CS8625

Multistage interconnection networks

CS8625 High Performance and Parallel Computing

Dr. Ken Hoganson

Class

Will

Start

Momentarily…

Multi-Stage Interconnects

- Multi-stage interconnection networks: are a class of networks that provide a switched connection path.
- Switching buses are expensive, so a completely connected architecture is not scalable.
- So connections are made in a fixed pattern that uses multiple stages of switches.
- Transfers will go through multiple switches from source to destination.
- Multiple stages of switches take additional time (but it is small, due to the speed of the hardware).
- Not all combinations of connections are possible, dut to conflicts at individual switches.

Dynamic and switched interconnection networks.

C,d,e, are multistage networks.

Cost of (a)?___________

Cost of (b)____________

Cost of (c) ____________

.

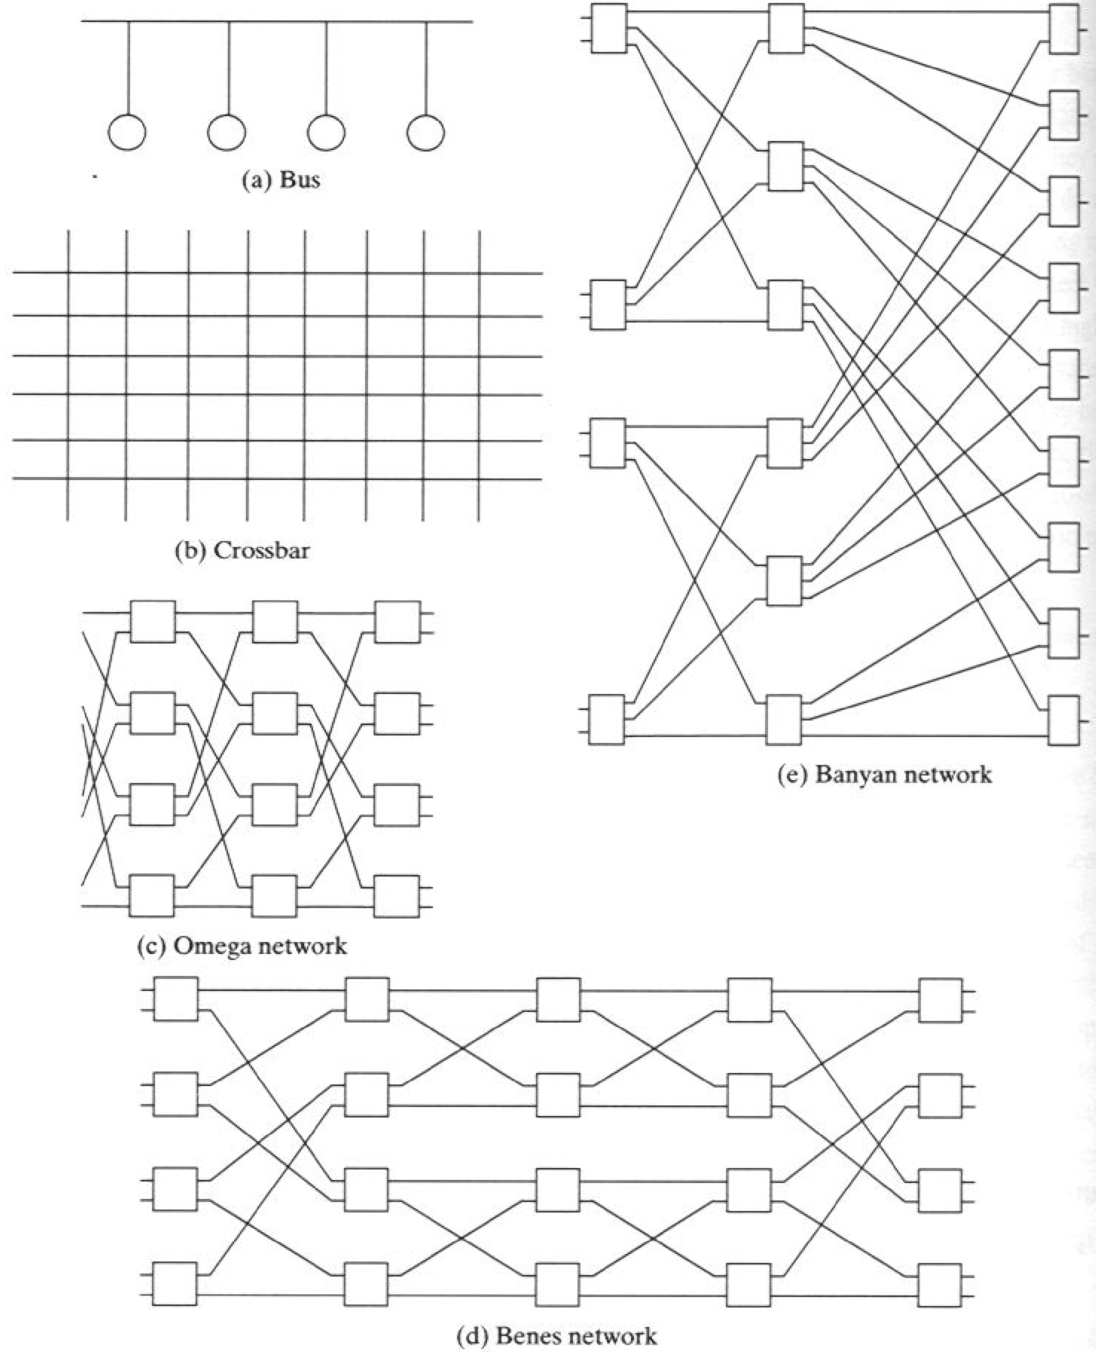

Figure 6-2 Dynamic interconnection networks.

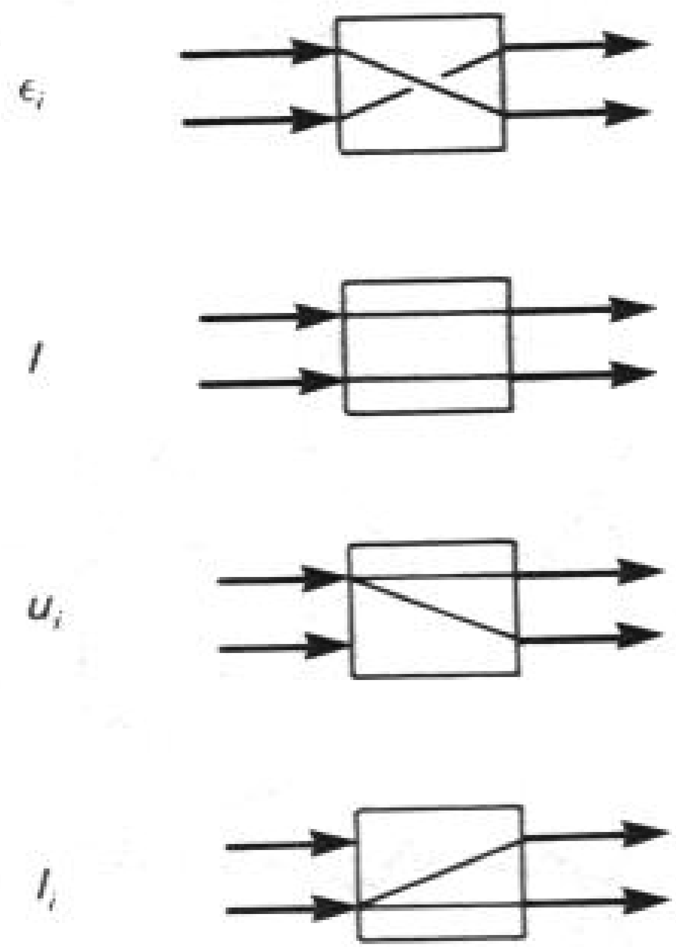
ε_i
I
u_i
l_i

X-bar switch too expensive.

Reduction (recursive) lowers cost.

Supports all N! permutations of inputs onto outputs.

