15-441 15-641 Computer Networking

Lecture 11: Virtual Circuits

Peter Steenkiste

Fall 2016

www.cs.cmu.edu/~prs/15-441-F16

Outline

- Circuit switching refresher
- Virtual Circuits - general
  - Why virtual circuits?
  - How virtual circuits? -- tag switching!
- Two modern implementations
  - ATM - teleco-style virtual circuits
  - MPLS - IP-style virtual circuits

Circuit Switching

- Source first establishes a connection (circuit) to the destination.
  - Each router or switch along the way may reserve some bandwidth for the data flow
- Source sends the data over the circuit.
  - No destination address needed - routers know the path
- The connection is torn down.
- Example: traditional telephone network.

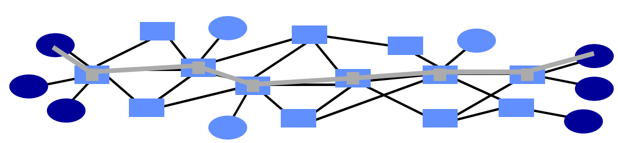

Circuit Switching

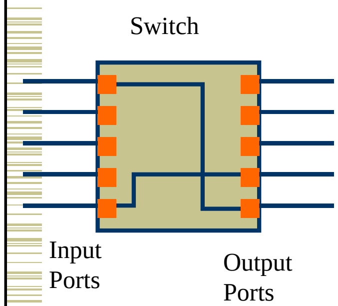

- Switch remembers how to forward data
  - No addresses!
- Many options
  - Between specific wires (circuit = wire)
  - Between timeslots (TDMA on each wire)
  - Between frequencies (FDMA on each wire)

Circuit Versus Packet Switching

Circuit Switching

- Fast switches can be built relatively inexpensively
- Inefficient for bursty data
- Predictable performance (e.g. hard QoS)
- Requires circuit establishment before communication

Packet Switching

- Switch design is more complex and expensive
- Allows statistical multiplexing
- Difficult to provide QoS guarantees
- Data can be sent without signaling delay and overhead

Can we get the benefits of both?

Virtual Circuits

- Each wire carries many "virtual" circuits.
- Forwarding based on virtual circuit (VC) identifier
  - IP header: src, dst, etc.
  - Virtual circuit header: just "VC"
  - A path through the network is set up when the VC is established
  - Can eue statistical multiplexing for efficiency
- Can support wide range of quality of service.
  - No guarantees: best effort service
  - Weak guarantees: delay < 300 msec, …
  - Strong guarantees: e.g. equivalent of physical circuit

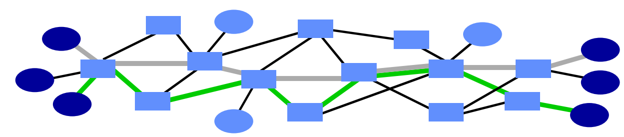

Virtual Circuits Versus Packet Switching

- Many similarities:
  - Forwarding based on "address" (VCID or dest address)
  - Statistical multiplexing for efficiency
  - Must have buffers space on switches
- Virtual circuit switching:
  - Uses short circuit identifiers to forward packets
  - Switches keep (hard) state for each circuit, so they can more easily implement features such as quality of service
  - Failures result in loss of virtual circuit
- Packet switching:
  - Use full destination addresses for forwarding packets
  - Can send data right away: no need to establish a connection first
  - Switches are stateless: easier to recover from failures
  - Adding QoS is hard

Virtual Circuit Forwarding

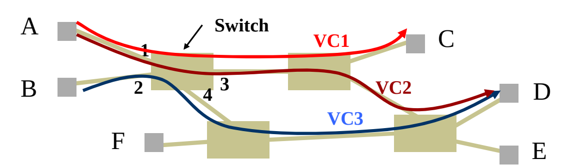

Address	Next Hop
VC1	3
VC2	3
VC3	4
VC4	?
VC5	?

