TSN: Lecture 8
Buffer Placement

Topics Covered

- Buffering
- Input buffering (input queueing)
- Dealing with HOL blocking
- Output queueing
- Buffered fabric

Buffering

- All packet switches need buffers to match input rate to service rate
  - or cause heavy packet loses
- Where should we place buffers?
  - input
  - in the fabric
  - output
  - shared

Input buffering (input queueing)

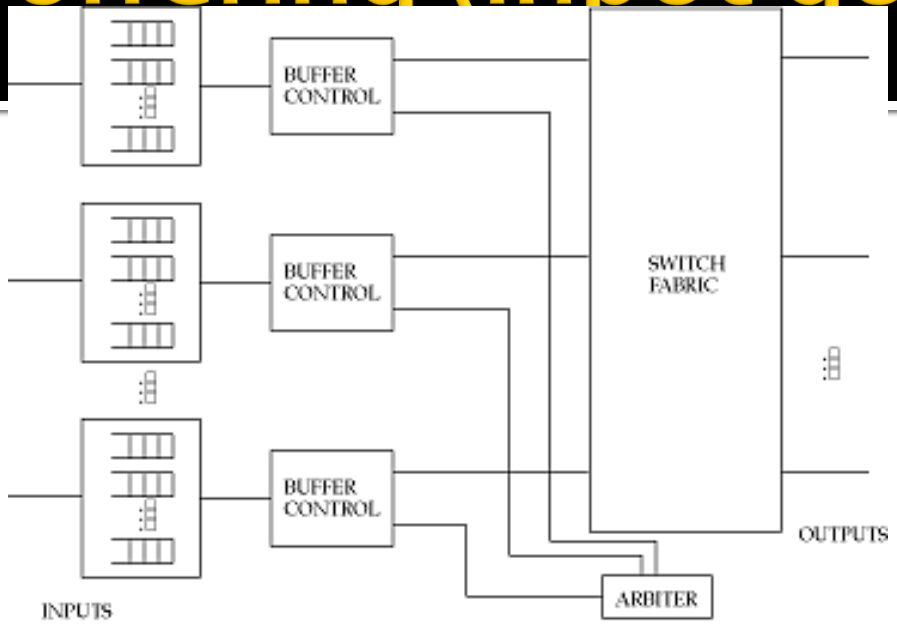

- No speedup in buffers or trunks (unlike output queued switch)
- Needs arbiter
- Problem: HOL (*head of line blocking)*
  - with randomly distributed packets, utilization at most 58.6%
  - worse with *hot spots*

