Karlsruher Institut für Technologie

4. Label Switching

Prof. Dr. Martina Zitterbart
Institut für Telematik

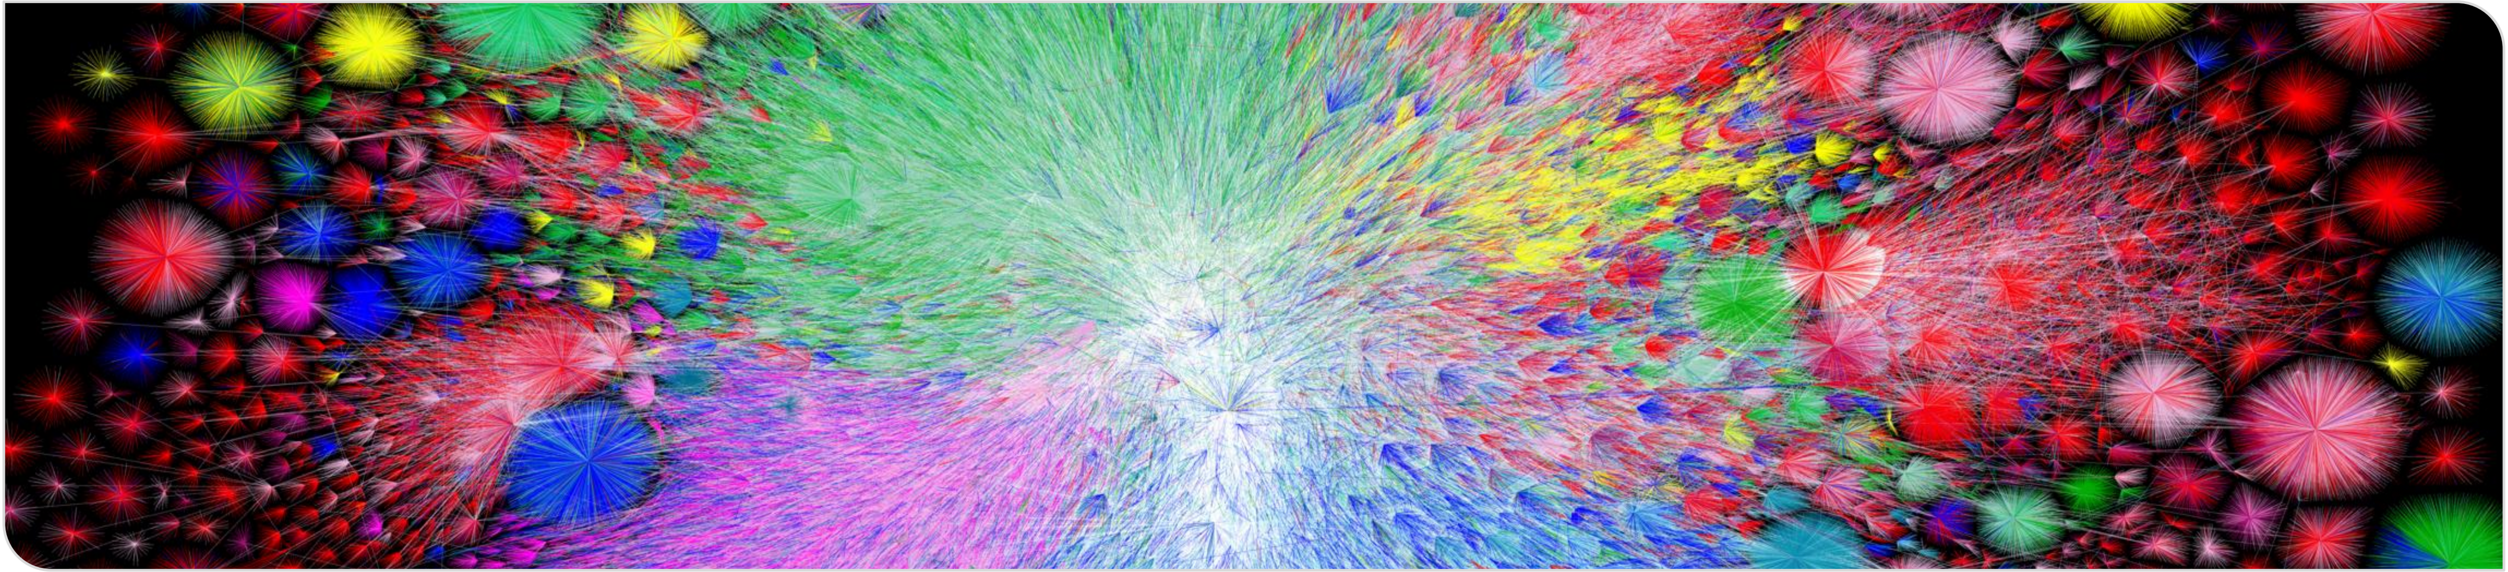

Label switching

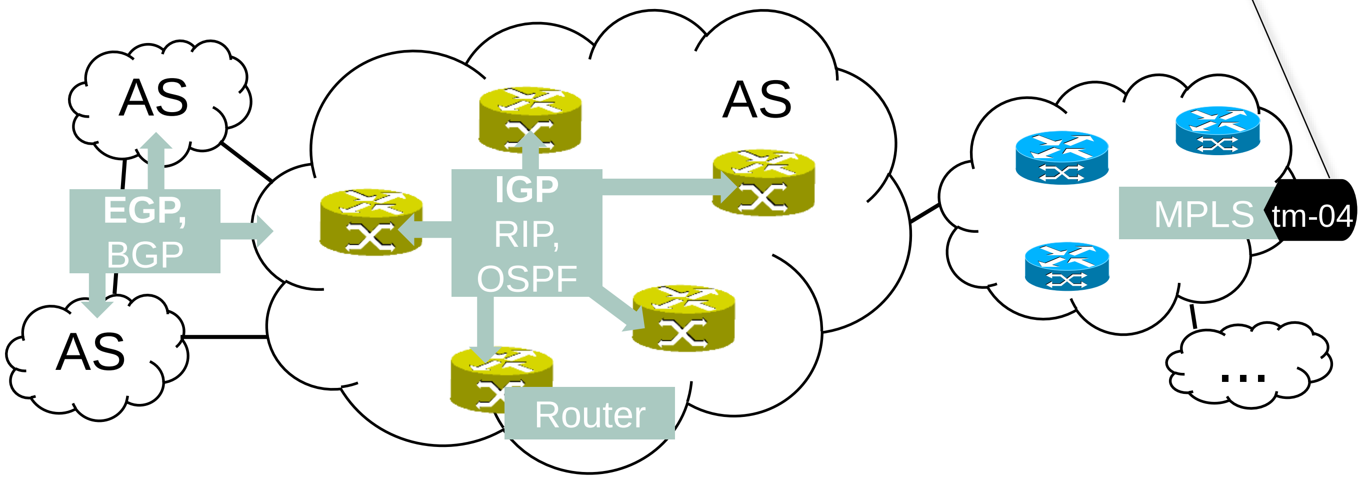

Karlsruher Institut für Technologie

4

Label Switching

4.1 Motivation

Observations

- Issues related to IP based routing as discussed previously
  - Lookup is rather complex
    - Longest matching prefix
      → high performance forwarding needed
  - Shortest path routing selects shortest path to destination
    - Multiple paths to destination can not be utilized concurrently
      → traffic engineering desirable
  - Strictly packet based
    - Each IP datagram is handled individually
      → no support for the concept of data streams (flows)

4.2 Flows

Flow

- Coarse grained definition
  - Packets that somehow belong together

- … a bit more technical
  - A flow is a sequence of packets traversing a network that share a set of header field values

- Different levels of granularity possible, e.g.,
  - All packets belonging to a particular TCP connection
  - HTTPS traffic
  - VoIP traffic
    - Of a particular sender
    - Within a network
  - …

[Stal17]

What is a „flow“?

- Examples

A particular TCP connection between Alice and Bob

VoIP traffic

Ethernet broadcasts

From a particular customer

Destination IP address in subnet X

In VLAN X

HTTPS traffic

TCP source port in the range X … Y

What is a „flow“?

- Combinations are also possible

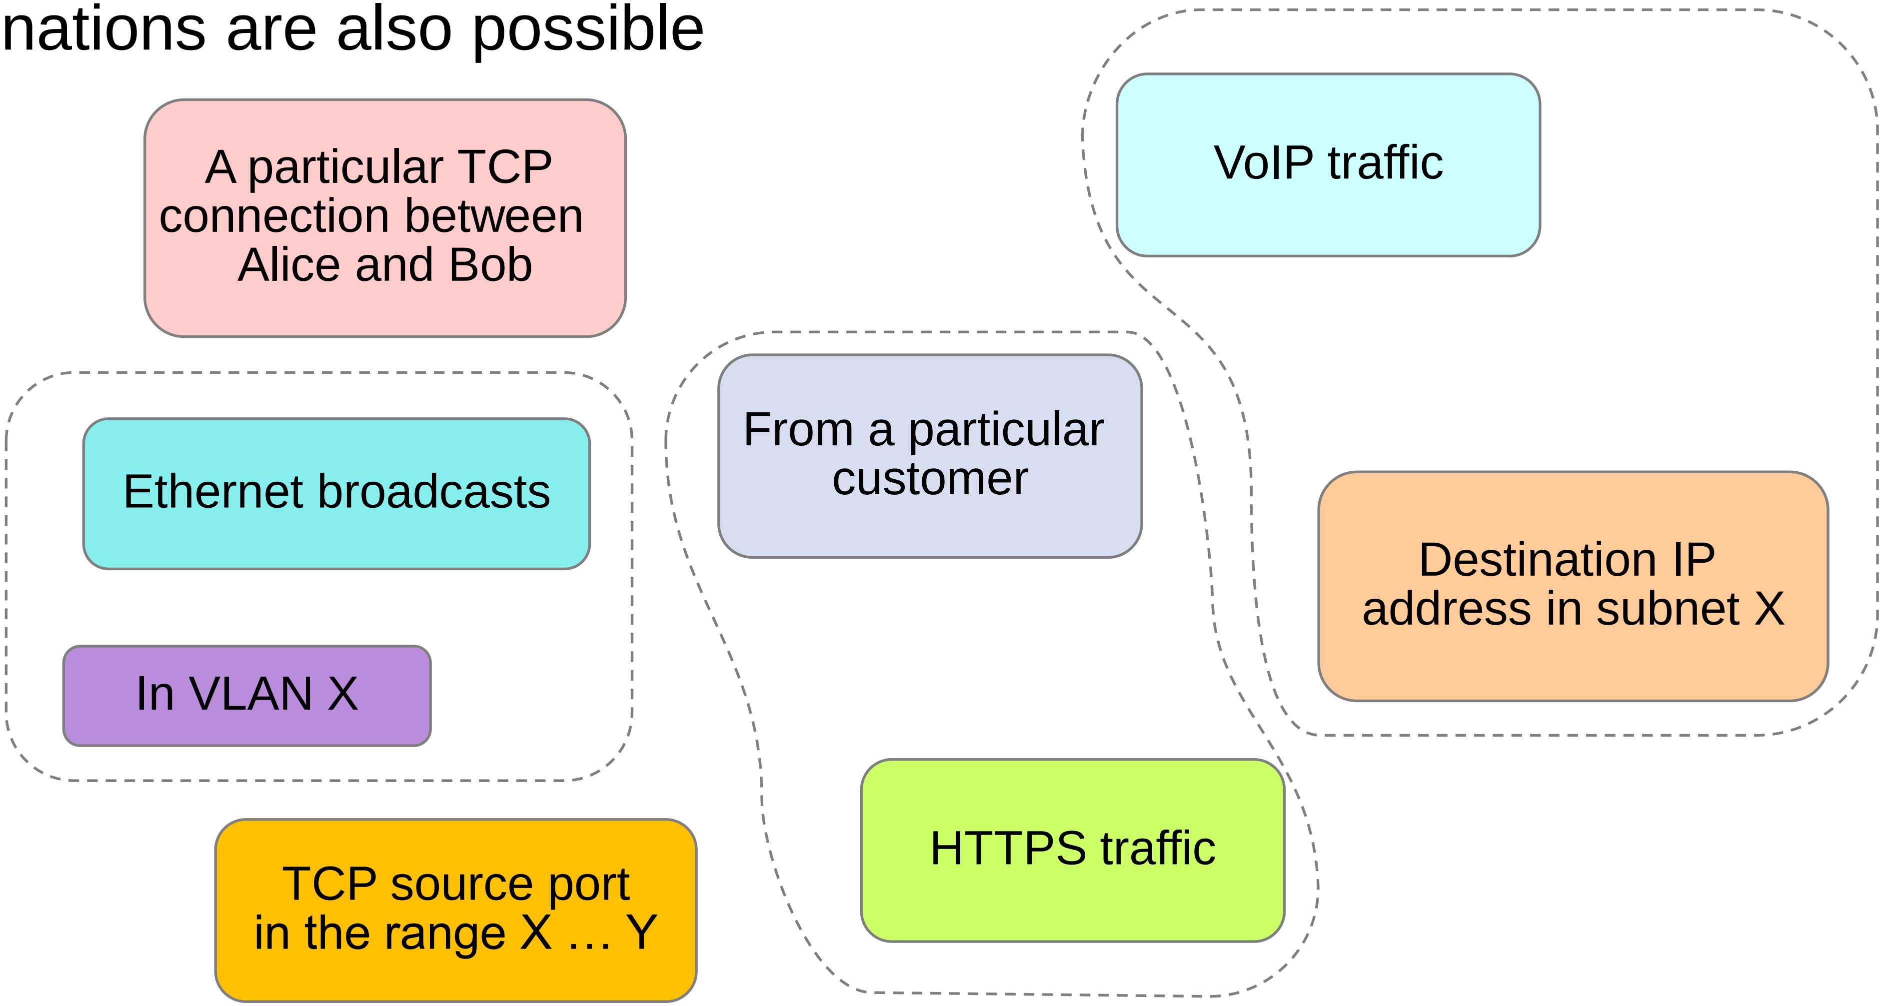

At which Layer are we?

- IP routing can be considered
  - special case of flow based forwarding
  - with quite coarse grained flows

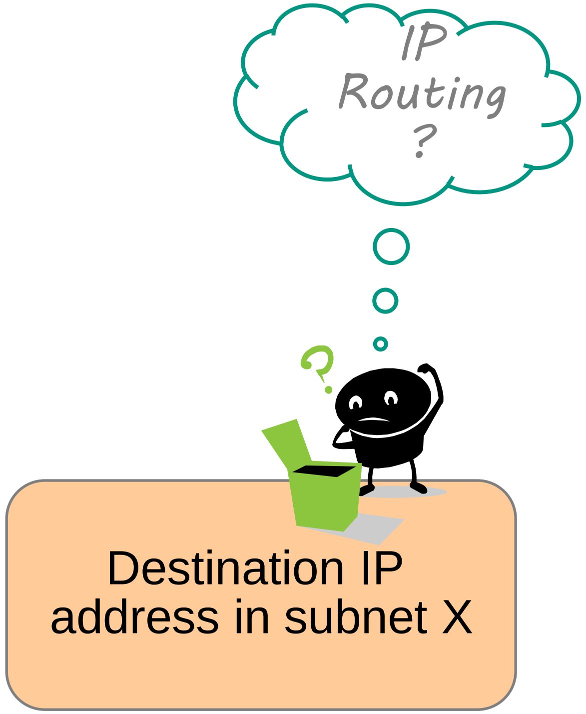

At which Layer are we?

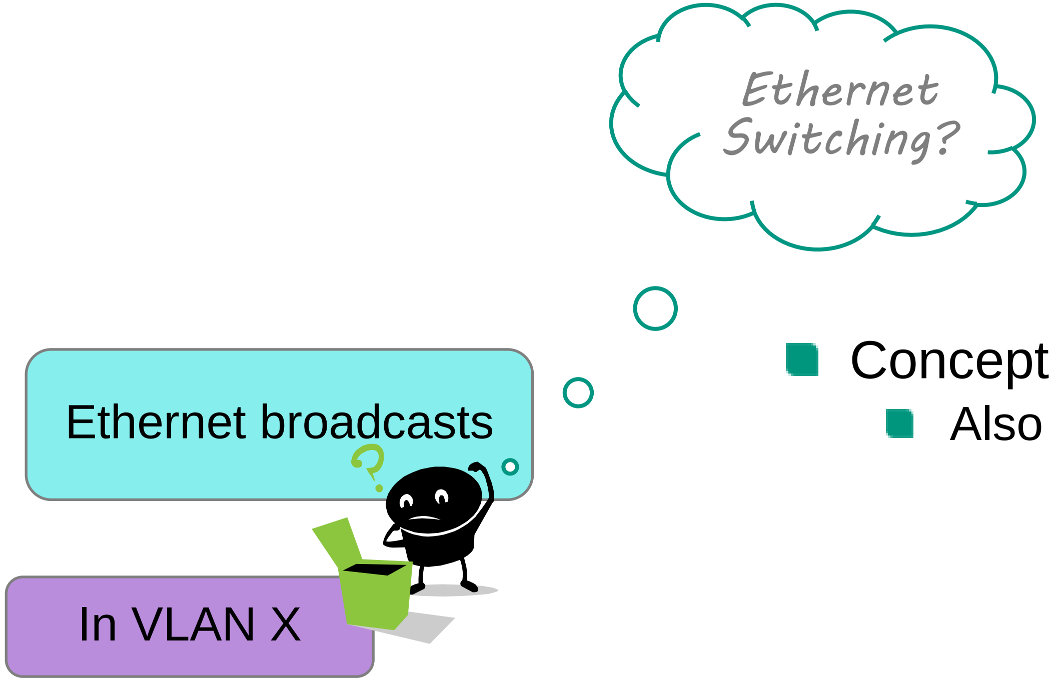

- Concept of flow based forwarding
  - Also applicable to Ethernet switching

At which Layer are we?