- Address used for look up is a virtual circuit identifier (VC id)
- Forwarding table entries are filled in during signaling
- VC id is often shorter than destination address

8

VC versus Packets: Control over Path

Payload | VCI   Payload | Dst

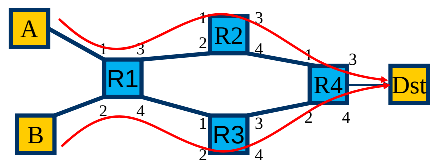

R1 packet forwarding table:

Dst R2

R1 VC table:

VC 1 R2

VC 2 R3

Different paths to same destination!

(useful for traffic engineering!)

How to Pick a VC Id?

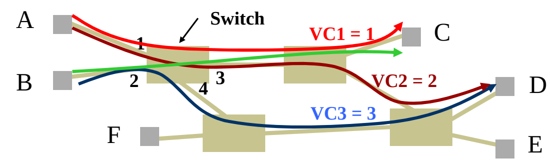

- When B establishes green virtual circuit, how does it know what VC ids are available?
- Even worse: every VC id may already be used on a link along the path to the destination
- Solution: VC id swapping

VC id Swapping

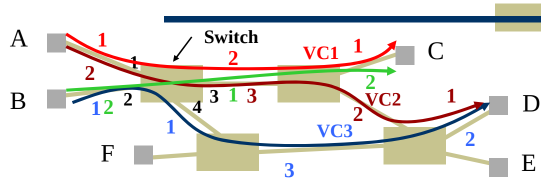

Address	Next Hop	Next id
VC1 = 1	3	2
VC2 = 2	3	3
VC3 = 1	4	1
VC4 = 2	3	1

- Look up is based on VC id in header + incoming port number
- Forwarding table specifies outgoing port and new VC id
- VC id conflicts can be resolved locally during signaling

Connections and Signaling

- Permanent vs. switched virtual connections (PVC/SVC)
  - PVCs last "a long time"
    - E.g., connect two bank locations with a PVC or for traffic engineering
    - PVCs administratively configured
  - SVCs dynamically set up on a "per-call" basis
    - A bit more like a phone call
    - Significant signaling overhead for each connection
- Topology
  - point to point, point to multipoint, multipoint to multipoint

SVC Connection Setup

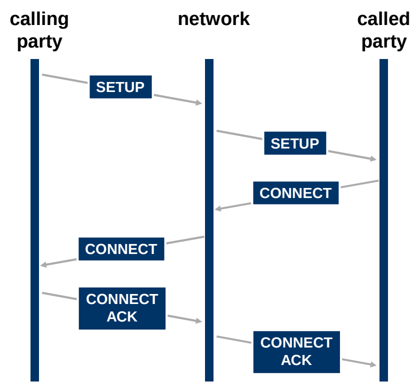

Virtual Circuits In Practice

- Older tech: Frame Relay
  - Only provided PVCs. Used for quasi-dedicated 56k/T1 links between offices, etc.
- Asynchronous Transfer Mode - ATM: Telco approach
  - Kitchen sink: voice, file transfer, video, etc.
  - Intended as IP replacement. That didn't happen. :)
  - Today: rarely used.
- MPLS: The "IP Heads" answer to ATM
  - Stole good ideas from ATM
  - Integrates well with IP
  - Today: Used inside many transit networks to provide traffic engineering, VPN support, simplify core.

Outline

- Circuit switching refresher
- Virtual Circuits - general
  - Why virtual circuits?
  - How virtual circuits? -- tag switching!
- Two modern implementations
  - ATM - teleco-style virtual circuits
  - MPLS - IP-style virtual circuits
- Virtual LANs
  - How do they differ?

ATM Background

- Telephone companies supported voice telephony: 4 kHz analog, 64 kbs digital
- They provided lines for data networking
  - ISDN: 64 kbps and faster channels
  - T1 (1.544 Mbps)
  - T3 (44.736 Mbps)
- Wanted to become the primary service provider for data networking services
  - Internet traffic, voice, multimedia, ..
- ATM is based on small, fixed sized cells
  - Voice needs small "packets" – reduces latency
  - Efficient technology: simple buffer management, fast forwarding based on flat VCIDs
  - Initial version ran at 155 Mbs (versus 10 Mbs Ethernet)

