

MPLS

□ Circuit switching

- Packet switching vs. circuit switching
- Virtual circuits

□ MPLS

- Labels and label-switching
- Forwarding Equivalence Classes
- Label distribution
- MPLS applications

Packet switching vs. circuit switching

□ Packet switching

- Data traffic divided into packets
 - Each packet contains its own header (with address)
 - Packets sent separately through the network
- Destination reconstructs the message
- Example: sending a letter through postal system

□ Circuit switching

- Source first establishes a connection to the destination
 - Each router on the path may reserve bandwidth
- Source sends data over the connection
 - No destination address, since routers know the path
- Source tears down the connection when done
- Example: voice conversation on telephone network

Advantages of circuit switching

- ❑ Guaranteed bandwidth
 - Predictable communication performance
 - Not “best-effort” delivery with no real guarantees
- ❑ Simple abstraction
 - Reliable communication channel between hosts
 - No worries about lost or out-of-order packets
- ❑ Simple forwarding
 - Forwarding based on time slot or frequency
 - No “longest prefix match” on each packet
- ❑ Low per-packet overhead
 - Forwarding based on time slot or frequency
 - No IP (and TCP/UDP) header on each packet

Disadvantages of circuit switching

❑ Wasted bandwidth

- Bursty traffic leads to idle connection during silent period
- Unable to achieve gains from statistical multiplexing

❑ Blocked connections

- Connection refused when resources are not sufficient
- Unable to offer "okay" service to everybody

❑ Connection set-up delay

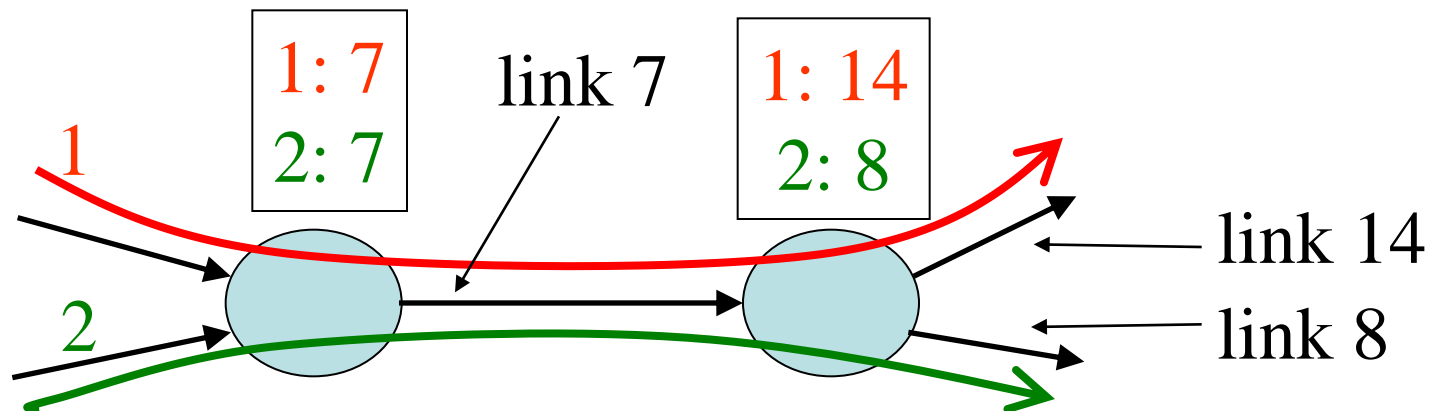
- No communication until the connection is set up
- Unable to avoid extra latency for small data transfers

❑ Network state

- Routers must store per-connection information
- Unable to avoid per-connection storage and state failover

