Lecture 13
ATM, SONET, and MPLS

Peter Steenkiste
School of Computer Science
Carnegie Mellon University

15-441 Networking, Spring 2006
http://www.cs.cmu.edu/~prs/15-441

Outline

- Circuit switching.
- ATM overview.
- SONET.
- MPLS.

Packet Switching

- Source sends information as self-contained packets that have a destination address.
  - » Source may have to break up single message in multiple
- Each packet travels independently to the destination host.
  - » Routers and switches use the destination address in the packet to determine how to forward the packets
- Destination recreates the message.
- Analogy: a letter in surface mail.

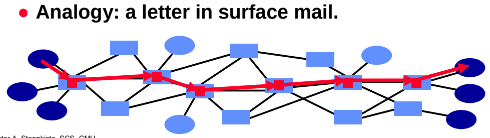

Circuit Switching

- 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 need to include the destination address with the data since the routers know the path
- The connection is torn down.
- Example: telephone network.

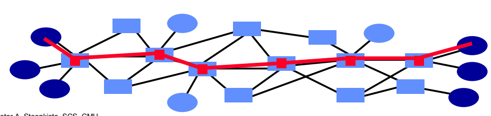

Connection Setup

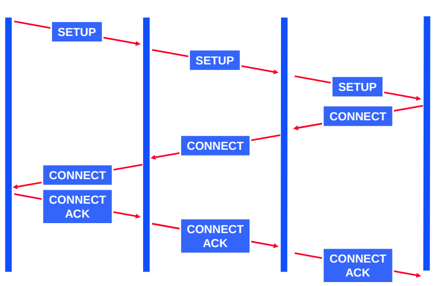

Circuit Switching Discussion

- Traditional circuits: on each hop, the circuit has a dedicated wire or slice of bandwidth.
  - » Physical connection - clearly no need to include addresses with the data
- Advantages, relative to packet switching:
  - » Implies guaranteed bandwidth, predictable performance
  - » Simple switch design: only remembers connection information, no destination address look up
- Disadvantages:
  - » Inefficient for bursty traffic (wastes bandwidth)
  - » Delay associated with establishing a circuit
- Can we get the advantages without (all) the disadvantages?

Virtual Circuits

- Each wire carries many "virtual" circuits.
  - » Forwarding based on virtual circuit (VC) identifier
  - » A path through the network is determined for each VC when the VC is established
  - » Use 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

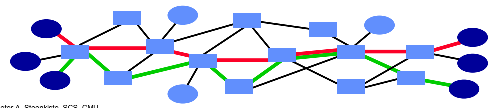

Packet Switching and Virtual Circuits: Similarities

- "Store and forward" communication based on an address.
  - » Address is either the destination address or a VC identifier
- Must have buffer space to temporarily store packets.
  - » E.g. multiple packets for some destination arrive simultaneously
- Multiplexing on a link is similar to time sharing.
  - » No reservations: multiplexing is statistical, i.e. packets are interleaved without a fixed pattern
  - » Reservations: some flows are guaranteed to get a certain number of "slots"

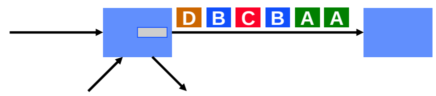

Virtual Circuits Versus Packet Switching

- Circuit switching:
  - » Uses short connection identifiers to forward packets
  - » Switches know about the connections so they can more easily implement features such as quality of service
  - » Virtual circuits form basis for traffic engineering: VC identifies long-lived stream of data that can be scheduled
  - » Requires "hard" state in the network
- 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
    - – Only have "soft" state
  - » Adding QoS is hard
  - » Traffic engineering is hard: too many packets!

How To Assign VC Ids?