..at lower cost, and only slightly longer latency

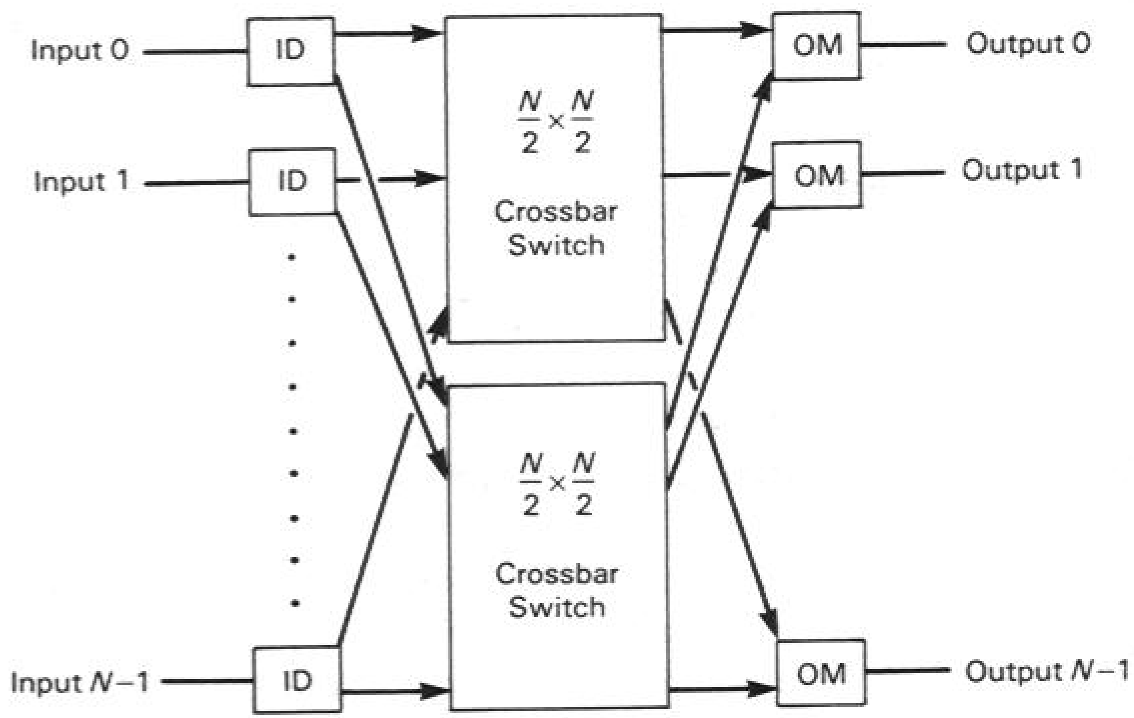

Figure 3.12 Beneš reduction of the cross-bar switch

Diagram ways to map 0→7

Cost:

$$\frac{n}{2}[2\log_2 N - 1]$$

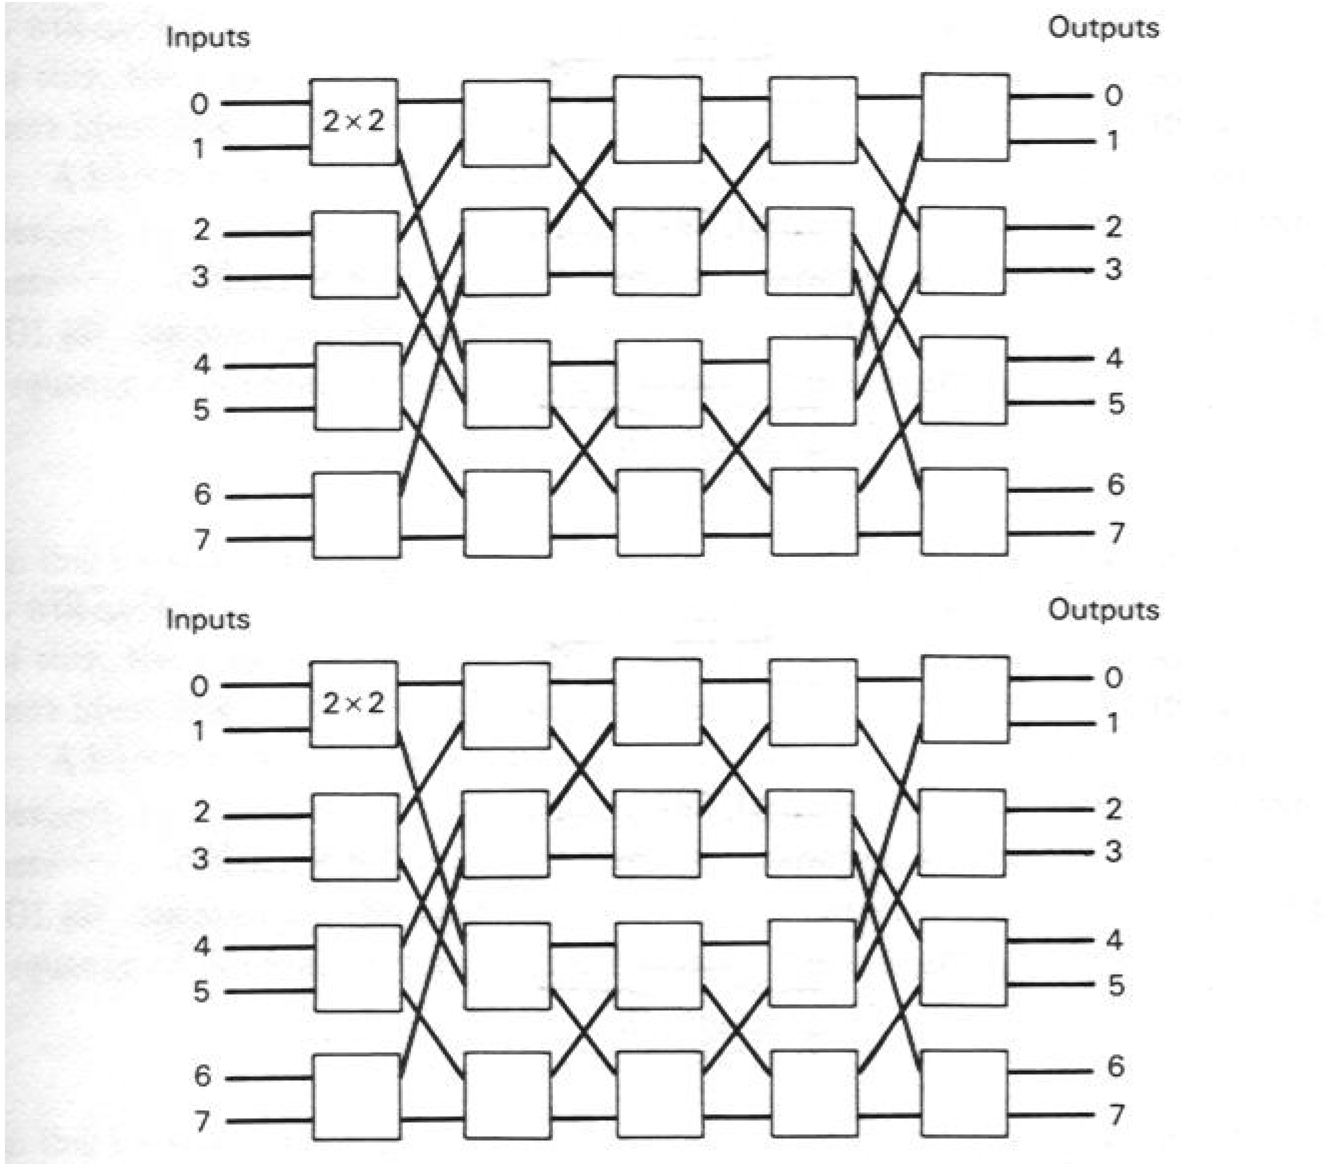

Figure 3.13 An 8-way Beneš network reduced to 2 × 2 cross-bar switches

Cost Comparison

- X-bar switch On^2
- Benes Network $$\frac{n}{2}[2\log_2 N - 1]$$

N=	**8**	**16**	**32**	**64**	**128**
Xbar:	64	256	1024	4096	16384
Benes:	20	56	144	352	832

Test a mapping:

0 → 7

1 → 5

2 → 4

3 → 0

4 → 1

5 → 6

6 → 3

7 → 2

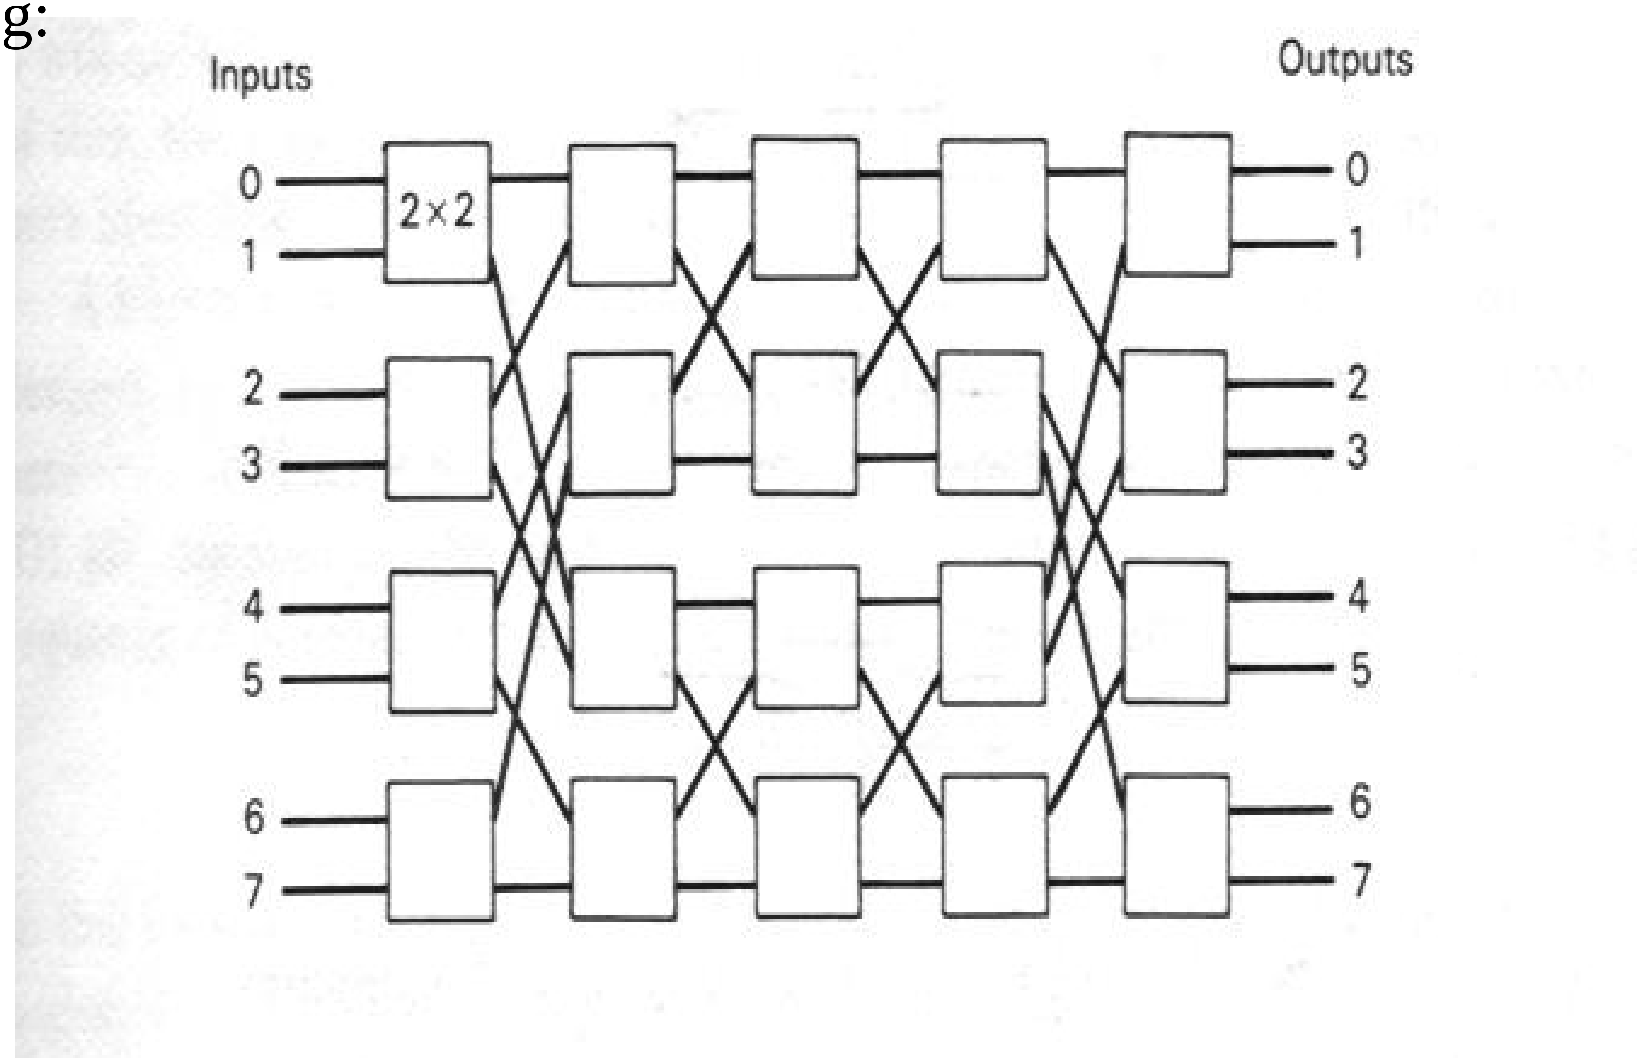

Figure 3.13 An 8-way Beneš network reduced to 2 × 2 cross-bar switches

- The design of the pattern of the switches is based on mathematics, that allows the switch settings to be described easily.
- Different settings of the switches (which define connections from source to destination) represent different "permutations", which can be mathematically described.

Cyclic Shifts

Cyclic shifts of 2^h places for $0 \le h < m$, where $N = 2^m$, can be done with $N\log_2 N$ switches and $\log_2 N$ control signals. Each