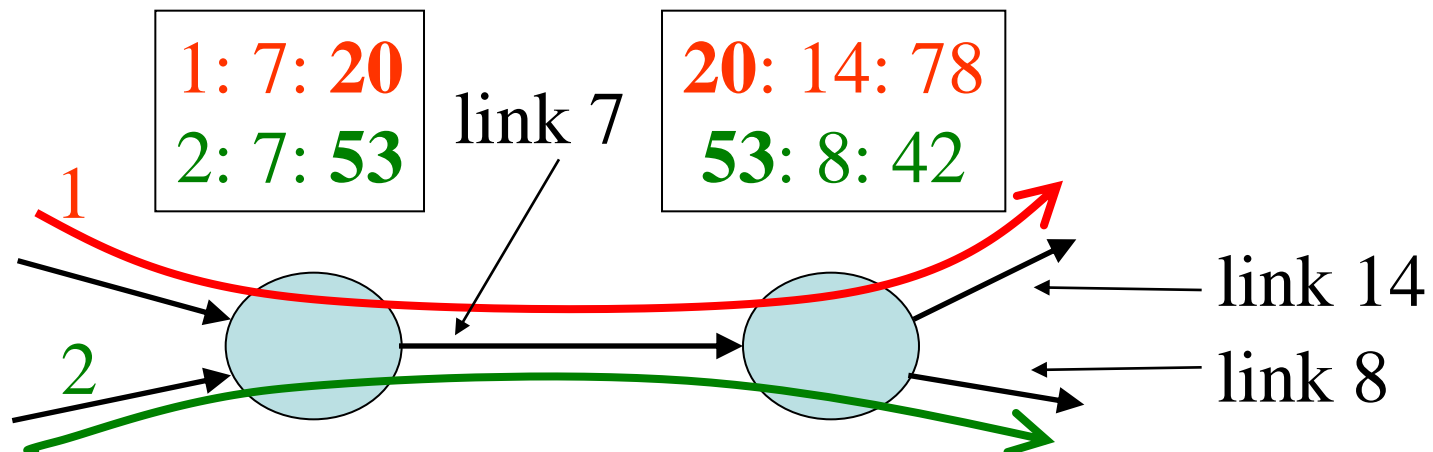
Virtual circuits

- ❑ Hybrid of packet and circuit switching
 - Logical circuit between a source and destination
 - Packets from different VCs multiplex on a link
- ❑ Virtual Circuit Identifier (VC ID)
 - Source set-up: establish path for the VC
 - Switch: mapping VC ID to an outgoing link
 - Packet: fixed length label in the header



Swapping the label at each hop

- ❑ Problem: Using VC ID along the whole path
 - Each virtual circuit consumes a unique ID
 - Starts to use up all of the ID space in the network
- ❑ Label swapping
 - Map the VC ID to a new value at each hop
 - Table has old ID, next link, and new ID
 - Allows reuse of the IDs at different links



Virtual circuits similar to IP datagrams

□ Data divided in to packets

- Sender divides the data into packets
- Packet has an address (e.g., IP address or VC ID)

□ Store-and-forward transmission

- Multiple packets may arrive at once
- Need buffer space for temporary storage

□ Multiplexing on a link

- No reservations: statistical multiplexing
 - Packets are interleaved without a fixed pattern
- Reservations: resources for group of packets
 - Guarantees to get a certain number of “slots”

Virtual circuits differ from IP datagrams

❑ Forwarding look-up

- Virtual circuits: fixed-length connection id
- IP datagrams: destination IP address

❑ Initiating data transmission

- Virtual circuits: must signal along the path
- IP datagrams: just start sending packets

❑ Router state

- Virtual circuits: routers know about connections
- IP datagrams: no state, easier failure recovery

❑ Quality of service

- Virtual circuits: resources and scheduling per VC
- IP datagrams: difficult to provide QoS

Multi-Protocol Label Switching

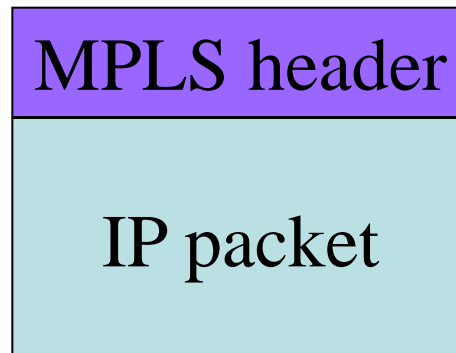
Multi-Protocol Label Switching

❑ Multi-Protocol

- Encapsulate a data packet
 - Could be IP, or some other protocol (e.g., IPX)
- Put an MPLS header in front of the packet
 - Actually, can even build a stack of labels...