- Easy solution: globally unique VC identifiers.
  - » But very large number of VCs makes this inefficient
    - – Requires large ids that have to be stored in header
    - – Centralized assignment of ids to avoid conflicts
- Better solution: use ids that have only local significance for a link.
  - » Smaller number of VCs on a link reduces the id space
  - » Allows local selection of VC identifiers
  - » But requires "VC id remapping"

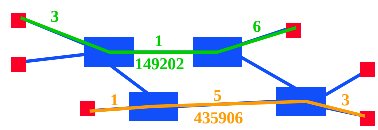

VC Packet Forwarding

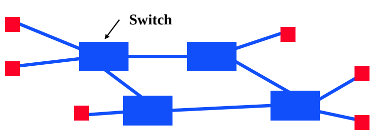

VC ID	Next Hop	Info
3	3	13
8	3	2
A21023C90590	0	-
128.2.15.3	1	(2,34)

- Lookup based on VC identifier.
  - » Short, fixed-sized ids
  - » Easier than MAC addresses (48 bits) or hierarchical IP addresses
- Next hop: output port for packet.
- Info: priority, VC id, ..
- Table is filled in by routing protocol.
  - » Similar to IP, but applies to setup request only

Outline

- Circuit switching.
- ATM overview.
- SONET.
- MPLS.

ATM History

- 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)
- They wanted to become the primary service provider for data networking services.
  - » file transfer: bursty, many Mbps peak
  - » database access: bursty, low latency
  - » Multimedia: synchronized
  - » Video: 6 MHz analog, 1.2-200 Mbps digital
- How?

ATM

- Asynchronous Transfer Mode
  - » Replace Synchronous Transfer Mode, which used slots in fixed frame structure (circuits)
- Instead of predefined TDM slots, tag each slot with a virtual connection ID.
  - » Bandwidth can change dynamically

- Small packets allow good real time behavior.
- Fixed sized packets (cells) support fast switching

ATM Features

- Fixed size cells (53 bytes).
- Virtual circuit technology using hierarchical virtual circuits (VP,VC).
- Physical layer processing delineates cells by frame structure, cell header error check.
- 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 Standard Protocol Layers

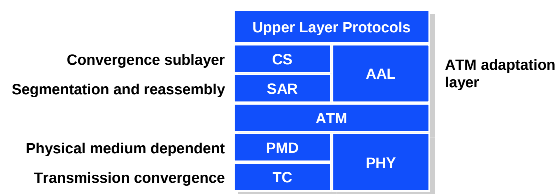

The ATM Cell (UNI)

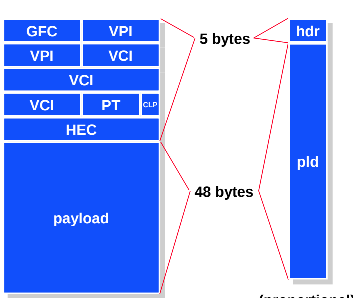

Why 53 Bytes?

- Small cells favored by voice applications
  - » delays of more than about 10 ms require echo cancellation
  - » each payload byte consumes 125 µs (8000 samples/sec)
- Large cells favored by data applications
  - » Five bytes of each cell are overhead
- France favored 32 bytes
  - » 32 bytes = 4 ms
  - » France is 3 ms wide

- USA, Australia favored 64 bytes
  - » 64 bytes = 8 ms
  - » USA is 16 ms wide

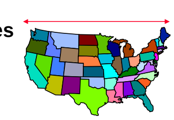

- Compromise

Virtual Circuit Switching: Label Swapping

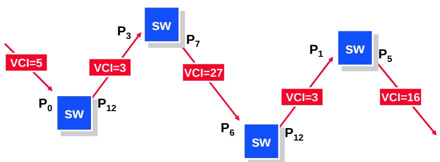

- Signaling establishes mapping from ($Port_{in}$, VCI_{in}) to ($Port_{out}$, VCI_{out}) at each switch on path.
  - » VCI remapping
- Cells in a VC arrive in order.

Virtual Paths

- Virtual path is a bundle of virtual circuits.
  - » VCs in a virtual path follow the same route
- Benefits:
  - » route and rerouting at the virtual path level
  - » fast connection set up
  - » bandwidth management

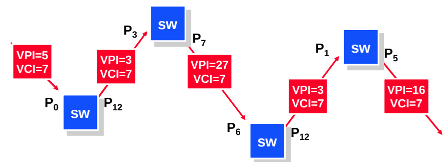

Virtual Path Trunking

- Allows aggregated resource management and fault recovery.

