universität wien

Chair of Future Communication
Faculty of Computer Science
Prof. Dr. K. Tutschku

050069

VO Netzwerktechnologie für Multimedia Anwendungen

Lecture 2: Review of Networking Concepts

Prof. K. Tutschku (kurt.tutschku@univie.ac.at)

Endowed by A1

Bachelor Informatik (Medieninformatik)

WS 2010/11

Review of Networking Concepts

- Overview:
- Protocol layering and Internet protocol stack
- Circuit switching vs. packet switching
- Connectionless vs. connection-oriented networks, routing, forwarding, and switching
- Transport layer protocols
- Application layer
  - Sockets
  - Client-server and peer-to-peer communication
- Web services

Überblick

- Das TCP/IP Referenzmodell
- Kritik an Referenzmodellen

Chair of
Future Communication

endowed by

universität wien

The Hourglass Model

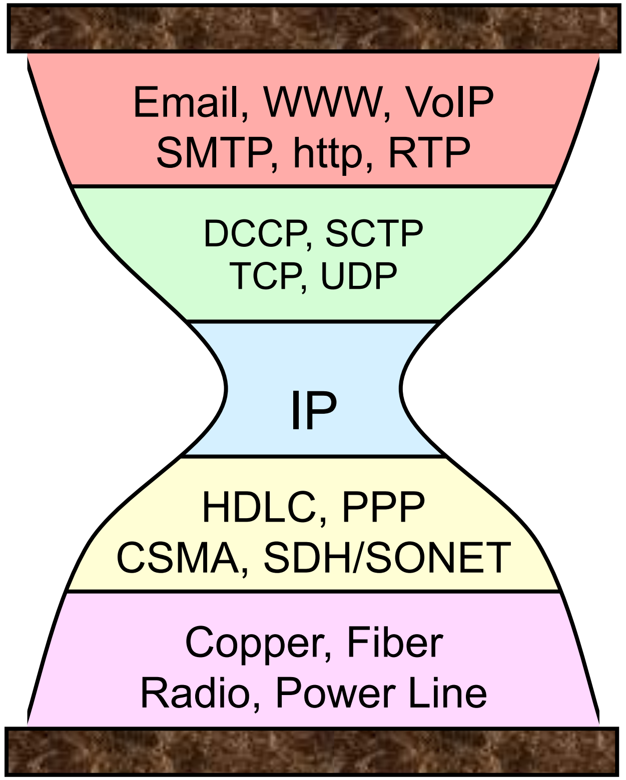

Everything over IP

IP over everything

Chair of **Future Communication**

endowed by A1

universität wien

Überblick

- Das TCP/IP Referenzmodell
- Kritik an Referenzmodellen

universität wien

Kritik an Referenzmodellen (1)

- ISO/OSI
  - Unausgeglichene Funktionsfülle der einzelnen Schichten: Presentation Layer, Session Layer, Application Layer
  - Hat bei Implementierung in der Praxis versagt
  - ITU-T bei der Standardisierung neuer Protokolle langsam
- TCP/IP
  - Schnelle Verbreitung über BSD/UNIX: gut implementiert, einfach zu benutzen, kostenfrei
  - IETF bei der Standardisierung neuer Protokolle schnell
  - Spezielle Beschreibung des Status Quo, nicht allgemein
- 5 Schichten-Modell
  - ISO/OSI ohne Session und Presentation Layer
  - Meist nützliche und häufig verwendete Taxonomie

Kritik an Referenzmodellen (2)

- Gefahr: schränken Denken ein
- Erlauben keinen Informationsaustausch zwischen den Schichten
  - Cross Layer Design (CLD): Austausch von Informationen über möglicherweise mehrere Layer
  - Beispiele für CLD
    - Ausnutzung von Physical Layer Information um Application Layer zu adaptieren bzw. zu optimieren
    - Location-aware Services
    - TCP reagiert auf ECN-bit (explicit congestion notification) in IP Header
- Einordnung von Protokollen in Schichten manchmal problematisch
  - TCP/IP/MPLS/SDH/WDM: MPLS gilt als Layer 2.5
  - TCP/IP1/UMTS-Schichten/IP2/ATM/SONET:
    IP kommt im Network und im Link Layer zum Einsatz

