

# ECSE-6600: Internet Protocols

## Informal Quiz #14

Shivkumar Kalyanaraman:  
GOOGLE: "Shiv RPI"  
shivkuma@ecse.rpi.edu



**High Speed Router Design  
(Slide set #16):  
Informal Quiz**

# Router Design

T F

- The trie data structure for IP forwarding lookup facilitates binary search in terms of prefix length.
- An overwhelmingly large number of prefixes in the global routing mesh are between 16 and 24 bits long
- Larger routers require lesser space, power and number of interfaces for the same capacity
- Processing and memory access become bottlenecks in the data plane of high speed routers
- Current PC software routers can easily handle multiple Gb/s interfaces at line rates simultaneously
- Aggregate packet forwarding rates are usually measured using the smallest packet sizes (eg: 64 bytes)
- The goal of network processors is to provide speeds approaching that of ASICs, but with programmability
- IP options are usually handled in the slow processing path
- Bridges and L2 switches use a simple hashing procedure for lookup.
- Longest prefix match is equivalent to a two-dimensional search problem
- At OC-192c rates, the time available for looking up a 64 byte packet is in the order of 100s of nanoseconds.
- A trie is a binary tree.

T F

# Router Design

- 24-bit prefixes are the most common prefixes in global routing tables
- Very few prefixes in global routing tables are larger than 24 bits
- First generation IP routers used switched backplanes to interconnect network interface cards
- Faster routers use asymmetric multi-processing and parallelism to push specialized computing functions to NICs and by using co-processors (ASICs)
- Call blocking in circuit switches leads to failure to connect calls, whereas blocking in packet switches leads to packet buffering
- Crossbars can arbitrarily scale in terms of number of cross-points and management of crosspoints
- Modern voice switches combine both space-division and time-division switching
- Multi-stage cross-bar designs save on crosspoints but may increase blocking probability
- Banyan switch fabric is an example of a self-routing recursive fabric
- Blocking in Banyan switches can be avoided by choosing the order in which packets appear at the input ports (I.e. using a batcher sorting stage)
- Output queueing is usually not used because it requires prohibitive maximum speeds of memory access or processing.
- Output queues face the head-of-line blocking problem

T F

# Router Design

- Virtual output queues are used at the output ports of routers
- The knockout principle distributes the pain of losing packets across inputs, thus facilitating lower speed output queues
- Packet classification is a generalization of the longest-prefix-match problem
- Micro-engines in network processors run their own operating systems
- Micro-engines in network processors are ASICs
- The slow-path processing in network processors is handled by the StrongARM/Xscale processor
- A CAM (content-addressible memory) module returns all elements that match a given key
- T-CAMs are useful for packet classification
- Intel IXA network processors use a variety of register structures (eg: ring buffers, next-neighbor register) to ease packet processing
- Intel IXA network processors use specialized modules for key functions (eg: CRC, hashing, random number generation etc)
- Various types of parallelism and pipelining (both at the hardware and software level) is used in network processors
- Power constraints determine the layout of chips and boards in ultra-fast routers
- Power constraints also lead to multi-rack packaging for Tb/s routers
- Multi-rack packaging increases the switch round-trip time and hence buffering requirements
- Power constraints are likely to lead to increasing use of optical switching inside next-generation ultra-fast routers