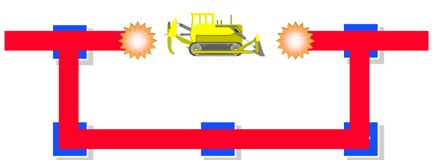

ATM Adaptation Layers

1	2	3	4	5
synchronous		asynchronous		
constant	variable bit rate			
connection-oriented			connectionless	

- AAL 1: audio, uncompressed video
- AAL 2: compressed video
- AAL 3: long term connections
- AAL 4/5: data traffic

AAL5 Adaptation Layer

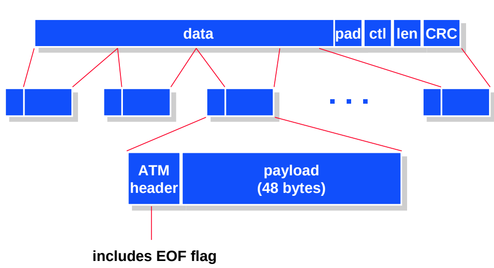

ATM Traffic Classes

- Constant Bit Rate (CBR) and Variable Bit Rate (VBR).
  - » Guaranteed traffic classes for different traffic types.
- Unspecified Bit Rate (UBR).
  - » Pure best effort with no help from the network
- Available Bit Rate (ABR).
  - » Best effort, but network provides support for congestion control and fairness
  - » Congestion control is based on explicit congestion notification
    - – Binary or multi-valued feedback
  - » Fairness is based on Max-Min Fair Sharing.

    (small demands are satisfied, unsatisfied demands share equally)

ABR: Max-Min Fair Sharing

- Flows are divided in two groups.
  - » Flows that are bottlenecked elsewhere
  - » Flows that are bottlenecked here
- The max-min fair share rate R_{fair} of a network link is defined such that
  - » Flows bottlenecked at the link have rate $r = R_{fair}$
  - » Flows bottlenecked elsewhere have rate r, where
    - – $r < R_{fair}$
    - – r is the max-main fair share rate of the bottleneck link
- Two implementations:
  - » Multi-valued feedback: switch returns rate
  - » Single bit feedback: use congestion bit in the header

Max-Min Fair Sharing Example

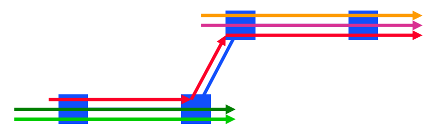

Connections and Signaling

- Permanent vs. switched virtual connections
  - » static vs. dynamic
  - » services often start with PVCs (Permanent Virtual Circuits)
- Call = bundle of connections, e.g. voice + video + data
- Topology
  - » point to point
  - » point to multipoint
  - » multipoint to multipoint
- Signaling VCs
  - » dedicated
  - » metasignaled, i.e. dynamically allocated

Q.SAAL: Signaling ATM Adaptation Layer

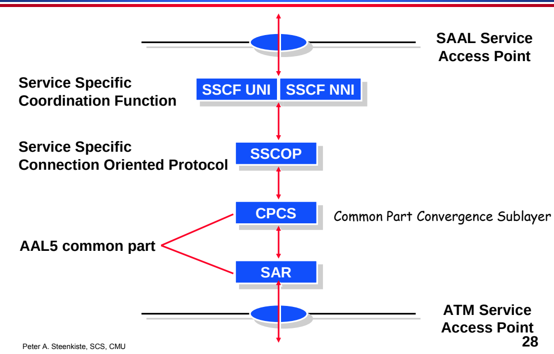

IP over ATM and SONET

- Many options!
- IP over ATM, with signaling support.
  - » I.e. establish new VC for each session
- IP over ATM, using statically configured ATM pipes.
  - » I.e. network manager sets up set of VCs
- IP over SONET (Packets over SONET).
- Differences in efficiency and flexibility in bandwidth management.

IP over ATM

- When sending IP packets over an ATM network, set up a VC to destination.
  - » ATM network can be end to end, or just a partial path
  - » ATM is just another link layer
- Virtual connections can be cached.
  - » After a packet has been sent, the VC is maintained so that later packets can be forwarded immediately
  - » VCs eventually times out
- Properties.
  - – Overhead of setting up VCs (delay for first packet)
  - – Complexity of managing a pool of VCs
  - + Flexible bandwidth management
  - + Can use ATM QoS support for individual connections (with appropriate signaling support)

LAN Emulation

- Motivation: making a non-broadcast technology work as a LAN.
  - » Focus on 802.x environments
- Approach: reuse the existing interfaces, but adapt implementation to ATM.
  - » MAC - ATM mapping
  - » multicast and broadcast
  - » bridging
  - » ARP
- Example: Address Resolution “Protocol” uses an ARP server instead of relying on broadcast.

