ECSE-6660: Broadband Networks

Homework 3

Please Submit Online in the WebCT dropbox

Deadline : 29th March (non-tape-delayed)

April 3rd (tape-delayed)

Rensselaer Polytechnic Institute

Shivkumar Kalyanaraman

Reading Assignment & Quick Questions

Reading assignments count for a substantial part of homework credit

Carefully review slide sets 5,6,7,8,9; Read Chapter 4,8,9,14,15 of S. Keshav's book, and Chap 6, 10, 13 of Ramaswami/Sivarajan's book.

Then answer the following quick true/false questions that test your knowledge. Please submit the electronic version of this powerpoint file with your answers. (Cut-and-paste the tick (√) over the appropriate boxes on the left)

[92 questions; 10/9 points per question (upto two mistakes ignored)]

T F

☐ ☐ Address resolution is a mapping (internetworking) problem that occurs in the data-plane (I.e. as packets are processed)

☐ ☐ IP-over-ATM and LANE approaches primarily use server-based approach to handle address resolution.

☐ ☐ Mapping OSPF over ATM involves in part dealing with adjacency issues (for LSA flooding, hello maintainance, and dijkstra computation)

☐ ☐ The NBMA and point-to-multipoint OSPF models are equivalent in the context of ATM networks

☐ ☐ Designated routers are not necessary in the point-to-multipoint model

☐ ☐ IP lookup involves longest-prefix-match whereas switching involves exact label match

- MPLS is aimed at small and medium-sized enterprise networks
- Ipsilon's IP switching is an example of using IP data plane with ATM control plane
- A key drawback of the Ipsilon model was its data-driven model of setting up VCs
- Tag switching and MPLS use a control-driven model of VC-setup; MPLS offers more flexibility in the choice of control-plane methods.
- MPLS is a layer 2 protocol in the eyes of IP, even though it provides "layer-3-like" functions.
- Label-stacking is a mechanism to set up nested tunnels in an MPLS network
- G-MPLS is a generalization of MPLS to optical networks where the "label" could mean an optical wavelength or TDM time-slot
- MPLS uses the shim label concept in mapping to ATM and frame-relay
- A vanilla LSP is actually part of a sink-tree rooted at the destination
- An explicitly routed LSP is one where the labels are assigned starting from the destination node, and moving backwards towards the source
- RSVP, though originally developed for int-serv, has now been extended to support MPLS signaling (I.e. LSP setup)
- A central problem (related to traffic engineering) in connectionless routing protocols today is that they are limited to a single path.
- The coupling of traffic mapping to route availability means that when we desire to remap traffic, we have to also incur the routing dynamics to update routes

- Forwarding equivalence classes (FEC) pre-defines the route for the traffic aggregate
- The MPLS traffic engineering approach may be easily extended across area and domain boundaries
- The explicit route object (ERO) in RSVP allows for strict as well as loose-source-routes
- Constrained LSPs essentially allow the incorporation of resource optimization and QoS goals to be incorporated into the routing selection process
- MPLS tunnels are lightweight and opaque to L3-traffic
- BANANAS offers a method of emulating the explicit routing feature of MPLS without signaling, by defining a label with global semantics for the PathID
- BANANAS forwarding is a mix of traditional IP-like longest-prefix-match plus an MPLS-like exact-match forwarding
- BANANAS-TE is restricted to only a single area or a single domain.
- QoS is meaningful only when the the service is specified a priori and quantitatively
- The essential ideas in QoS are to offer some degree of service isolation and service differentiation, both of which are not possible with just a FIFO queuing discipline
- The delay conservation law holds only for the FIFO scheduling scheme
- QoS involves two zero-sum games (assuming constant capacity): capacity and delay
- Traffic conditioning is an example of a control-plane QoS mechanism
- QoS routing and traffic engineering signaling are examples of data-plane QoS schemes
- The progress of end-to-end QoS mechanisms is coupled with that of network-based QoS mechanisms
- A leaky or token bucket ensures that incoming traffic conforms to a linear-bounded traffic envelope specification

