

Asynchronous Operation of Bufferless Crossbars

Georgios Passas and Manolis Katevenis

Foundation for Research and Technology (FORTH),
Inst. Of Computer Science (ICS), Greece

&

Computer Science Dept., Univ. of Crete, Greece


May 31, 2007

Topic

Switching of variable-size packets in crossbar switches

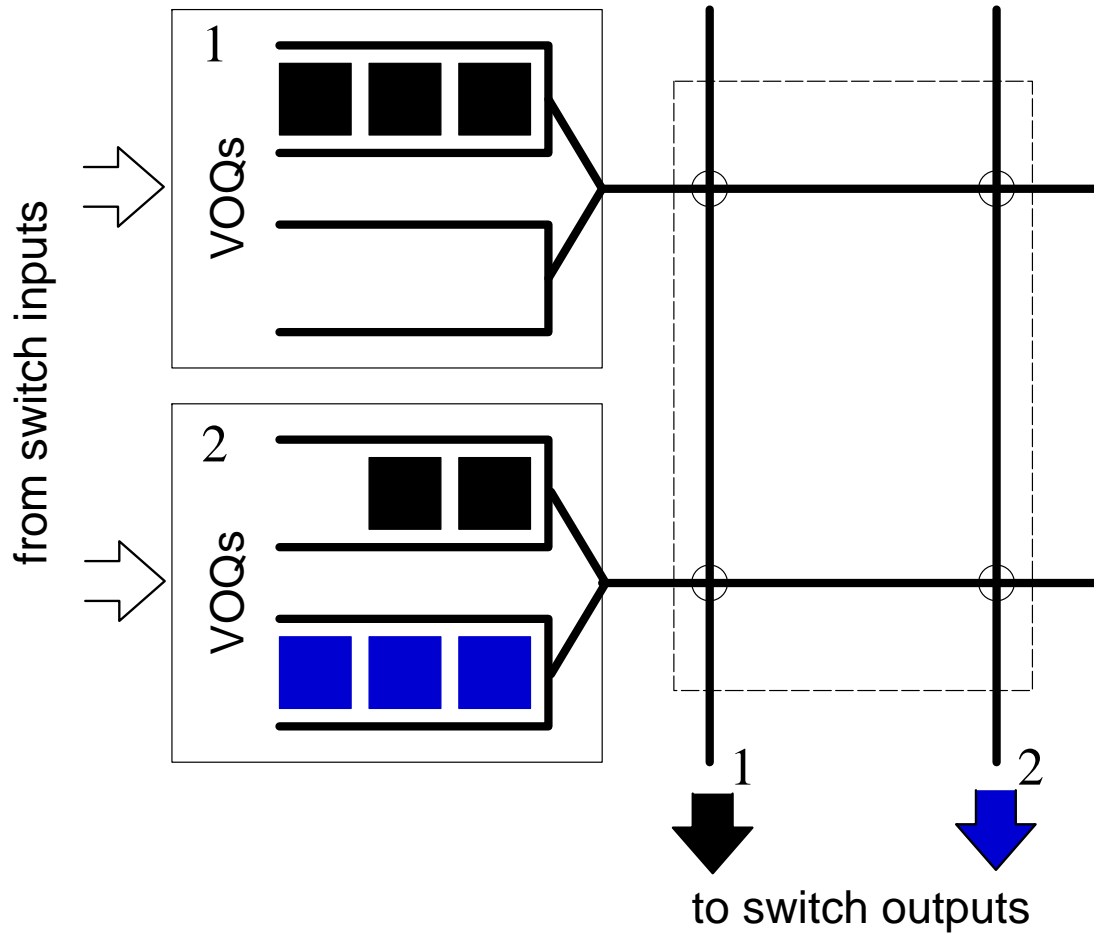
Outline

- Background
- Problem statement and Contribution
- Proposed operation and Results
- Summary




Example of a 2x2 VOQ switch with a bufferless crossbar

- VOQ eliminates HOL blocking
 - HOL blocking may limit switch throughput to 59% of link capacity
- Scheduling is needed to configure the crossbar
 - each input can be connected to at most one output and vice versa
 - bipartite graph matching problem
- Crossbar schedulers usually run a 3-phase matching algorithm

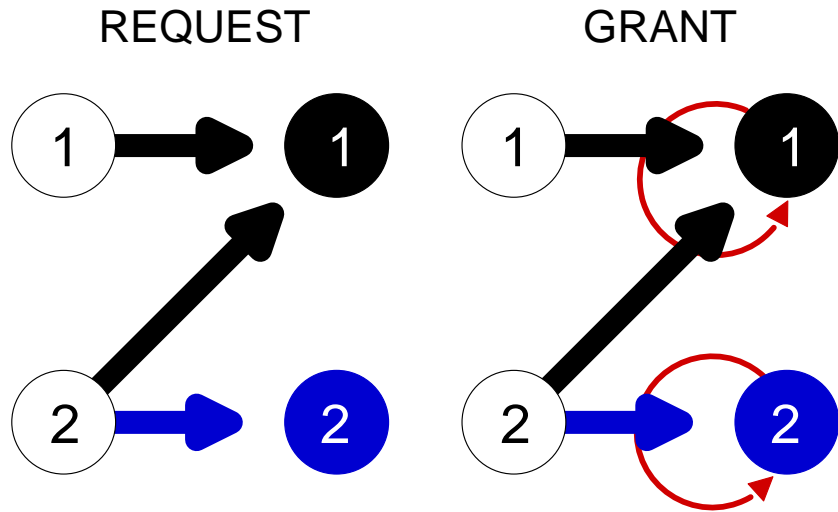


Traffic scenario

REQUEST




- Each input requests the outputs for which it has traffic
- Inputs submit requests in synchrony