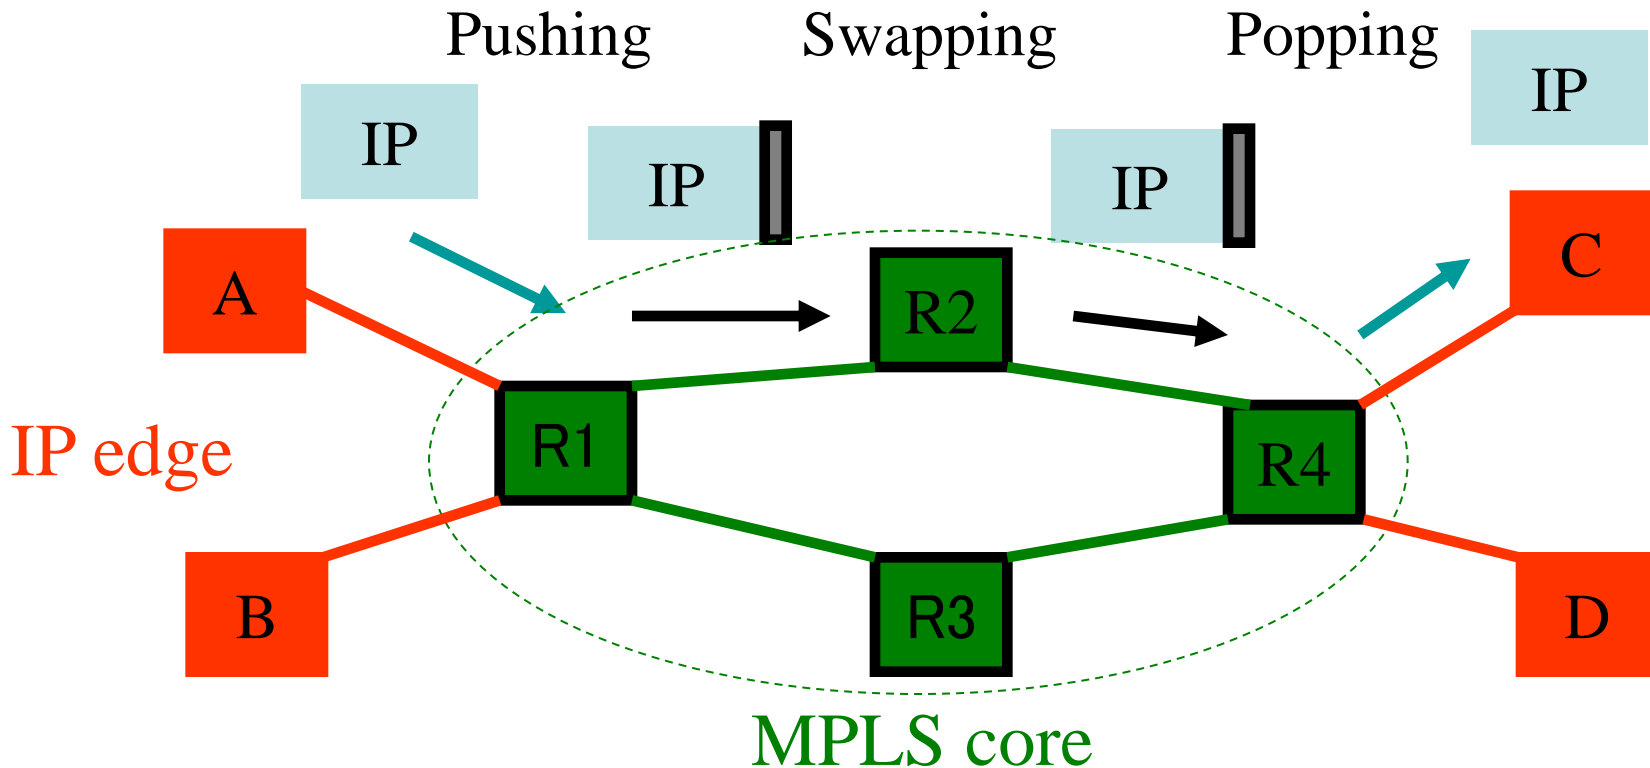
❑ Label Switching

- MPLS header includes a label
- Label switching between MPLS-capable routers



Pushing, swapping, and popping

- ❑ Pushing: add the initial "in" label
- ❑ Swapping: map "in" label to "out" label
- ❑ Popping: remove the "out" label



Forwarding Equivalence Class (FEC)

❑ Rule for grouping packets

- Packets that should be treated the same way
- Identified just once, at the edge of the network

❑ Example FECs