IP over ATM Static VCs

- Establish a set of “ATM pipes” that defines connectivity between routers.
- Routers simply forward packets through the pipes.
  - » Each statically configured VC looks like a link
- Properties.
  - – Some ATM benefits are lost (per flow QoS)
  - + Flexible but static bandwidth management
  - + No set up overheads

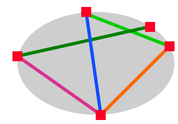

ATM Discussion

- At one point, ATM was viewed as a replacement for IP.
  - » Could carry both traditional telephone traffic (CBR circuits) and other traffic (data, VBR)
  - » Better than IP, since it supports QoS
- Complex technology.
  - » Switching core is fairly simple, but
  - » Support for different traffic classes
  - » Signaling software is very complex
  - » Technology did not match people’s experience with IP
    - – deploying ATM in LAN is complex (e.g. broadcast)
    - – supporting connection-less service model on connection-based technology
  - » With IP over ATM, a lot of functionality is replicated
- Currently used as a datalink layer supporting IP.

Outline

- Circuit switching.
- ATM overview.
- SONET.
- MPLS.

SONET

- SONET is the Synchronous Optical Network standard for data transport over optical fiber.
- One of the design goals was to be backwards compatible with many older telco standards.
- Beside minimal framing functionality, it provides many other functions:
  - » operation, administration and maintenance (OAM) communications
  - » synchronization
  - » multiplexing of lower rate signals
  - » multiplexing for higher rates

Standardization History

- Process was started by divestiture in 1984.
  - » Multiple telephone companies building their own infrastructure
- SONET concepts originally developed by Bellcore.
- First standardized by ANSI T1X1 group for the US.
- Later picked up by CCITT and developed its own version.
- SONET/SDH standards approved in 1988.

A Word about Data Rates

- Bandwidth of telephone channel is under 4KHz, so when digitizing:

  8000 samples/sec * 8 bits = 64Kbits/second
- Common data rates supported by telcos in North America:
  - » Modem: rate improved over the years
  - » T1/DS1: 24 voice channels plus 1 bit per sample

    (24 * 8 + 1) * 8000 = 1.544 Mbits/second
  - » T3/DS3: 28 T1 channels:

    7 * 4 * 1.544 = 44.736 Mbits/second

Synchronous Data Transfer

- Sender and receiver are always synchronized.
  - » Frame boundaries are recognized based on the clock
  - » No need to continuously look for special bit sequences
- SONET frames contain room for control and data.
  - » Data frame multiplexes bytes from many users
  - » Control provides information on data, management, …

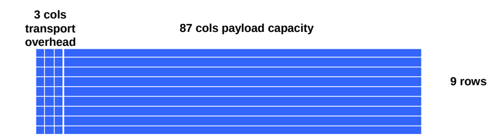

SONET Framing

- Base channel is STS-1 (Synchronous Transport System).
  - » Takes 125 μsec and corresponds to 51.84 Mbps
  - » 1 byte/frame corresponds to a 64 Kbs channel (voice)
  - » Also called OC-1 = optical carrier
- Standard ways of supporting slower and faster channels.
  - » Support both old standards and future (higher) data rates
- Actual payload frame "floats" in the synchronous frame.
  - » Clocks on individual links do not have to be synchronized

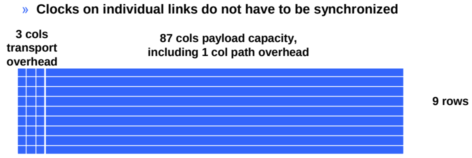

How Do We Support Lower Rates?