ATM Features

- Fixed size cells (53 bytes)
  - 5 byte header, 48 byte payload
- Virtual circuit technology using hierarchical virtual circuits (VP,VC).
- Support for multiple traffic classes by adaptation layer.
  - E.g. voice channels, data traffic
- Elaborate signaling stack.
  - Backwards compatible with respect to the telephone standards
- Standards defined by ATM Forum.
  - Organization of manufacturers, providers, users

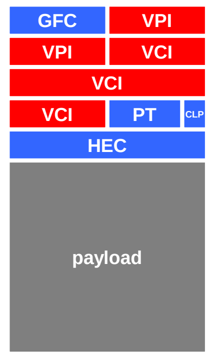

ATM Discussion

- Vision: ATM is a replacement for IP.
  - Could carry both traditional telephone traffic (constant bit rate circuits) and other traffic (data, variable BR)
  - Simple switching core: forwarding based on VC identifiers
  - Better than IP, since it supports QoS, traffic engineering
- But: the technology is very complex
  - Signaling software is very complex
  - Supporting connection-less service model on connection-based technology is painful
  - Deploying ATM in LAN is complex (e.g. broadcast)
  - With IP over ATM, a lot of functionality is replicated
- Reality: Too little benefit compare with IP, although traffic engineering benefits were attractive
  - Fast VCI lookup became less critical over time

MPLS – Multi-Protocol Label Switching

- Bringing virtual circuit concept into IP
- Driven by multiple forces
  - Traffic engineering
  - High performance forwarding
  - VPN
  - Quality of service

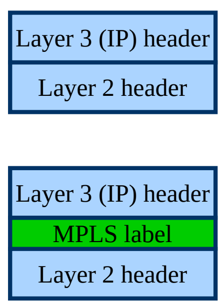

Some MPLS slides from H. Zhang

Label Switched Path (LSP)

- Simplex path through interior network
- Packets forwarded by label-switched routers(LSR)
  - Performs LSP setup and MPLS packet forwarding
  - Label Edge Router (LER): LSP ingress or egress
  - Transit Router: swaps MPLS label, forwards packet

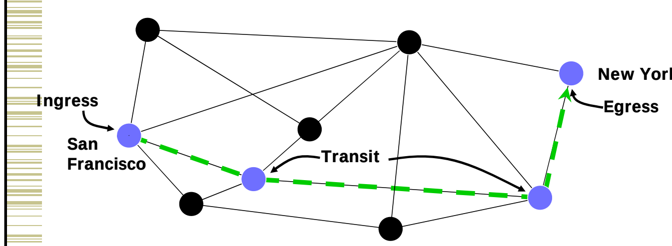

MPLS Header

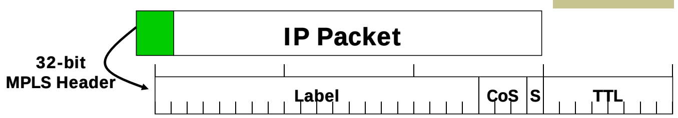

- IP packet is encapsulated in MPLS header
  - Label
  - Class of service
  - Stacking bit: if next header is an MPLS header
  - Time to live: decremented at each LSR, or pass through
- IP packet is restored at end of LSP by egress router
  - TTL is adjusted, transit LSP routers count towards the TTL
- MPLS is an optimization – does not affect IP semantics

Forwarding Equivalence Classes

FEC = “A subset of packets that are all treated the same way by a LSR”

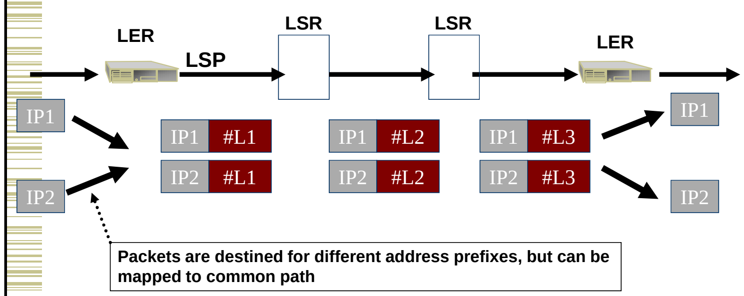

Establishing LSPs

- Use a Label Distribution Protocol (LDP) to establish paths based on IP forwarding tables
  - Simple
  - MPLS packets follow the same path as IP
  - But offers only limited benefits
- Explicitly route circuits
  - More work
  - Provides finer grain control over how traffic is distributed throughout the network
  - Important tool for traffic engineering