- Destination prefix
 - Longest-prefix match in forwarding table at entry point
 - Useful for conventional destination-based forwarding
- Src/dest address, src/dest port, and protocol
 - Five-tuple match at entry point
 - Useful for fine-grain control over the traffic
- Sent by a particular customer site
 - Incoming interface at entry point
 - Useful for virtual private networks

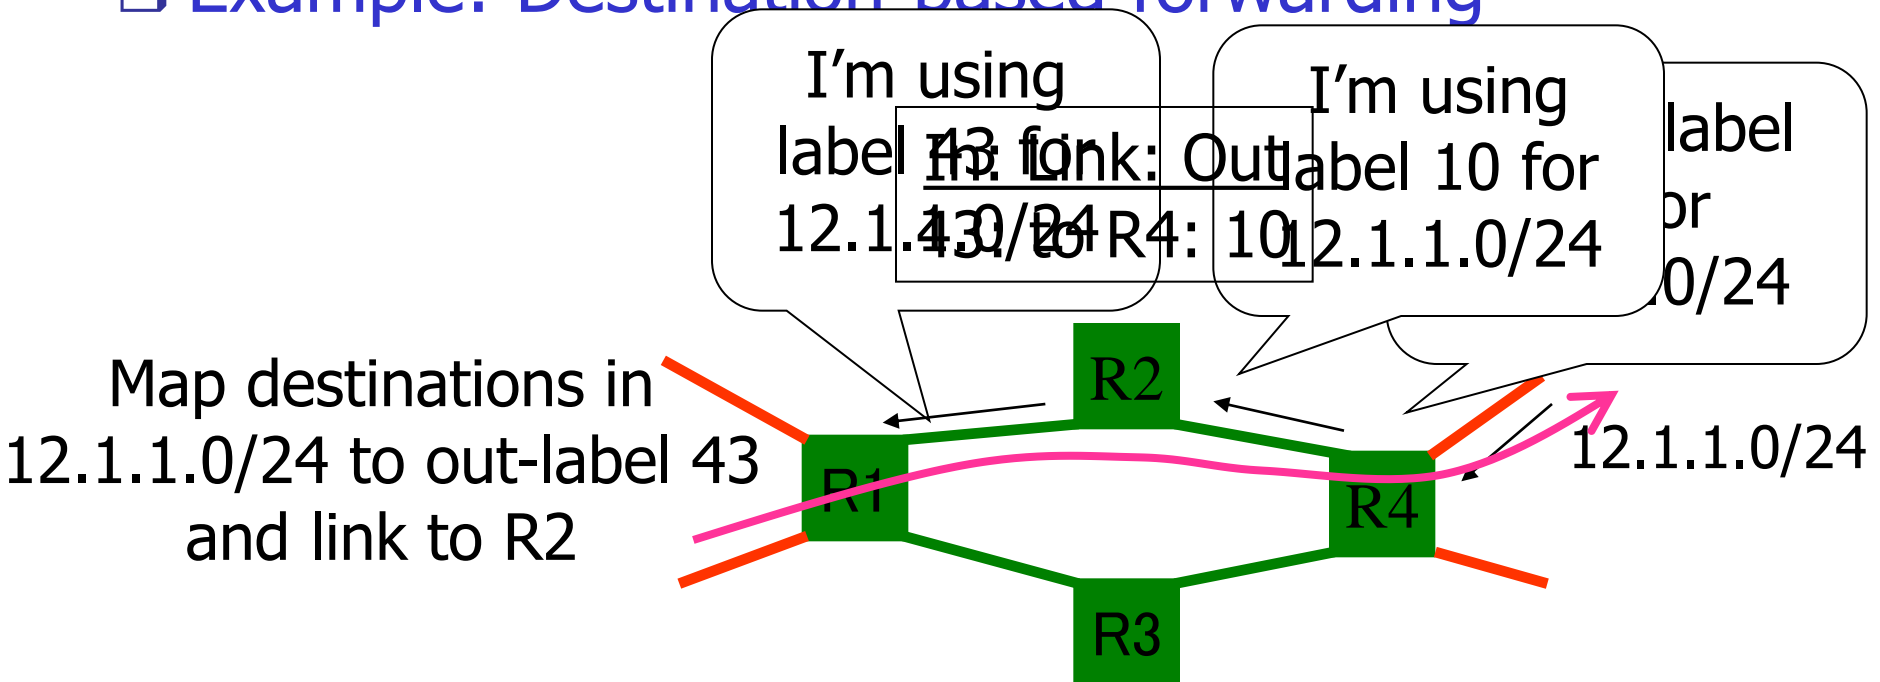
A label is just a locally-significant identifier for a FEC