- Goes beyond classical routing / forwarding technologies

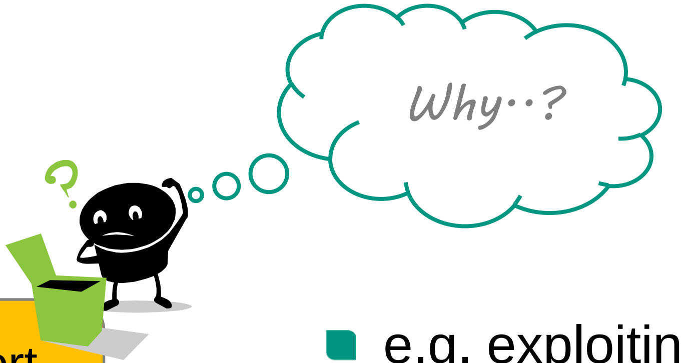

TCP source port
in the range X … Y

- e.g. exploiting the source port, for
  - Load balancing
  - Simple firewall functionalities

Flow Based Forwarding

- Fundamental concept, independent of certain layers
  - Can span multiple layers
- Incorporates classic routing/forwarding concepts
- Goes beyond classic concepts

Which Header Fields Need to be Evaluated?

Headers of layers 2 - 4

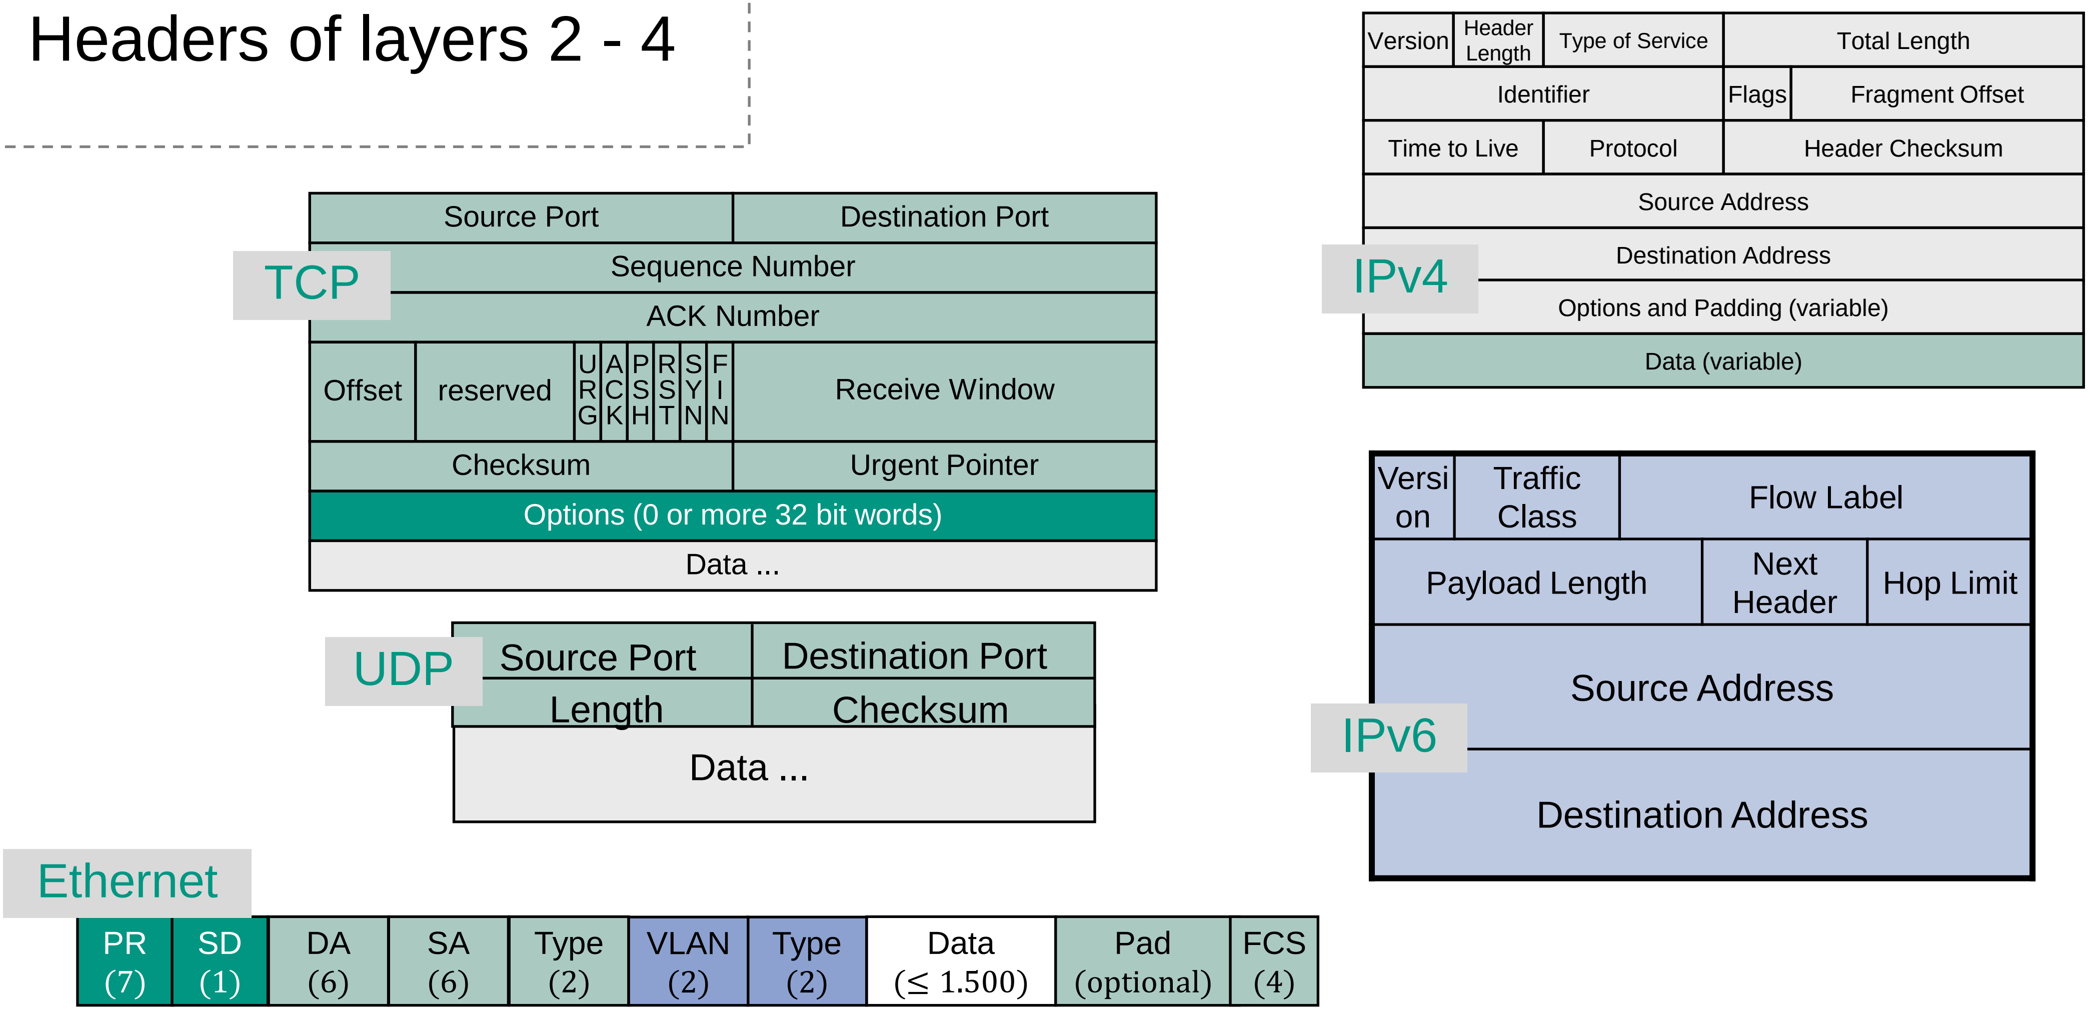

Which Header Fields Need to be Evaluated?

Karlsruher Institut für Technologie

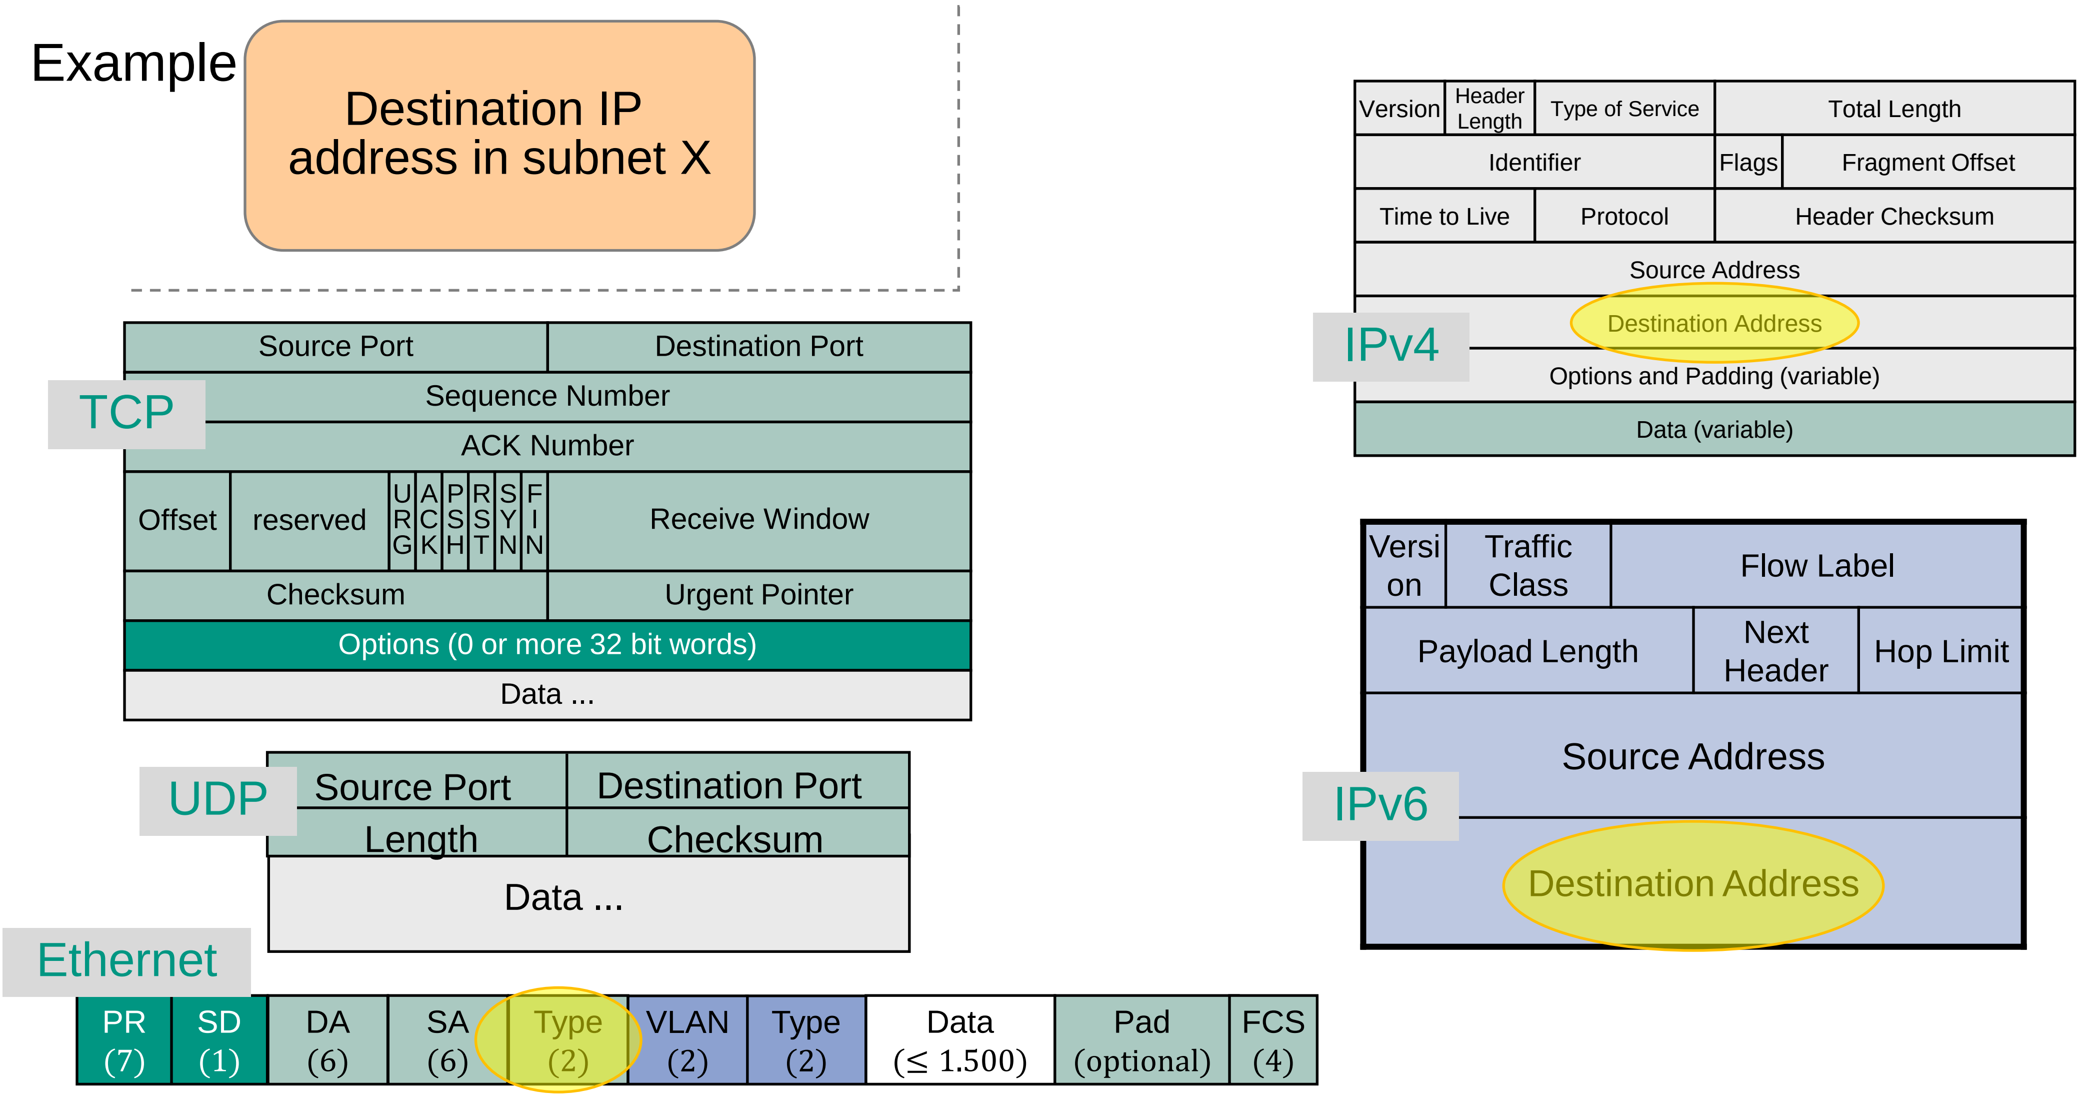

TeleMatics

Which Header Fields Need to be Evaluated?