Head of Line blocking

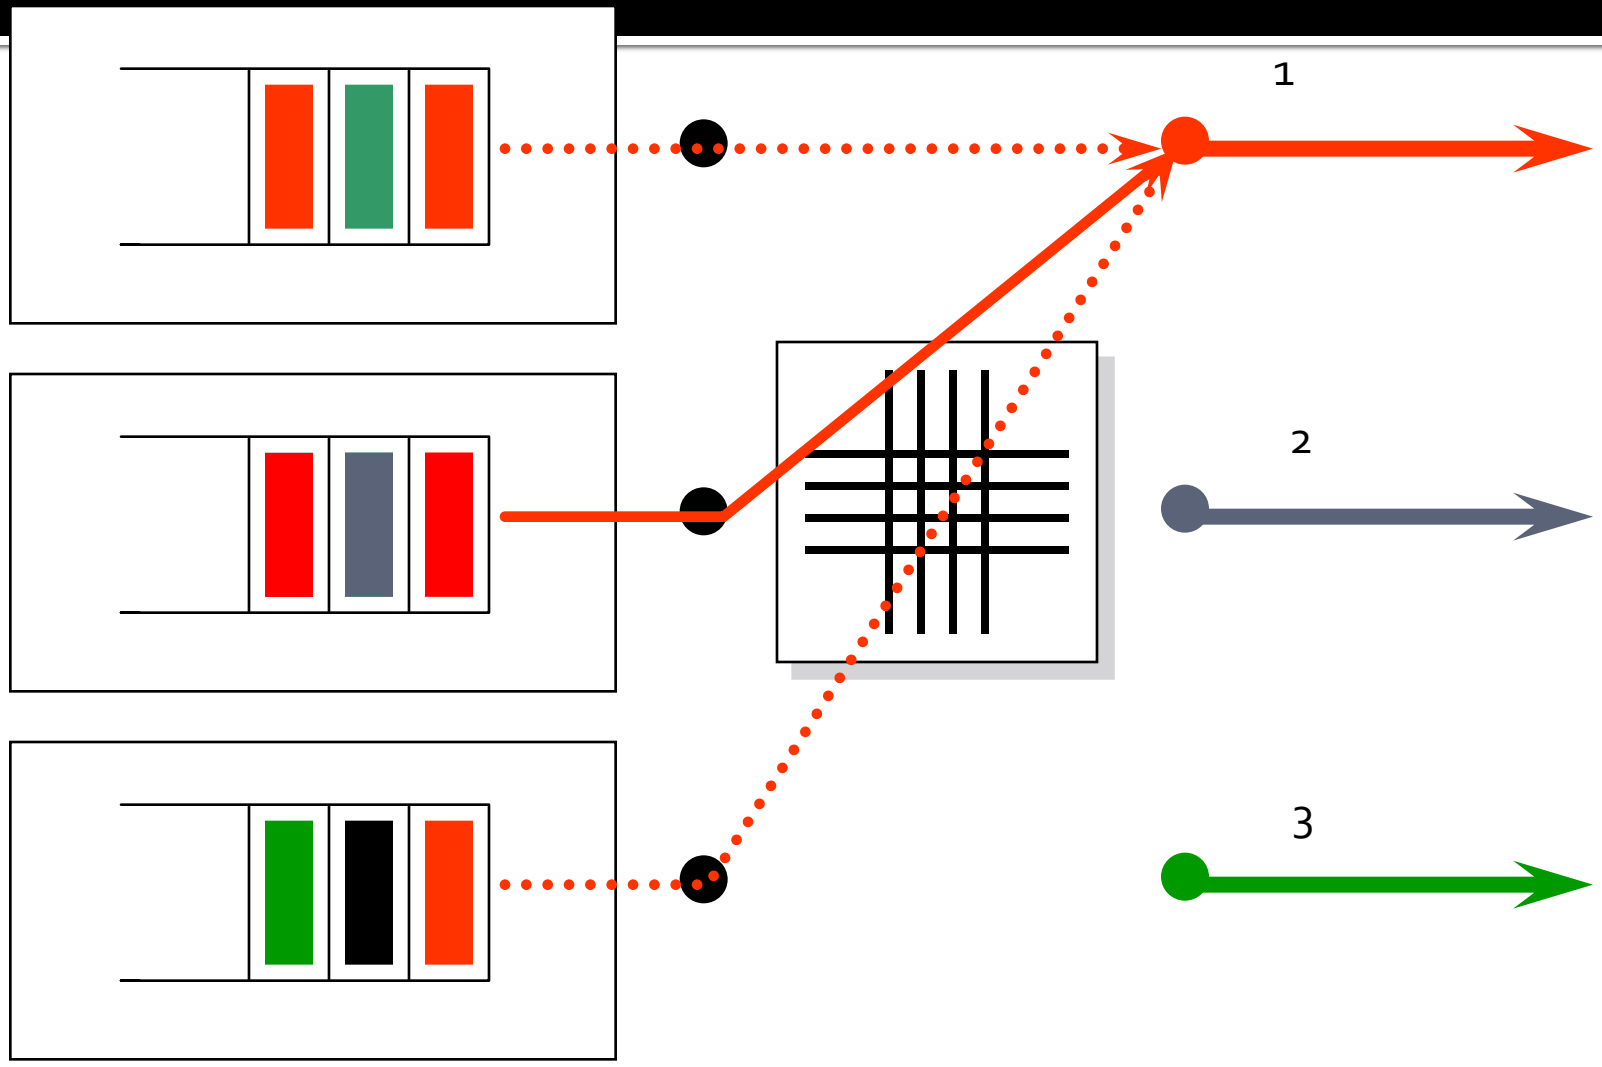

Dealing with HOL blocking

- Per-output queues at inputs (Virtual Input Queueing)
- Arbiter must choose one of the input ports for each output port
- How to select?
- Parallel Iterated Matching
  - inputs tell arbiter which outputs they are interested in
  - output selects one of the inputs
  - some inputs may get more than one *grant*, others may get none
  - if >1 grant, input picks one at random, and tells output
  - losing inputs and outputs try again
- Used in many large switches

