady function
  - Recursion by *continuous fractal space-filling curve* using **Hilbert space-filling curve**

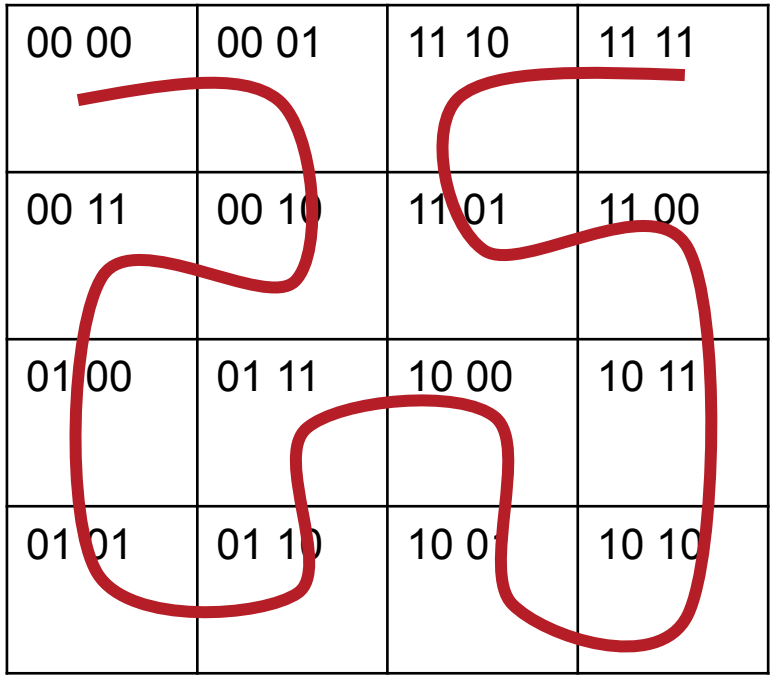

Room-Filling Curves

- **Approach**
  - Map curve to n-dimensional space
  - Requirement: complete filling of space with steady function
  - Recursion
    - base curve partitions room into 4 areas
    - rotation of base curve
    - continue up to needed depth

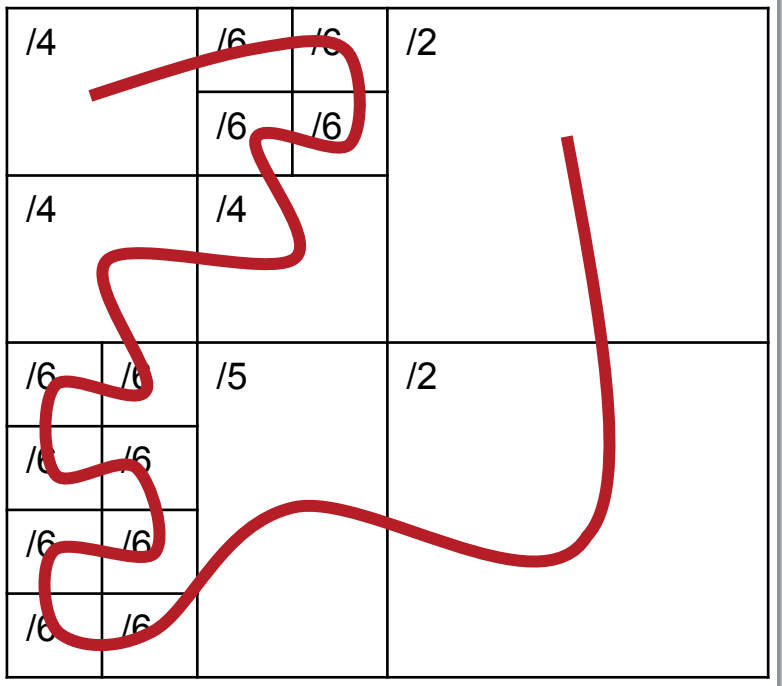

Room-Filling Curves

- **Hilbert curve** for 2D representation of IPv4 address space