Introduction

Advanced Computer Networks

Summer Semester 2016

Prof. Dr. Xiaoming Fu

Organizational Information

- https://wiki.net.informatik.uni-goettingen.de/wiki/Advanced_Computer_Networks_(Summer_2016)

- Slides are available online
- Course is held in English

- 5 ECTS credits
- AI: M.Inf.1222.Mp, M.Inf.1223.Mp
- ITIS: 3.17

Organizational Information

- 100% final exam
- Check the FlexNow registration deadlines

Course Overview

Date	Topic	Lecturer	
14.04.2015	Introduction	Prof. X. Fu	
21.04.2015	P2P Networking	Prof. X. Fu	
28.04.2015	NO LECTURE (Girls Day)		
05.05.2015	Data Analytics	Prof. X. Fu	
12.05.2015	NO LECTURE (Public Holiday)		
19.05.2015	Wireless I	Dr. X. Chen	
26.05.2015	Wireless II	Dr. X. Chen	
02.06.2015	Wireless III	Dr. X. Chen	
09.06.2015	Software-defined Networking I	Dr. D. Koll	
16.06.2015	Software-defined Networking II	Dr. M. Arumaithurai	
23.06.2015	Datacenter Networks	Dr. D. Koll	
30.06.2015	Information Centric Networks I	Dr. M. Arumaithurai	
07.07.2015	Information Centric Networks II	Dr. M. Arumaithurai	
14.07.2015	Written Examination		

Course Materials - Basics

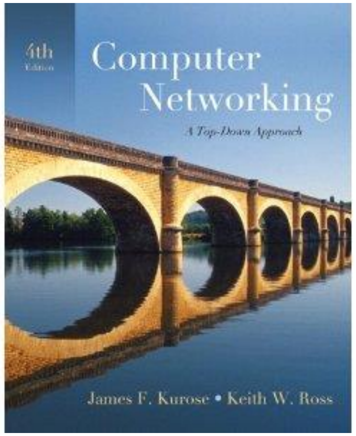

- J. Kurose and K. Ross, "Computer Networking: A Top-Down Approach Featuring the Internet“ (some slides are based on the book)
- A. S. Tanenbaum, "Computer Networks“

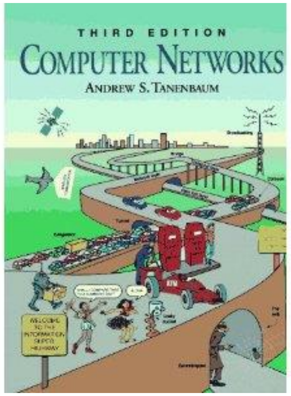

- Further materials are released on the wiki

Other related courses in SS16

- Practical Course on Advanced Computer Networking: Data Science Edition
- Seminar on Advanced Topics in Mobile Communications: Social Network in Mobile Big Data
- Seminar on Internet Technology

Course Materials - Advanced

- Peer-to-Peer Networks
  - Mahlmann and Schindelbauer, Peer-to-Peer Netzwerke: Methoden und Algorithmen, Springer, 2007, german

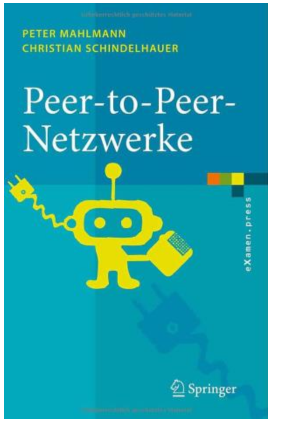

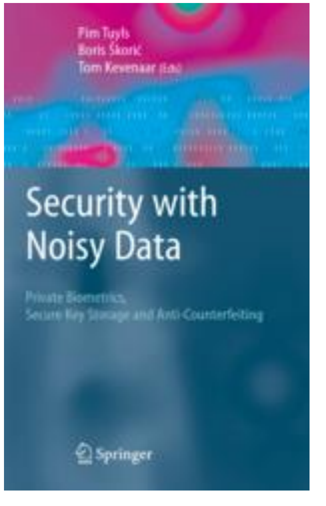

- Wireless Networks
  - Pim Tuyls et al, Security with Noisy Data, Springer, 2007

Course Materials - Advanced

- New Networking Technologies
  - Paul Goransson and Chuck Black, Software Defined Networks: A Comprehensive Approach (1st edition), Morgan Kaufmann, 2014.

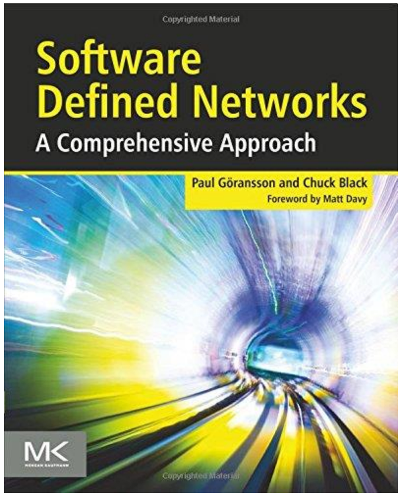

Recap of Basics

- The following slides repeat fundamentals for advanced networks:
  - Protocols / Layering
  - The hierarchical structure of the Internet
  - Addressing on the different layers
  - Routing and IP subnet aggregation
  - Layer 4 services such as reliable data transfer

- If something is new for you, please review the Computer Networks slides at: https://wiki.net.informatik.uni-goettingen.de/wiki/Computer_Networks_(Winter_2015/2016)

What is the Internet?

What is the Internet?

- millions of connected computing devices: *hosts, end-systems*
  - PCs, workstations, servers
  - PDAs, phones, toasters
  - running *network apps*
- *communication links*
  - fiber, copper, coax, radio, satellite
  - transmission rate = *bandwidth*
- *routers:* forward packets (chunks of data)

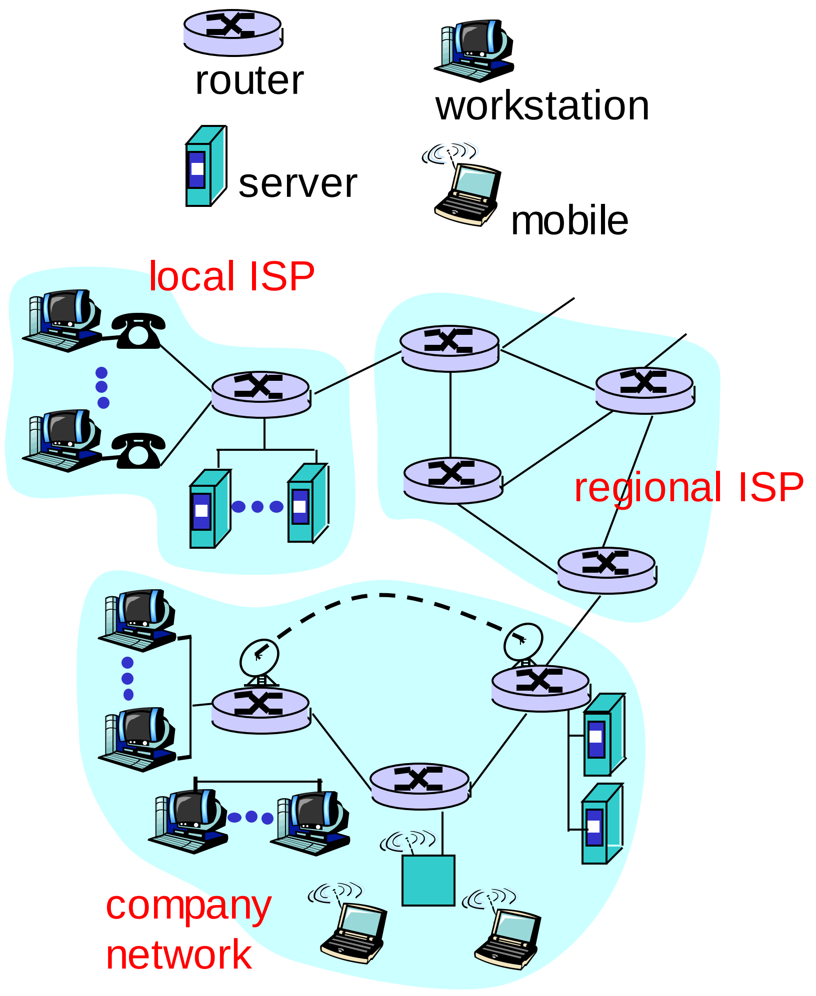

What is the Internet?

