

Exercise Set 1: Internal Blocking, Q'ing Arch. Throughput

Assigned: Mon. 2 Feb. 2009 (week 1) - Due: Wed. 11 Feb. 2009 (week 2)


1.1 Introductory Examples on Internal Blocking

Consider the two versions of 4x4 networks shown in the figure, below, made of 2x2 switching elements. All boxes shown are 2x2 switching elements with **no** internal blocking, each. The 4+4=8 I/O links (the links connected to the ports P0 through P3) have a transmission capacity of 1, each; similarly, all internal links have a capacity of 1, each.

(a1) Network (a) is called a *banyan*. Is this a *single-path* or a *multi-path* network, i.e. are there input-output port pairs (flows) for which multiple paths through the network exist, or is it that a single possible path exists for any packet, given its input and output ports?

(a2) Banyan networks, (a), exhibit *internal blocking*: Show a *feasible* (for the overall network, as observed "from the outside") traffic pattern (Chapter 1, slide 7) that **cannot** be routed by this network, due to internal blocking (Chapter 1, slide 8).

(a3) The fact that internal blocking exists for *some* traffic patterns does **not** mean that *all* ("demanding") traffic patterns will encounter internal blocking. Show a non-trivial traffic pattern for the banyan network (a) that does not encounter internal blocking; to qualify as "non-trivial", make this traffic pattern such that all (external) network ports are *fully loaded*, i.e. the sum of all flow rates at each port is 1.



(b1) Network (b) consists of two banyans connected back-to-back, with their middle stages merged; this is a *Benes* network. Is this a single-path or multi-path network? How many different routes are there from a given input port to a given output port, for the different input-output port pairs (flows)?

(b2) Try finding a feasible traffic pattern that cannot be sustained by this network. You will find none, because Benes networks are internally **non-blocking**. Show some "difficult" traffic patterns (e.g. some for which the banyan (a) exhibited internal blocking) and how the Benes network can route them.

(b3) When all packets of a flow follow a same, single path, show that flows in the Benes network (b) may need *rerouting* in order for the non-blocking property to be satisfied. Consider, for simplicity, flows of rate 1 each; in this case, there is a single flow per input port and a single flow per output port (flows among identically numbered ports are allowed). Start setting up flows, one by one; in some cases, the flows that have already been established will dictate the routing choice for the new flow. Give a sequence of such flow set-up's and routing choices such that when a request for a new flow arrives (between two non-busy ports), this new flow can only be satisfied if some previously established flow(s) are rerouted, i.e. previous routing decisions are revised. We call these *rearrangeably non-blocking* networks.

1.2 Queueing Architectures under Buffer Memory Throughput Limitations

Slide 19 of chapter 1 motivated us for the fact that limitations exist in buffer memory throughput. Accurate limitations will be discussed in next chapters of this course, but, for now, consider the following, simplified, set of limitations for a switch built inside a chip:

- Aggregate throughput of all input ports plus all output ports of the switch (i.e. aggregate off-chip bandwidth) cannot exceed 640 Gbits/s.
- No single buffer memory can be built with an I/O throughput (sum of input plus output throughput) exceeding 240 Gbits/s.
- The aggregate I/O throughput of *all* buffer memories on the chip cannot exceed 16 Tbits/s (16 Tera b/s).

Consider the four (4) queueing architectures seen in chapter 1 (e.g., consider all rows but the last, in the table on the last slide (#26) of chapter 1). Assume that external switch links have a throughput of 10 Gbits/s each (10 Gb/s per input port, plus 10 Gb/s per output port). Given the above throughput limitations, what is the *largest possible* switch that can be built inside this chip for these 4 different queueing architectures? Consider symmetric $N \times N$ switches, i.e. with equal numbers, N , of input and output ports, and find the maximum possible N for each architecture. *Optional additional question:* if we allow non-symmetric $M \times N$ switches, with different numbers of input and output ports, will we be able to increase the switch size for any of the above architectures?

[Up to the Home Page of CS-534](#)

© copyright University of Crete, Greece.
Last updated: 2 Feb. 2009, by [M. Katevenis](#).