Überblick

- Das TCP/IP Referenzmodell
- Kritik an Referenzmodellen

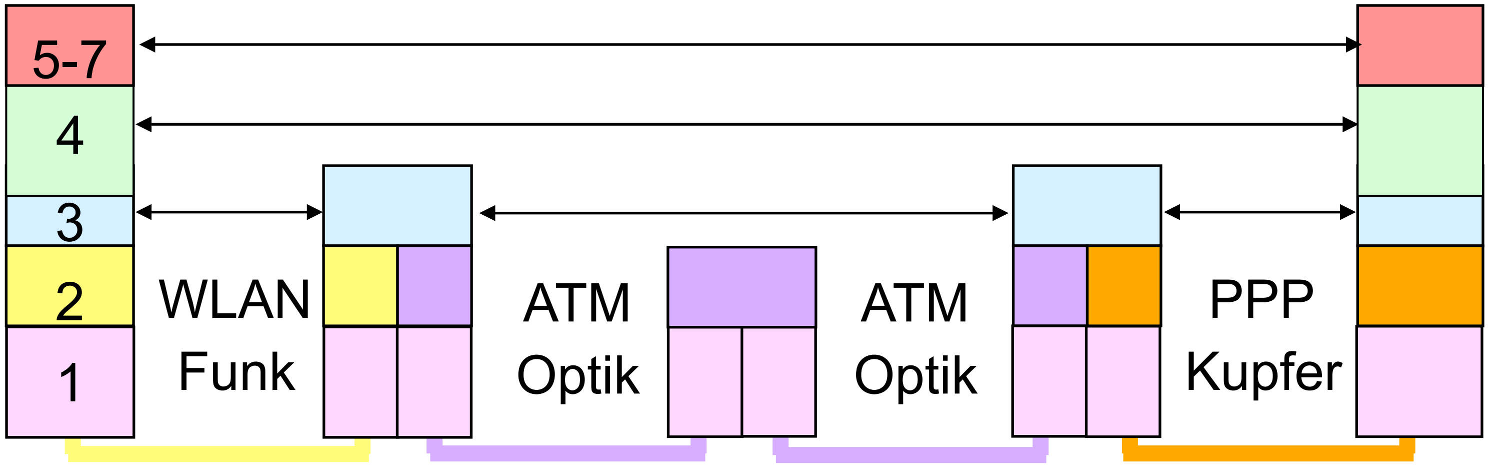

universität wien

Revision of Networking Concepts

- <u>Overview:</u>
- Protocol layering and Internet protocol stack
- Circuit switching vs. packet switching
- Connectionless vs. connection-oriented networks, routing, forwarding, and switching
- Transport layer protocols
- Application layer
  - Sockets
  - Client-server and peer-to-peer communication
- Web services

Circuit Switching: FDM and TDM

- Network resources (e.g., bandwidth) divided into "pieces"
  - Frequency division multiplex
  - Time division multiplex
  - Pieces allocated to calls
  - Resource piece *idle* if not used by owning call *(no sharing)*