- *protocols* define format, order of msgs sent and received among network entities, and actions taken on msg transmission, receipt
- *Internet:* "network of networks"
  - loosely hierarchical
  - public Internet versus private intranet
- *Internet standards*
  - RFC: Request for Comments
  - IETF: Internet Engineering Task Force

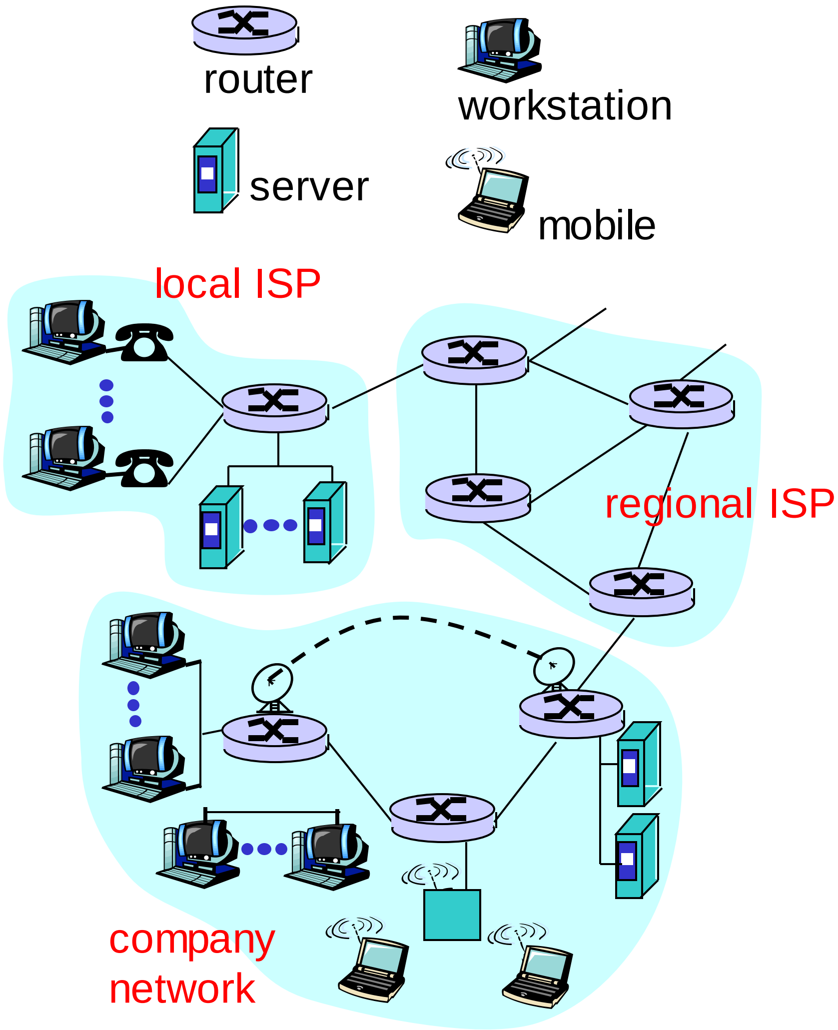

What is the Internet?

- The network is typically divided into the network core and the network edge
  - Edge: end systems, access networks, links
  - Core: circuit switching, packet switching, network structure, routers etc.

- Users access the Internet from the edge: e-mail, http web access, applications, social networks

At the edge

- client/server model
  - client host requests, receives service from always-on server
  - e.g. web browser/server; email client/server
  - **Is that optimal for large scale?**
- peer-peer model:
  - minimal (or no) use of dedicated servers
  - e.g. Skype, BitTorrent

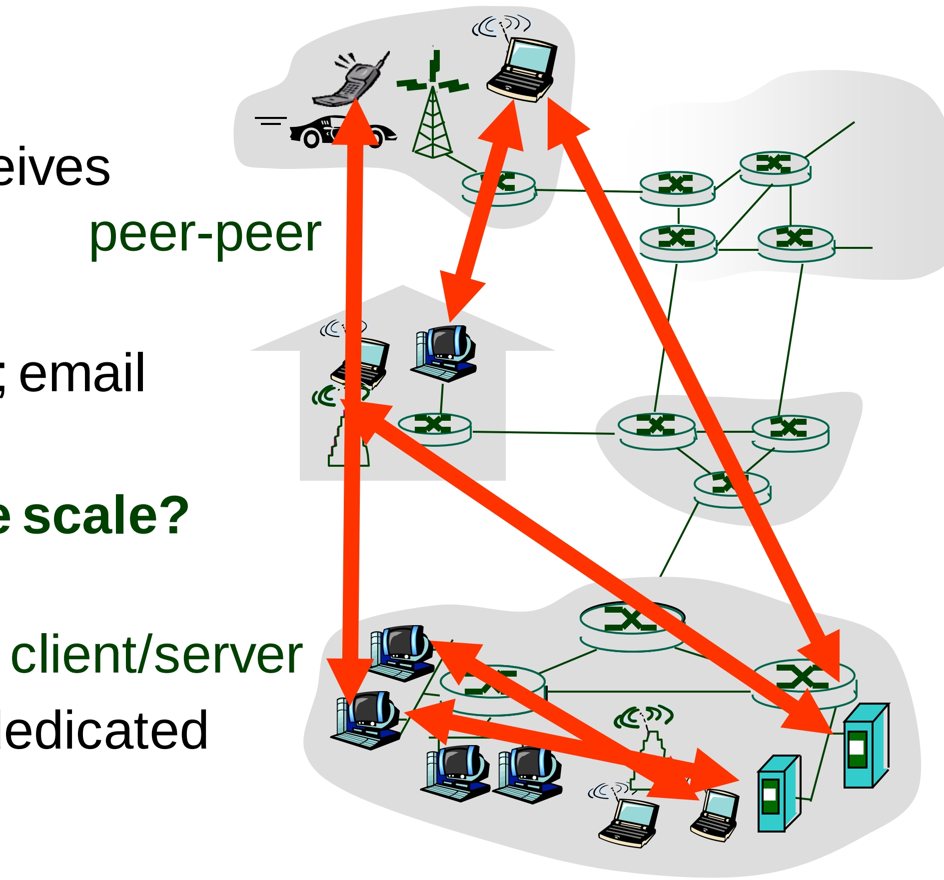

Packet Switching Implications

- End data stream is divided into „packets“
- Multiple users share network resources
  - Sequence of packets does not have fixed pattern, this effect is called statistical multiplexing
  - Each packet uses full bandwidth
  - Resources are used as needed
- No dedicated resource allocation!
- Demand may exceed resources:
  - Congestion, delay, loss

Hierachical Structure

- A packet traverses multiple networks from source to destination!