Virtual Input (Output) Queueing

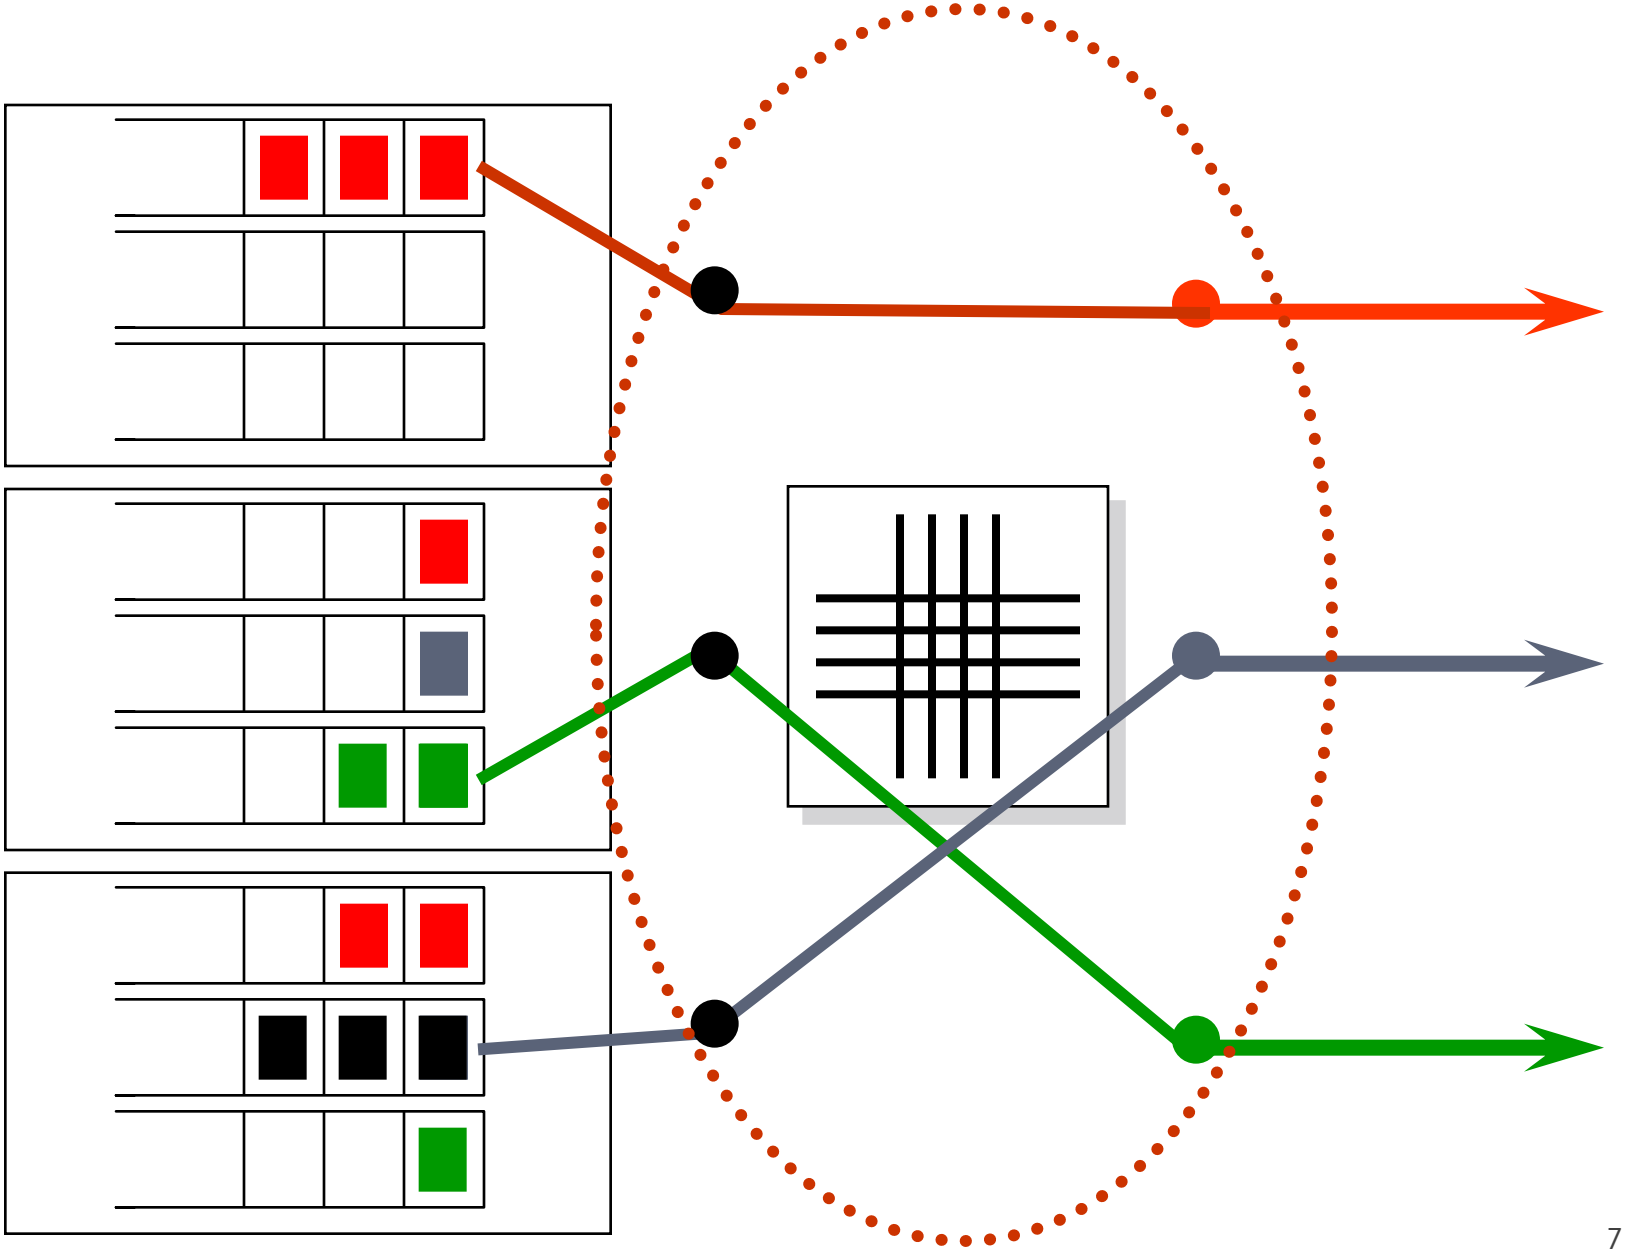

Output queueing

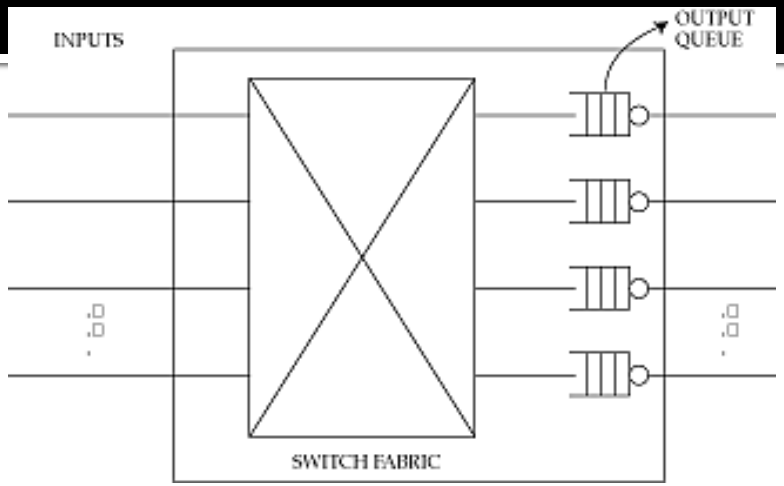

- Doesn't suffer from head-of-line blocking
- But output buffers need to run much faster than trunk speed (why?)
- Can reduce some of the cost by using the *knockout* principle
  - unlikely that all N inputs will have packets for the same output
  - drop extra packets, fairly distributing losses among inputs