Karlsruher Institut für Technologie

Example

A particular TCP connection between Alice and Bob

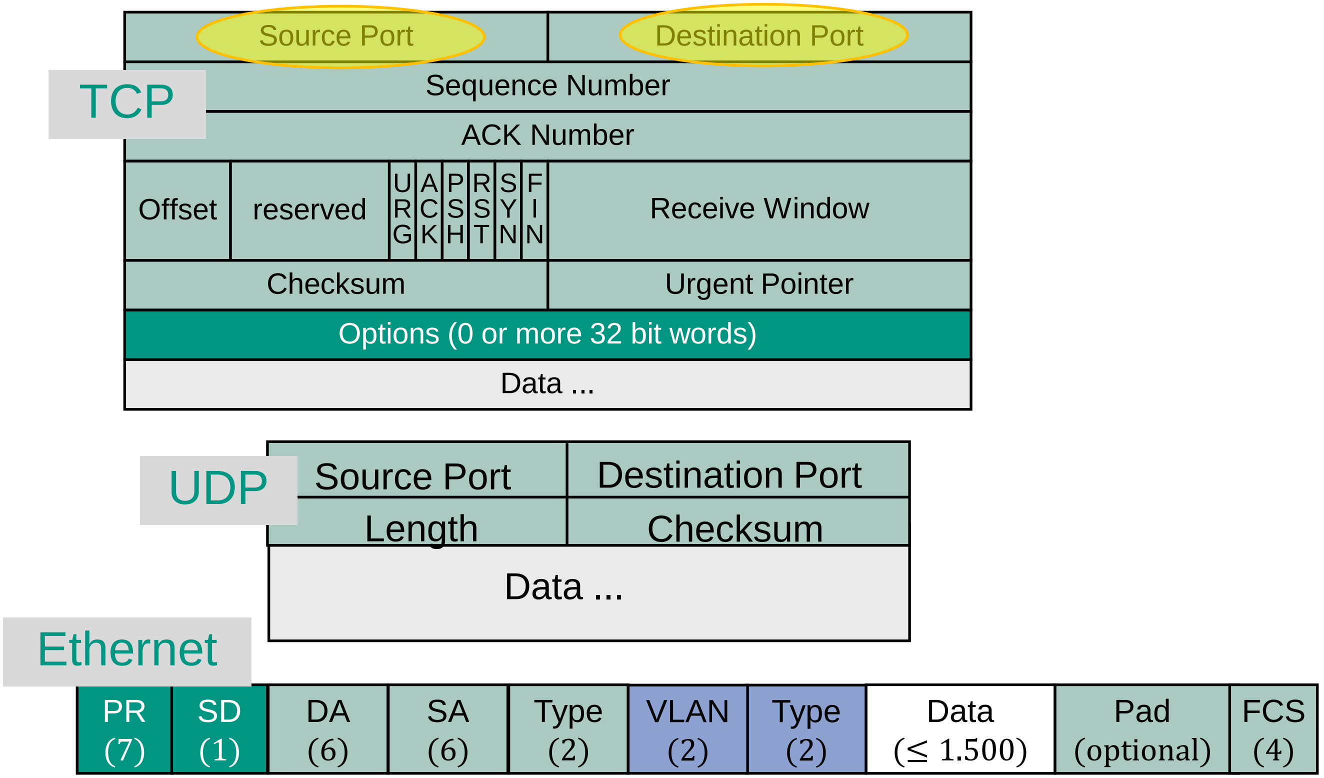

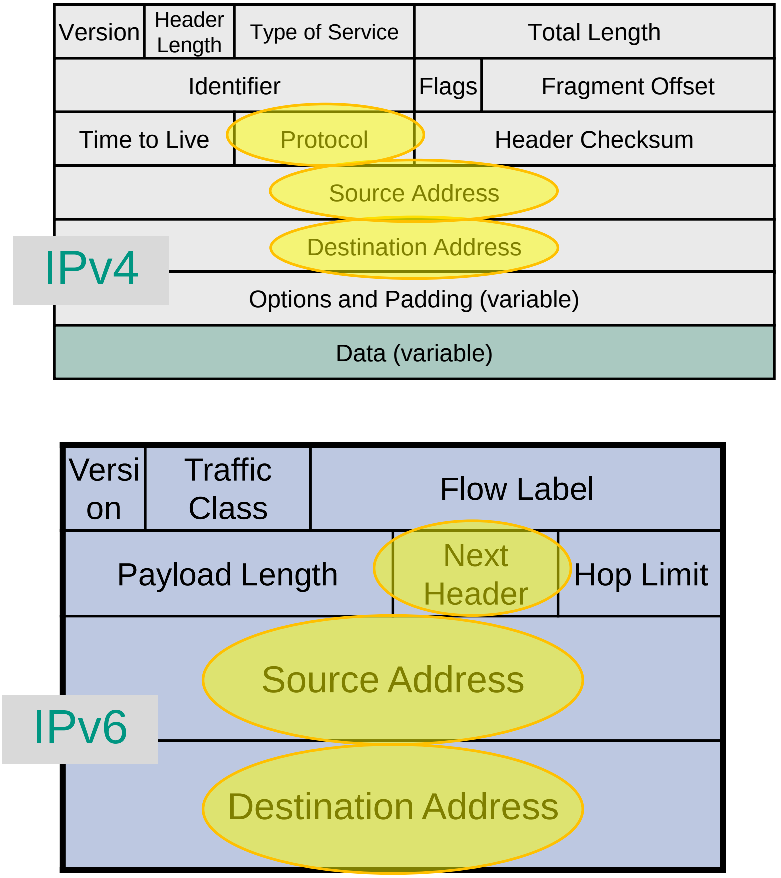

Aggregation

- Micro-flows
  - Consider a single "connection"
    - e.g., a TCP connection
  - Fine grained control
  - High number of flows possible

A particular TCP connection between Alice and Bob

- Macro-flows
  - Higher level of aggregation
  - Aggregation of several "connections"
    - e.g., IP destination address in specific subnet
  - Lower number of flows

Destination IP address in subnet X

4.3 Label Switching

Types of Communication Networks

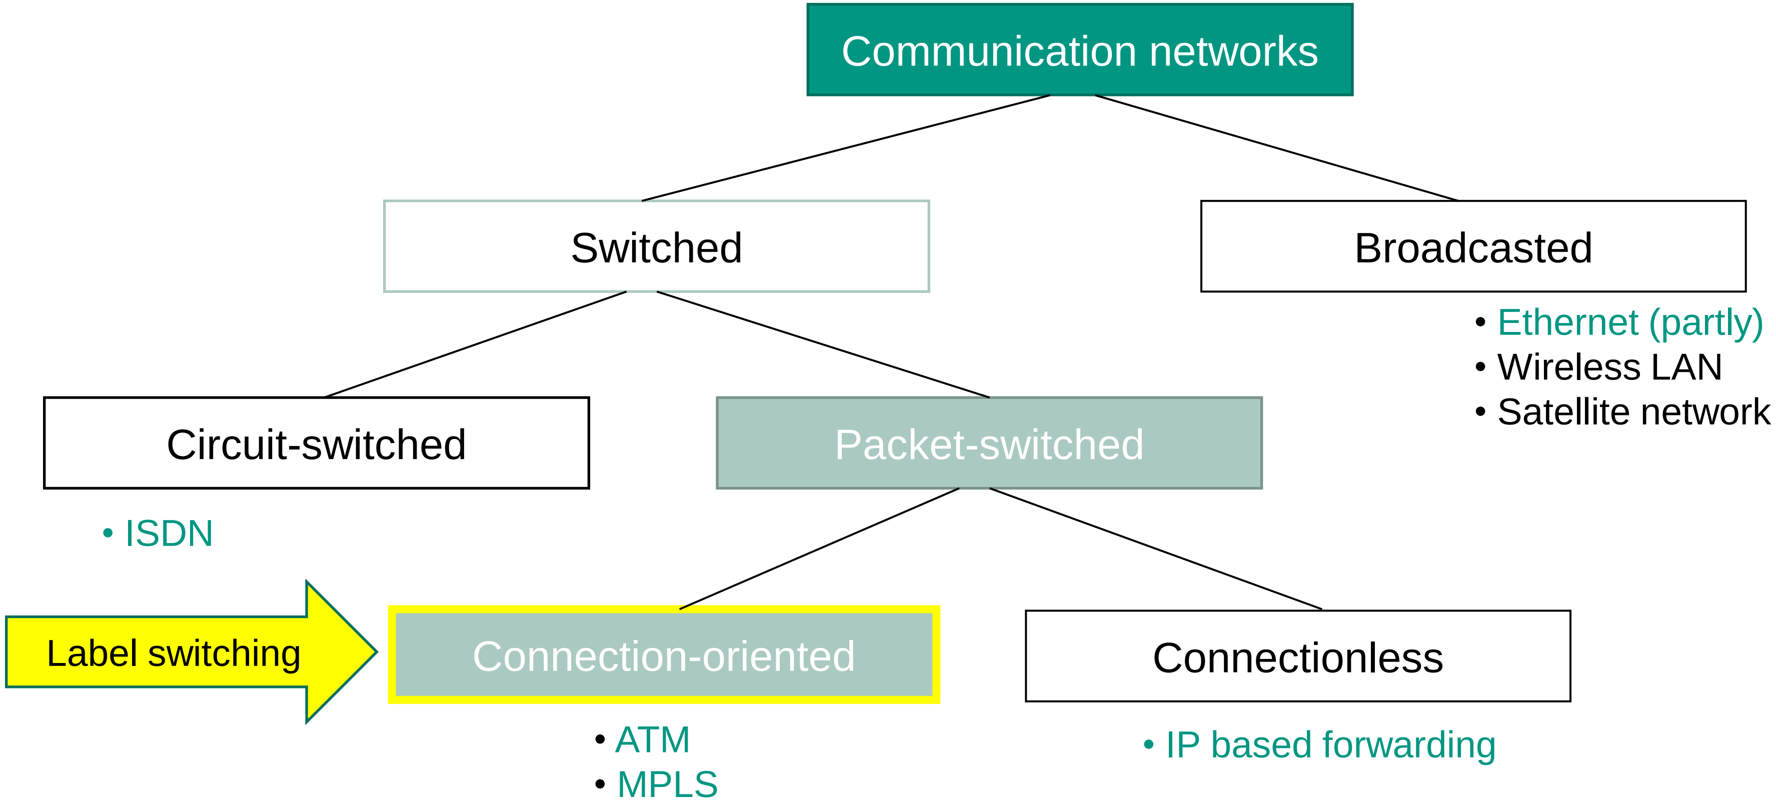

ATM: Asynchronous Transfer Mode
ISDN: Integrated Services Digital Network

Types of Communication Networks

Karlsruher Institut für Technologie

Type	Usage	Characteristics
Packet switching / Datagrams	IP	Packets are forwarded independent of each other, meta data required, expensive forwarding decision
Circuit switching	ISDN	Connection with fixed resource reservations, no meta data within the data stream required
Virtual circuits	MPLS	Connections without fixed resource reservation, packets are efficiently forwarded on the same path, inexpensive label-based forwarding decision

Label Switching

- Combination of
  - Packet switching
    - Packets are forwarded individually
    - Packets include metadata needed for forwarding decision

  - Circuit switching
    - Paths established for flows through the network
    - Simple forwarding decision
    - Differentiation of flows possible
      - Load balancing
      - Quality of service (QoS)

Label Switching

- Implementation
  - Switching at layer 2
    - Instead of routing at layer 3
  - Labels
    - Identification which is only locally valid
  - Virtual circuits
    - Sequence of labels

Label

- Short unstructured identification of fixed length
  - Does not carry any layer-3-information
  - Unique: only locally at the corresponding switch
  - Label swapping
    - Mapping from input label to output label
- Virtual circuit
  - Identified through sequence of labels on the path

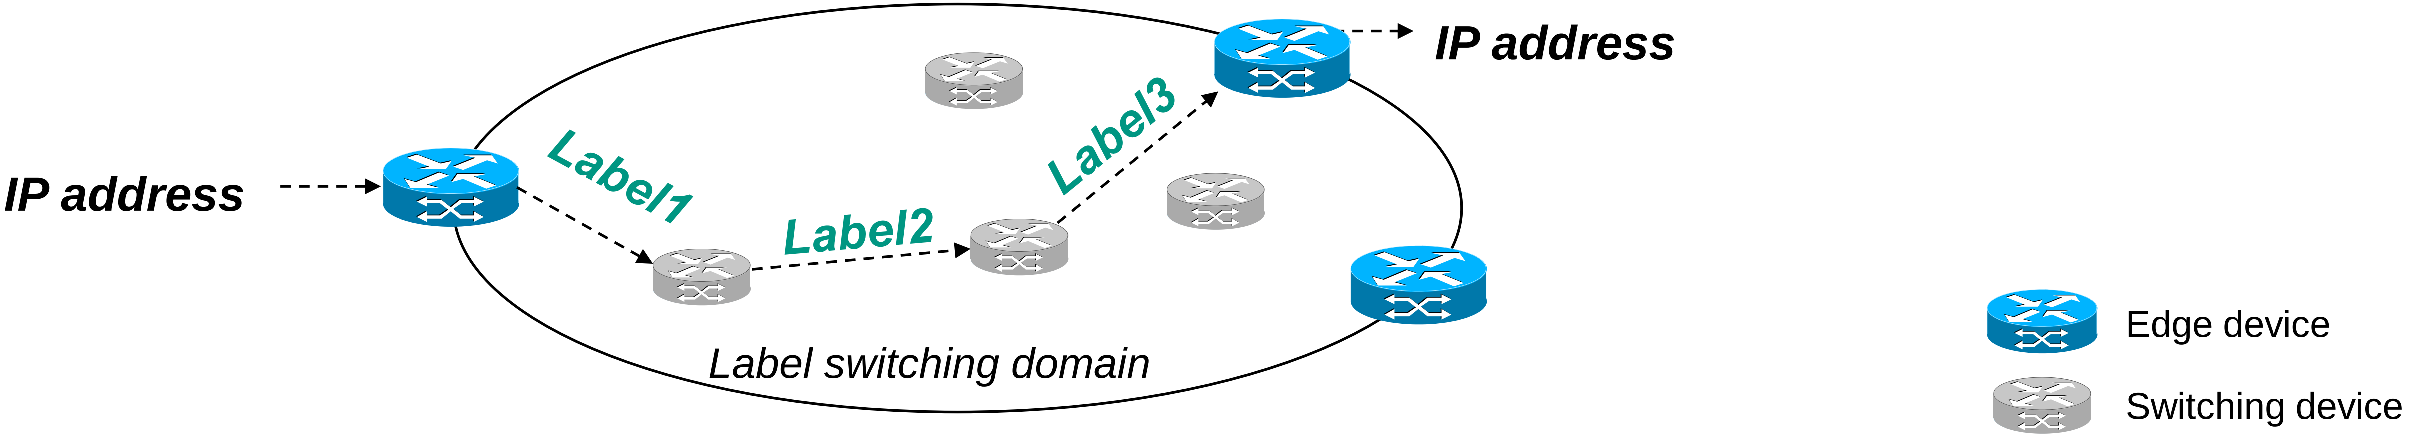

Transport of Label

- Label must be transported within the packet
  - Additional „header“ in the packet
    - Between headers of layer 2 and layer 3

      → *layer 2,5*

  - Alternative
    - In specialized fields within existing packet headers
      - IPv6: flow label
        - 20 bit field in IPv6 header, to identify micro flows more easily

Label Switching Domain

- Basic architecture
  - Border of the domain (edge devices)
    - Add / remove label
    - Map flow to forwarding class
    - Access control
    - …
  - Within the domain (switching device)
    - Forward packets based on label information
    - Label swapping

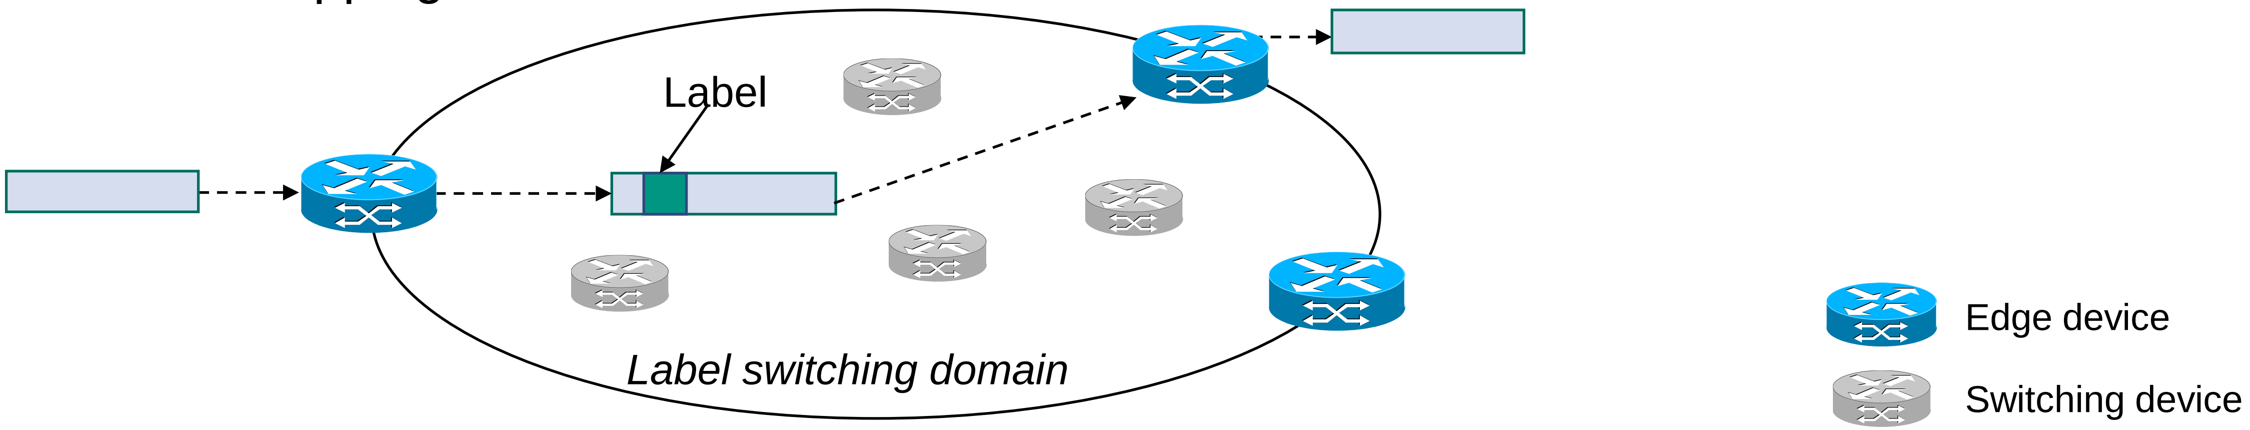

Label Forwarding Information Base

- Forwarding table in case of label switching
  - Efficient access through label (no longest prefix matching needed)