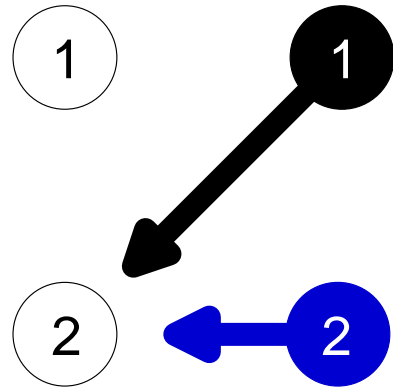


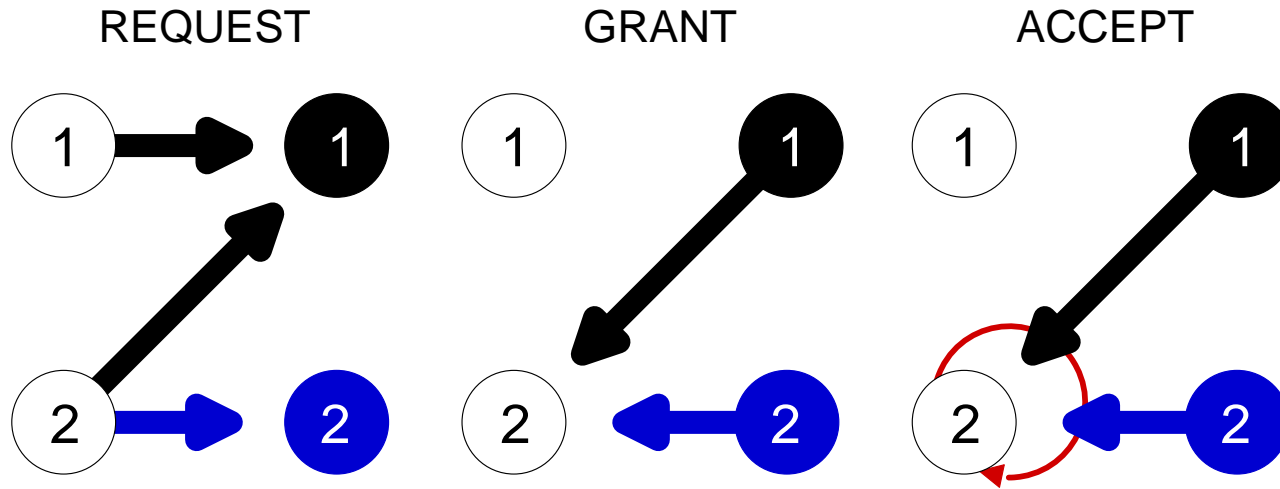
- Each output selects among requests the one to grant
- Independent but synchronized output decisions

REQUEST





GRANT





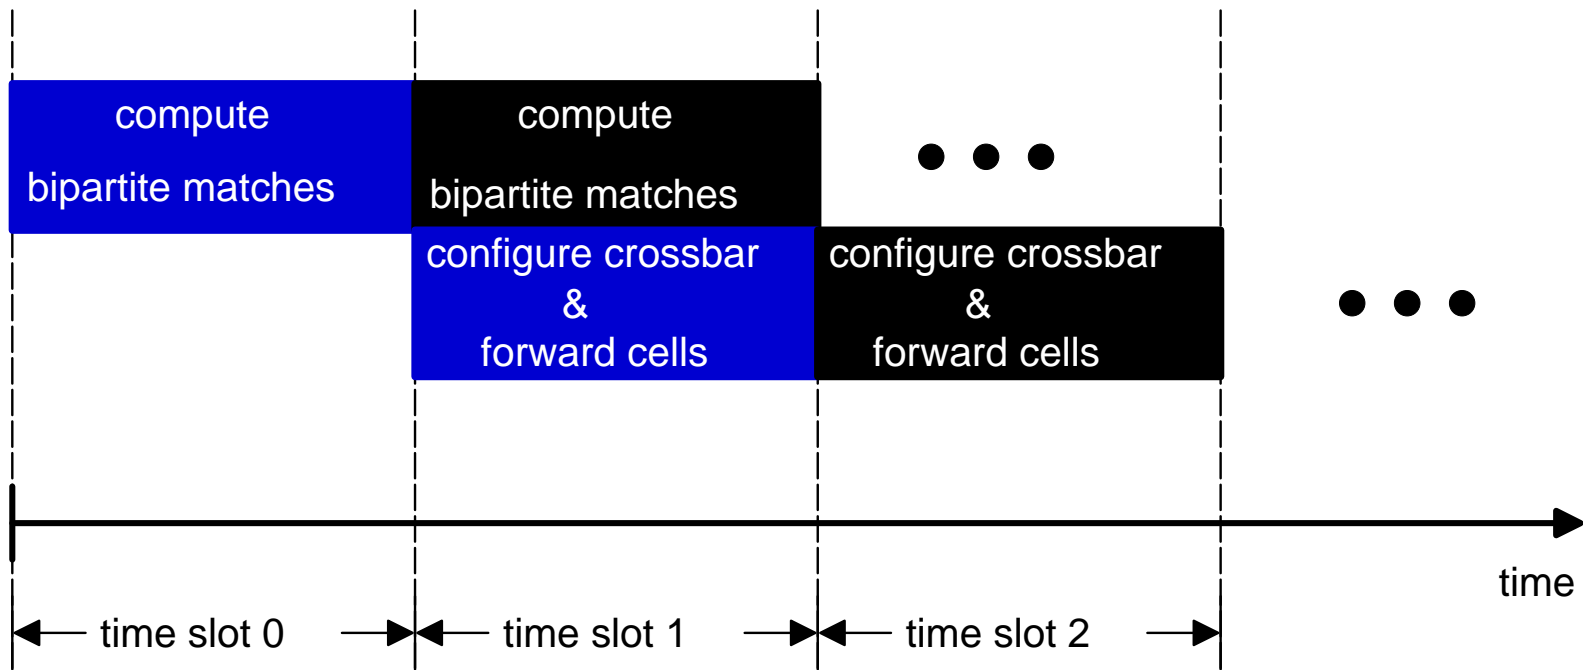
- Each input selects among grants the one to accept
- Independent but synchronized input decisions