FDM

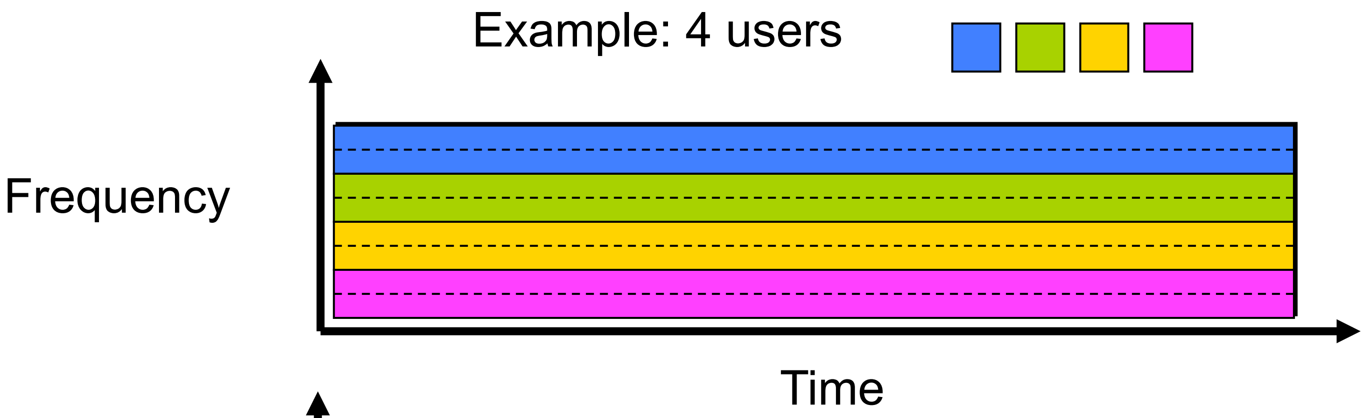

TDM

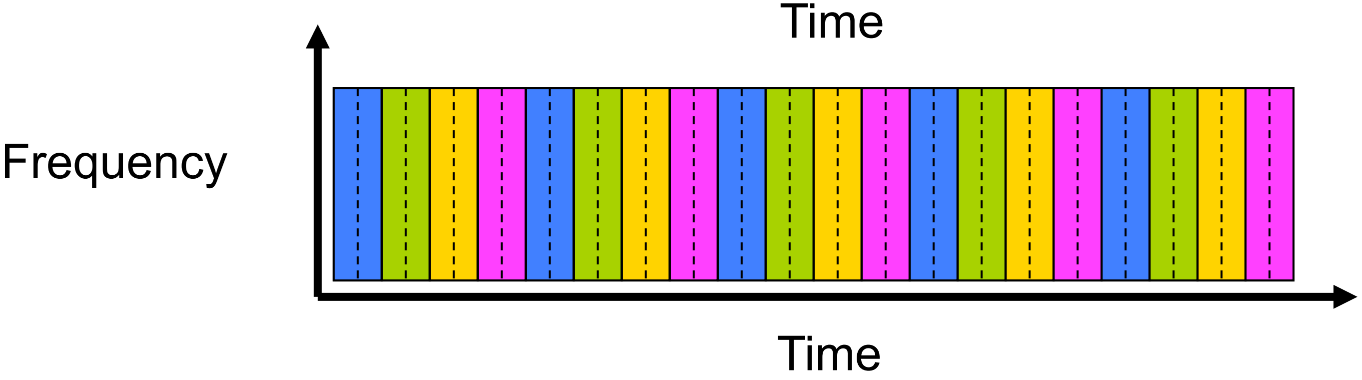

Chair of **Future Communication**

endowed by

universität wien

Packet Switching

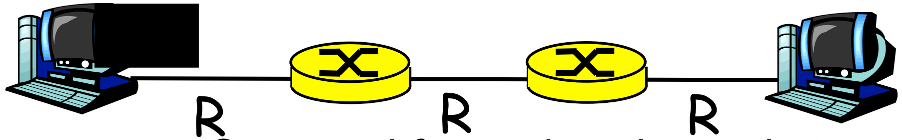

- Data stream divided into packets
  - Packet streams of different flows share network resources
  - Each packet uses full bandwidth
  - Resources used as needed
- Resource contention
  - Aggregate demand rate can exceed available capacity
  - Congestion: packets queue, wait for link use

- Store-and-forward: entire packet must arrive at router before it can be transmitted on next link
- Example
  - Link bandwidth R=1.5 Mbit/s
  - Msg size L=7.5 Mbit
  - Takes L/R=5 sec to transmit packet
  - 3 hops ⇒ overall delay = 15 sec

Bandwidth division into "pieces"
Dedicated allocation
Resource reservation

Chair of **Future Communication**

endowed by A1

Packet Switching: Statistical Multiplexing

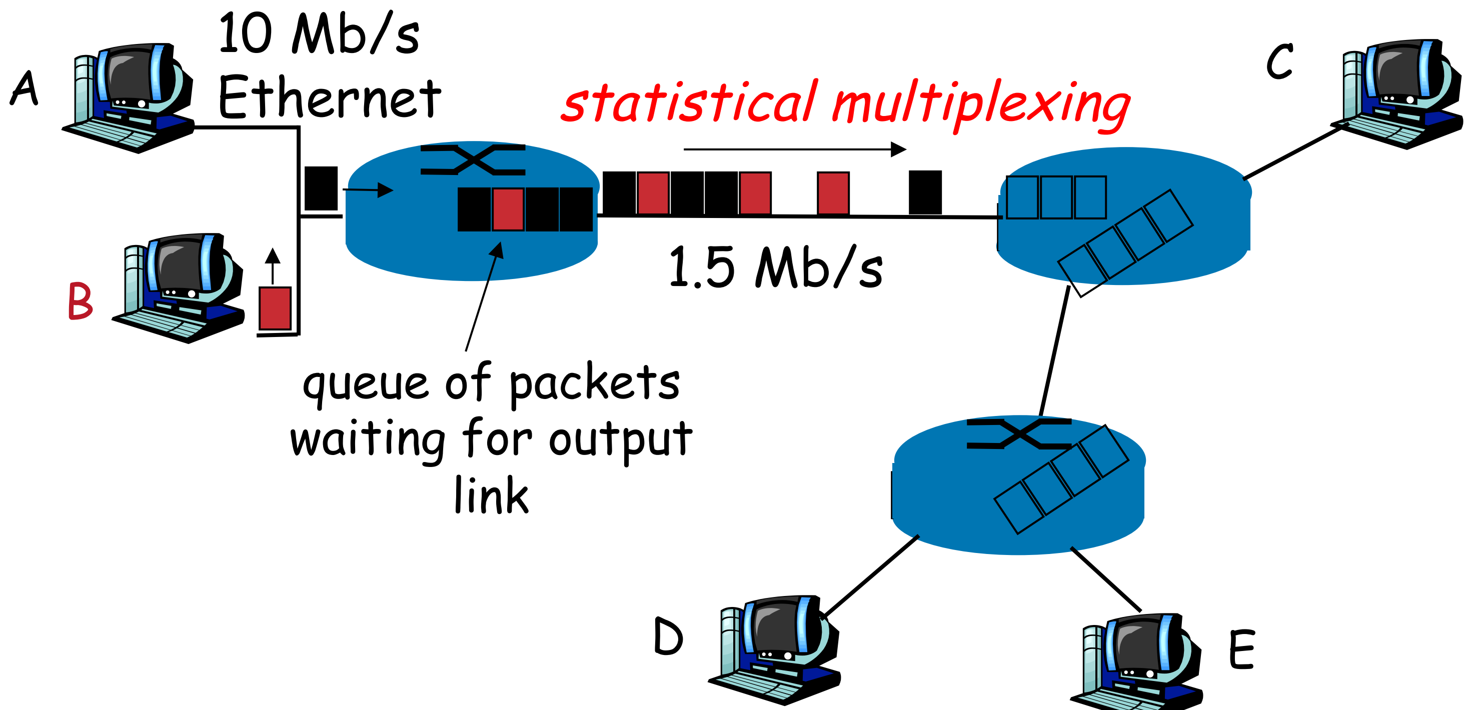

- Sequence of blue and green packets does not have fixed pattern ➡ *statistical multiplexing*
- In TDM each host gets same slot in revolving TDM frame.

Chair of **Future Communication**

endowed by A1

universität wien

Packet Switching vs. Circuit Switching

Chair of Future Communication

endowed by A1

- Comparison
  - Link: 1 Mbit/s
  - User
    - 100 kbit/s when “active”
    - Active 10% of time
  - Circuit switching: 10 users
  - Packet switching: with 35 users, probability > 10 active less than .0004
- Packet switching allows more users to use network!