Switch	Label in	Label out	Next hop	Interface out
A	–	62	B	0
B	62	15	D	0
C	–	15	D	2
D	15	60	E	2
E	60	–	–	0

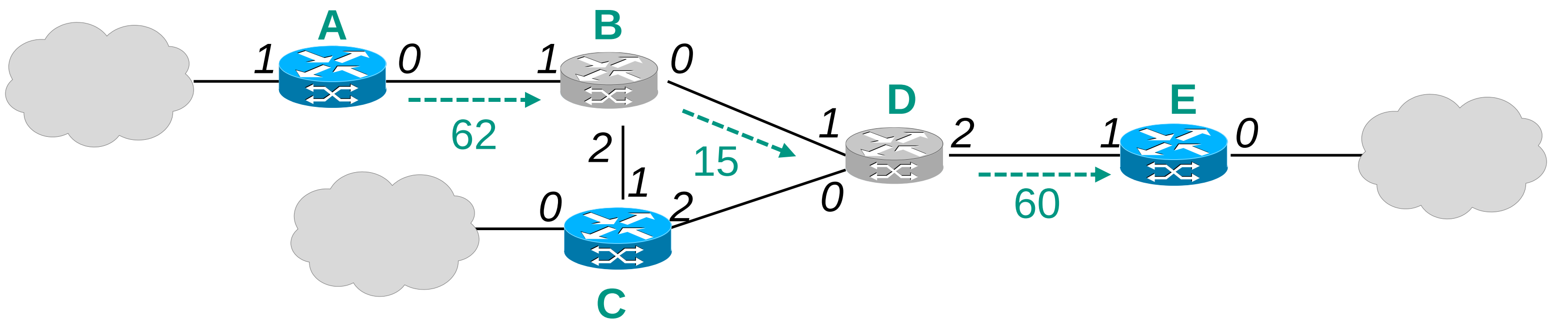

4.4 MPLS

4.4.1 General Aspects

Multiprotocol Label Switching (MPLS)

Karlsruher Institut für Technologie

- Based on label switching
- Originally: data plane optimization
  - Still requires control plane
  - Allows more sophisticated control plane techniques (e.g., traffic engineering)
- Standardized within the IETF
  - Key specification: RFC 3031
    - … and many additional RFCs
- Increasingly applied in larger autonomous systems
  - Adapted much quicker than IPv6

[RFC3031]

Main Features

- Fast forwarding due to reduced amount of packet processing
- QoS support
  - Guarantees on latency and capacity, e.g., for voice traffic
  - …
- Traffic engineering
  - Supports load balancing in order to optimize network utilization
  - …
- Virtual private networks
  - Isolate traffic from other packets on the Internet
  - …
- Multiple networks support
  - Usable on different network technologies, e.g., IP, ATM …

Rapid Acceptance of MPLS

- Reasons for the growth of MPLS
  - Connection-oriented communication on top of IP
  - Clear separation of forwarding (label switching) and control (manipulation of label binding)
    - Data plane – control plane
  - Not limited to IP
    - Ethernet, SONET, ATM …
  - Support of metrics
  - Versatile concept
    - Original goal: speed-up lookup of next hop
    - Applied over time to new applications: traffic engineering, virtual private networks (layer 3, layer 2) …
  - Scales
    - Applied in many global networks

4.4.2 Architecture, Components and Basic Operation

Architecture of an MPLS Network

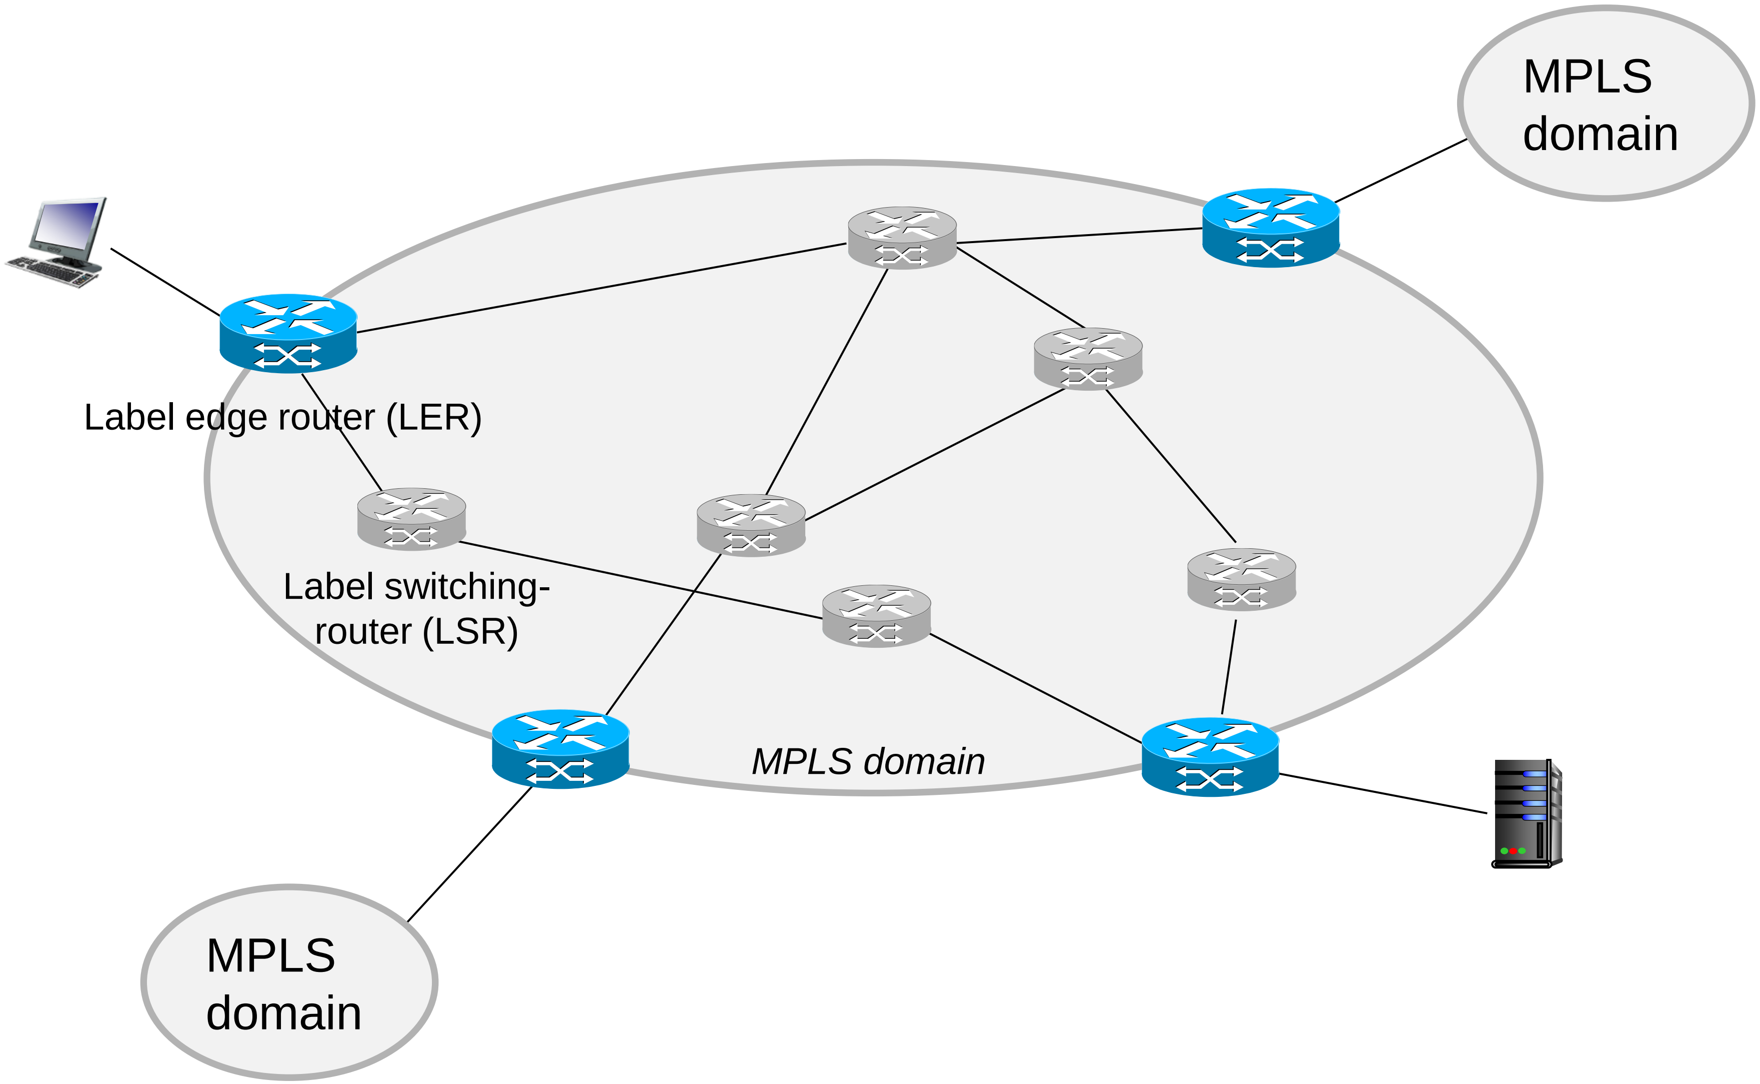

Components

- Label-switching router (LSR)
  - MPLS-capable IP router
    - Can forward packets based on both, IP prefixes and MPLS labels
    - Typically: IP for control plane and MPLS for data plane
- Label edge router (LER)
  - Router at the edge of an MPLS domain
    - Each LSR with a non-MPLS capable neighbor is an LER
    - Also called: label ingress router resp. label egress router
  - Classifies packets that enter the MPLS domain
    - Forwarding equivalency class (FEC)
- MPLS-Node
  - General term for MPLS-capable intermediate systems, like LSRs

Architecture of a Label-switching Router

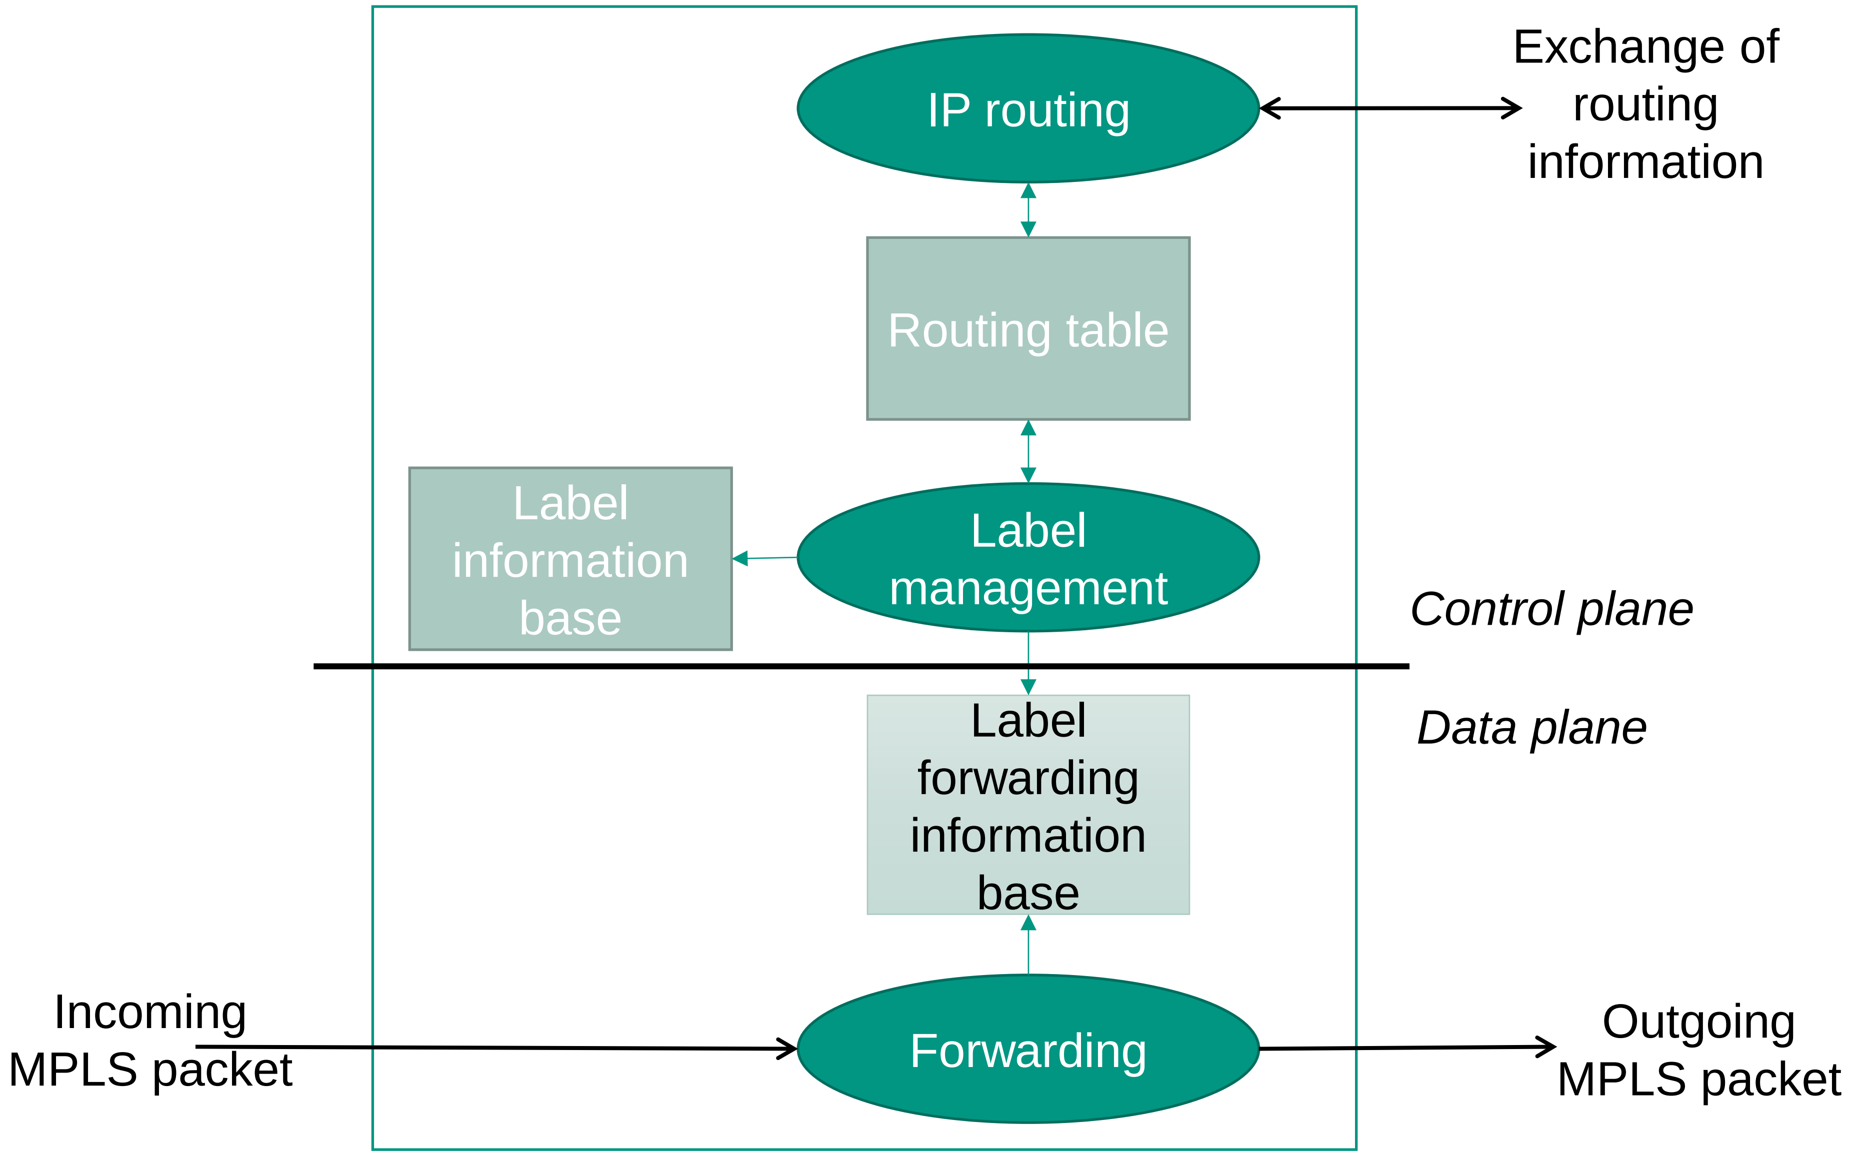

Forwarding Equivalence Class

- Forwarding equivalence class (FEC)
  - Class of packets that should be treated equally
    - Same path through the network
    - Same QoS properties
  - Basis for label assignment
  - MPLS-specific term, roughly comparable to „flow“

- Example for forwarding equivalence classes
  - Same address prefix and same type-of-service field
  - Same IP addresses and same port numbers
  - VoIP traffic with destination address in subnet X