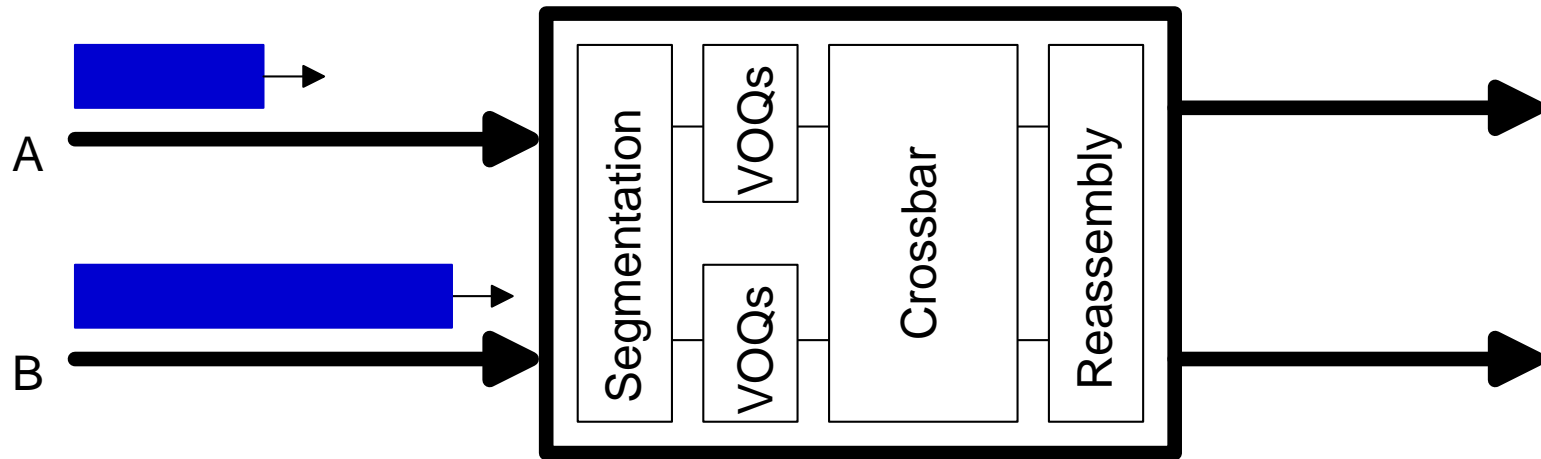


- Bipartite matches emerge after the accept phase is completed
- The crossbar is configured accordingly

Background: synchronous xbar reconfiguration



- Crossbar configuration changes synchronously
 - Configuration of time slot $n+1$ specified by matches computed in time slot n
- Operation on fixed-size internal cells
 - 1 time slot = 1 cell time

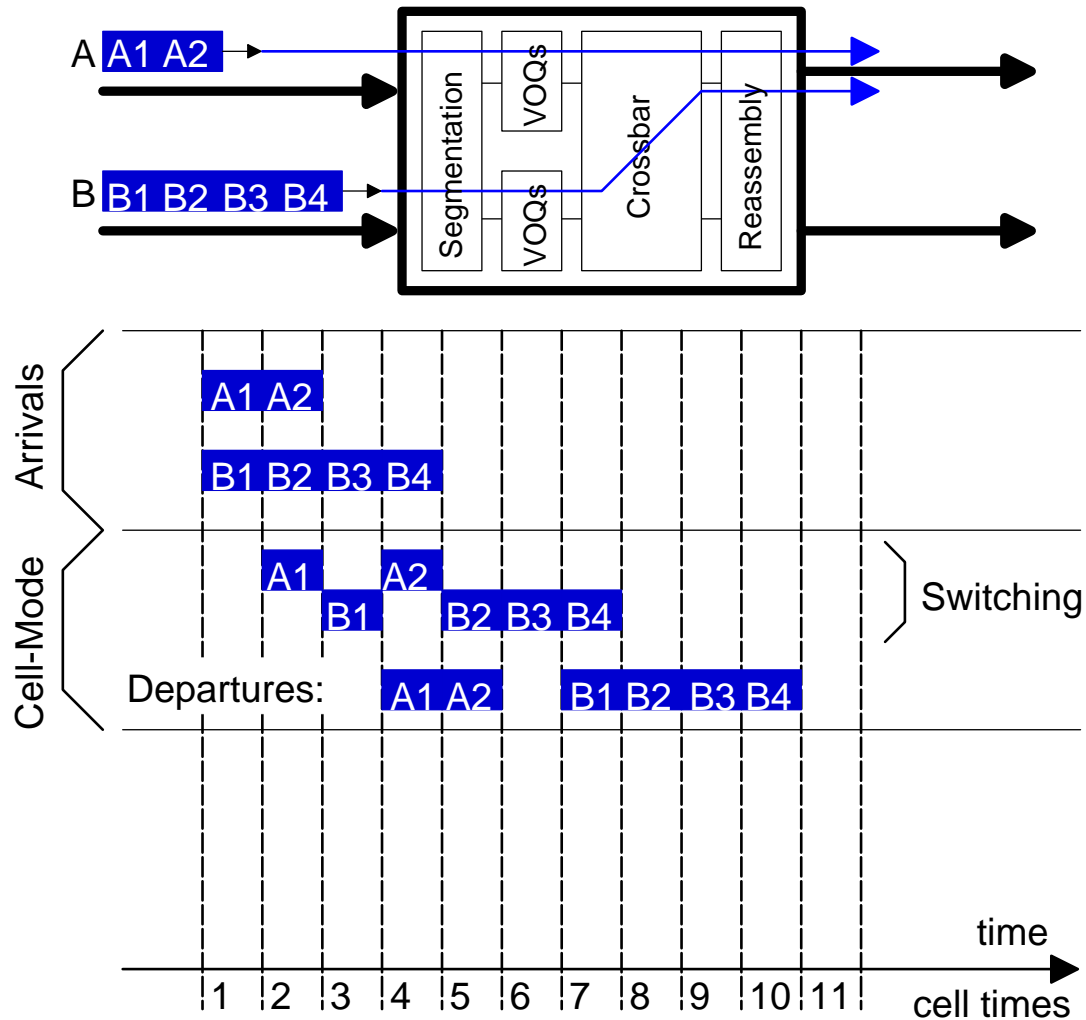


- External packets have variable size (e.g. IP packets)
- Segmentation to cells, switching, reassembly, transmission