Shared memory

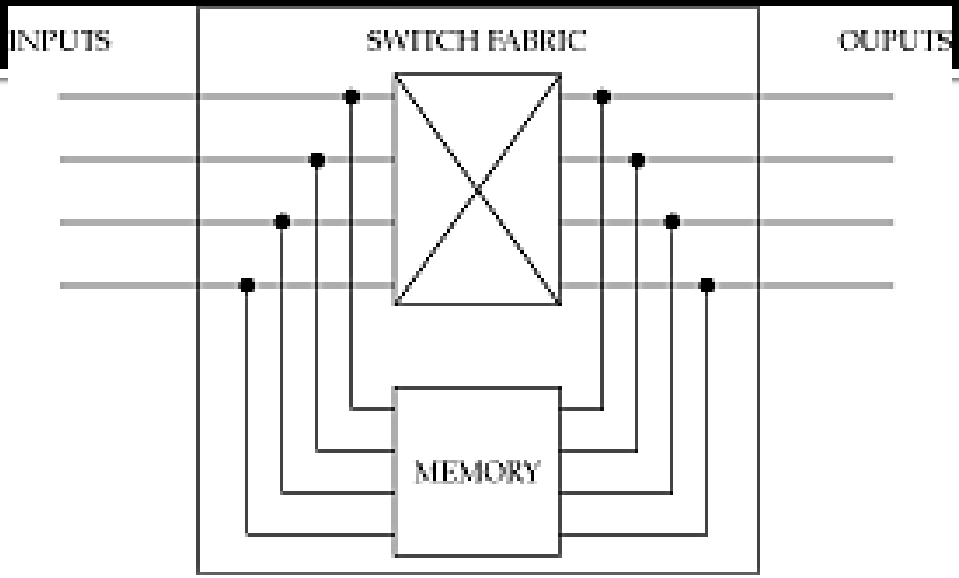

- Route only the header to output port
- Bottleneck is time taken to read and write multiported memory
- Doesn't scale to large switches
- But can form an element in a multistage switch

Buffered fabric

- Buffers in each switch element
- Pros
  - Speed up is only as much as fan-in
  - Hardware backpressure reduces buffer requirements
- Cons
  - costly (unless using single-chip switches)
  - scheduling is hard

Summary of Buffer Placement

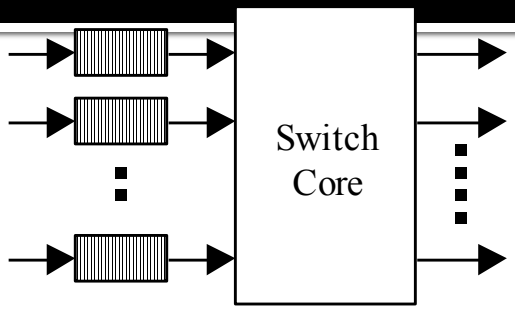

a) Input Queueing

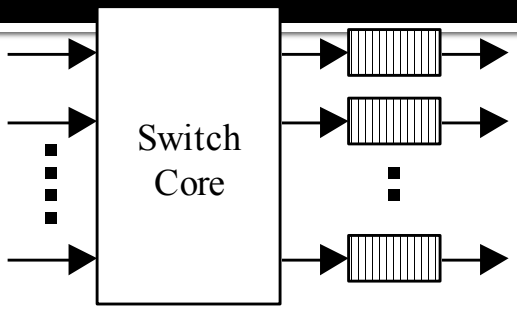

c) Output Queueing

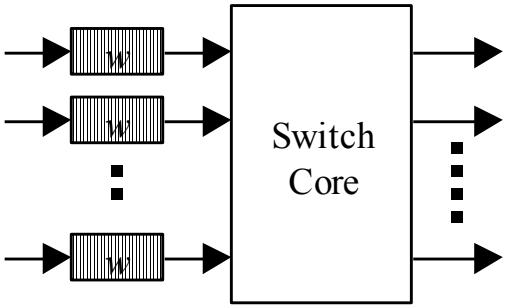

b) Window Selection

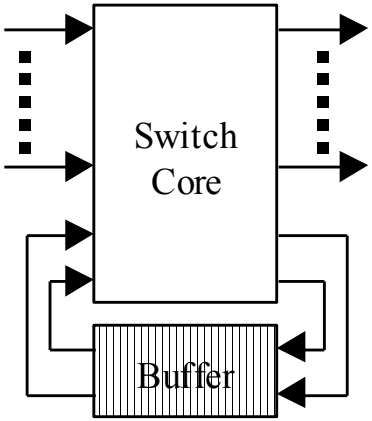

d) Shared Buffering