- A service curve can intersect and cross over the arrival curve
- The horizontal distance between the arrival and service curves for a given bit indicates the queuing delay for that bit
- The effect of sending a non-conforming packet into the network with a bit set (eg: DE bit) is exactly the same as if the packet were delayed at the edge till it obtained a token
- Priority dropping is an example of a scheduling strategy
- Round robin is an example of a buffer management strategy
- Unlike priority scheduling, round-robin offers protection (I.e. isolation) for flows against the misbehavior of other flows.
- The GPS scheduler offers proportional fairness among flows
- Weighted fair queuing is an exact bit-by-bit emulation of GPS scheduling
- Virtual time measure the service that a backlogged flow with weight of 1 would receive in a GPS schedule
- The Parekh-Gallagher (P-G) theorem offers rate and delay bounds assuming token-bucket shapers at the edge and per-flow WFQ schedulers at the core
- Int-serv represents a stateless architecture that achieves the QoS using P-G theorem
- The real challenge in stateful approaches is the configuration, establishment and maintainence of fine-grained state
- Diff-serv represents a stateless architecture that achieves a broad range of QoS.
- Per-hop behavior (PHB) is a generalization of scheduling, buffer management and forwarding support for QoS at internal nodes.
- AF forwarding can be used to emulate leased line behavior

- Dynamic packet state (DPS) approaches to QoS simplify the core network (I.e. minimal state, and some computation) but use state in packets that may change at every hop
- Unequal error protection (UEP) is a method of getting service differentiation at the application level even though the network offers an undifferentiated best-effort service
- RTP follows the application-layer framing philosophy and hence can be extended by the the application-level protocols.
- RTP specifies buffer management and coding mechanisms
- H.323 specifies unequal error control and buffer management mechanisms.
- The key problem in inter-domain QoS is that ISPs do not have a clear incentive and lack mechanisms to coordinate the management of traffic crossing their boundaries
- Closed-loop building blocks for QoS can provide service differentiation in packet-by-packet time-scales.
- Content delivery networks represent an application-level approach to offering QoS.
- The key economic factors in high-speed routers are speed, space, power and interface cost.
- Most routers today are output queued routers
- Both DRAM and Moore's law trends underperform the demand growth for router speeds
- Software routers that support multiple (> 2) Gigabit Ethernet ports can be built on 1 Ghz Pentium platforms.
- Routers have seen a trend towards more specialized processing power in NICs
- Cross-bar is the simplest example of a time-division switch
- Time-slot interchangers with demultiplexors and multiplexors form a switch
- In a multi-stage Banyan fabric, switching elements are shared
- Layer 2 lookups involve longest-prefix match and trie data structures

- Longest-prefix match can be modeled as a two-dimensional lookup problem (in the space of prefix length and prefix values)
- Forwarding minimum sized packets at 10 Gbps gives a lookup time budget of about 100 ns
- Multiway tries waste memory but are efficient to lookup (I.e. they reduce access times)
- A vast majority of prefixes are less than or equal to 24 bits in size
- Fourth generation switches have the interfaces placed on separate racks and connected to a single rack that contains the switching fabric
- Call-blocking minimization is the goal in large circuit switches.
- Large CO switches use a combination of time- and space-division switching techniques
- Output buffered switches suffer from the head-of-line (HOL) blocking problem
- Banyan is a self-routing recursive fabric.
- The Batcher sorting stage is placed in front of a banyan to make it non-blocking
- With head-of-line blocking, the switch cannot operate beyond a load of about 58.6%
- Virtual output queueing (VOQ) is a method used to handle head-of-line blocking
- Packet classification is a multi-dimensional version of the longest-prefix-match lookup problem
- Network processors attempt to achieve the programmability of general purpose processors, and also the speed and specialization of ASICs
- Methods used in network processors include multi-processors, hardware threading, pipelining, and offering a general purpose processor for control/slow-path operations
- Given a key, a ternary CAM finds all the entries matching the key, subject to a mask.
- Requirements like multicast and in-order packet delivery place hurdles to achieve multi-tera-bit router capacities.

- Five nines implies 50 minutes of downtime per year
- Most SONET topologies use linear protection
- In 1+1 dedicated protection, traffic is sent both on the working and protect lines simultaneously
- UPSR allows 1:n survivability schemes
- The UPSR and BLSR schemes use the APS feature in SONET for their signaling functions
- BLSR uses 2 or 4 fibers and allows bi-directional transmission, and m:n protection schemes
- UPSR rings are employed in the access topologies and BLSR in metro/core topologies
- Squelching is done only in the case of node-failures
- Fast re-route and mesh restoration operate at longer time-scales but offer more efficiency in the protection architecture.