Two alternatives to switch the internal cells

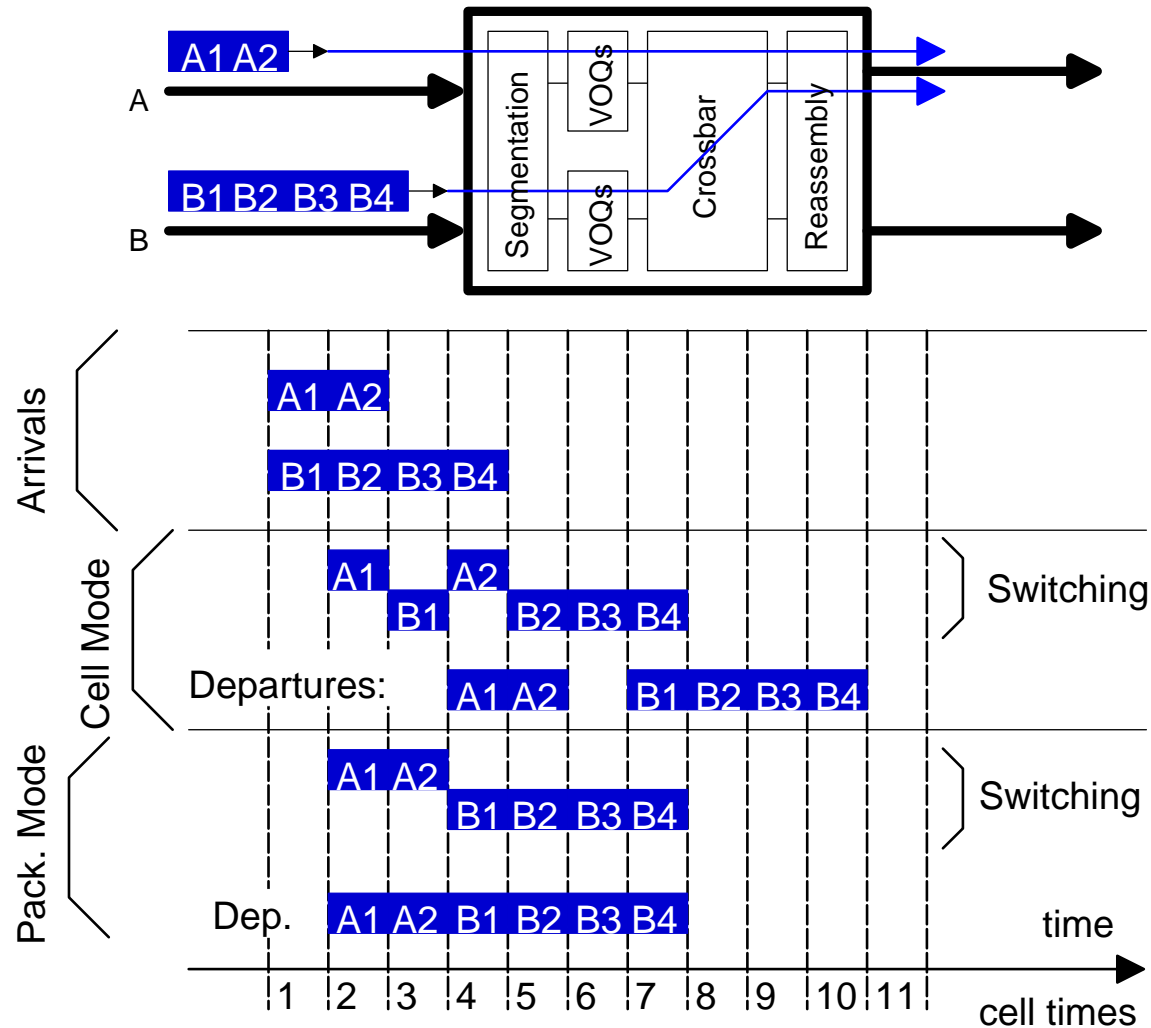
- cell mode:
ignore which cell belongs to which packet and reconfigure the crossbar from scratch every cell time
- packet mode:
keep cells of the same packet consecutive by tearing down connections only when whole packets have been fwded