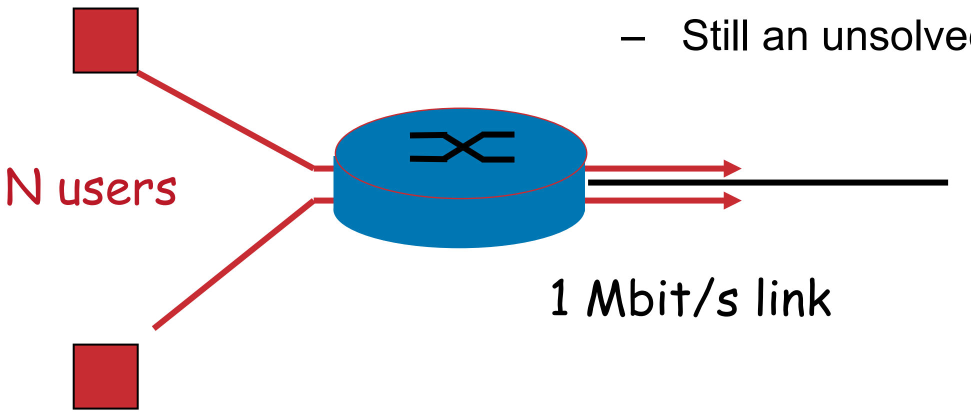

- Is packet switching better than circuit switching?
- Great for bursty data
  - resource sharing
  - simpler, no call setup
- Excessive congestion
  - Packet delay and loss
  - Protocols needed for reliable data transfer, congestion control
- How to provide circuit-like behavior?
  - Bandwidth guarantees needed for audio/video apps
  - Still an unsolved problem

Review of Networking Concepts

- Overview:
- A picture of the Internet
- Protocol layering and Internet protocol stack
- Circuit switching vs. packet switching
- Connectionless vs. connection-oriented networks, routing, forwarding, and switching
- Transport layer protocols
- Application layer
  - Sockets
  - Client-server and peer-to-peer communication
- Web services

Connectionless (Datagram) Networks

- No call setup at network layer
  - No network-level concept of “connection”
  - Routers: no per-flow state
- Packets forwarded using destination host address
  - But: packets between same source-dest pair may take different paths

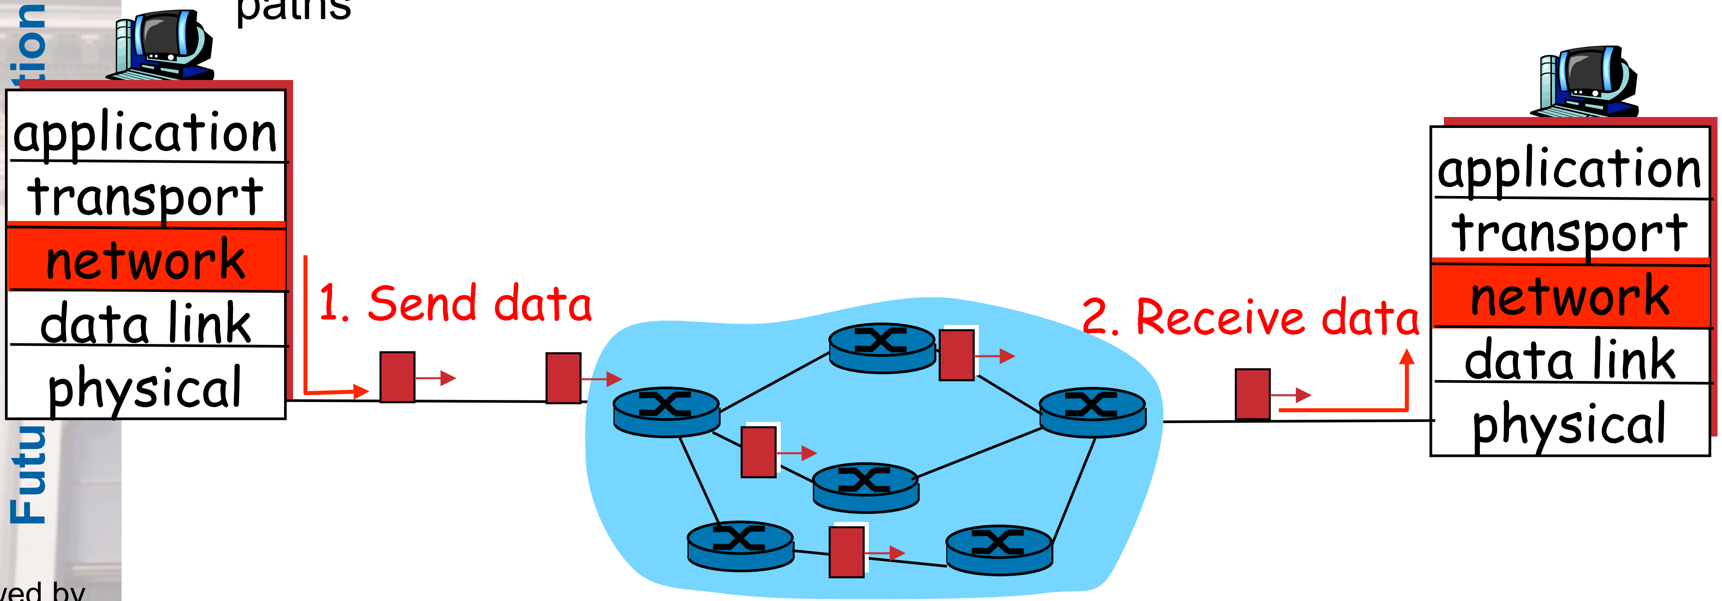

Connectionless Networks: Routing and Forwarding

Chair of Future Communication
endowed by A1

- *Routing*
  - Composes forwarding table
  - Distributed routing algorithms
  - Determines route taken by packets from source to dest
- *Forwarding*
  - Moves packets from router's input to appropriate output interface
  - Uses forwarding table
- *Analogy*
  - Routing: process of planning trip from source to dest
  - Forwarding: process of getting through single interchange