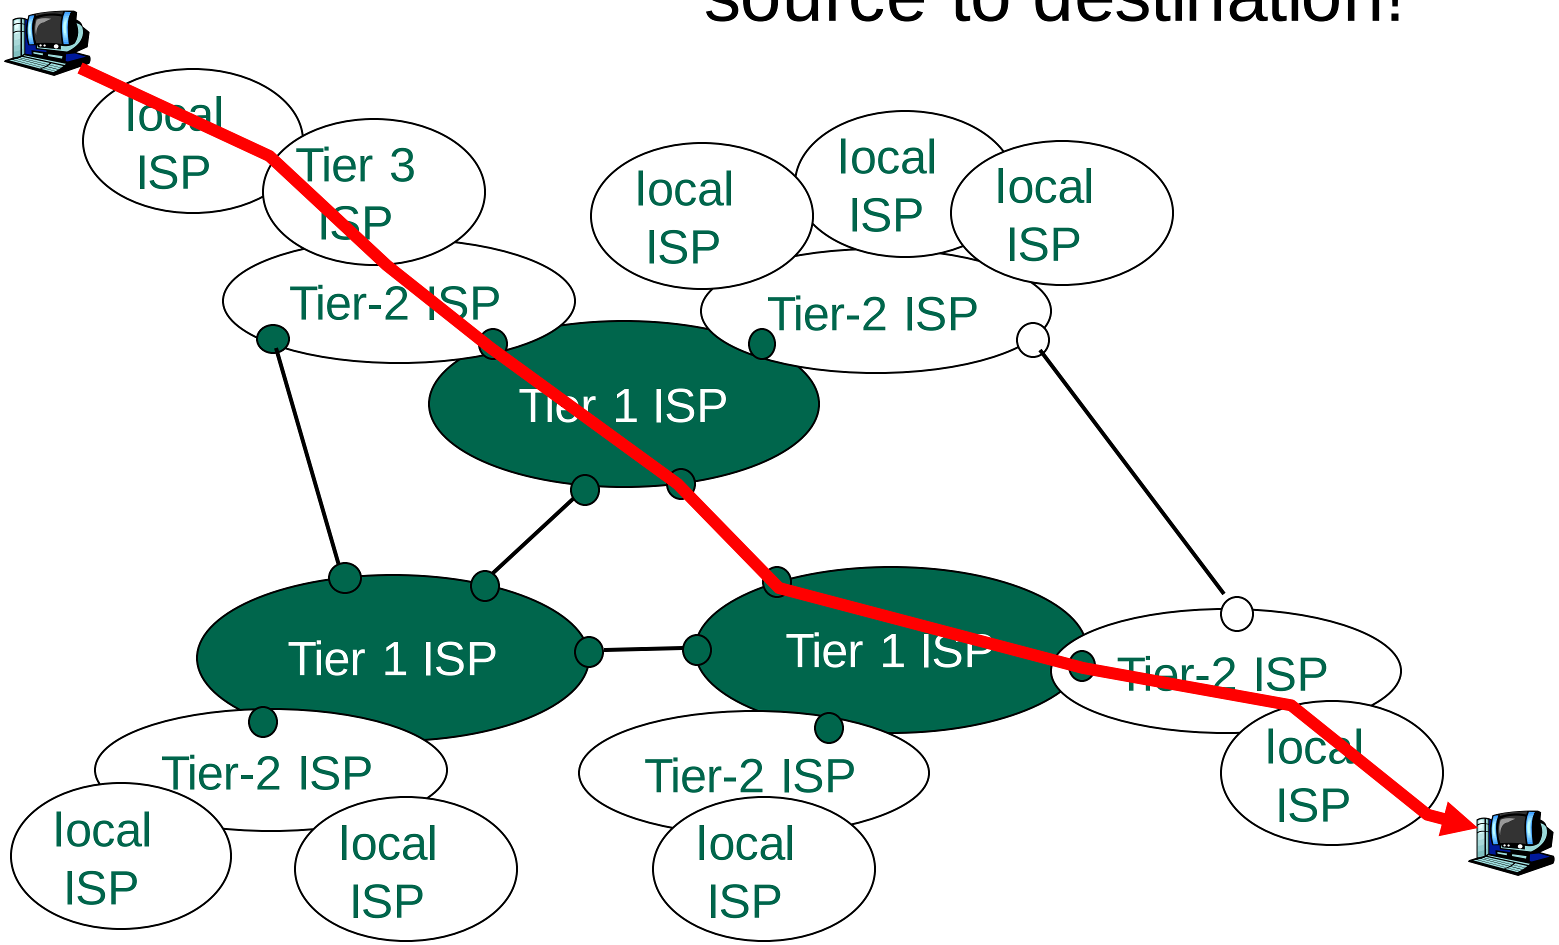

Impact of Hierarchy

- Main goal accomplished:
  - Efficient, works!
- But:
  - Providers have to buy resources from higher tiers or exchange directly on low tier level (peering)
  - Transferring data through external network is costly
  - Many modern technologies make it relevant HOW a packet traverses through the net
    - File sharing consumes large of resources, P2P is not business model of providers! Providers have an incentive to control P2P flows!
  - Efficient content delivery of large amounts of data (IP-Television, Video etc…)

computer NETWORKS

Protocol Layering

- Networks are complex and consist of many pieces: hosts, routers, links, applications, protocols et cetera...

- Idea: Simplify by using layers to distinguish organizational parts of networking.

- Inspired by other, real-world „layered“ processes

Internet Protocol Stack

- Application: network apps
  - FTP, SMTP, HTTP
- Transport: process-to-process data transfer
  - TCP, UDP
- Network: routing of datagrams from source to destination
  - IP, other routing protocols
- Link: data transfer between neighboring network elements
  - PPP, Ethernet
- Physical: bits „on the wire“

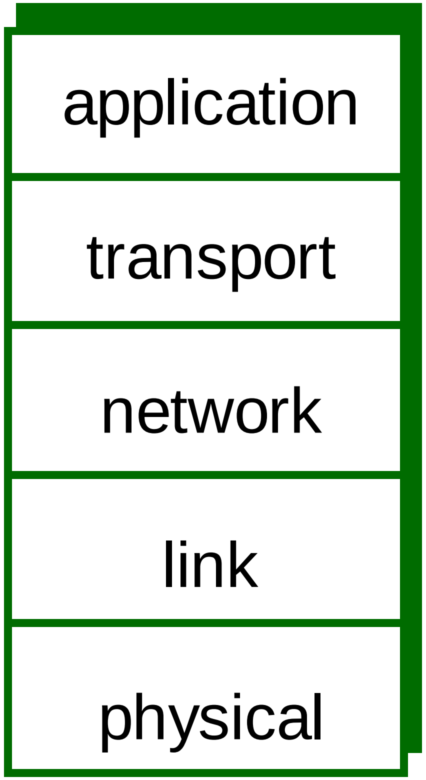

ISO/OSI Reference Model

- Presentation: allow applications to interpret meaning of data, e.g., encryption, compression, machine-specific conventions
- Session: synchronization, checkpointing, recovery of data exchange
- Internet stack "missing" these layers!
  - these services, *if needed,* must be implemented in application
  - needed?

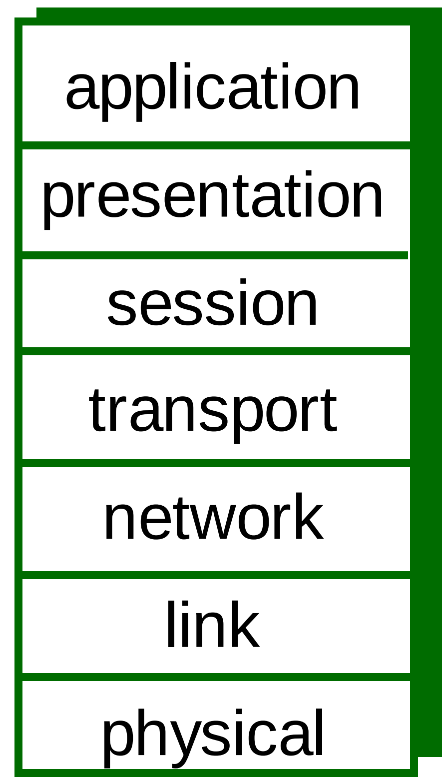

Layering Challenges

- Many modern networks are application level networks
  - Social Networks such as Facebook, Twitter etc.

- P2P networks are a mixture of different layers
  - Sometimes implementing application level data lookup, packet (chunk) forwarding etc.

- Large scale applications demand optimizations on various levels:
  - Content delivery of large data streams!