Background: cell-mode operation



Reassembly complexity is needed and store&fwd is enforced

Background: packet-mode operation



Packet reassembly is trivial and cut-through is allowed

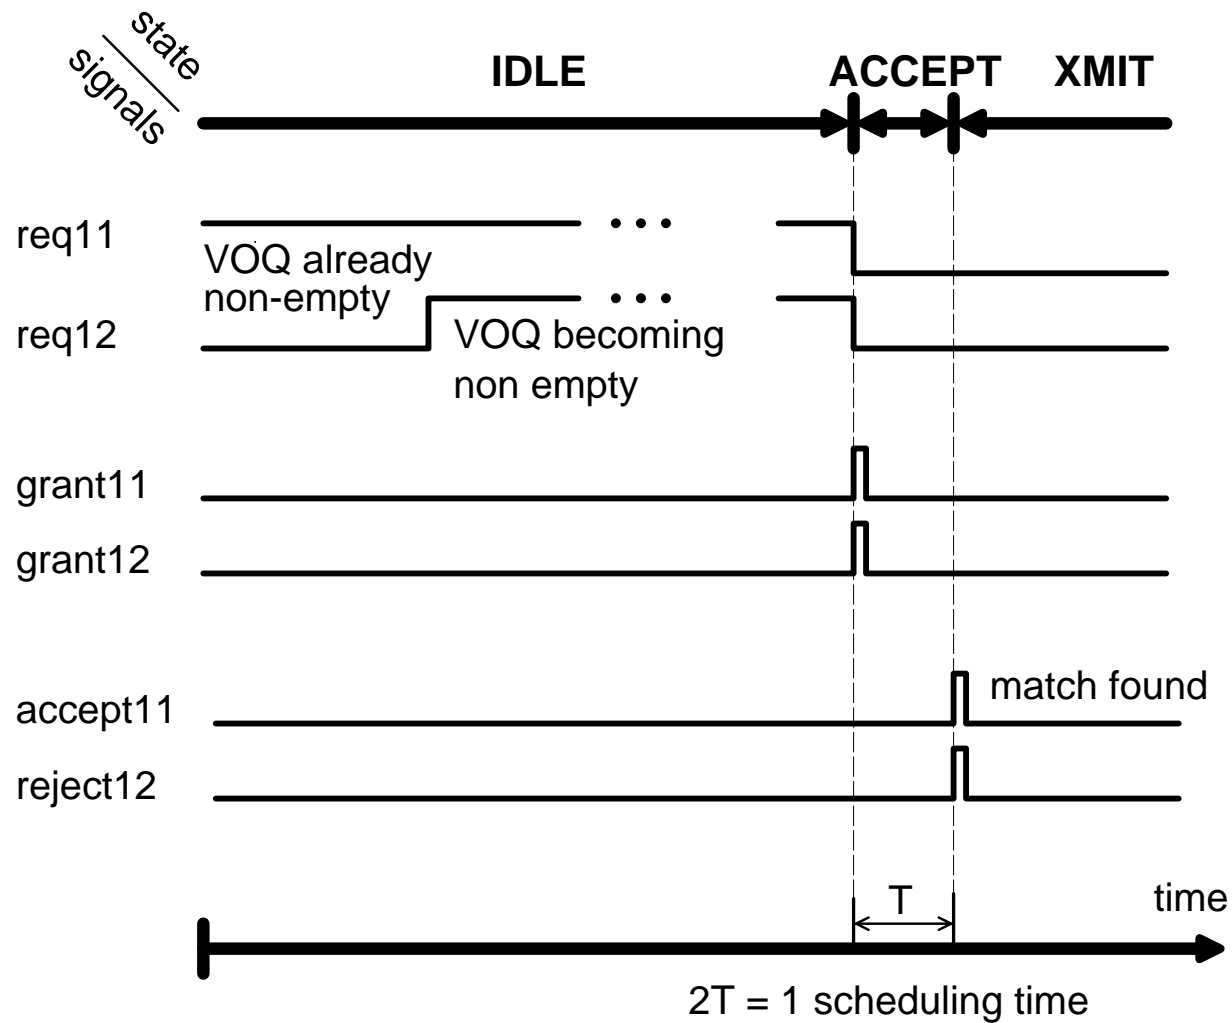
Problem statement

- Synchronous operation presumes cell size for packet-size granularity
 - 64 Bytes is a typical cell size in research and production switches
- External packets are not always integral multiples of internal cells
 - the last cell of a packet is padded
 - internal speed-up is required for padding overheads
- Internal speed-up for padding overheads may be ~ 2 in the worst case
 - e.g. with 65-Byte packets in a 64-Byte cell switch

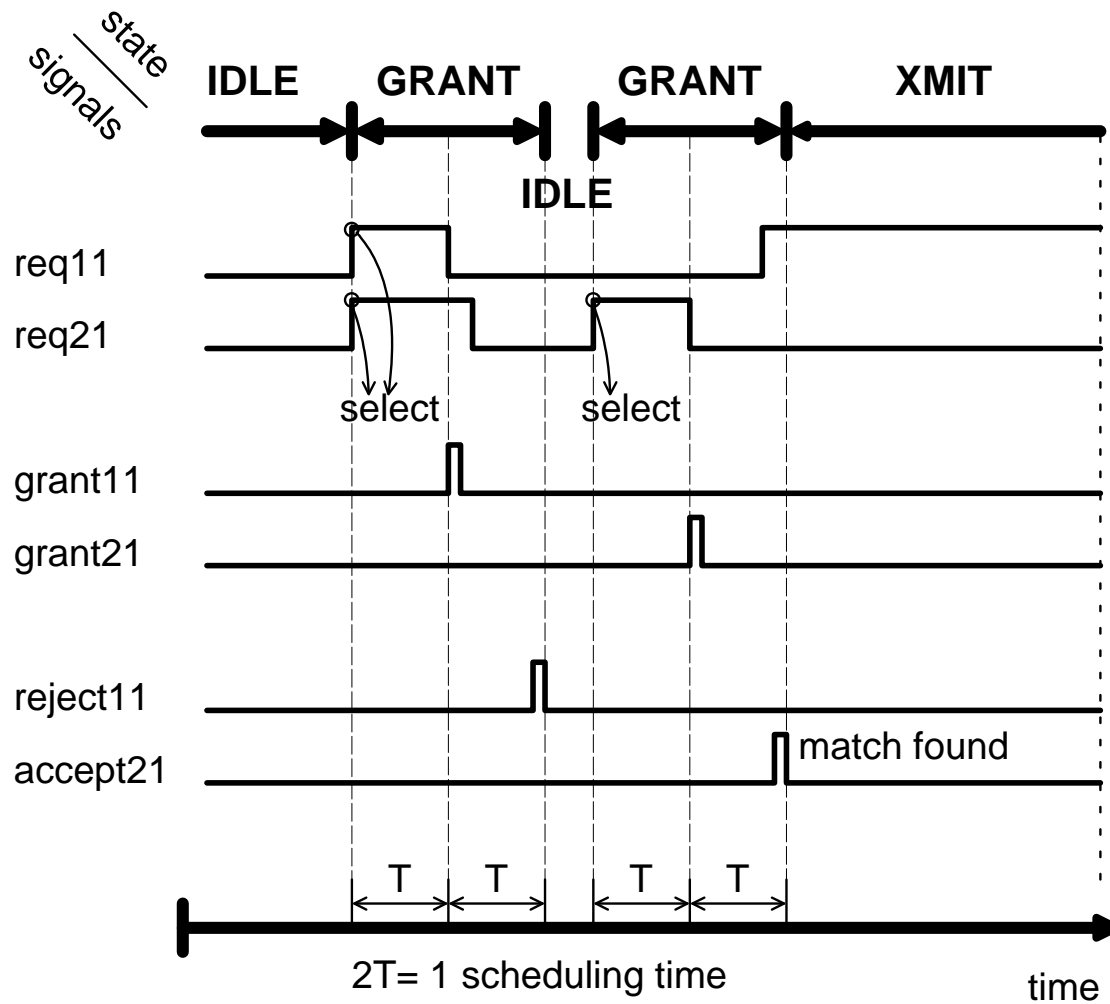
Bufferless crossbars operating directly on variable-size packets