- Granularity of forwarding equivalence classes
  - Coarse-grained: Important for quick forwarding and scalability
  - Fine-grained: Important for differentiated treatment of packets or flows
  - see „macro-flow“, „micro-flow“

Forwarding Equivalence Class

- Example
  - Very fine granular FEC (“micro flow”)
    - A single TCP connection, identified by 5-tuple

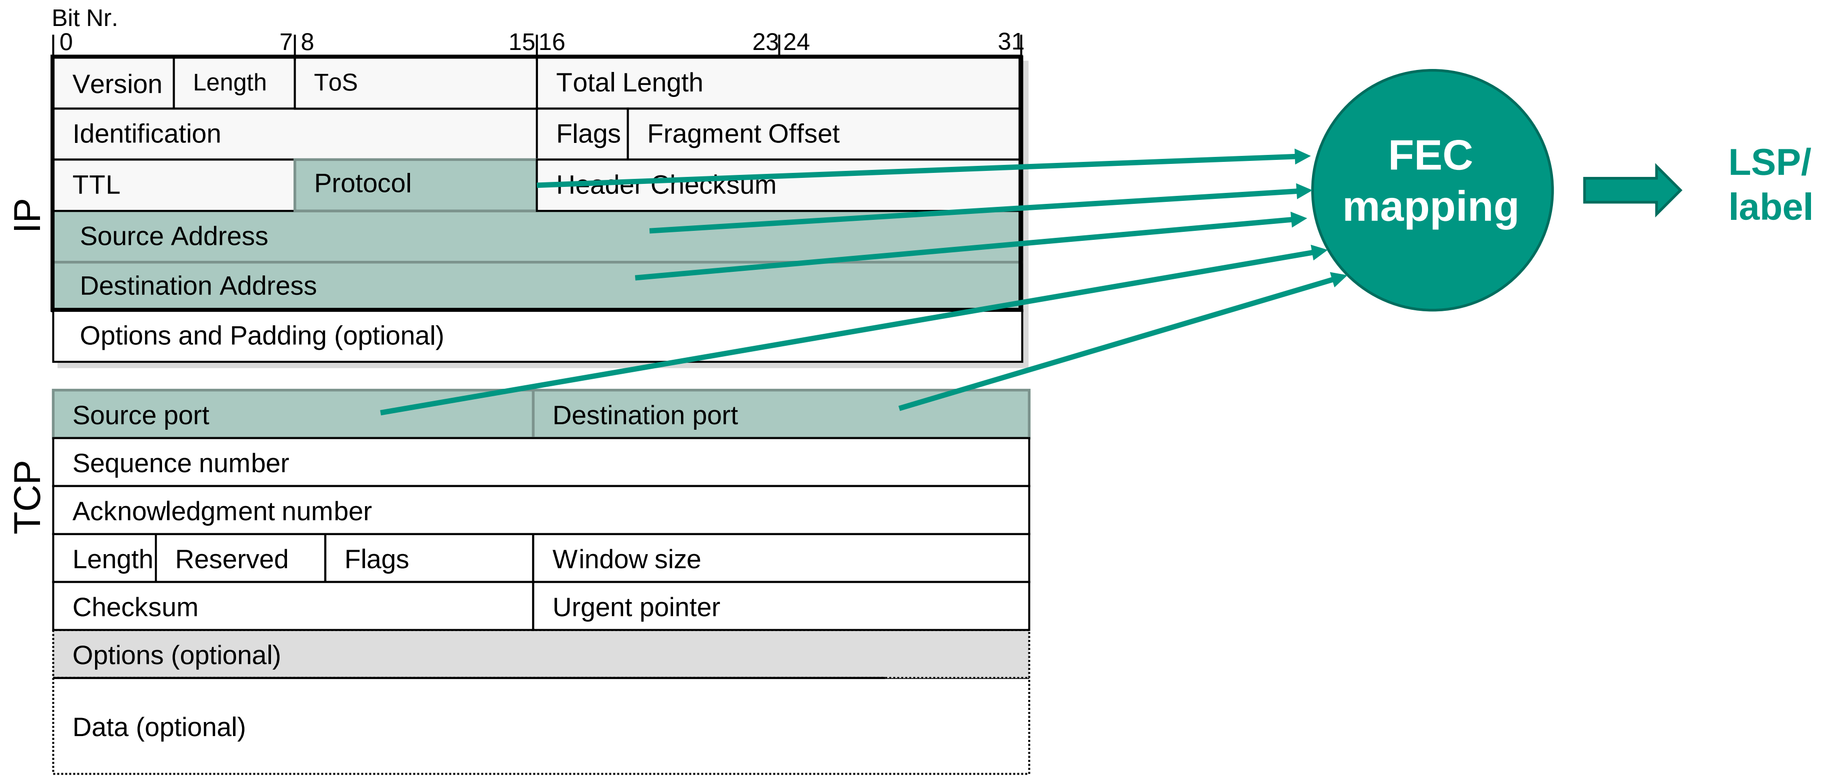

Forwarding Equivalence Classes: Example

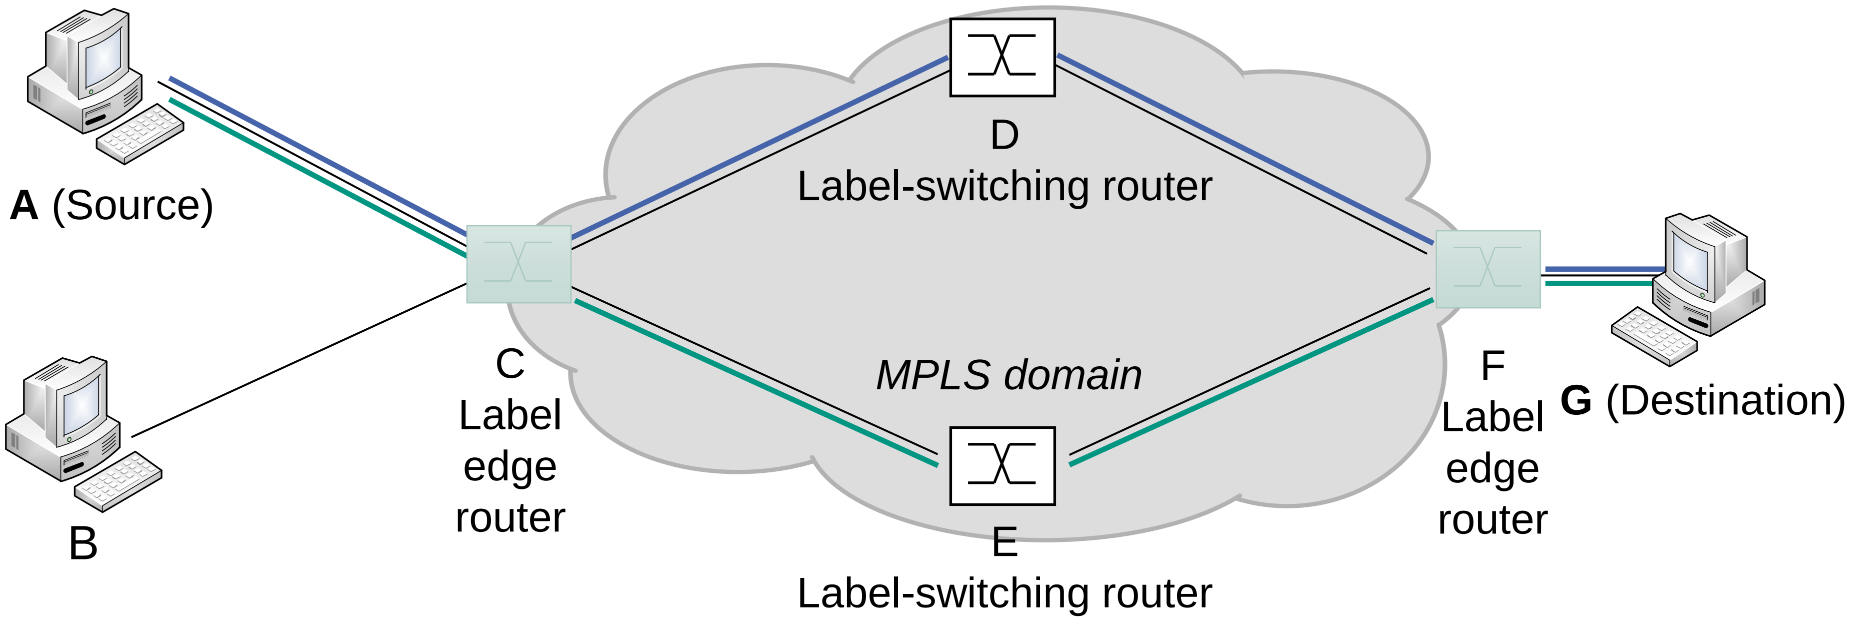

Source A sends two data streams (high priority and low priority traffic) towards destination G. Possible paths

- Path 1: A → C → *D* → F → G
- Path 2: A → C → *E* → F → G

Label edge router assigns the data streams to different forwarding equivalence classes → *differentiation*

Differentiation of Data Streams

- Traffic engineering
  - Usage of different paths
    - Not possible (to the same extent) with IP routing
  - Goals
    - Load balancing
    - Utilization of all available resources
    - Prioritization of individual data streams

  … realized through separate virtual connections

- Support of quality of service
  - Different quality of service for different data streams

Label Switched Path

- Virtual connection
  - Sequence of labels on a path through MPLS domain

Interface in	Label in	Interface out	Label out
3	50	1	40

Interface in	Label in	Interface out
3	40	1

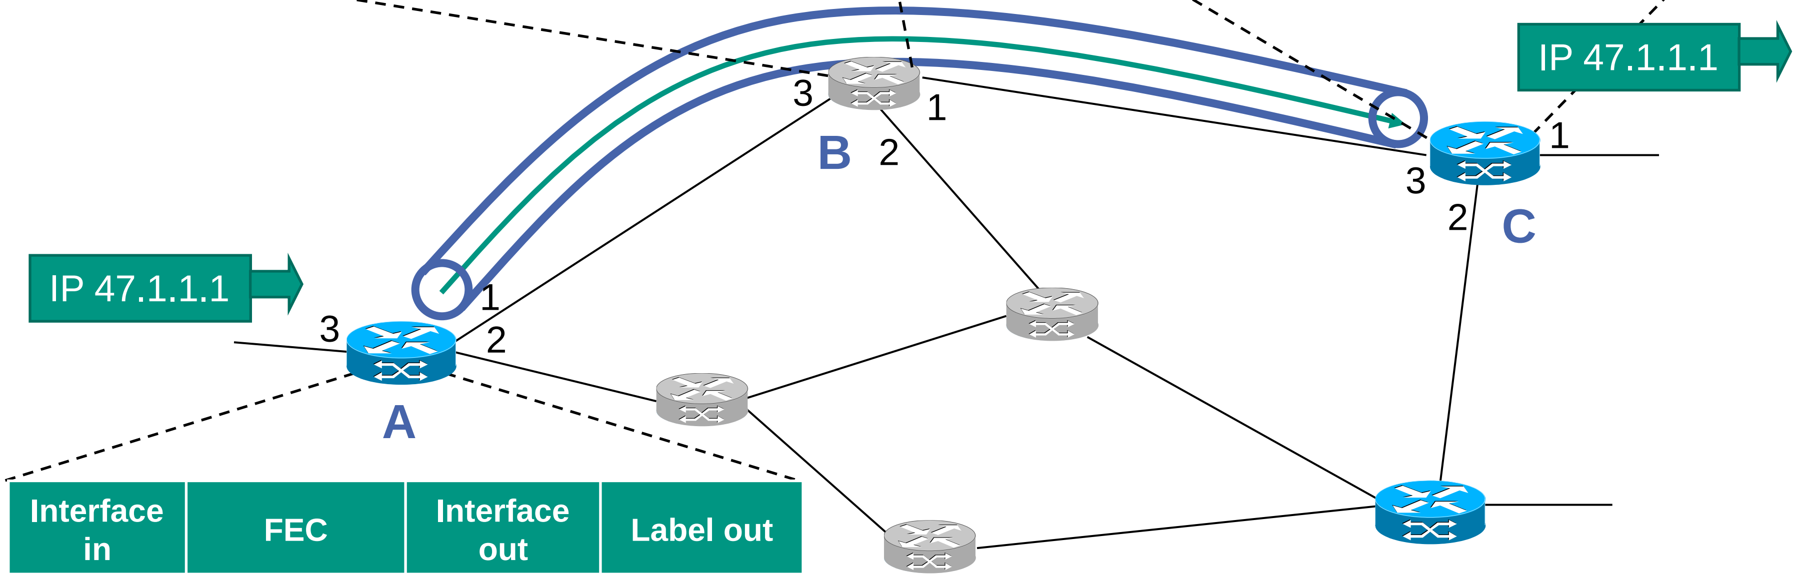

Interface in	FEC	Interface out	Label out
3	xyz	1	50

MPLS-Label

- Encapsulation
  - Between headers of layer 2 and layer 3

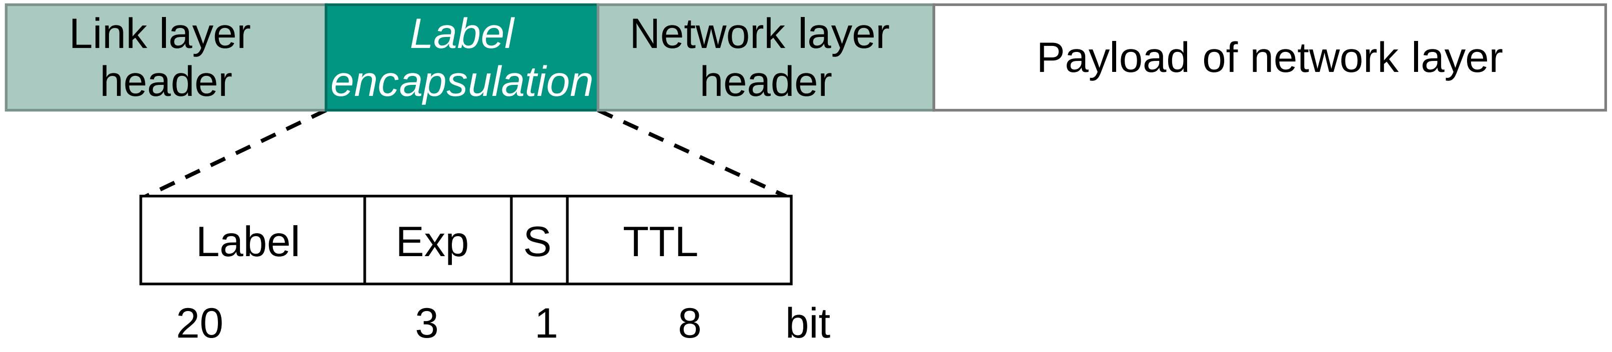

- Structure
  - Label: The label itself
  - Exp: Bits for experimental usage
  - S: Stack-bit
  - TTL: Time-to-live

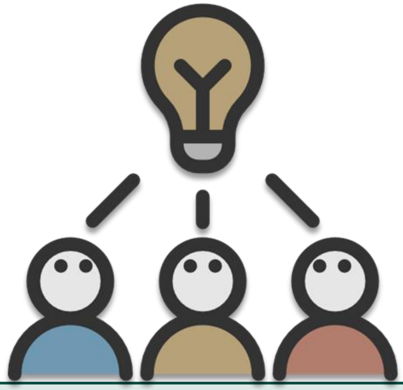

Task 04-1

Label Switching Path

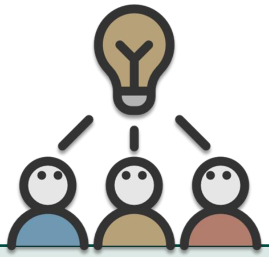

Task 04-1

4.4.3 Label Distribution

Label Binding and Distribution

- Label binding
  - Associate specific label to FEC
  - Stored in *label forwarding information base*
    - Used as incoming label

- Label distribution
  - Label binding is distributed to neighboring routers
  - Stored in *label forwarding information base*
    - Used as outgoing label

Types of Label Distribution

- "Roles" of a label-switching router
  - Downstream LSR
    - In direction of data flow
  - Upstream LSR
    - Against direction of data flow

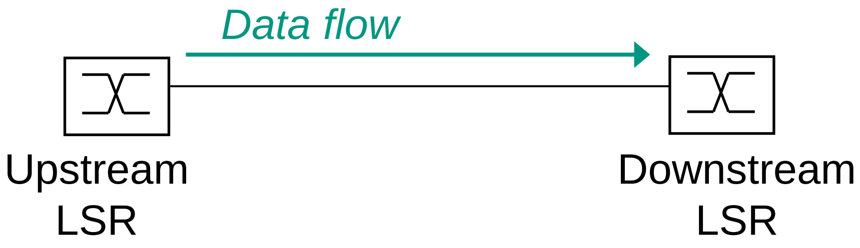

- Unsolicited downstream
  - Router generates label bindings as soon as it is ready to forward MPLS packets of the respective FEC
    - Upstream neighbors (according to IP routing) update forwarding tables
      - Label used as outgoing label
    - Non-upstream neighbors can store label for later use
      - Quicker reactions on route changes

- Downstream on demand
  - Downstream router generates label binding on demand
  - Upstream router has to request label binding for FEC

Label Distribution Protocol

- Originally
  - Label distribution protocol defined along with MPLS
- De-facto
  - RSVP-TE is used mostly