Layer 2 - Principles

- Switch processes up to layer 2
- Principles:
  - Reliable transmission of data over a link
  - Error detection, correction
  - Sharing a broadcast channel: multiple access
  - Link layer addressing
- Layer 2 packet is referred to as frame

data-link layer has responsibility of transferring datagram from one node to adjacent node over a link

Layer 2 – Error Detection

- Implemented in layer 2 adaptors (network interface cards etc.)
- Idea: Combine data payload with a Checksum to detect transmission errors

- Implemented using Cyclic Redundancy Checksum

- Typically frame is dropped if error occured. Re-transmit (reliability) has to be implemented at other layer

Layer 2 – Multiple Access Prot.

- In general two types of „links“:
  - Point-to-Point (e.g., PPP for dial-up, link between Ethernet switch and host)
  - Broadcast medium
    - Shared wire in old Ethernet
    - 802.11 wireless LAN

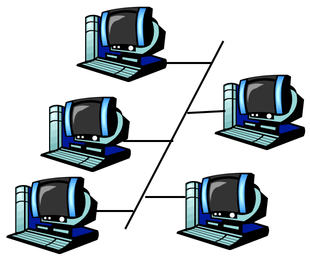

shared wire (e.g., cabled Ethernet)

- In broadcast:
  - Distributed algorithms determine how nodes share channel and transmit
  - No „out-of-band“ channel for coordination

Layer 2 – Addressing

- 48 BIT MAC address
  - Burned in NIC Rom (sometimes software settable)
  - Should be unique (vendors have MAC-ranges)
  - No hierachical or aggregatable information! Also called flat address.
- IP to MAC address resolved using the Address Resolution Protocol (ARP)
  - Idea: Broadcast query: „Who has IP ...?“. Host with queried IP replies unicast „IP ... is at MAC ...“.

computer NETWORKS

Layer 2 – Forwarding

- Solved by switches
  - Learns in a plug and play mode, no configuration necessary!
  - Question: On what switch-port to forward incoming dataframe
    - If destination MAC is known (in forwarding table) and associated to a port, use that port
    - If destination MAC is unknown, broadcast the dataframe on all ports
    - If a packet with an unknown source MAC is observed, create an entry in the forwarding table that binds port to MAC
- Only works in the broadcast domain! No routing!

Layer 3 – Principles

- Internet Protocol dominant networking technology
- Network layer transports datagrams from sending to receiving host
- Protocol runs in every host and router (not in switches!)

Hourglass architecture

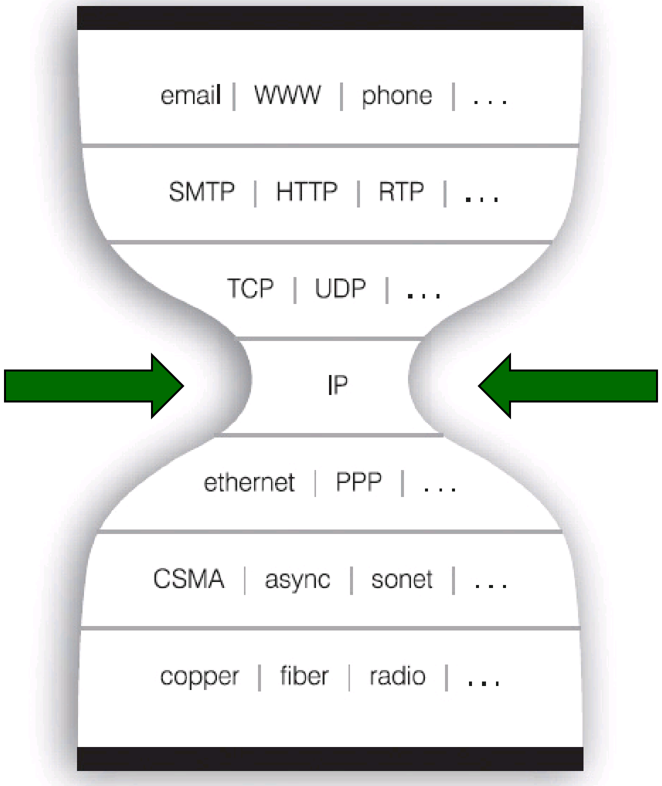

Layer 3 – Routing

- Routing: determine the route taken by packts from source to destination
  - Analogy: planning a road trip from one city to another
- Functions of a router:
  - Forward datagram from incoming to outgoing link
  - Select outgoing link by running a routing algorithm
    - Examples: RIP, OSPF, BGP
- Routers buffer packets (buffer size rule of thumb: Round-Trip-Time times link capacity
  - Buffering source of delay, overflowing buffer causes loss

Layer 3 – Addressing

- IP address: 32 bit identifier for host, router interface
- Interface: connection between host/router and physical link
  - Routers typically have multiple interfaces
  - Hosts typically have one interface
  - IP address associated with each interface

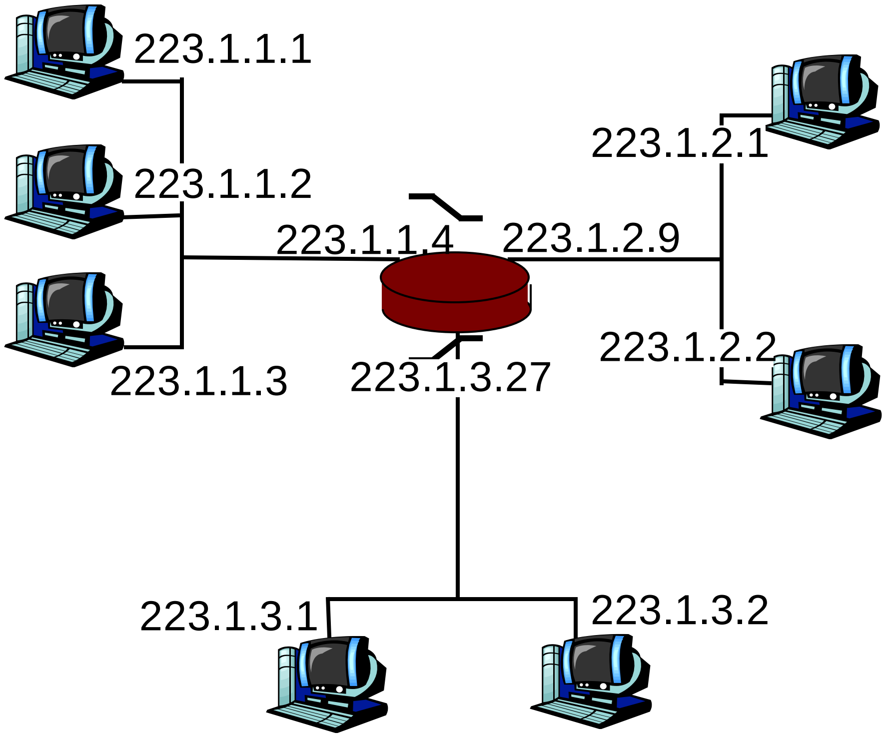

223.1.1.1 = 11011111 00000001 00000001 00000001

223 1 1 1

Layer 3 – Subnets

- IP address not flat
- Hierachical structure of „subnets“
- Inside a subnet:
  - All device interfaces share part of IP address
  - Devices can reach each other without intervening a router