- No segmentation & no reassembly
 - advantages of packet mode operation are inherited
- No segmentation to fixed-size cells
 - no padding overheads, no internal speed-up for them

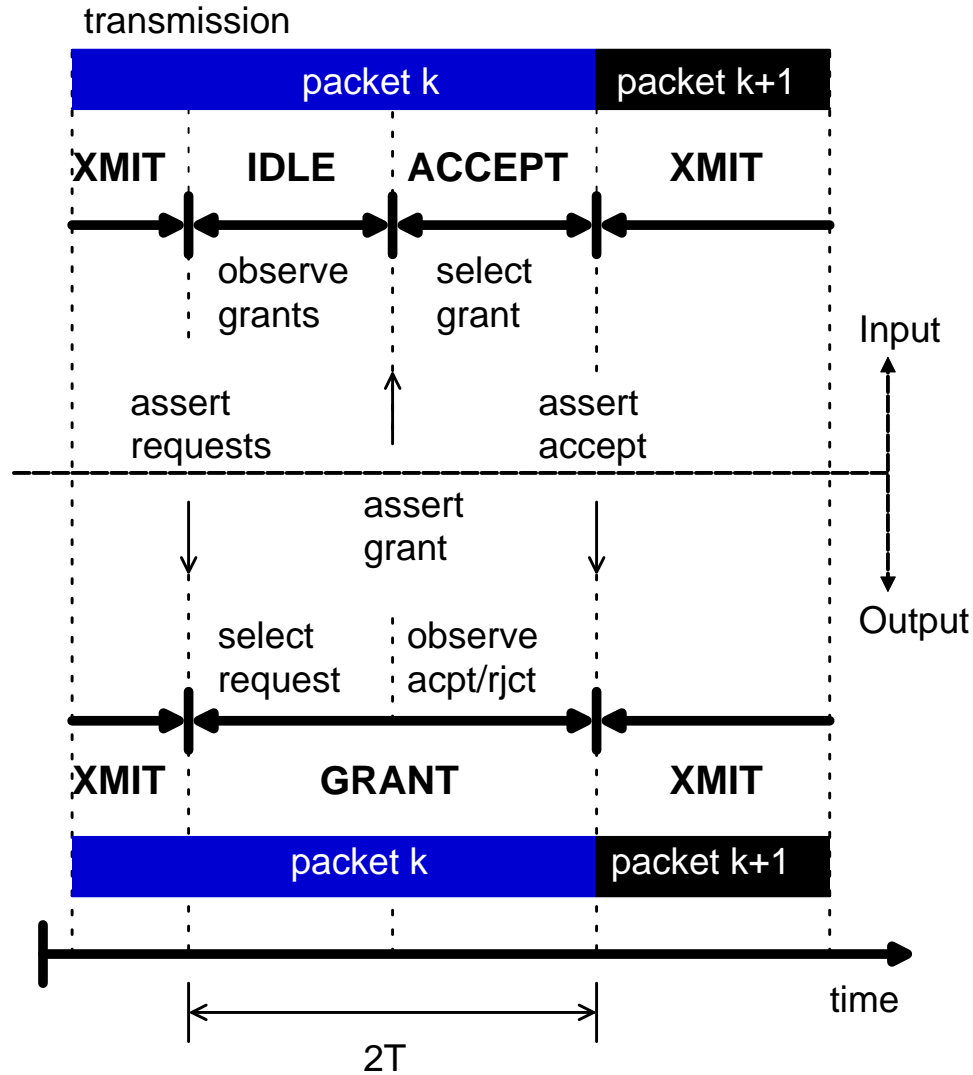
- Bipartite matches computed in three phases
- Inputs and outputs do not synchronize their scheduling decisions
- Crossbar configuration may change at any moment in time
 - infinitely small packet-size granularity