TE: traffic engineering

[Medh17]

RSVP (Resource ReserVation Protocol)

- Goal
  - Bandwidth reservation for end-to-end data streams
- Soft state principle
  - Establish a session and periodically signal that session is still alive
  - In case of failure, state is automatically removed after some time

- Signaling
  - Path message
    - From sender to receiver
    - Find path to receiver
    - Each hop is recorded in the message
  - Resv message
    - From receiver to sender
    - Bandwidth reservation on return path

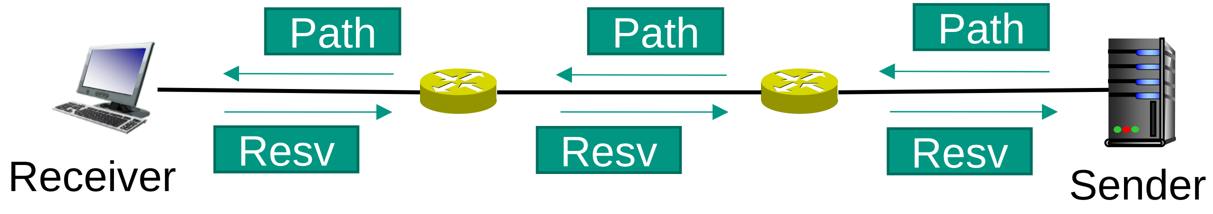

[RFC2205]

RSVP-TE (Traffic Engineering)

- Extension to RSVP to support label distribution
  - Many additional fields and functionality, e.g., fast reroute

- Label distribution type
  - Downstream on demand

- Signaling
  - Path message
    - From upstream LER to downstream LER
    - Label request
    - Source route (“explicit route”) [optional]
  - Resv message
    - In response to path message
    - From downstream LER to upstream LER
    - Label binding (hop-per-hop)

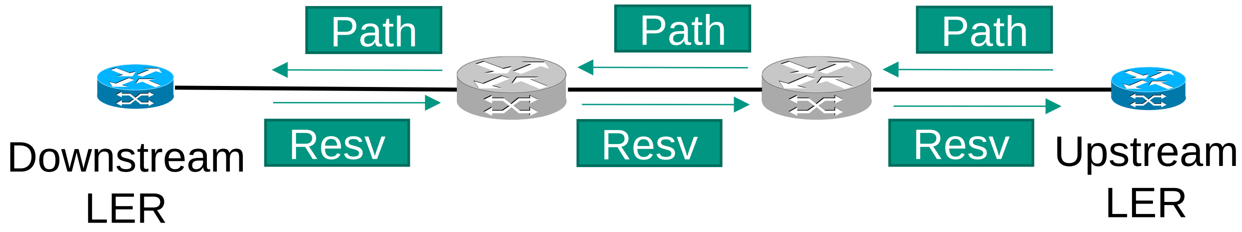

[RFC3209]

Virtual Private Networks

- MPLS is useful for virtual private networks (VPNs)

- Use case: VPN traffic engineering
  - Customer with sites at different locations (e.g., different cities) wants to lease seamless “network” service

  - Requirements
    - Connect physically remote locations
    - Carry IP-based intranet traffic
    - Each customer has obtained an IP address block
    - Guaranteed bandwidth / SLAs

  - Options
    - “Dark fibre” provider or
    - VPN backbone provider

[Medh17]

Example: Private Networks over "Dark Fibre"

- Three companies have sites at remote locations
  - Company A: Karlsruhe, Paris, Zürich
  - Company B: Karlsruhe, Paris
  - Company C: Karlsruhe, Paris
- Each company runs a private network
  - Different subnet for each site from customers IP address space
  - Router connects site to other site(s)
  - Data is transported over leased fiber optic cables ("dark fibre")
    - Capacity of each cable: 155 Mbit/s, utilization marked in graph

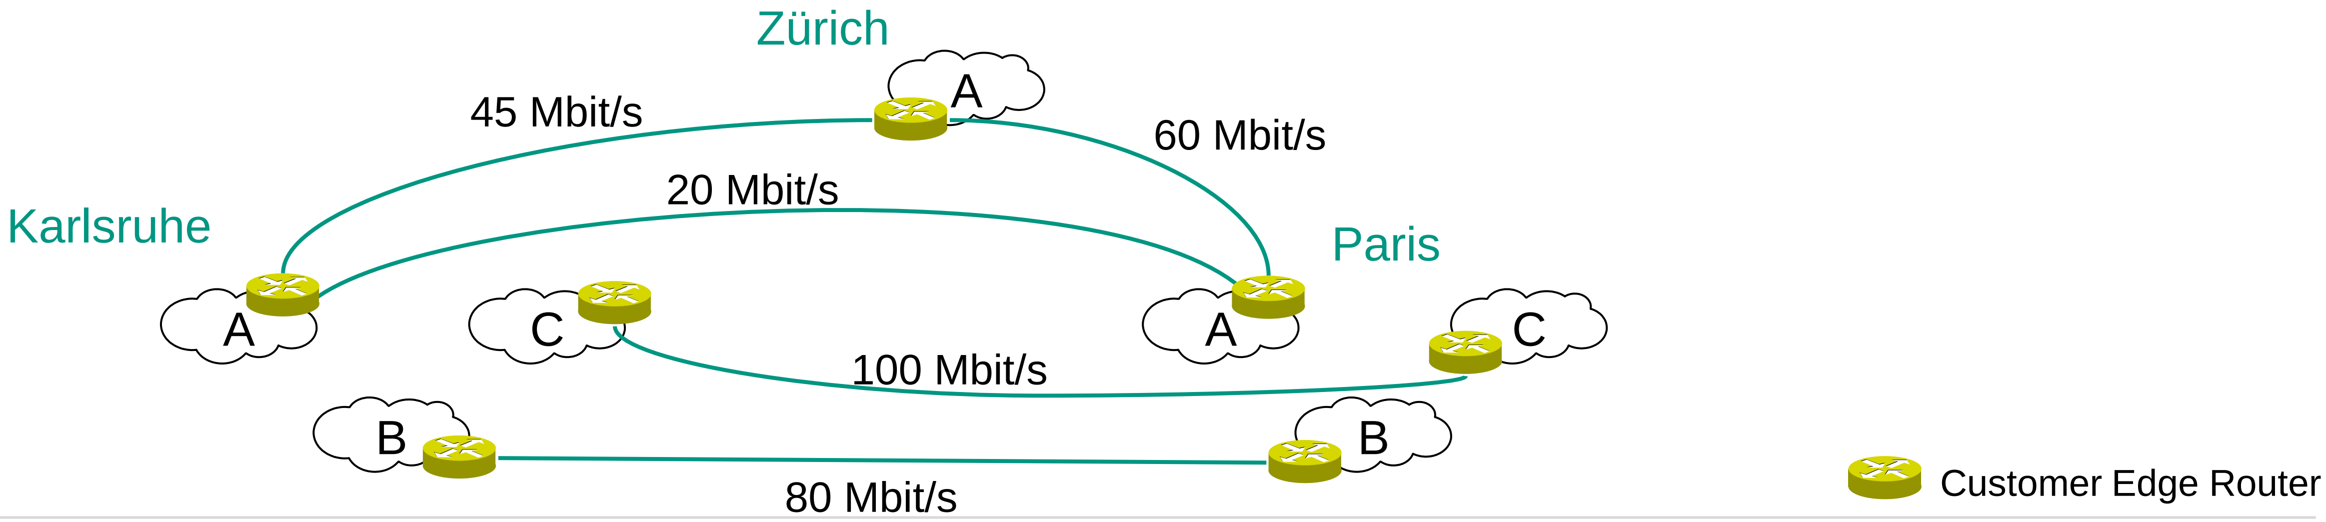

Example: Virtual Private Networks

- A provider uses MPLS to offer virtual private networks
  - Has „points of presence (PoP)“ in all three cities
  - Offers bandwidth at arbitrary rates
  - Is cheaper than leasing fiber optic cables

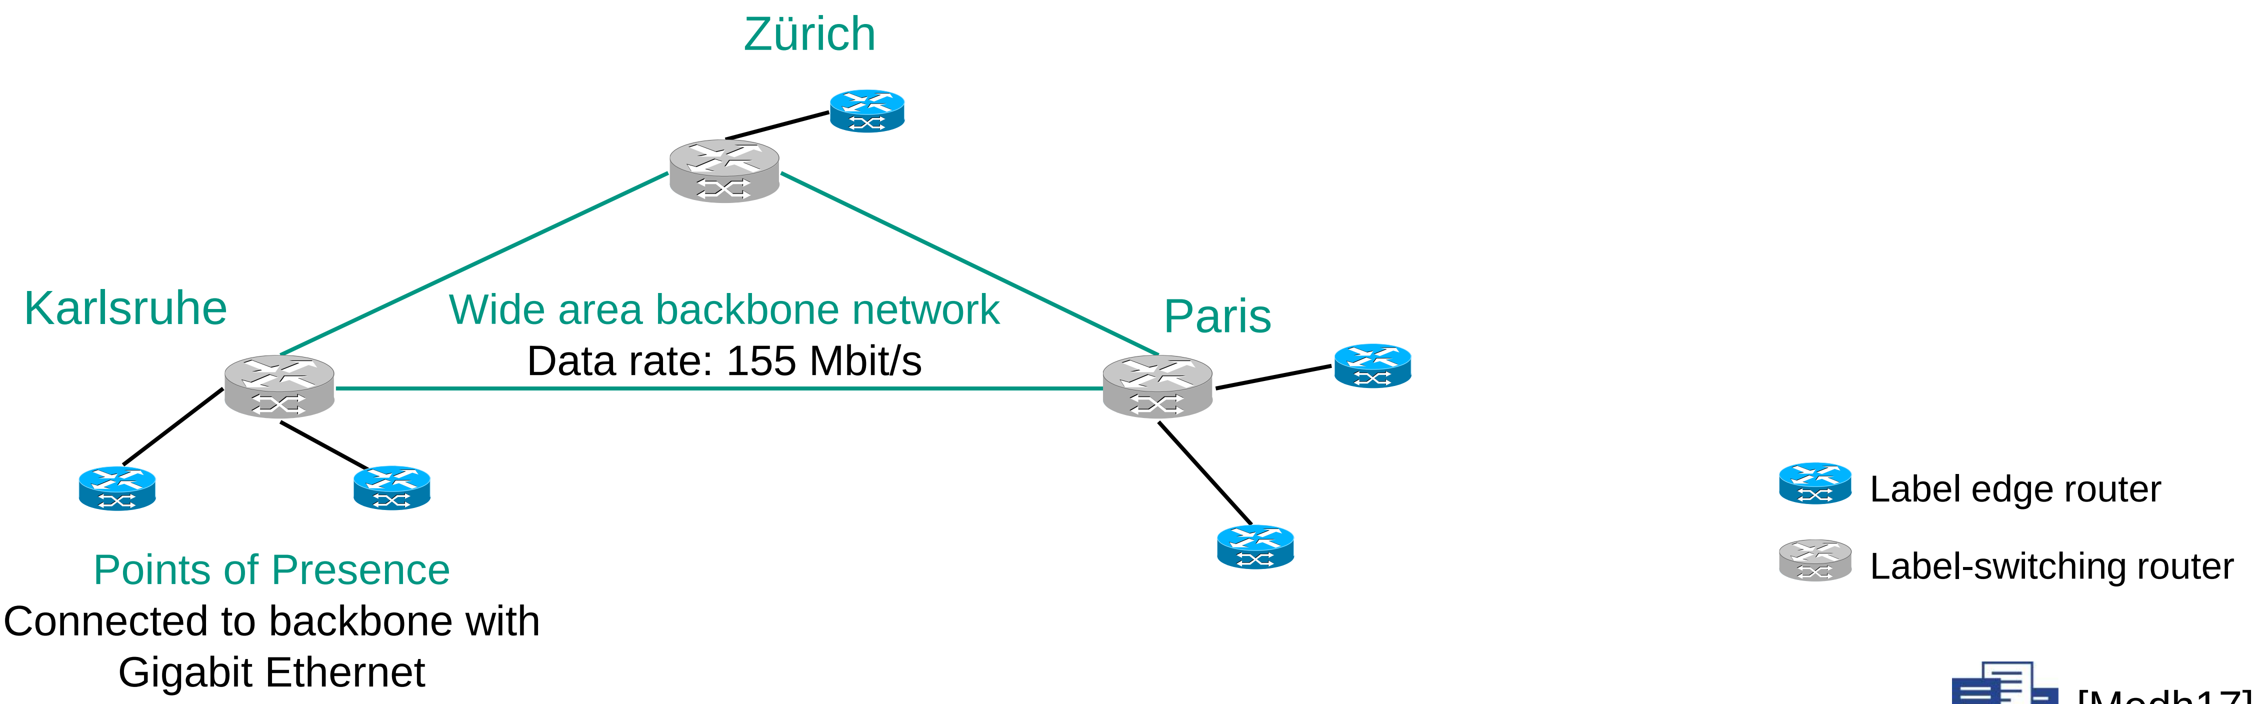

[Medh17]

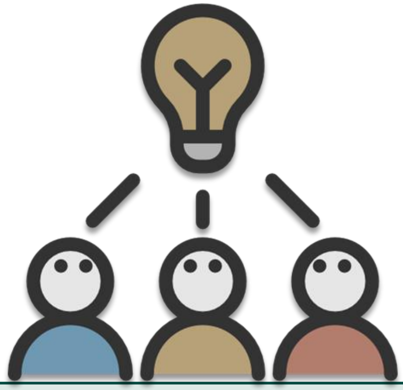

Task 04-2

VPN

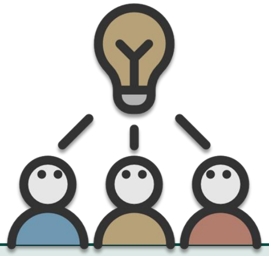

Task 04-2

VPNs Implemented by Label Stacking

- Outer label identifies path to LER
- Inner label identifies VPN instance / customer

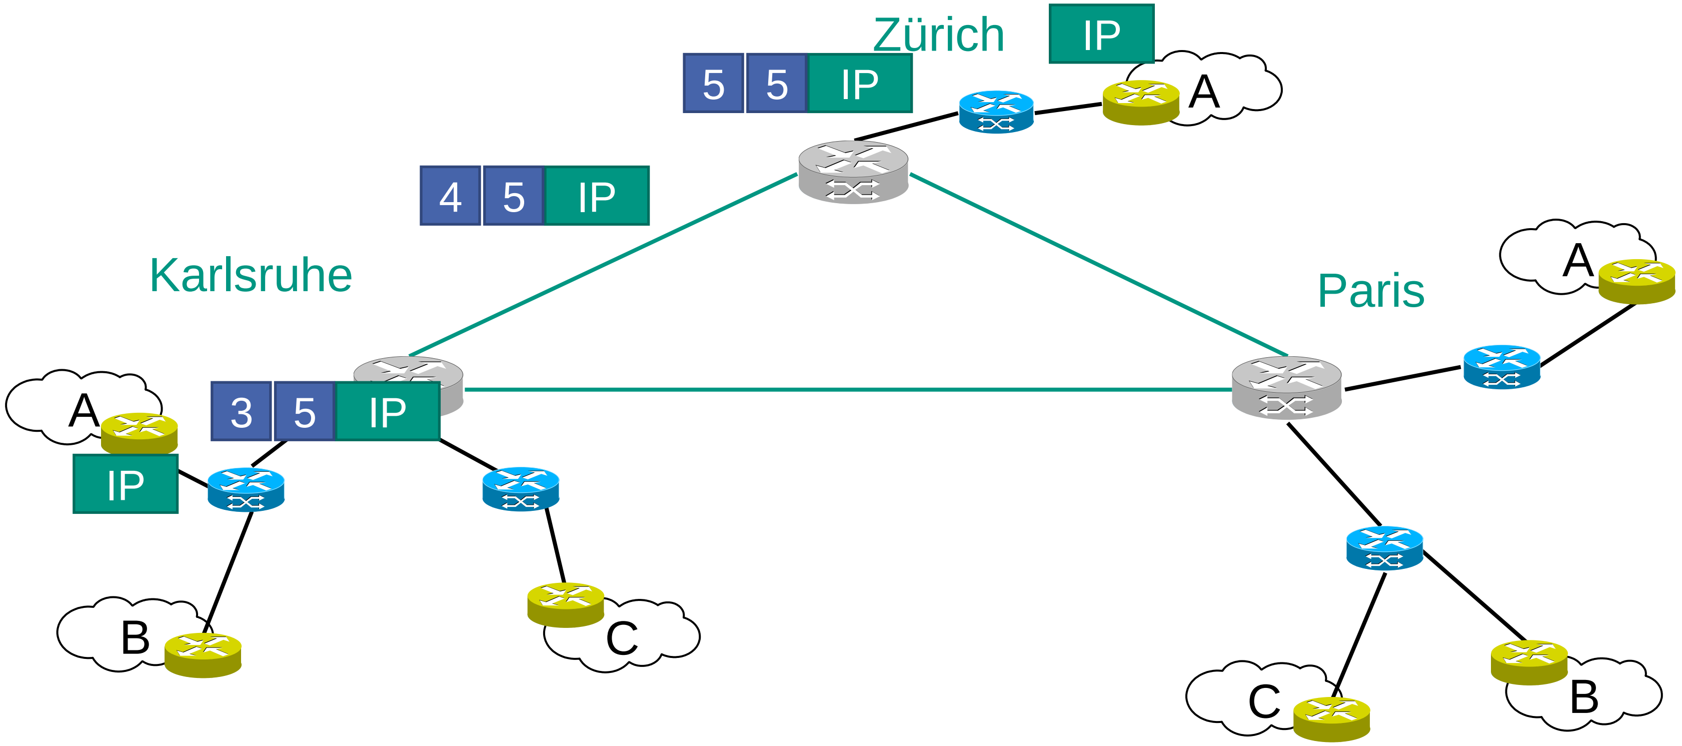

VPNs Implemented by Label Stacking

- Outer label identifies path to LER
- Inner label identifies VPN instance / customer