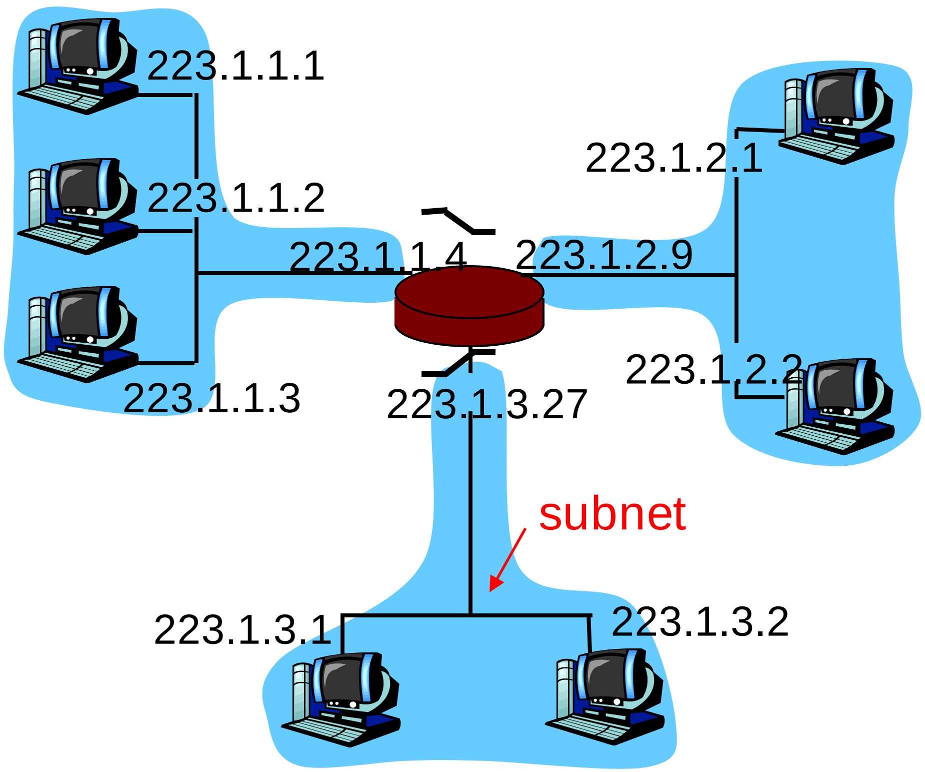

network consisting of 3 subnets

Layer 3 – How to get an IP?

- MAC address was set by vendor

- IP needs to be obtained from ISP
  - Either in a fixed configuration
  - Or using the Dynamic Host Configuration Protocol (DHCP) (server offers service and hosts request IPs)

- How does the ISP get address space?
  - Allocated by higher Tier! (remember the hierachy)
  - Highest assignment done by Internet Assigned Numbers Authority (IANA) via ICANN in the US

Layer 3 – Aggregation

- Hierachical addresses allow to aggregate some prefixes:

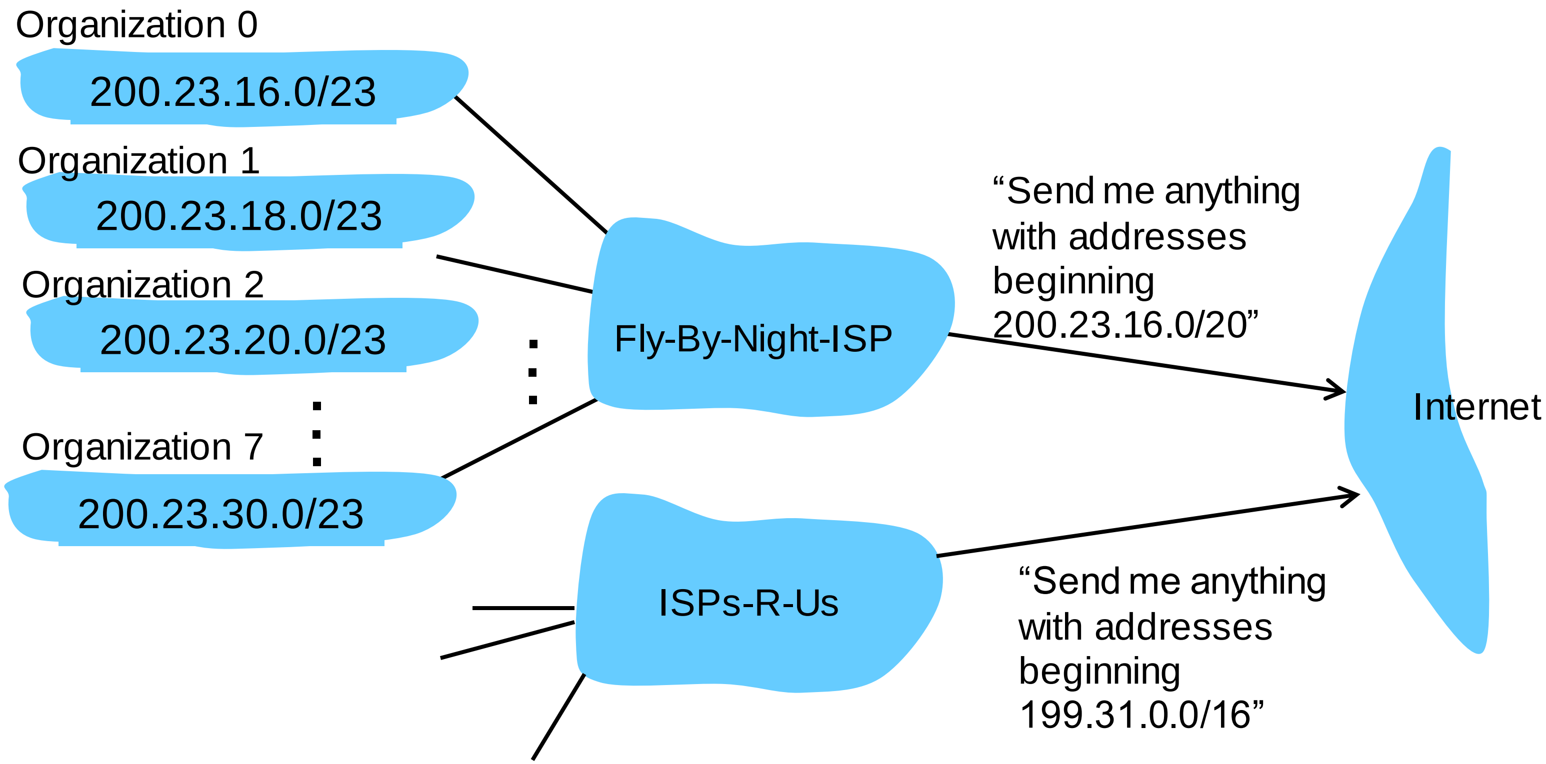

Layer 3 – NAT

- Network address translation (NAT) allows local network to have more users than assigned IP addresses!
  - Use one IP address for all devices
  - Addresses inside the network can change without need to notify the outside world
  - The ISP may change without changes necessary inside the network
  - Devices inside the local network are not explicitly addressable by the outside world!
    - Security plus but makes connections more difficult!