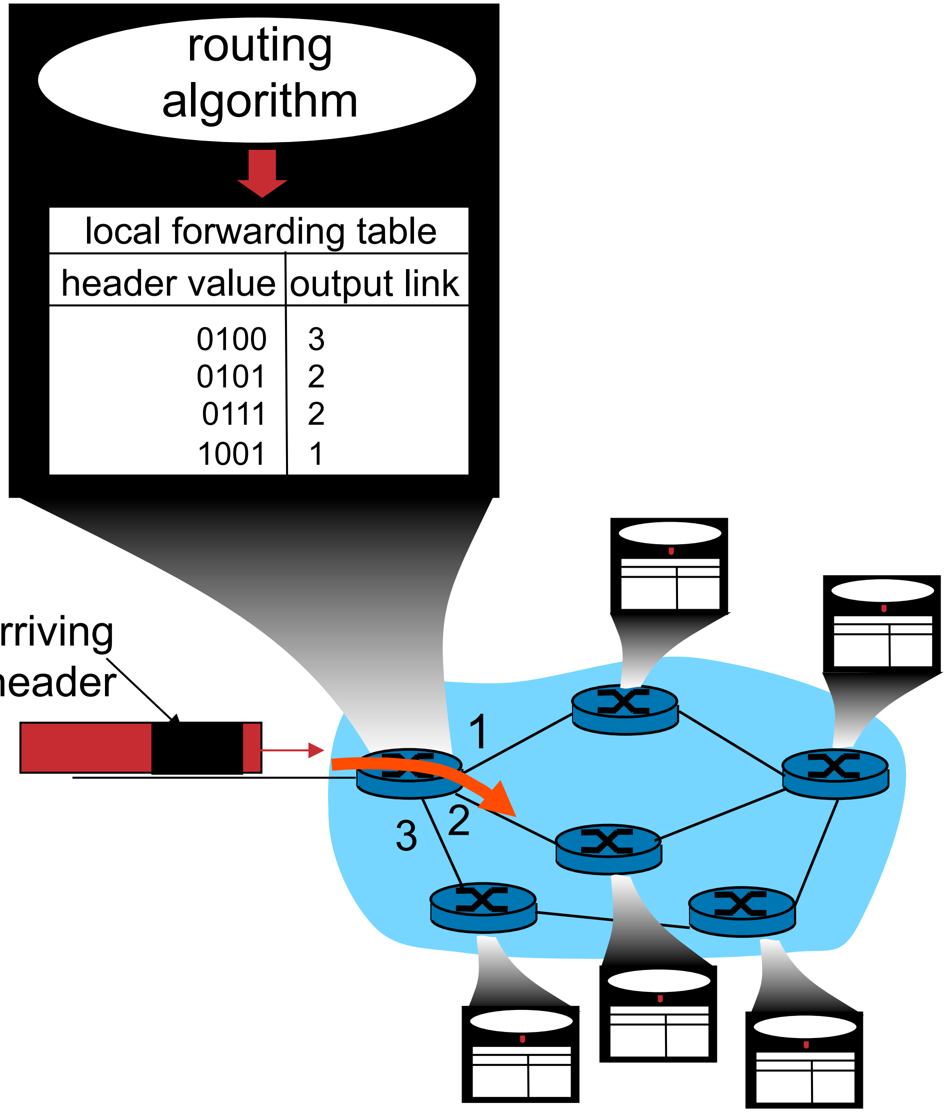

universität wien

Connectionless Networks: Longest Prefix Matching

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

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

4 billion possible entries in forwarding table

Examples:

DA: 11001000 00010111 00010110 10100001

DA: 11001000 00010111 00011000 10101010

Which interface?

Chair of Future Communication

endowed by A1

Connectionless Networks: Intra-AS Routing

- AS: autonomous system
- Also known as Interior Gateway Protocols (IGP)
- Most common Intra-AS routing protocols:
  - RIP
    - Routing Information Protocol
    - Distance vector routing protocol based on Bellman-Ford equation
  - OSPF
    - Open Shortest Path First
    - Link state routing protocol, Dijkstra's shortest path algorithm
  - IGRP
    - Interior Gateway Routing Protocol
    - Cisco proprietary
- IGPs follow usually the shortest paths with regard to a link cost metric
  - Hop count
  - Latency

universität wien

Connection-Oriented Networks (Virtual Circuits)

- Characteristics of a virtual circuit (VC)
  - Fixed path from source to destination
  - Packets belonging to VC carry a VC number
  - Forwarding tables along the path keep entry for each VC
- Connection setup: 3rd important function (next to routing and forwarding) in connection-oriented networks
- Signaling protocols for VC setup
  - Used to setup, maintain, and teardown VC
  - Used in ATM, frame-relay, X.25, in MPLS: RSVP-TE or LDP