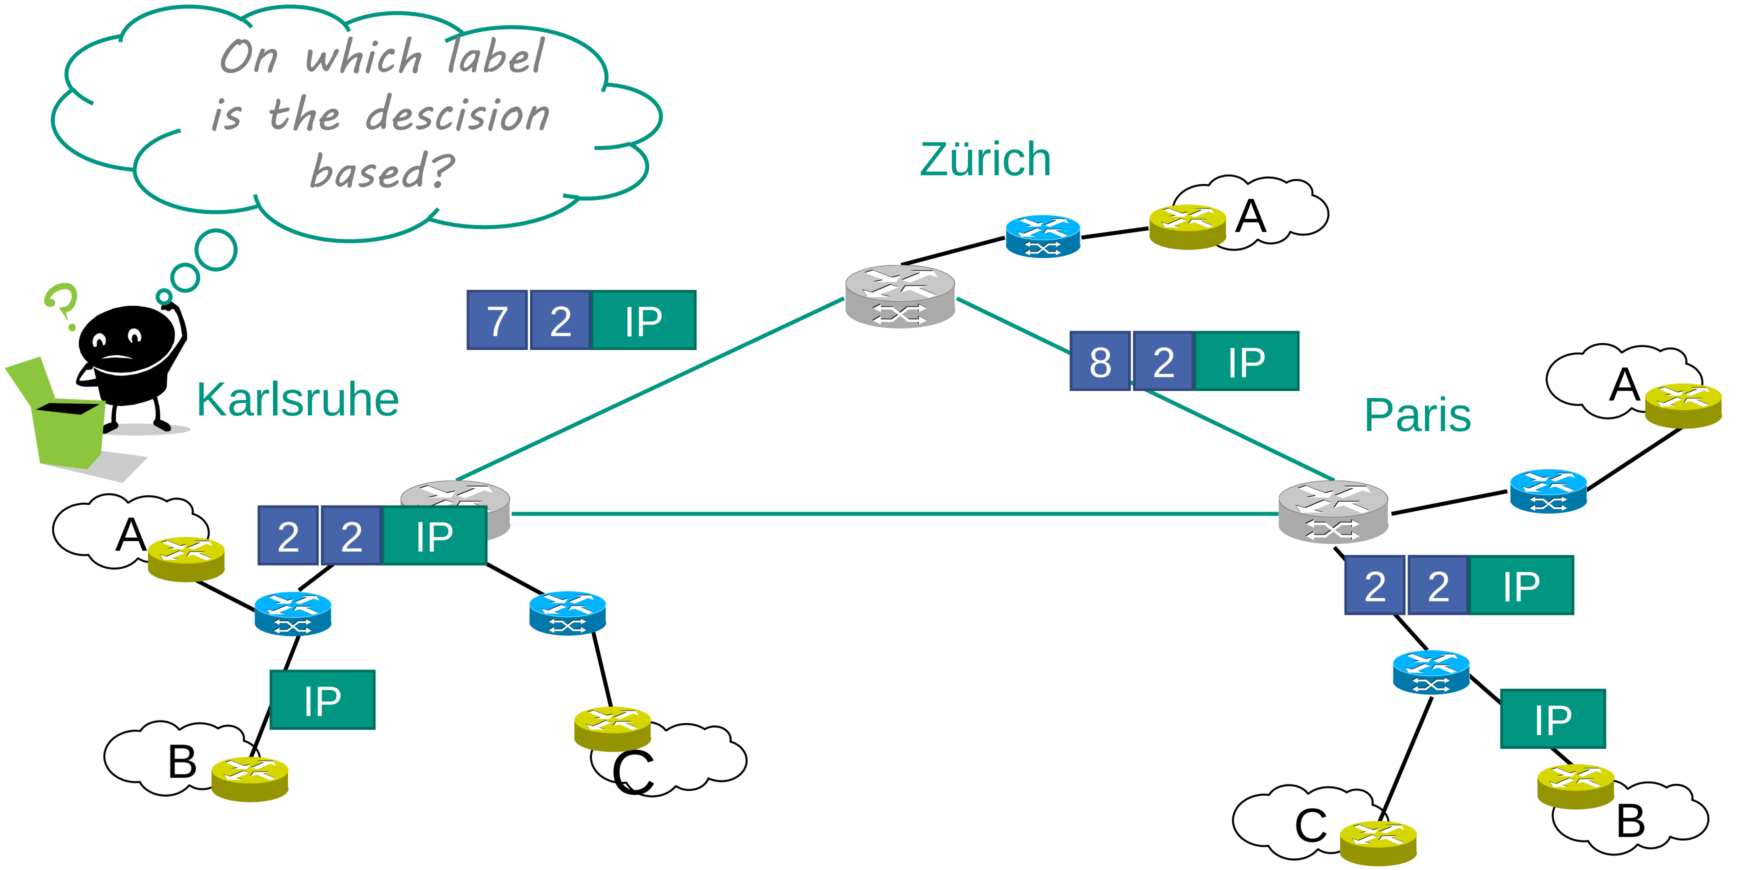

VPNs Implemented by Label Stacking

- Outer label identifies path to LER
- Inner label identifies VPN instance / customer

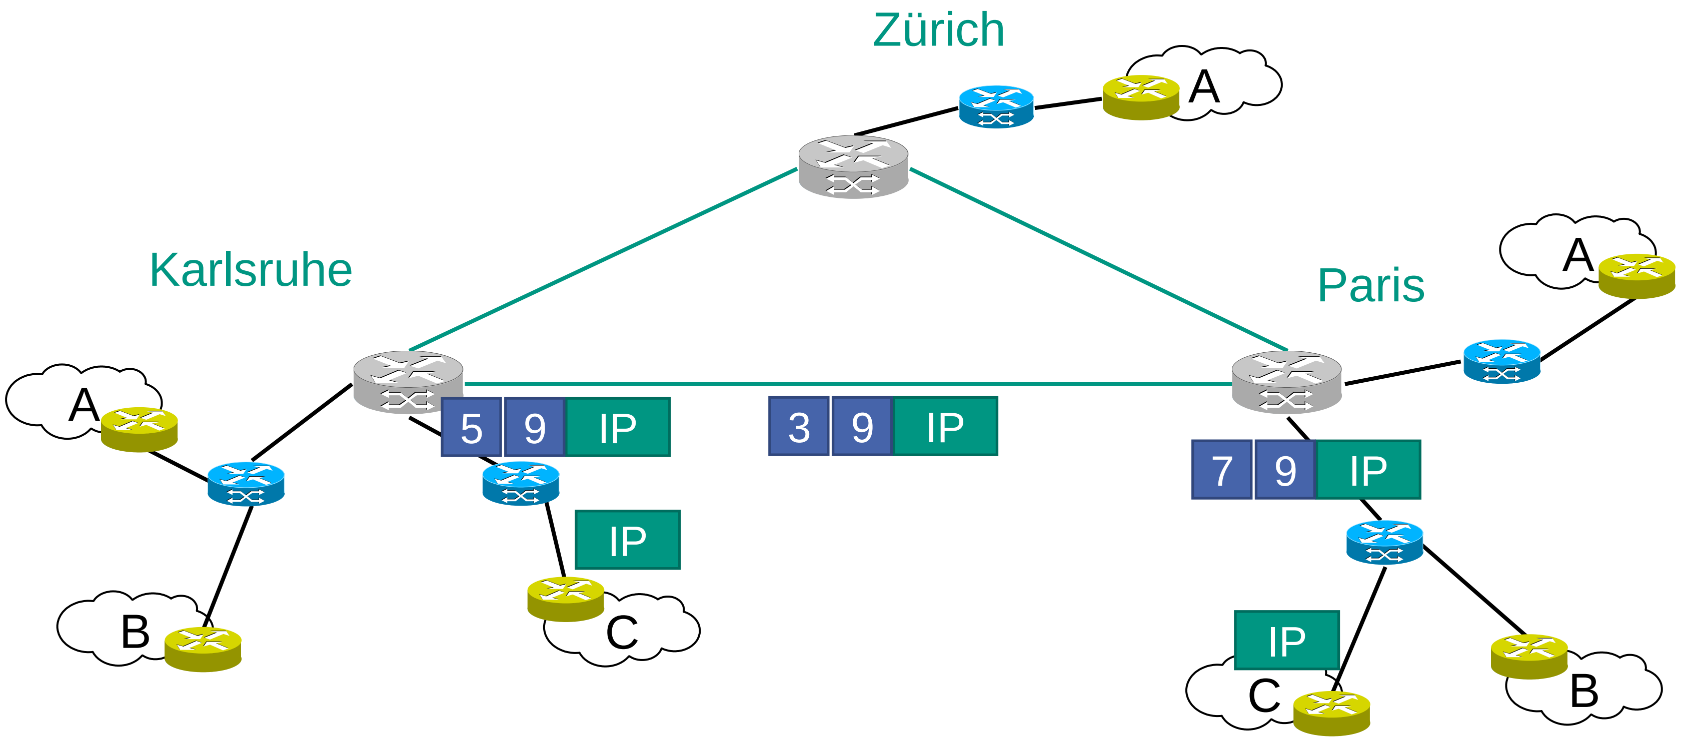

Label Distribution

- Recall VPN example from above
  - LSP for customer B (Karlsruhe → Paris) should take a “detour” over Zürich) to match bandwidth requirements
  - Setup of LSPs over explicitly given route with RSVP-TE

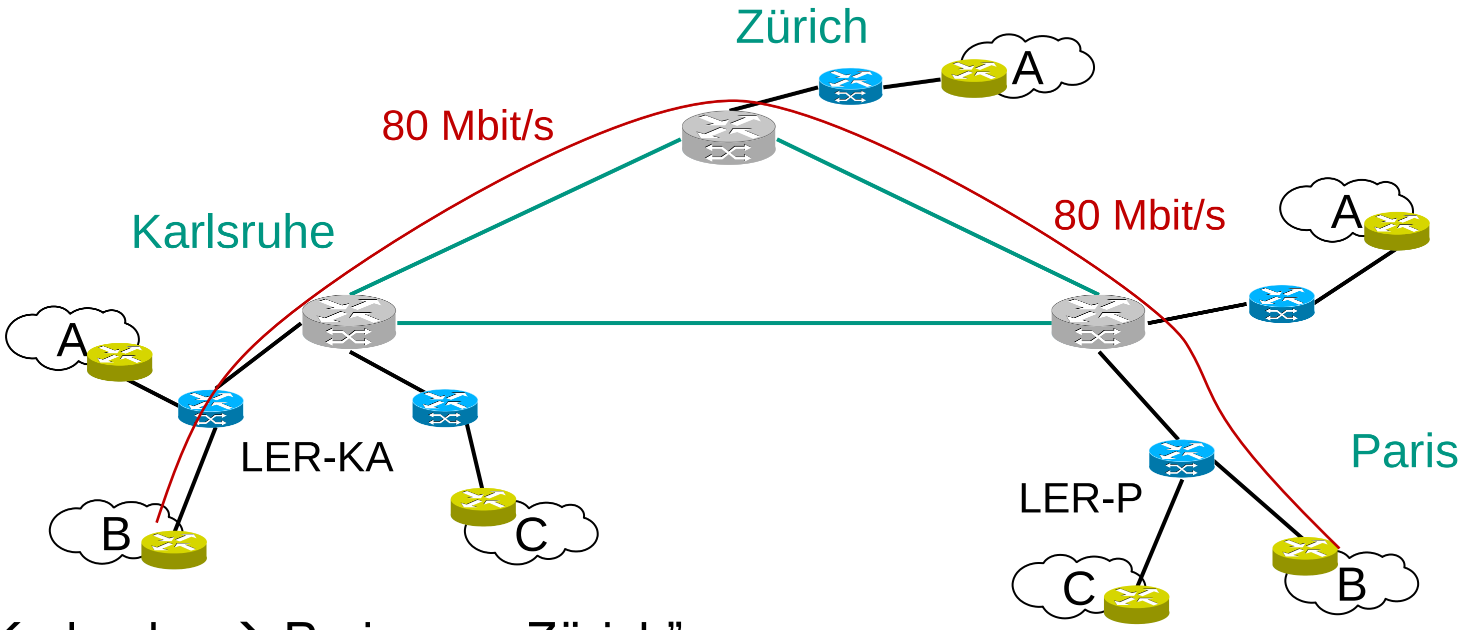

- Example: LSP “Karlsruhe → Paris over Zürich”
  - RSVP-TE signaling initiated at upstream LER (LER-KA)
  - Note: LSPs are unidirectional!

Label Distribution

1 Path Message	
Explicit route	Backbone-KA, Backbone-Z, Backbone-P, LER-P
Record-route	LER-KA
Reservation	80 Mbit/s
Label	Request new label

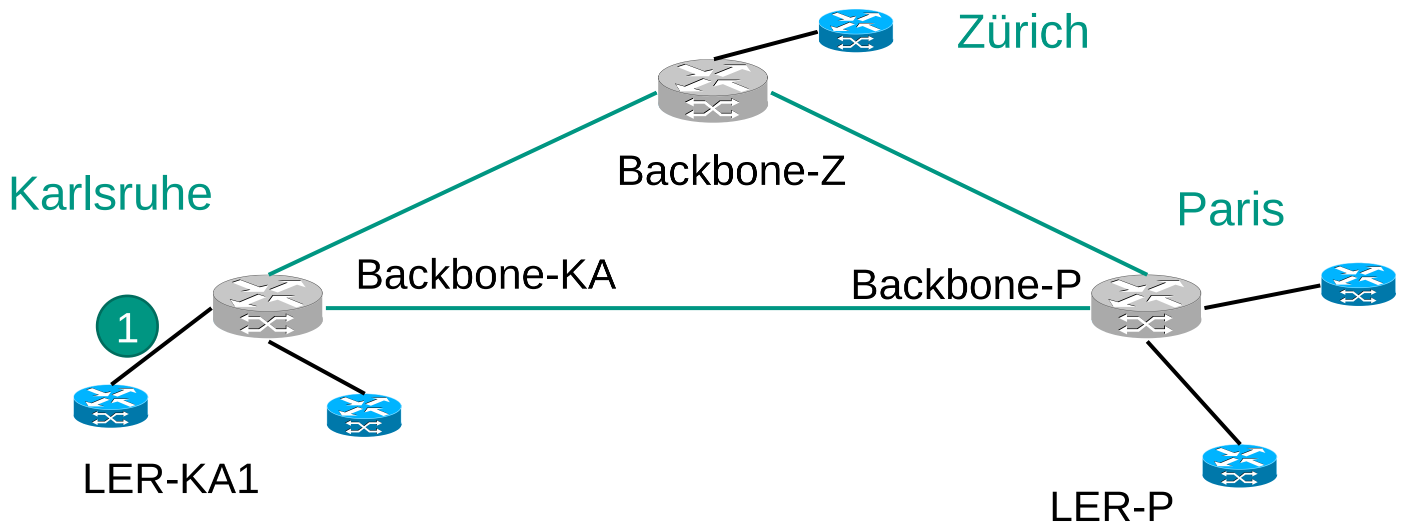

Label Distribution

2 Path Message	
Explicit route	Backbone-Z, Backbone-P, LER-P
Record-route	LER-KA1, Backbone-KA
Reservation	80 Mbit/s
Label	Request new label

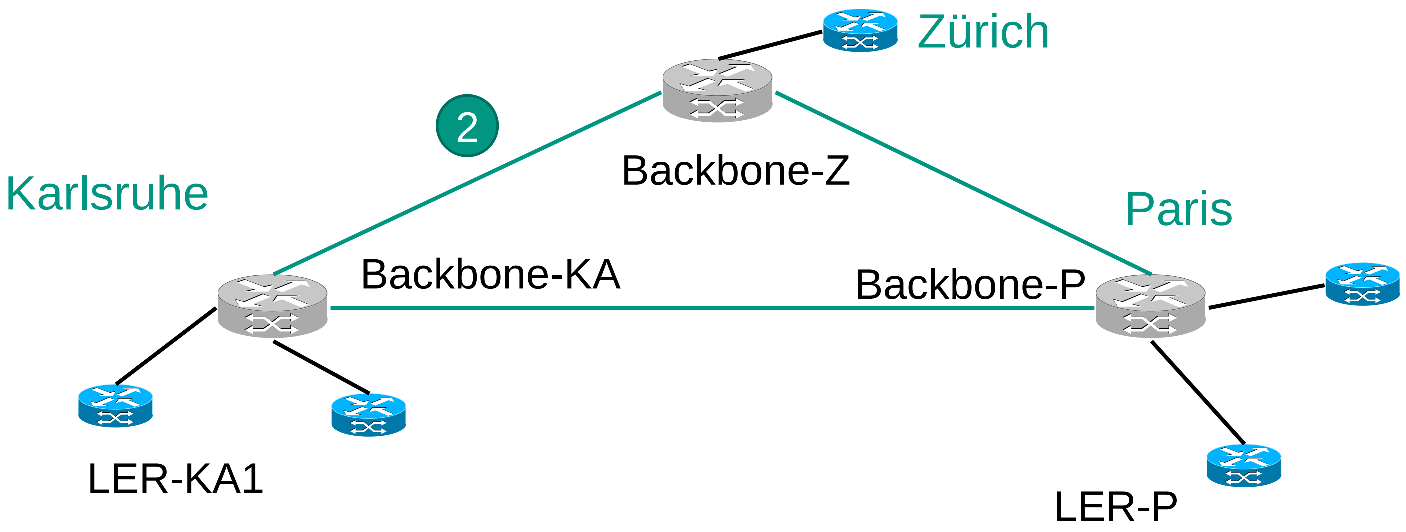

Label Distribution

3 Path Message	
Explicit route	Backbone-P, LER-P
Record-route	LER-KA1, Backbone-KA, Backbone-Z
Reservation	80 Mbit/s
Label	Request new label