Layer 3 – NAT

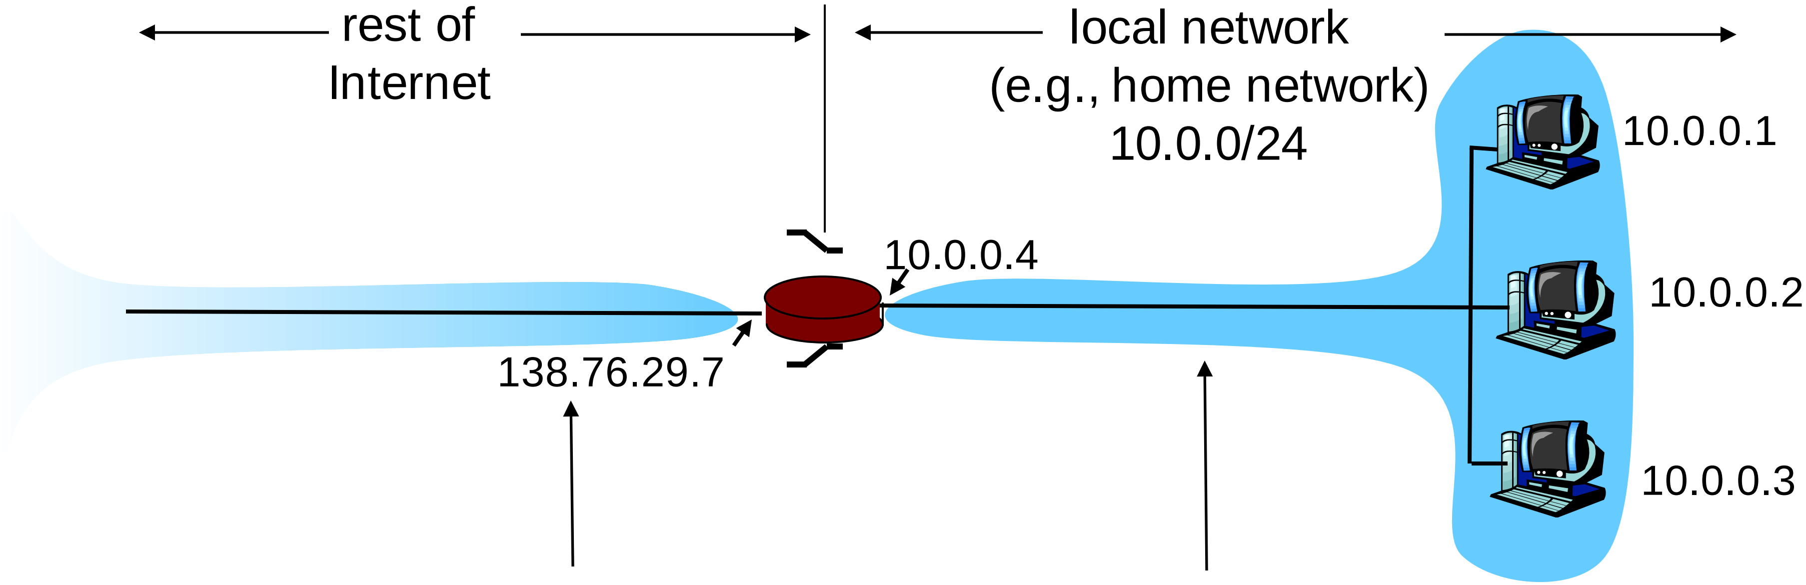

All datagrams *leaving* local network have same single source NAT IP address: 138.76.29.7, different source port numbers

Datagrams with source or destination in this network have 10.0.0/24 address for source, destination (as usual)

Layer 3 – NAT Challenges

- NAT violates end-to-end argument
- NAT traversal hard to solve for certain applications, for example P2P
- Address shortage should instead be solved by IPv6

Layer 3 – Routing

- Internet divided into Autonomous Systems (ASes)
- AS internal routing = intra-AS routing
  - E.g., OSPF
- Routing between ASes = inter-AS routing
  - E.g., BGP

Layer 3 – Intra-AS: OSPF

- "open": publicly available
- uses Link State algorithm
  - LS packet dissemination
  - topology map at each node
  - route computation using Dijkstra's algorithm
- OSPF advertisement carries one entry per neighbor router
- advertisements disseminated to entire AS (via flooding)
  - carried in OSPF messages directly over IP (rather than TCP or UDP

Layer 3 - OSPF features

- Security: all OSPF messages authenticated (to prevent malicious intrusion)
- multiple same-cost paths allowed
- For each link, multiple cost metrics for different TOS
- Integrated uni- and multicast support:
  - Multicast OSPF (MOSPF) uses same topology data base as OSPF
- Hierarchical OSPF in large domains.

Layer 3 - Hierarchical OSPF

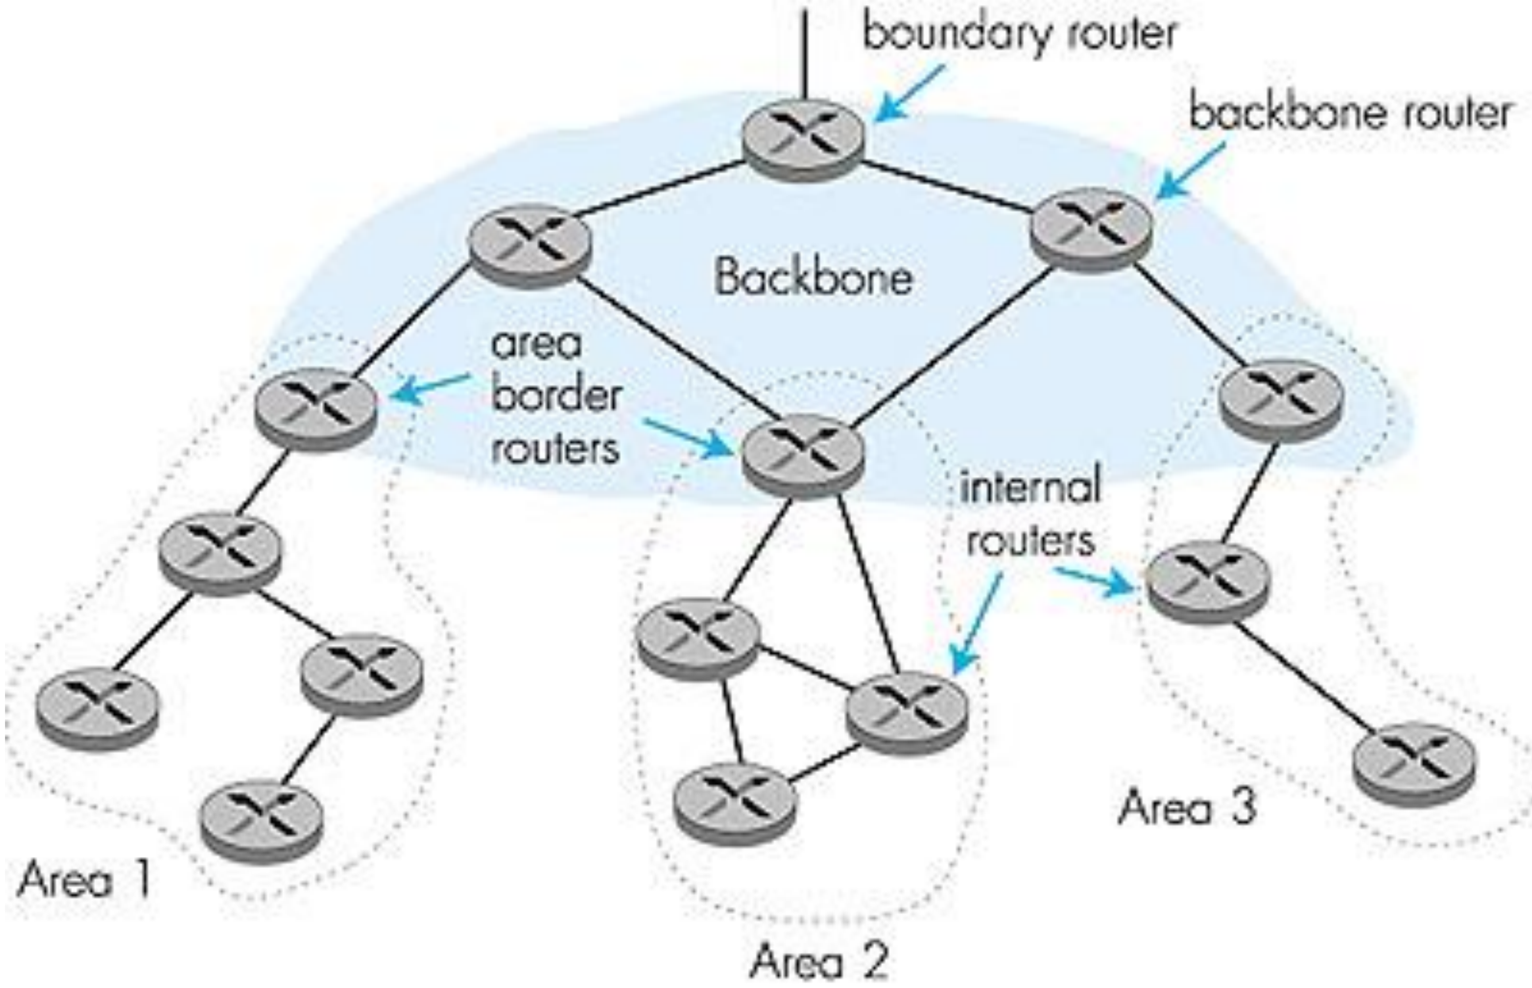

Layer 3 - Hierarchical OSPF

- two-level hierarchy: local area, backbone.
  - Link-state advertisements only in area
  - each nodes has detailed area topology; only know direction (shortest path) to nets in other areas.
- area border routers: “summarize” distances to nets in own area, advertise to other Area Border routers.
- backbone routers: run OSPF routing limited to backbone.
- boundary routers: connect to other AS’s.