Label Distribution Protocol

□ Distributing labels

- Learning the mapping from FEC to label
- Told by the downstream router

□ Example: Destination-based forwarding



Applications of MPLS

TE with constraint-based routing

□ Path calculation

- Constrained shortest-path first
- Compute shortest path based on weights
 - But, exclude paths that do not satisfy constraints
 - E.g., do not consider links with insufficient bandwidth

□ Information dissemination

- Extend OSPF/IS-IS to carry the extra information
 - E.g., link-state attributes for available bandwidth

□ Path signaling

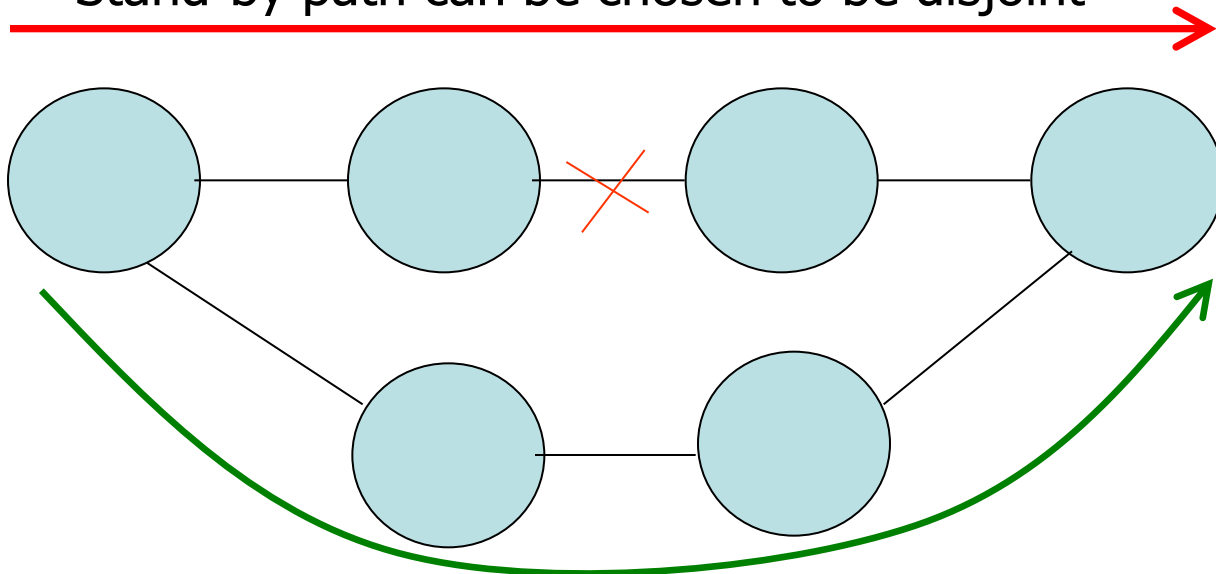
- Establish label-switched path on explicit route