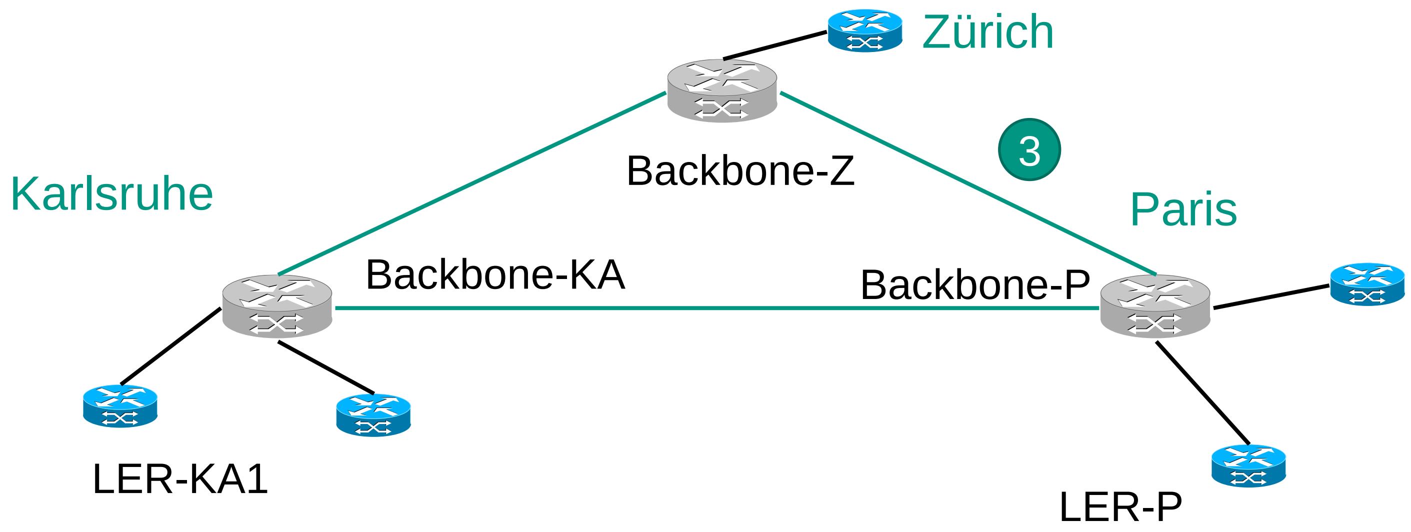

Label Distribution

4 Path Message	
Explicit route	LER-P
Record-route	LER-KA1, Backbone-KA, Backbone-Z, Backbone-P
Reservation	80 Mbit/s
Label	Request new label

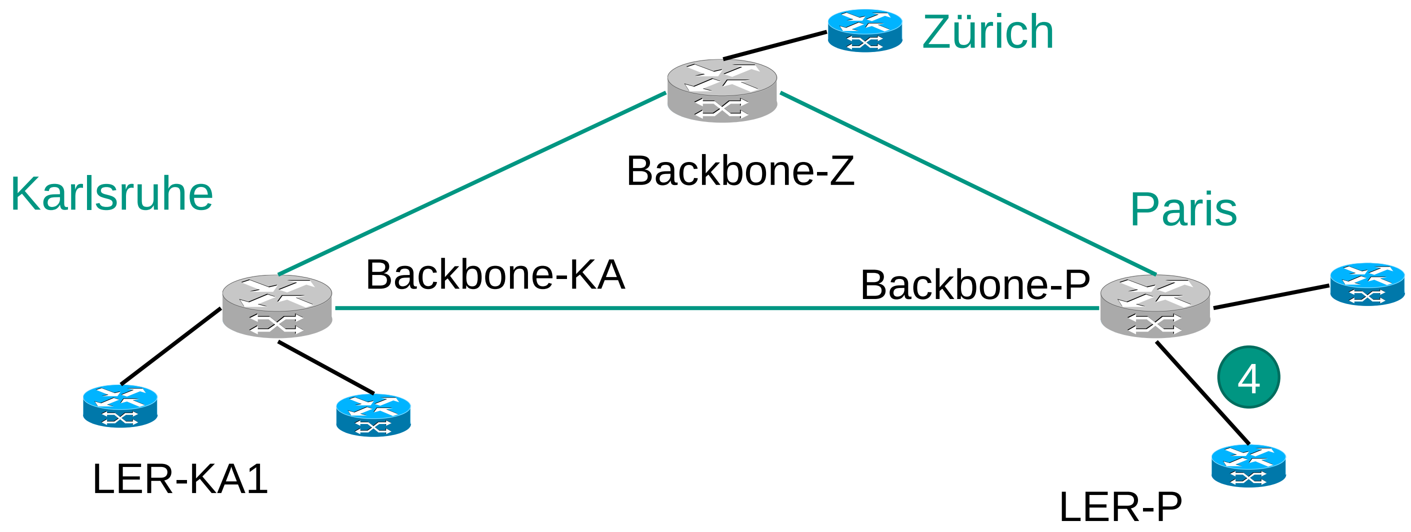

Label Distribution

5 Resv Message	
Record-route	LER-KA1, Backbone-KA, Backbone-Z, Backbone-P, LER-P
Label	2

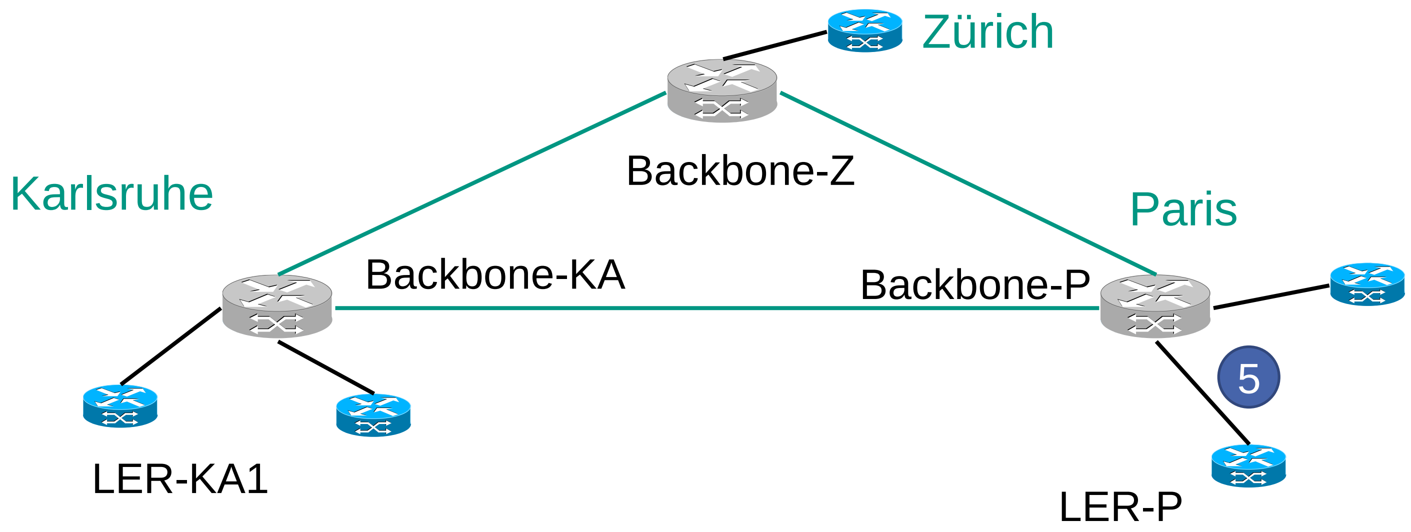

Label Distribution

6 Resv Message	
Record-route	LER-KA1, Backbone-KA, Backbone-Z, Backbone-P, LER-P
Label	8

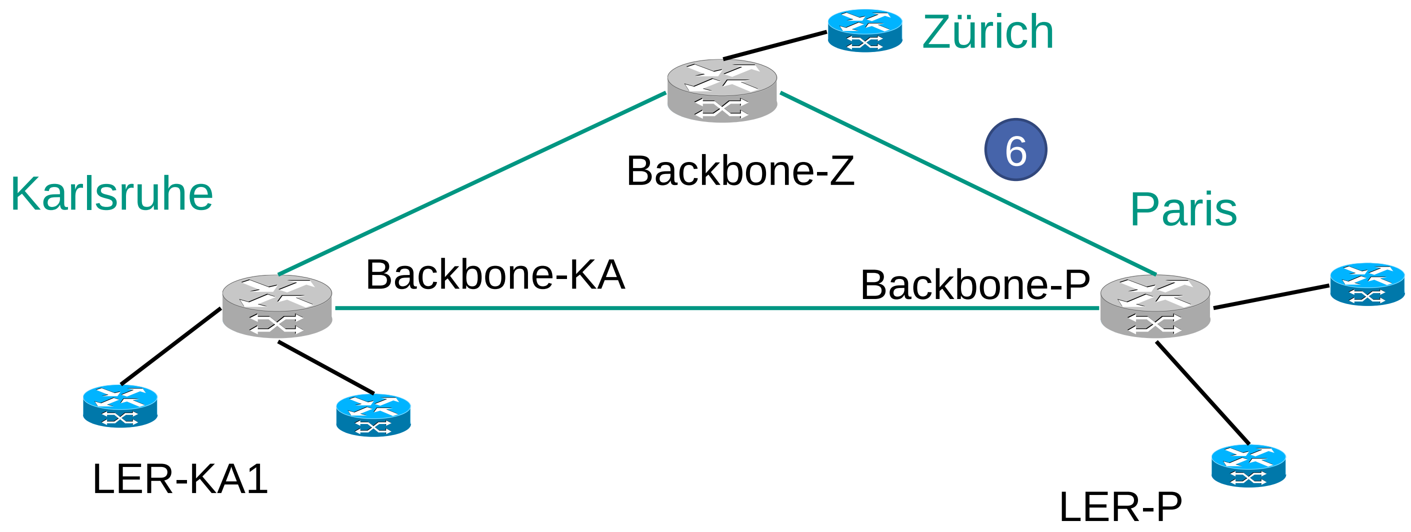

Label Distribution

7 Resv Message	
Record-route	LER-KA1, Backbone-KA, Backbone-Z, Backbone-P, LER-P
Label	7

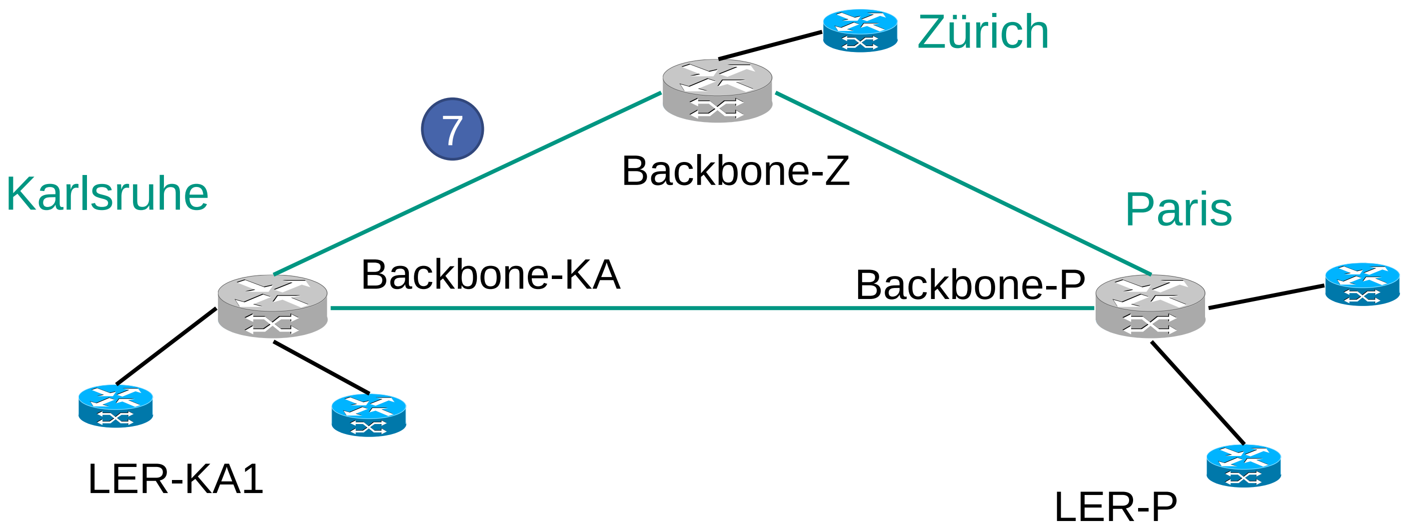

Label Distribution

8 Resv Message	
Record-route	LER-KA1, Backbone-KA, Backbone-Z, Backbone-P, LER-P
Label	2

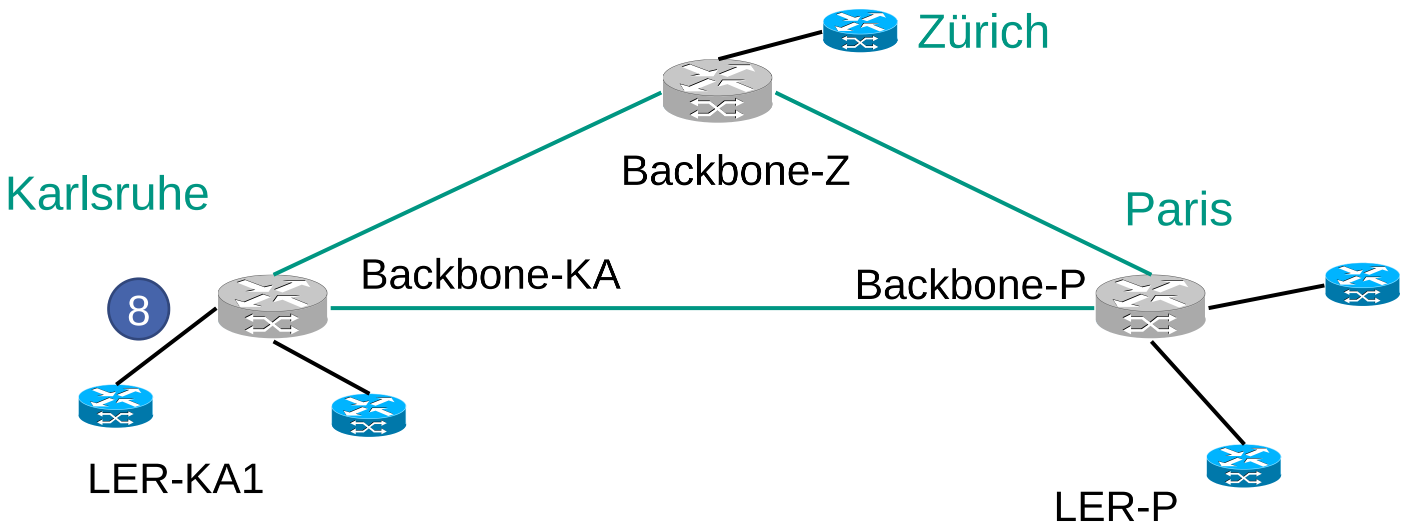

PROBLEMS

1) Explain the basic concepts of label switching.
2) Which goals are pursued by label switching?
3) How can data streams be aggregated by MPLS?
4) How can the IP data stream be mapped to different MPLS tunnels within the label edge router?
5) Which operations can be performed on labels?
6) Explain the label operations that happen when a packet travels through an MPLS domain.

LITERATURE

Literature

Karlsruher Institut für Technologie

[Medh17] D. Medhi, K. Ramasamy; Network Routing; Morgan Kaufmann, 2nd Edition, 2017

[RFC2205] R. Braden, L. Zhang, S. Berson, and S. Herzog, S. Jamin; Resource ReSerVation Protocol (RSVP); RFC 2205, IETF, September, 1997

[RFC3031] E. Rosen, A. Viswanathan, R. Callon; Multiprotocol Label Switching Architecture; RFC 3031, IETF, Januar, 2001

[RFC3209] D. Awduche, L. Berger, D. Gan, T. Li, V. Srinivasan, G. Swallow; RSVP-TE: Extensions to RSVP for LSP Tunnels; RFC 3209, IETF, December, 2001

[Stal17] W. Stallings; Data and Computer Communications; Prentice Hall, 10th Edition, 2017