Layer 3 - Inter-AS routing: BGP

- BGP (Border Gateway Protocol): the de facto standard
- allows subnet to advertise its existence to rest of Internet: “I am here”

Layer 3 - BGP basics

- pairs of routers (BGP peers) exchange routing info over semi-permanent TCP connections: BGP sessions
  - BGP sessions need not correspond to physical links.
- when AS2 advertises a prefix to AS1:
  - AS2 *promises* it will forward datagrams towards that prefix.
  - AS2 can aggregate prefixes in its advertisement

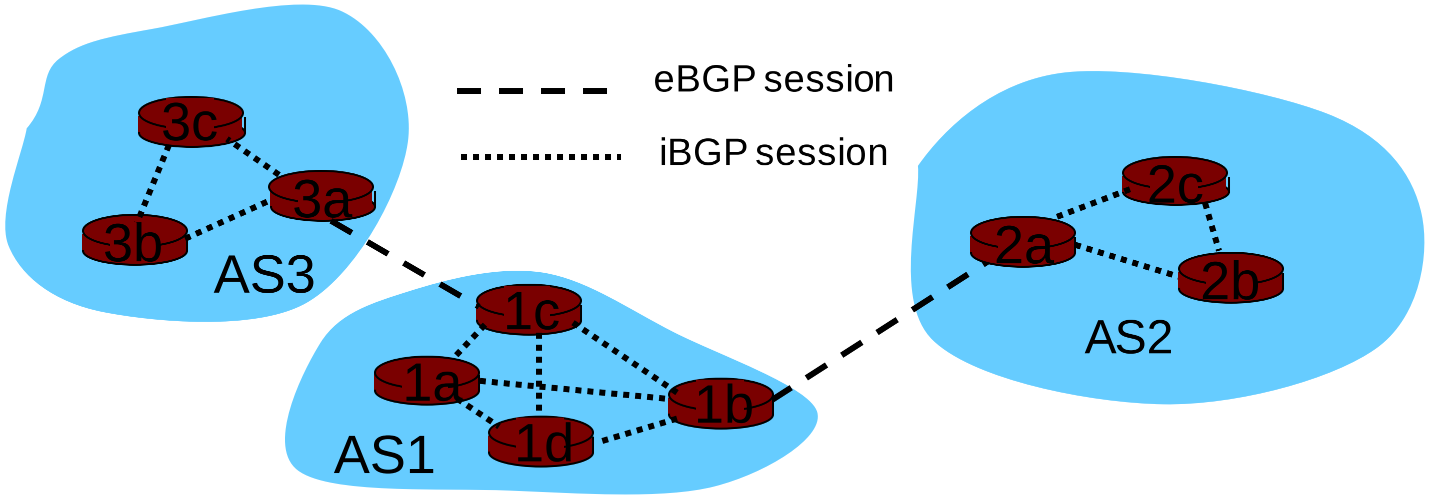

Layer 3 – BGP Reachability

- using eBGP session between 3a and 1c, AS3 sends prefix reachability info to AS1.
  - 1c can then use iBGP do distribute new prefix info to all routers in AS1
  - 1b can then re-advertise new reachability info to AS2 over 1b-to-2a eBGP session
- when router learns of new prefix, it creates entry for prefix in its forwarding table.

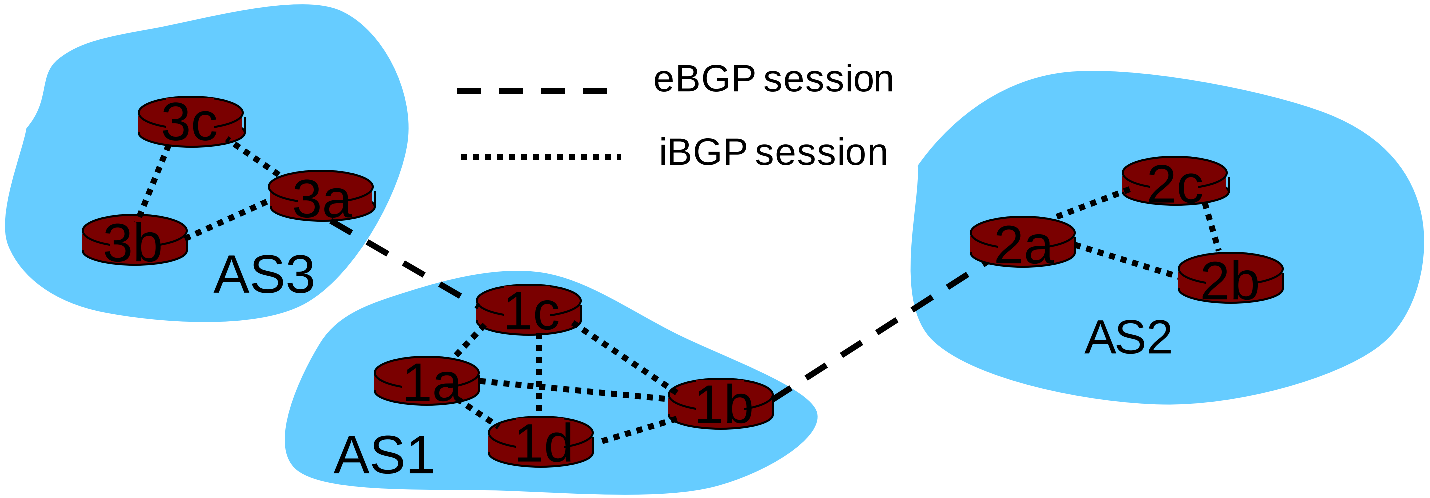

Layer 3 – BGP Paths and Routes

- advertised prefix includes BGP attributes.
  - prefix + attributes = “route”
- two important attributes:
  - AS-PATH: contains ASs through which prefix advertisement has passed: e.g, AS 67, AS 17
  - NEXT-HOP: indicates specific internal-AS router to next-hop AS. (may be multiple links from current AS to next-hop-AS)
- when gateway router receives route advertisement, uses import policy to accept/decline.

Layer 3 - BGP route selection

- Router may learn about more than 1 route to same prefix:
  - Router must select route
- Elimination rules:
  - Local preference value attribute: policy decision
  - Shortest AS-PATH
  - Closest NEXT-HOP router: hot potato routing
  - Additional criteria

Layer 3 – BGP Policies

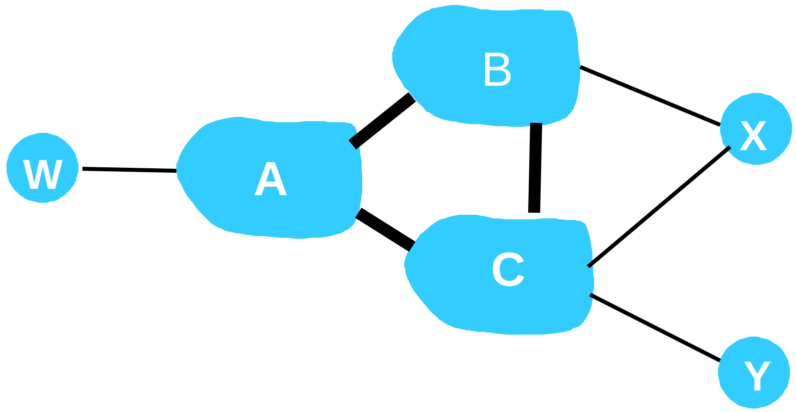

legend:

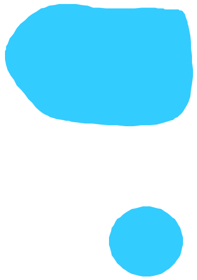

- A,B,C are provider networks
- X,W,Y are customer networks
- X is dual-homed: attached to two networks
  - X does not want to route from B via X to C
  - .. so X will not advertise to B a route to C

Layer 3 – BGP Policies

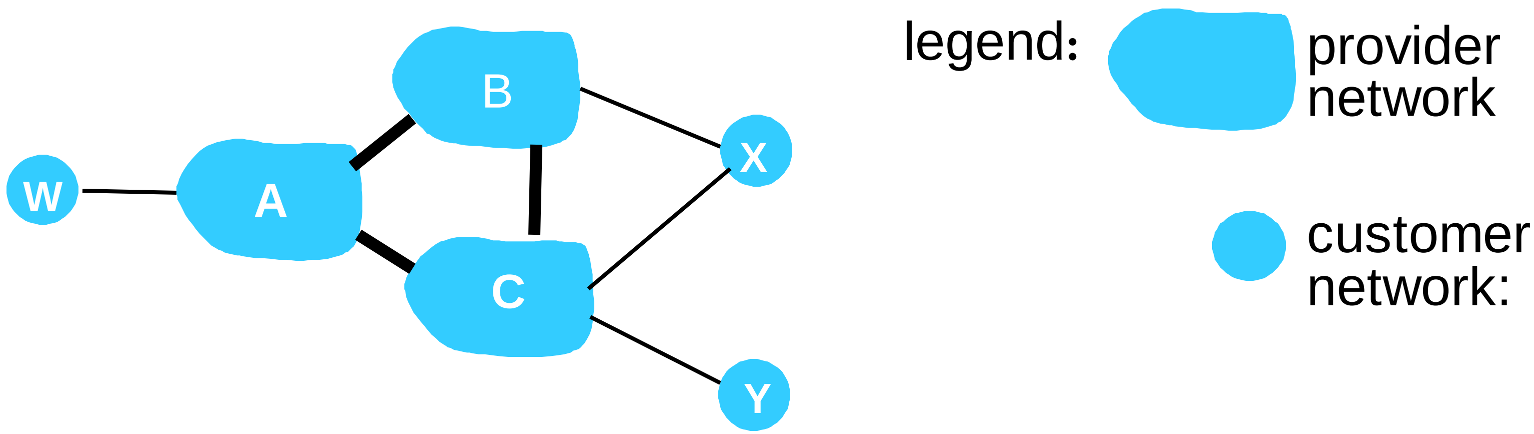

- A advertises path AW to B
- B advertises path BAW to X
- Should B advertise path BAW to C?
  - No way! B gets no “revenue” for routing CBAW since neither W nor C are B’s customers
  - B wants to force C to route to w via A
  - B wants to route *only* to/from its customers!

Layer 4 – Principles

- Provide logical communication between app processes running on different hosts
- Relies on and enhances network layer services
- Analogy: Network = postal service, Transport = secretaries at sender and receiver

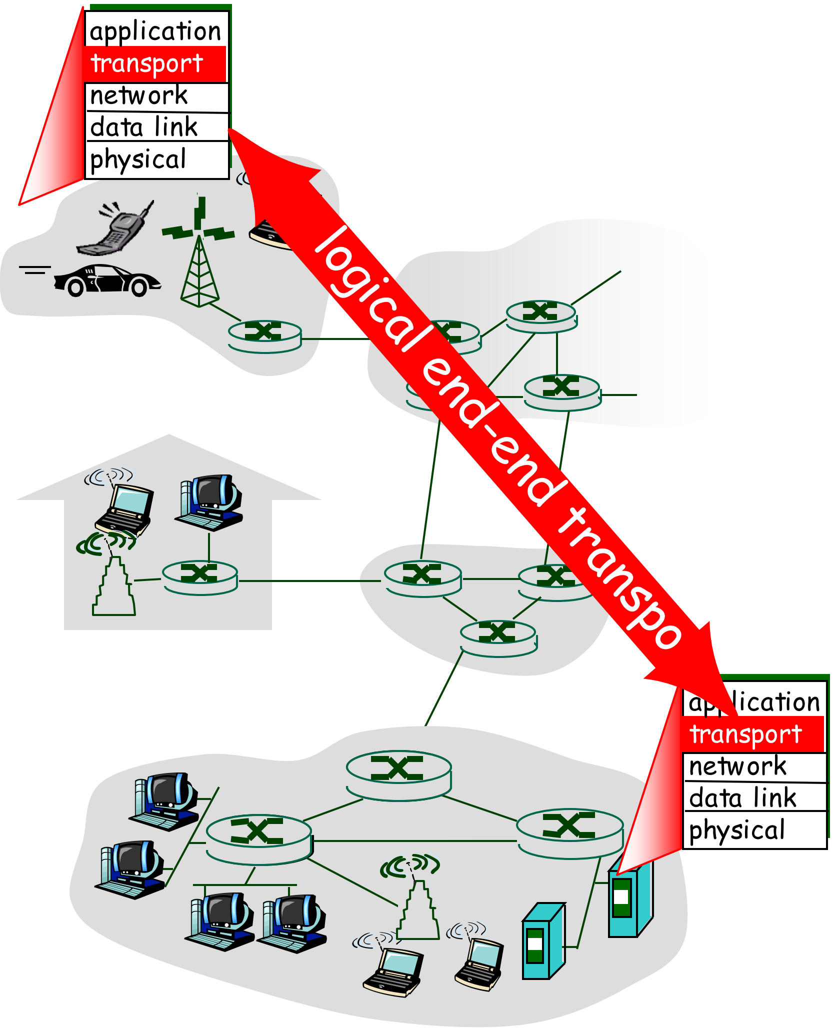

Layer 4 – UDP and TCP

- UDP provides unreliable, unordered delivery
  - Conceptually minimal transport layer protocol with „best-effort“ IP
- TCP provides
  - Reliable data transfer
  - In-order delivery
  - Congestion control
  - Flow control
  - Connection setup
- Both protocols do not guarantee certain delays or bandwidths

Layer 4 – UDP and TCP

- UDP is minimalistic, therefore
  - Used in time-critical applications (Voice over IP)
  - Used when loss or out-of-order delivery is not critical
  - When the concept of connection is not of high relevance
- TCP ensures reliable data transfer:
  - Used when correctness is important
  - Typically used in all long-term non voice/video connections such as file transfers

Layer 4 – Addressing

- TCP socket is identified by 4-tuple:
  - Source IP address
  - Source port number
  - Destination IP address
  - Destination port number
- This allows multiple apps to communicate simulataniously with same source and destination host but, for example, different source port

Layer 4 – Data Transfer

- Reliable data transfer (TCP only) is ensured by
  - Maintaining state for each packet or group of packets (timeout timers)
  - Resend upon loss
    - Loss is inferred by missing Acknowledgments
  - Only send new packets if previous ones got acked
- Speeding things up: Pipelining
  - Idea: Maintain a large group of packets „on the fly“ using a window. Examples: Go-back-N, Selective Repeat

Layer 4 – TCP Properties

- Point-to-Point - one sender, one receiver
- Reliable, in-order delivery
- Pipelined – TCP congestion and flow control sets window sizes
- Send & receive buffers
- Full-duplex data transfer
- Connection oriented
- Flow-Control – Sender prevents to overwhelm receiver

Application Layer

- Applications and application layer protocols:
  - HTTP/HTTPs for Web-Access
  - SMTP for e-mail
  - FTP/SFTP for data transfer
  - Many more…

- Challenge:
  - Utilize social information, peer-to-peer services, intelligent data storage and services (cloud) etc. to optimize content delivery!

Outlook

- In the next lectures:
  - Introduction to p2p and data analytics;
  - Problems with the traditional networking paradigms;
  - What can be done to alleviate these problems