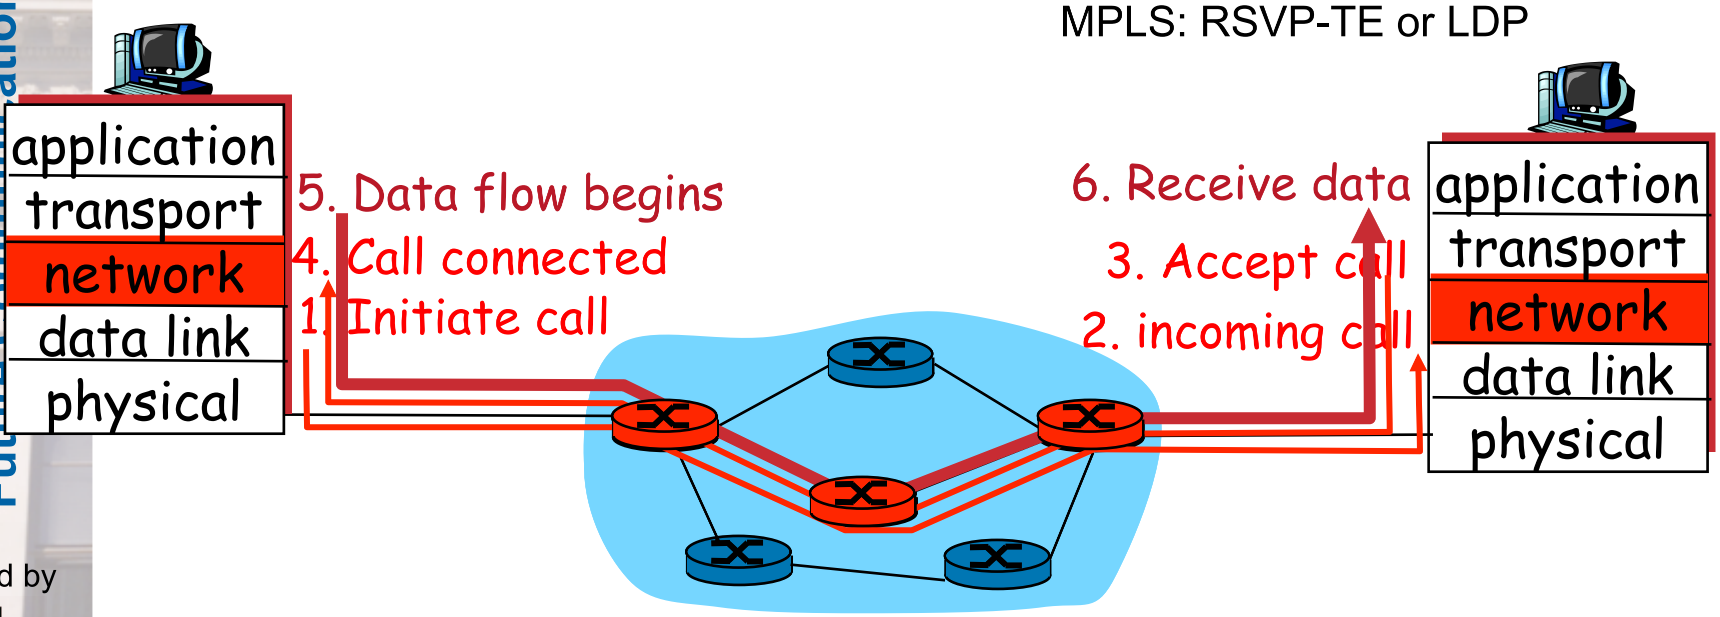

Connection-Oriented Networks: Forwarding Table

- Routers maintain connection state information
- VC number may change on each link

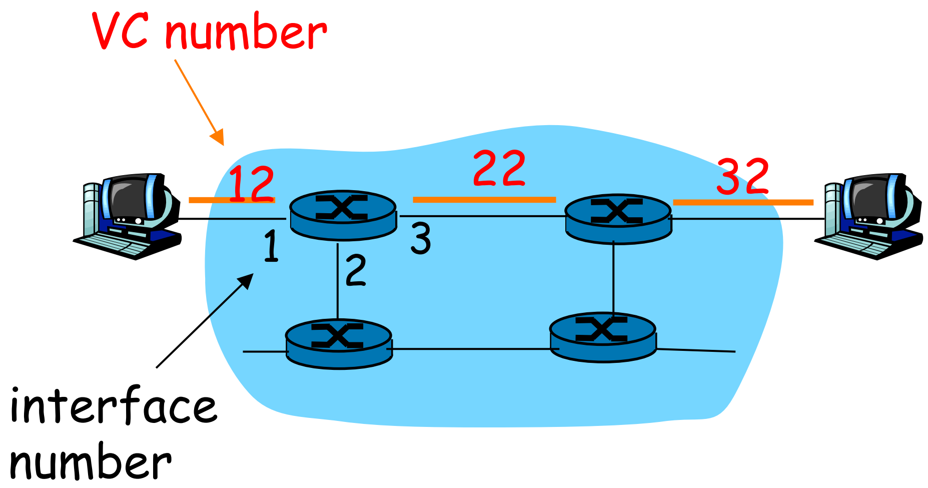

Forwarding table in northwest router:

Incoming interface	Incoming VC #	Outgoing interface	Outgoing VC #
1	12	3	22
2	63	1	18
3	7	2	17
1	97	3	87
…	…	…	…

Connectionless vs. Connection-Oriented Networks

- Connectionless networks
  - Routers are not flow-aware
  - Packets are routed solely based on destination adress
  - Example: IP datagrams
  - Simple operations
  - Difficult to add quality of service (QoS)

- Connection-oriented networks
  - Connections: setup, data transmission, teardown
  - Routers keep per connection state
  - Explicit paths
    - Deviation from shortest path routing possible
    - Example: label switched paths (LSPs) in MPLS
    - Used for traffic engineering
  - Easier support of QoS since flows are known

universität wien

Review of Networking Concepts

- <u>Overview:</u>
- Protocol layering and Internet protocol stack
- Circuit switching vs. packet switching
- Connectionless vs. connection-oriented networks, routing, forwarding, and switching
- Transport layer protocols
- Application layer
  - Sockets
  - Client-server and peer-to-peer communication
- Web services

TCP: Overview RFCs: 793, 1122, 1323, 2018, 2581

- Point-to-point
  - One sender, one receiver
- Flow control
  - Sender will not overwhelm receiver
- Congestion control
  - Sender reduces its rate in case of congested network
- Connection-oriented
  - Handshaking (exchange of control msgs) inits sender, receiver state before data exchange
- Pipelined
  - TCP congestion and flow control set window size
- Send & receive buffers
- MSS: maximum segment size
- Reliable, in-order byte stream
  - No "message boundaries"
- Full duplex data
  - Bi-directional data flow in same connection
- Does not provide: timing, minimum bandwidth guarantees

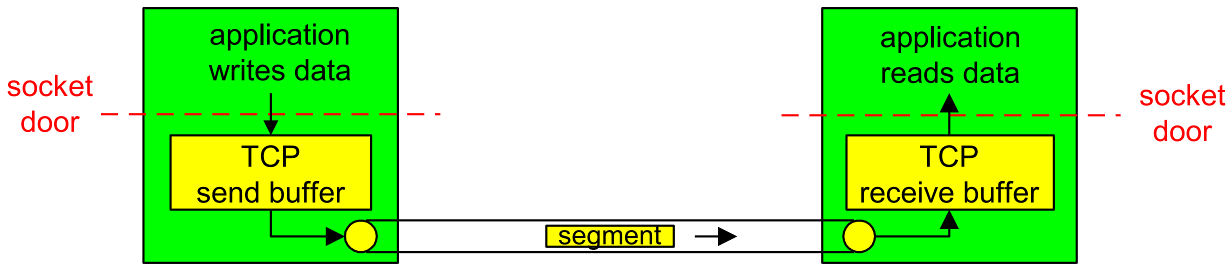

UDP: User Datagram Protocol [RFC 768]

- Very simple, connectionless transmission protocol
  - No handshaking between UDP sender, receiver
  - Each UDP segment handled independently of others
  - Multicast possible
- Unreliable data transfer between sending and receiving process
  - Packet loss
  - Packets delivered out of order to app
- Does not provide
  - Connection setup
  - Reliability
  - Flow control
  - Congestion control
  - Timing
  - Bandwidth guarantee
- Why is there a UDP?
  - No connection establishment
    - Fast transmission
  - Simple
    - No connection state at sender, receiver
    - Small segment header
  - No congestion control: UDP can blast away as fast as desired

What Transport Service is Needed?

Data loss

- some apps (e.g., audio) can tolerate some loss
- other apps (e.g., file transfer, telnet) require 100% reliable data transfer

Timing

- some apps (e.g., Internet telephony, interactive games) require low delay to be "effective"

Bandwidth

- some apps (e.g., multimedia) require minimum amount of bandwidth to be "effective"
- other apps ("elastic apps") make use of whatever bandwidth they get

Transport Service Requirements of Common Apps

Application	Data loss	Bandwidth	Time sensitive
file transfer	no loss	elastic	no
e-mail	no loss	elastic	no
Web documents	no loss	elastic	no
real-time audio/video	loss-tolerant	audio: 5kbps-1Mbps video:10kbps-5Mbps	yes, 100's msec
stored audio/video	loss-tolerant	same as above	yes, few secs
interactive games	loss-tolerant	few kbps up	yes, 100's msec
instant messaging	no loss	elastic	yes and no

universität wien

Internet apps: application, transport protocols

Application	Application layer protocol	Underlying transport protocol
e-mail	SMTP [RFC 2821]	TCP
remote terminal access	Telnet [RFC 854]	TCP
Web	HTTP [RFC 2616]	TCP
file transfer	FTP [RFC 959]	TCP
streaming multimedia	proprietary (e.g. RealNetworks)	TCP or UDP
Internet telephony	proprietary (e.g., Dialpad)	typically UDP

Chair of Future Communication

endowed by A1

Stream Control Transmission Protocol (SCTP, RFC2960)