□ Forwarding: MPLS labels

Surviving failures: Path protection

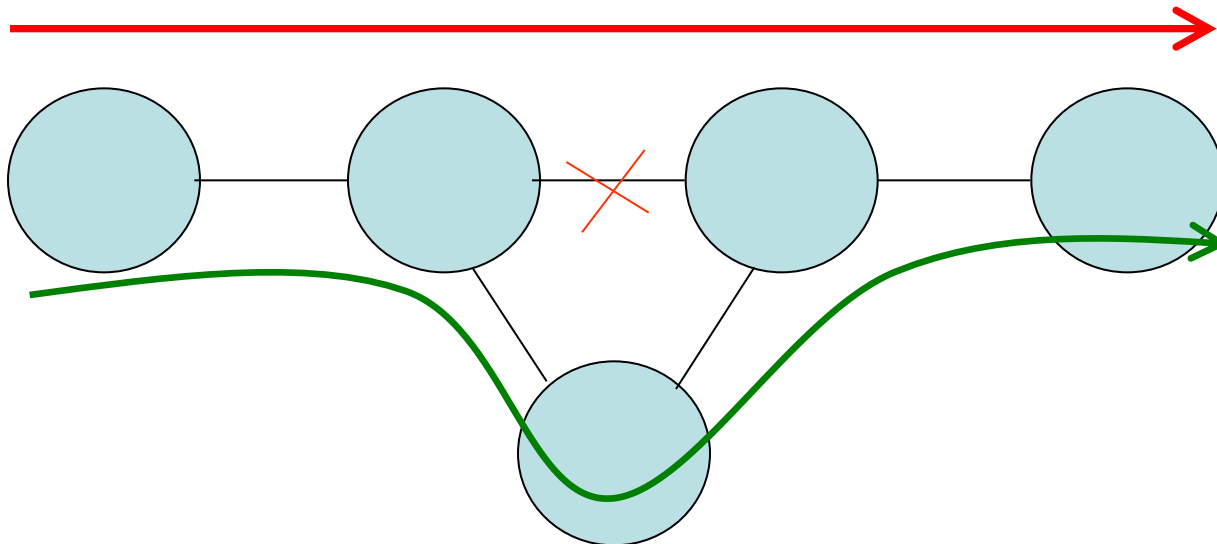
□ Path protection

- Reserve bandwidth on an alternate route
 - Protect a label-switched path by having a stand-by
- Much better than conventional IP routing
 - Precise control over where the traffic will go
 - Stand-by path can be chosen to be disjoint

