CS-534: Packet Switch Architecture
Spring 2008

Department of Computer Science


# Exercise Set 4:
# Switch Generations, Cut-through

Assigned: Fri. 14 Mar. 2008 (week 4) -- Due: Wed. 19 Mar. 2008 (week 5)

## 4.1 PCI Bus in a First/Second Generation Router

An (old) PCI (or PCI-X) bus can be 32- or 64-bit wide, and can use a 33, or 67, or 100, or 133 MHz clock. Assume that each data transfer over PCI(-X) takes a fixed 4-clock-cycle overhead (arbitration, turn-around, framing, addressing, etc), plus 1 additional clock cycle for each (32- or 64-bit) data word being transferred; thus, a 1-word transfer takes 5 clock cycles, while an 8-word-burst transfer takes 12 clock cycles.

(a) What is the PCI throughput, in Mbps, in the following 4 cases?

- (i) 32-bit 67 MHz PCI, transferring individual words all the time;
- (ii) 64-bit 133 MHz PCI, transferring individual words all the time;
- (iii) 32-bit 67 MHz PCI, transferring 64-byte bursts all the time;
- (iv) 64-bit 133 MHz PCI, transferring 64-byte bursts all the time.

**(b)** If a *first* generation switch *without* DMA is built around a PCI bus, and the system bottleneck is the PCI bus, what would be the maximum switch throughput (= aggregate incoming throughput = aggregate outgoing throughput), in Mbps, in the four cases of part (a)? If this is a 4-port switch (4x4 switch), what is the maximum link rate?

**(c)** Same question as (b), but for a first generation switch *with* DMA.

**(d)** Same question as (b), but for a *second* generation switch.

## 4.2 Cut-through when Port Rates differ

Consider a switch with some OC-12 links and some gigabit Ethernet links. The OC-12 (actually "OC-12c") links carry ATM traffic, with a peak throughput equal to what you calculated in exercise 2.2. The gigabit Ethernet links carry packets at the rate indicated in exercise 2.3. We wish the switch to provide cut-through, but, because of rate mismatches among its ports, cut-through transmission cannot always start "right away"; we wish to calculate the worst such required delay.

Consider the situation where a packet destined to a gigabit Ethernet output arrives through an OC-12 input. The maximum-sized Ethernet packet is 1518 bytes long, and, when arriving through an ATM port, it arrives segmented into 32 ATM cells (1518 bytes / 48 payload-bytes/cell = 31.63 cells). The first 1518 bytes in the (48-byte) payloads of the 32 cells are precisely the 1518 bytes that constitute the Ethernet packet. Assume that we are guaranteed that these 32 cells will always arrive back-to-back, without any idle cells in-between them (when ATM traffic is carried over SONET links, cells are always transmitted back-to-back, without any "spacing" between them other than the SONET overhead bytes; however, in general, not all cells need to be valid --idle cells can be freely injected into the SONET payload in the general case, but not in our case between same-packet cells).

How soon, in nanoseconds, after the arrival of the first payload byte of the first ATM cell of the above 32-cell train can we transmit the first byte of the 8-byte gigabit Ethernet preamble? Obviously, right after the 8-byte preamble is transmitted, the 1518 bytes of the Ethernet packet must be transmitted, back-to-back, at the gigabit Ethernet rate, without any "hiccups". Our switch circuits are able to transmit any given byte within 100 ns at the earliest, after that byte is received at an input port.

---

Up to the Home Page of CS-534


Last updated: 14 Mar. 2008, by M. Katevenis.