- 1 Byte in every consecutive frame corresponds to a 64 Kbit/second channel.
  - » 1 voice call.
- Higher bandwidth channels hold more bytes per frame.
  - » Multiples of 64 Kbit/second
- Channels have a "telecom" flavor - circuits.
  - » Fixed bandwidth
  - » Just data – no headers
  - » SONET multiplexers remember how bytes on one link should be mapped to bytes on the next link

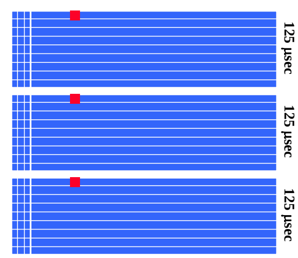

How Do We Support Higher Rates?

- Send multiple frames in a 125 μsec time slot.
- The properties of a channel using a single byte/ST-1 frame are maintained!
  - » Constant 64 Kbit/second rate
  - » Nice spacing of the byte samples
- Rates typically go up by a factor of 4.
- Two ways of doing interleaving.
  - » Frame interleaving
  - » Column interleaving
    - – concatenated version, i.e. OC-3c

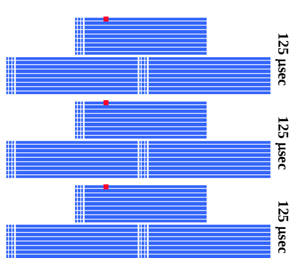

The SONET Signal Hierarchy

Signal Type	line rate	# of DS0
DS0 (POTS)	64 Kbs	1
DS1	1.544 Mbs	24
DS3	44.736 Mbs	672
OC-1	51.84 Mbs	672
OC-3	155 Mbs	2,016
OC-12	622 Mbs	8,064
OC-48	2.49 Gbs	32,256
OC-192	9.95 Gbs	129,024
OC-768	39.8 Gbs	516,096

Using SONET in Networks

Add-drop capability allows soft configuration of networks, usually managed manually.

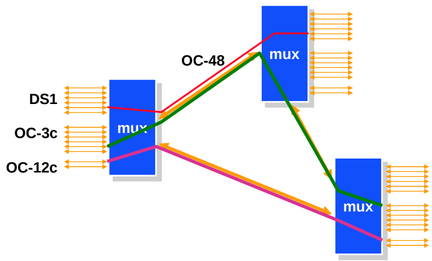

Self-Healing SONET Rings

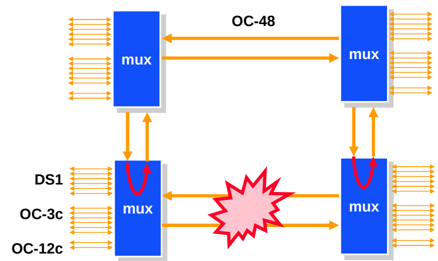

SONET as Physical Layer

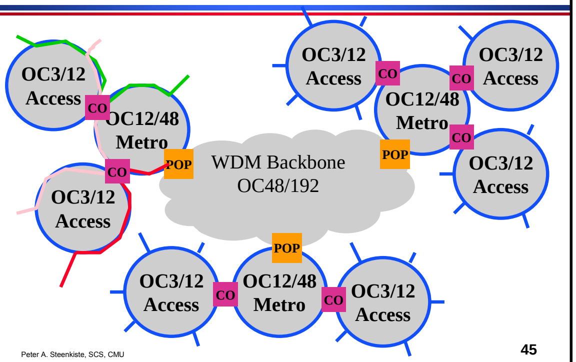

Packets over SONET

- Same as statically configured ATM pipes, but pipes are SONET channels.
- Properties.
  - Bandwidth management is much less flexible
  - Much lower transmission overhead (no ATM headers)

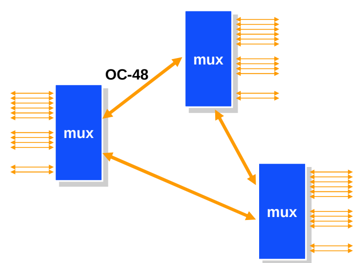

Outline

- Circuit switching.
- ATM overview.
- SONET.
- MPLS.