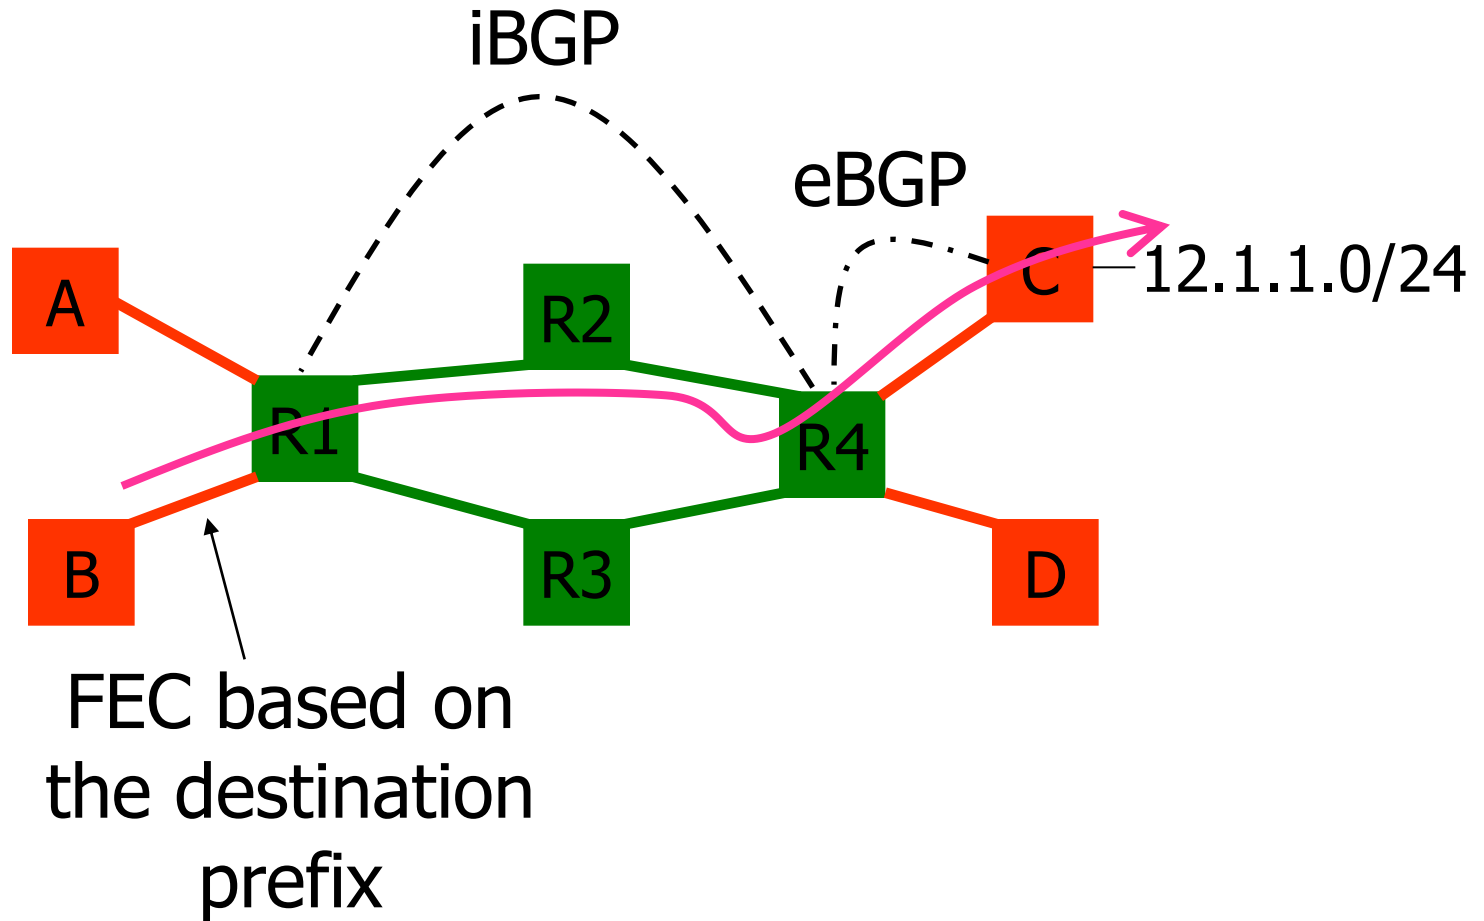


Surviving failures: Fast reroute

- ❑ Ensure fast recovery from a link failure
 - Protect a link by having a protection sub-path
- ❑ Much faster recovery than switching paths
 - Affected router can detect the link failure
 - ... and start redirecting to the protection sub-path