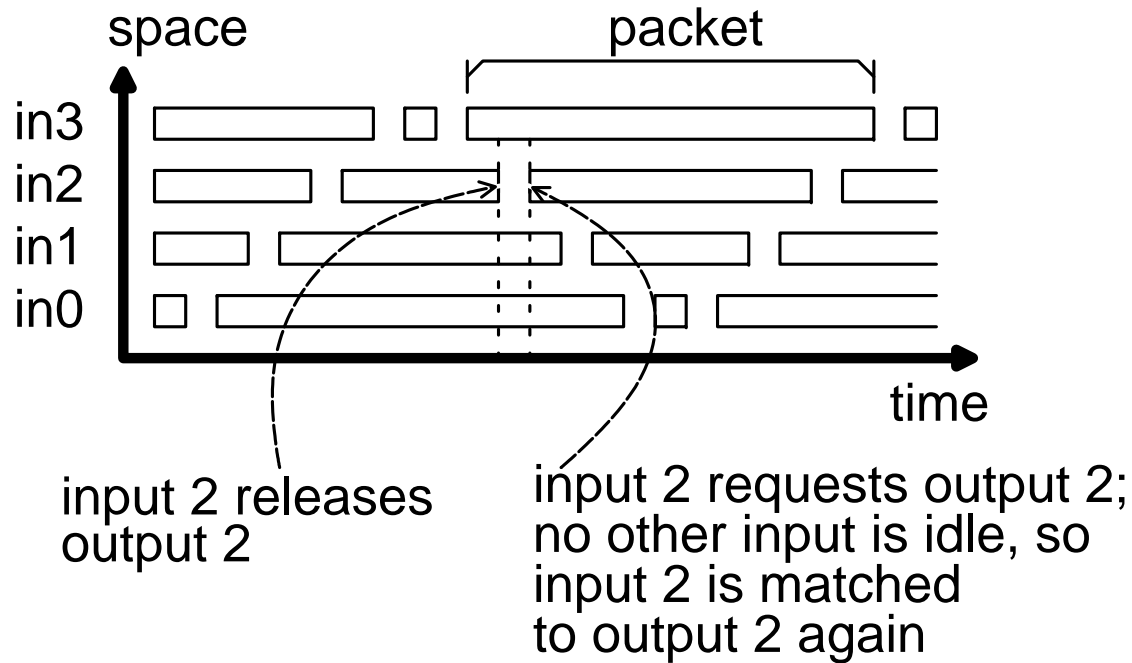
Example of input state transitions



Example of output state transitions



Example of joint input-output state transitions



- Scheduler does not consider candidacies for inputs or outputs which are in the middle of a packet forwarding
→ some flows may be locked-out
- Packet-mode scheduling similarly suffers from starvation

Simulations

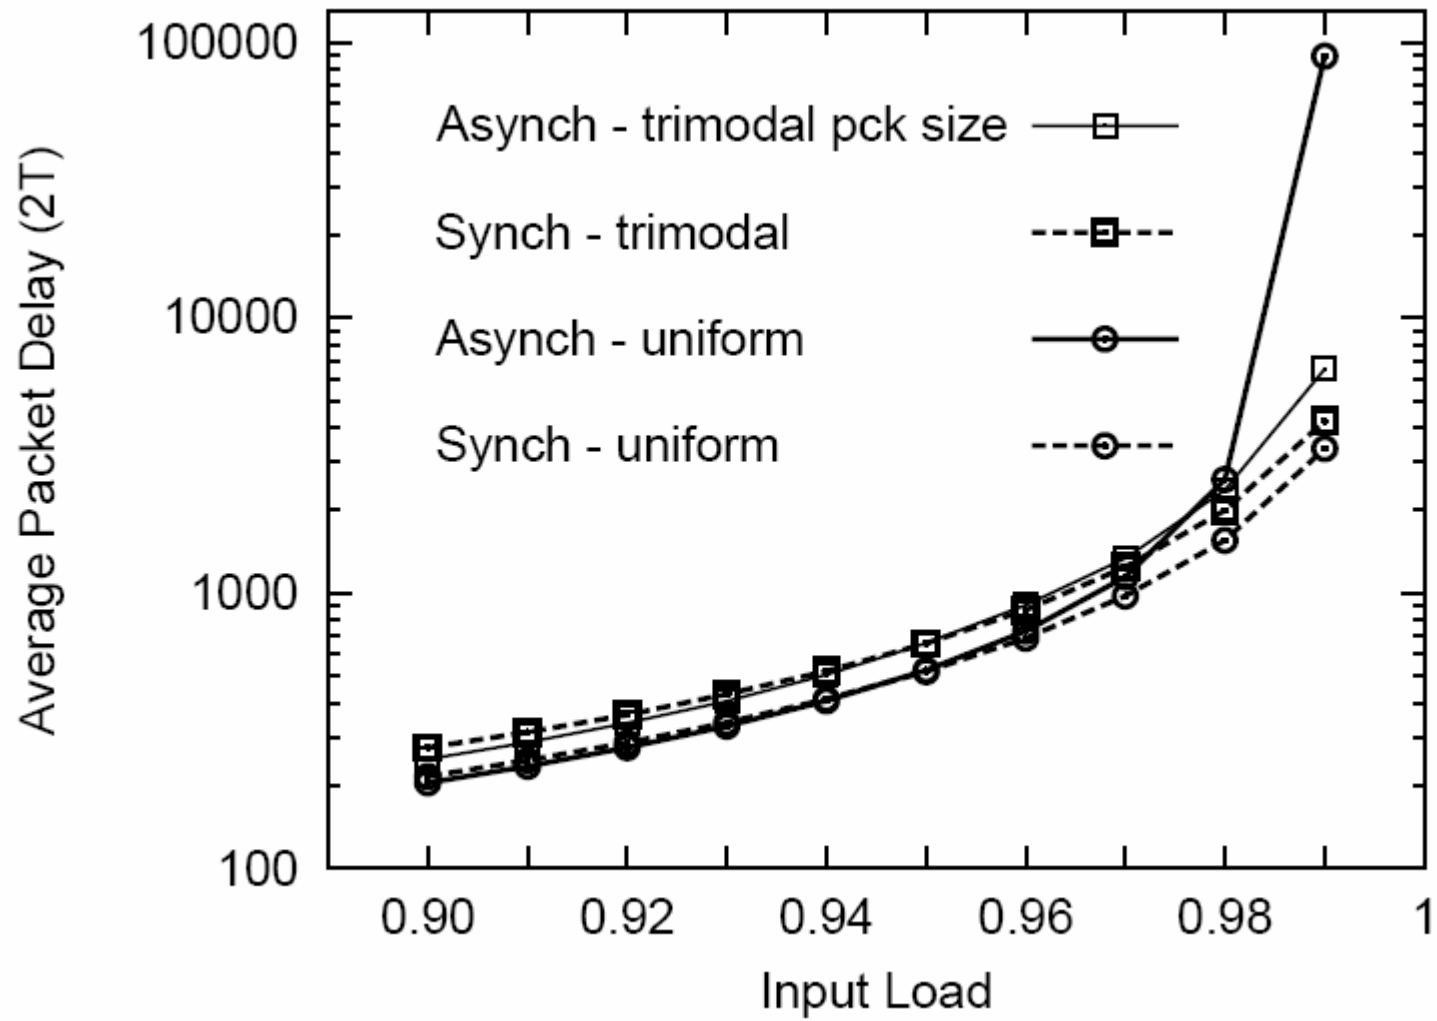
- Event-driven simulation for asynchronous operation
- Compared against packet-mode operation
- Packet-size granularity
 - asynchronous: 40 time smaller than the scheduling time
 - synchronous: cell-size
- Synthetic traffic
 - constant, uniform, trimodal packet-size distributions
 - uniform, unbalanced (Zipf) destination distributions
- Result: *Asynchronous operation does not degrade performance compared to packet-mode operation*