IP Switching

- How to use ATM hardware without the software.
  - ATM switches are very fast data switches
  - software adds overhead, cost
- The idea is to identify flows at the IP level and to create specific VCs to support these flows.
  - flows are identified on the fly by monitoring traffic
  - flow classification can use addresses, protocol types, ...
  - can distinguish based on destination, protocol, QoS
- Once established, data belonging to the flow bypasses level 3 routing.
  - never leaves the ATM switch
- Interoperates fine with “regular” IP routers.
  - detects and collaborates with neighboring IP switches

IP Switching Example

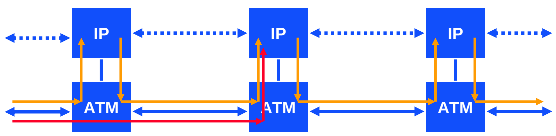

Peter A. Steenkiste, SCS, CMU

49

IP Switching Example

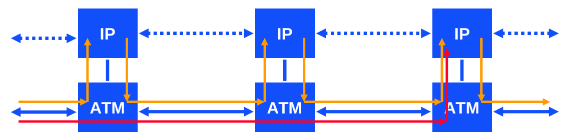

Peter A. Steenkiste, SCS, CMU

50

IP Switching Example

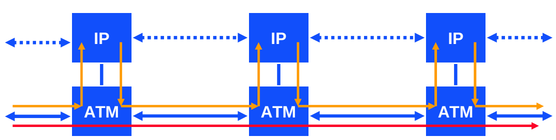

Peter A. Steenkiste, SCS, CMU

51

IP Switching Discussion

- IP switching selectively optimizes the forwarding of specific flows.
  - » Offloads work from the IP router, so for a given size router, a less powerful forwarding engine can be used
  - » Each data unit carries two addresses: IP and fast path
  - » Can fall back on traditional IP forwarding if there are failures
- IP switching couples a router with an ATM switching using the GSMP protocol.
  - » General Switch Management Protocol
- IP switching can be used for flows with different granularity.
  - » Flows belonging to an application .. Organization
  - » Controlled by the classifier
- Introduces a notion of flows/connections in IP.

Peter A. Steenkiste, SCS, CMU

52

An Alternative: Tag Switching

- Instead of monitoring traffic to identify flows to optimize, use routing information to guide the creation of “switched” paths.
  - » Switched paths are set up as a side effect of filling in forwarding tables
- Generalize to other types of hardware.
- Also introduced stackable tags.
  - » Made it possible to temporarily merge flows and to demultiplex them without doing an IP route lookup
  - » Requires variable size field for tag

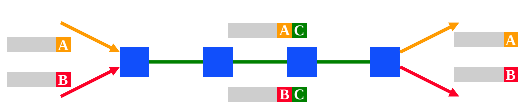

Peter A. Steenkiste, SCS, CMU

53

IP Switching versus Tag Switching

- Flows versus routes.
  - » tags explicitly cover groups of routes
  - » tag bindings set up as part of route establishment
  - » flows in IP switching are driven by traffic and detected by “filters”
    - – Supports both fine grain application flows and coarser grain flow groups
- Stackable tags.
  - » provides more flexibility
- Generality
  - » IP switching focuses on ATM
  - » not clear that this is a fundamental difference

Peter A. Steenkiste, SCS, CMU

54

Multi-Protocol Label Switching MPLS

- Map packet onto Forward Equivalence Class (FEC) based on its header.
  - » Simple case: longest prefix match of destination address
  - » More complex if QoS of policy routing is used
- In MPLS, a label is associated with the packet when it enters the network and forwarding is based on the label in the network core.
  - » Label is swapped (as ATM VCIs)
- Potential advantages.
  - » Packet forwarding can be faster
  - » Routing can be based on ingress router and port
  - » Can use more complex routing decisions
  - » Can force packets to followed a pinned route

MPLS Mechanisms

- Implementation of the label is technology specific.
  - » Could be ATM VCI or an extra header
- Label Distribution Protocols distributes information on label/FEC bindings.
  - » Extensions of existing protocols (routing, RSVP) or stand-alone protocols
  - » Can be upstream or downstream
- Supports stacked labels.

MPLS Discussion

- Original motivation for MPLS was efficiency and per-flow QoS.
  - » Efficiency: tag lookup instead of IP address lookup
  - » QoS: use ATM QoS support for IP traffic
- Today: traffic engineering.
  - » Manage bandwidth distribution in the network
- Idea: use MPLS to force traffic over certain paths through the network.
  - » Not possible in IP network using traditional packet forwarding
  - » Ingress router decides which “Label Switched” path to use
  - » Many uses: load balancing, VPNs, QoS, ...