- High level introduction in RFC3286
- Transport protocol
- Similarities with UDP
  - Message-orientation (no byte stream as in TCP)
- Similarities with TCP
  - Reliable, in-order delivery
  - Congestion control
- Multi-streaming
  - Transmission of several streams over a single SCTP connection, e.g. two images
- Uses per-stream sequence numbers for messages
  - If a packet of a specific stream is lost, only this stream suffers from retransmission (no head of the line blocking)
- Multihoming support
  - Both sender and receiver may have multiple IP addresses
  - Transparent failover if one of these addresses fails
- Path selection and monitoring

Datagram Congestion Control Protocol (DCCP, RFC4340)

universität wien

Chair of **Future Communication**

endowed by A1

- Transport protocol
- Similarities with UDP
  - Message-orientation
  - No reliable in-order delivery
- Similarities with TCP
  - Connection-orientation
  - Congestion control
- Makes use of ECN
- More appropriate than TCP for realtime data
  - No retransmissions for in-order delivery
  - Better timeliness for remaining data

Network Taxonomy

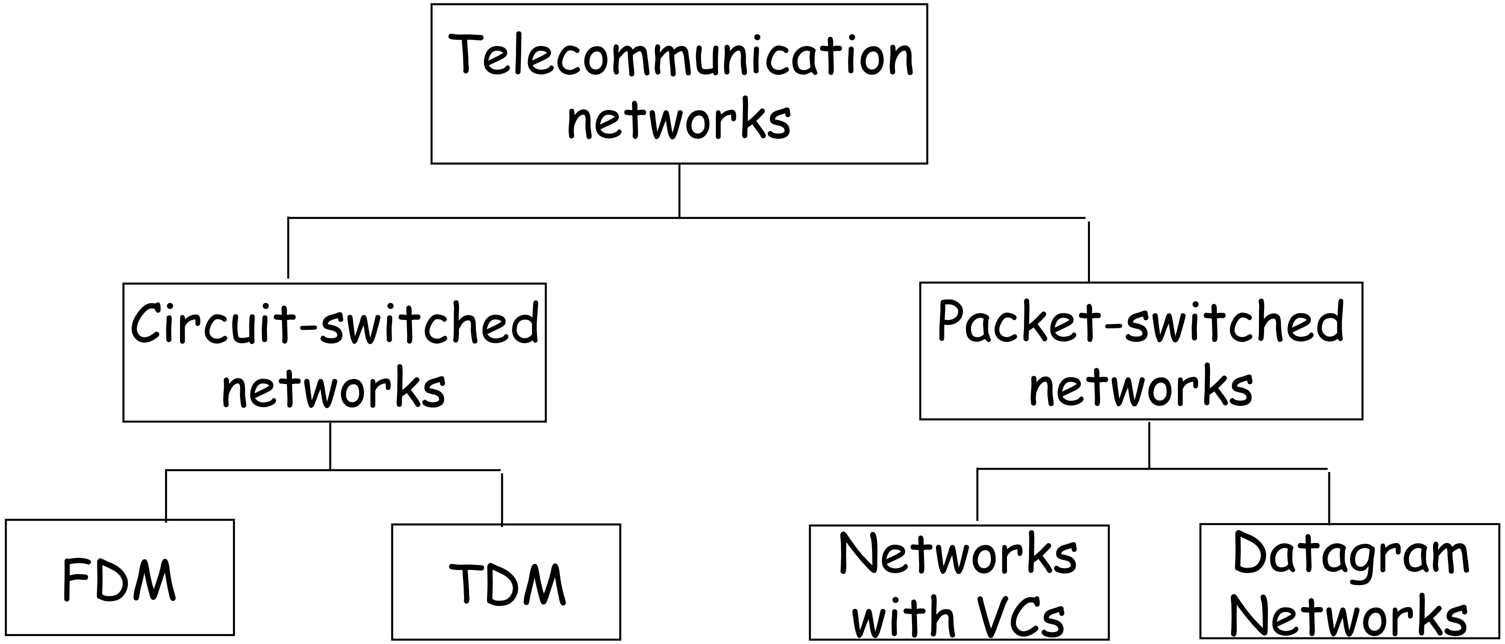

- Networks do *not* follow either the connection-oriented or connectionless principle
  - Internet provides both connection-oriented (TCP) and
  - Connectionless services (UDP) to apps
- Combinations on different layers possible
  - Packet-switched over circuit-switched
    - IP over optical
  - Circuit-switched over packet-switched
    - ISDN over MPLS

universität wien

Where Do Which Principles Apply?

Chair of **Future Communication**

endowed by A1

Internet

- Connectionless IP datagram forwarding
- Transport layer
  - Connectionless UDP
  - Connection-oriented TCP
- “Smart” end systems (computers)
  - Can adapt, perform control, error recovery
  - Simple inside network, complexity at “edge”
- Consequence of TCP
  - “Elastic” service
  - No strict timing
- Many link types
  - Different characteristics
  - Uniform service difficult

Asynchronous Transfer Mode (ATM)

- Connection-oriented ATM cell forwarding
  - Fixed size cells (48+5 bytes)
  - Virtual path connections (VPCc)
  - Virtual channel connections (VCCs)
  - Today used as link layer below IP
- Evolved from telephony
  - Human conversation
  - Guaranteed service needed
  - Strict timing, reliability requirements
- “Dumb” end systems (telephones)
  - complexity inside network

Multiprotocol Label Switching (MPLS)

- Connection-oriented packet forwarding
- Variable size packets
- Often used as link layer below IP
- Simple end-to-end measurements
- Hides network topology from traceroute