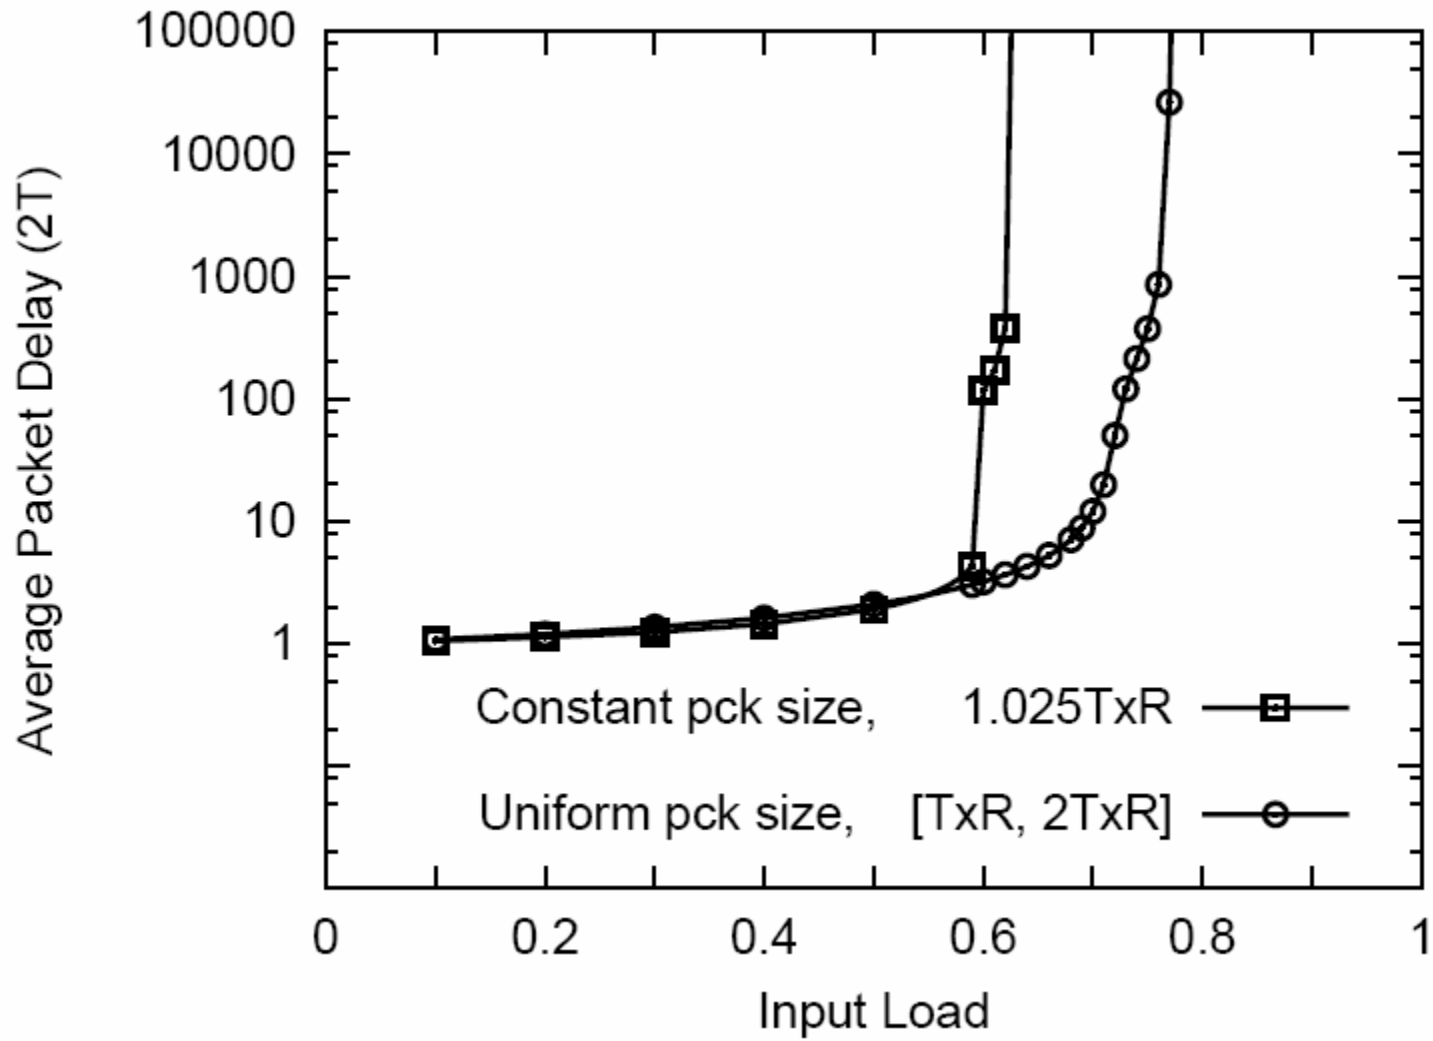
Summary


- We proposed asynchronous bufferless crossbar operation for direct switching of variable-size packets
- Like packet-mode operation:
 - no packet reassembly is needed and cut-through is allowed
- Unlike packet-mode operation:
 - no internal speed up for cell padding overheads is needed
- Simulations showed almost no performance degradation compared to packet-mode scheduling

THANK YOU

BACKUP SLIDES







Crossbar Reconfiguration Probability