source, S, has m switches connecting it to destinations, $(S + 2^h) \bmod N$, for $0 \le h < m$, see Figure 6-3. To maintain a readable diagram, output connections are only shown for nodes 0, 3, and 5. A shift of any length from 1 to N can be done by $\log_2 N$ or fewer of these shifts. The set of permutations for different values of h constitutes the *m-dimensional hypercube routing* functions that are supported by the hypercube networks.

For instance, connections at node 3 are :

If data at 3, needs to get to 1?

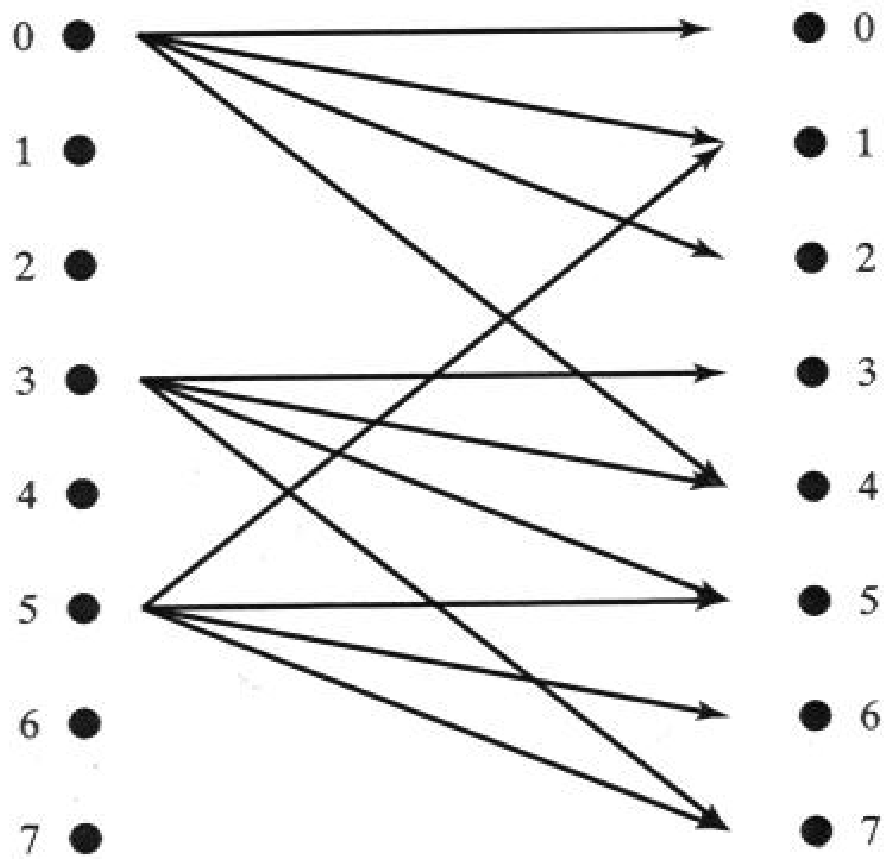

End of Lecture

End

Of

Today's

Lecture.

Blank