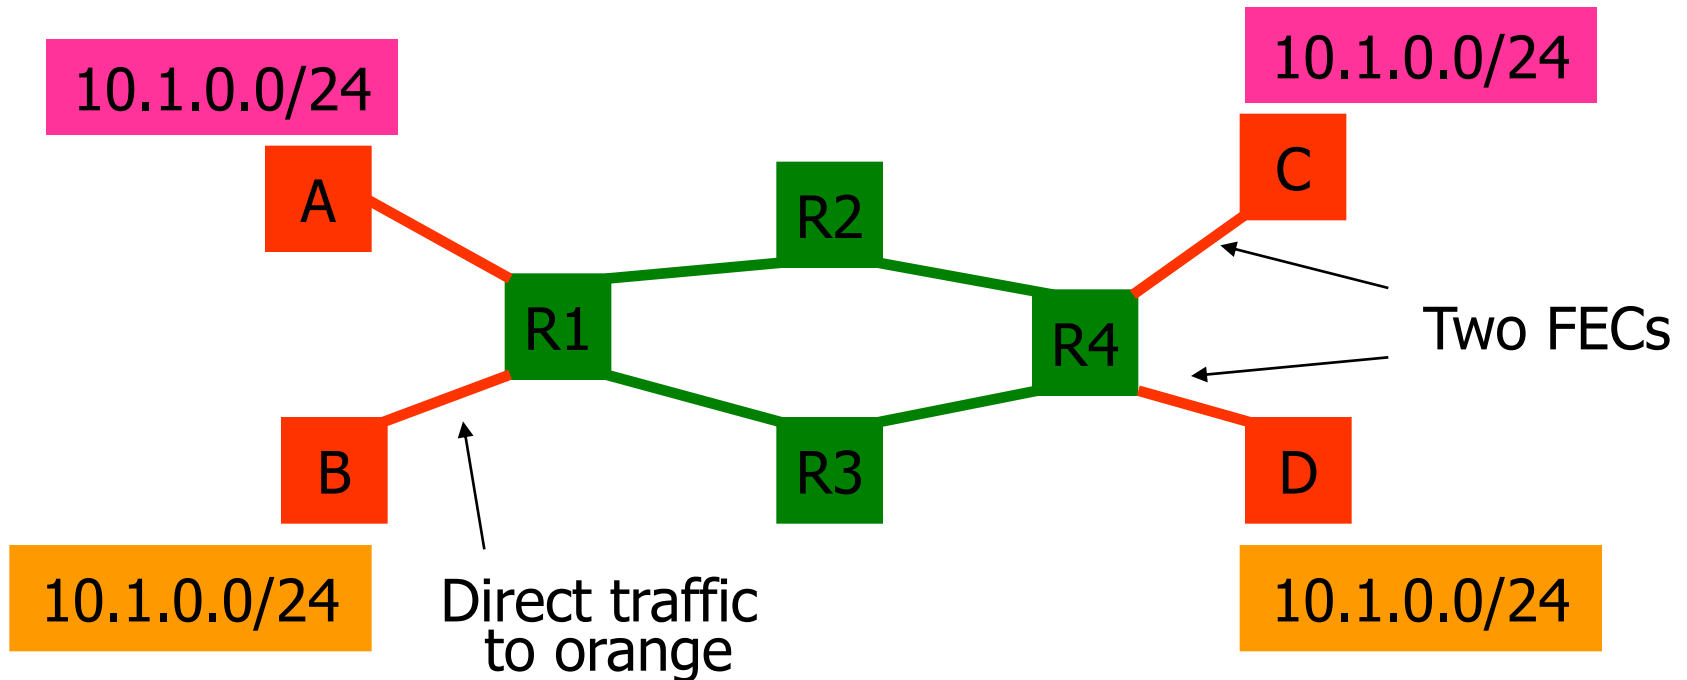


BGP-Free core



Routers R2 and R3 don't need to speak BGP

VPNs with private addresses



MPLS tags can differentiate green VPN from orange VPN.

Status of MPLS

❑ Deployed in practice

- BGP-free core
- Virtual Private Networks
- Traffic engineering

❑ Challenges

- Protocol complexity
- Configuration complexity
- Difficulty of collecting measurement data

❑ Continuing evolution

- Standards
- Operational practices and tools