23

LSPs Driven by IP Routing

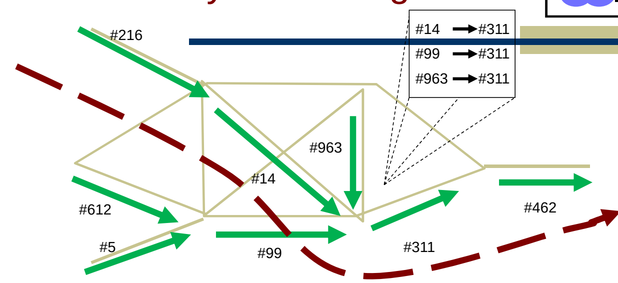

- A LSP is actually part of a tree from all sources to a destination (unidirectional).
- A control protocol (e.g. Label Distribution Protocol, LDP) builds the tree based on the IP forwarding tables.

MPLS Builds on Standard IP

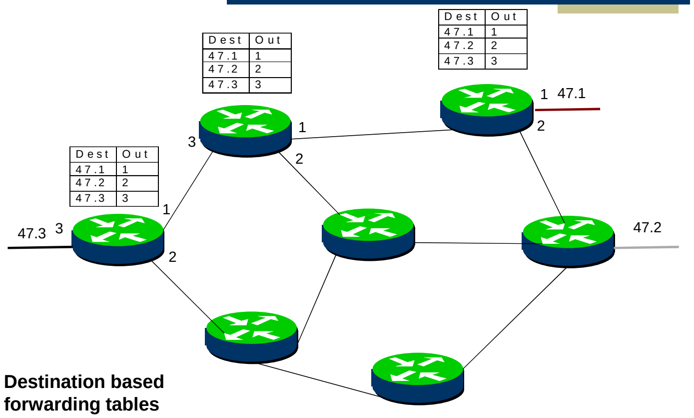

Destination based forwarding tables as built by OSPF, IS-IS, RIP, etc.

Label Switched Path (LSP)

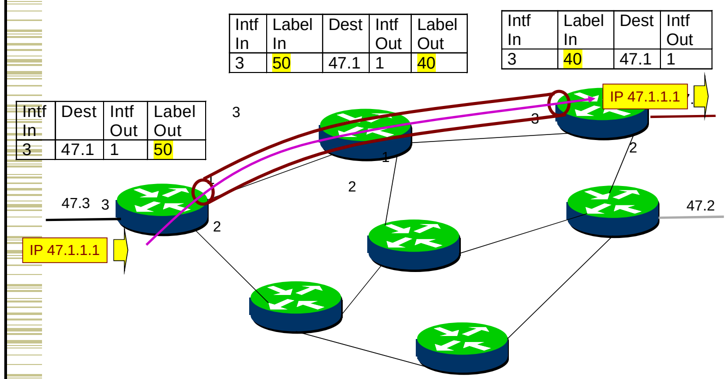

Explicitlyly Routed - ER-LSP

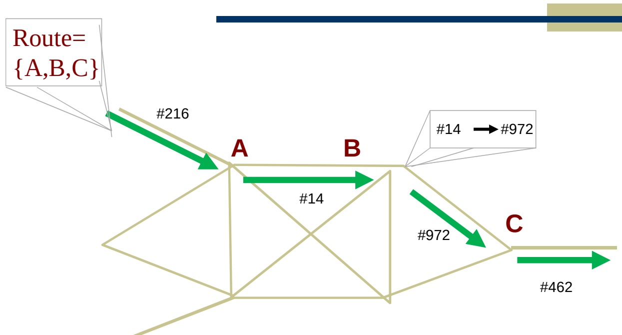

ER-LSP follows route that source chooses. In other words, the control message to establish the LSP (label request) is ***source routed***.

Explicitly Routed LSP - Example