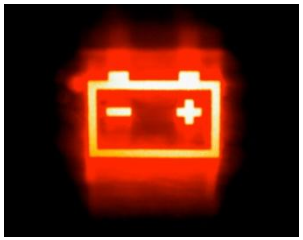

Order is Power: Selective Packet Interleaving for Energy Efficient Networks-on-Chip

Amit Berman

Joint work with Idit Keidar and Ran Ginosar

Department of Electrical Engineering,
Technion – Israel Institute of Technology

IEEE/IFIP International Conference on VLSI and System-on-Chip 2010 (VLSI-SoC'10)

The power wall

➢**Power density continues to get worse**

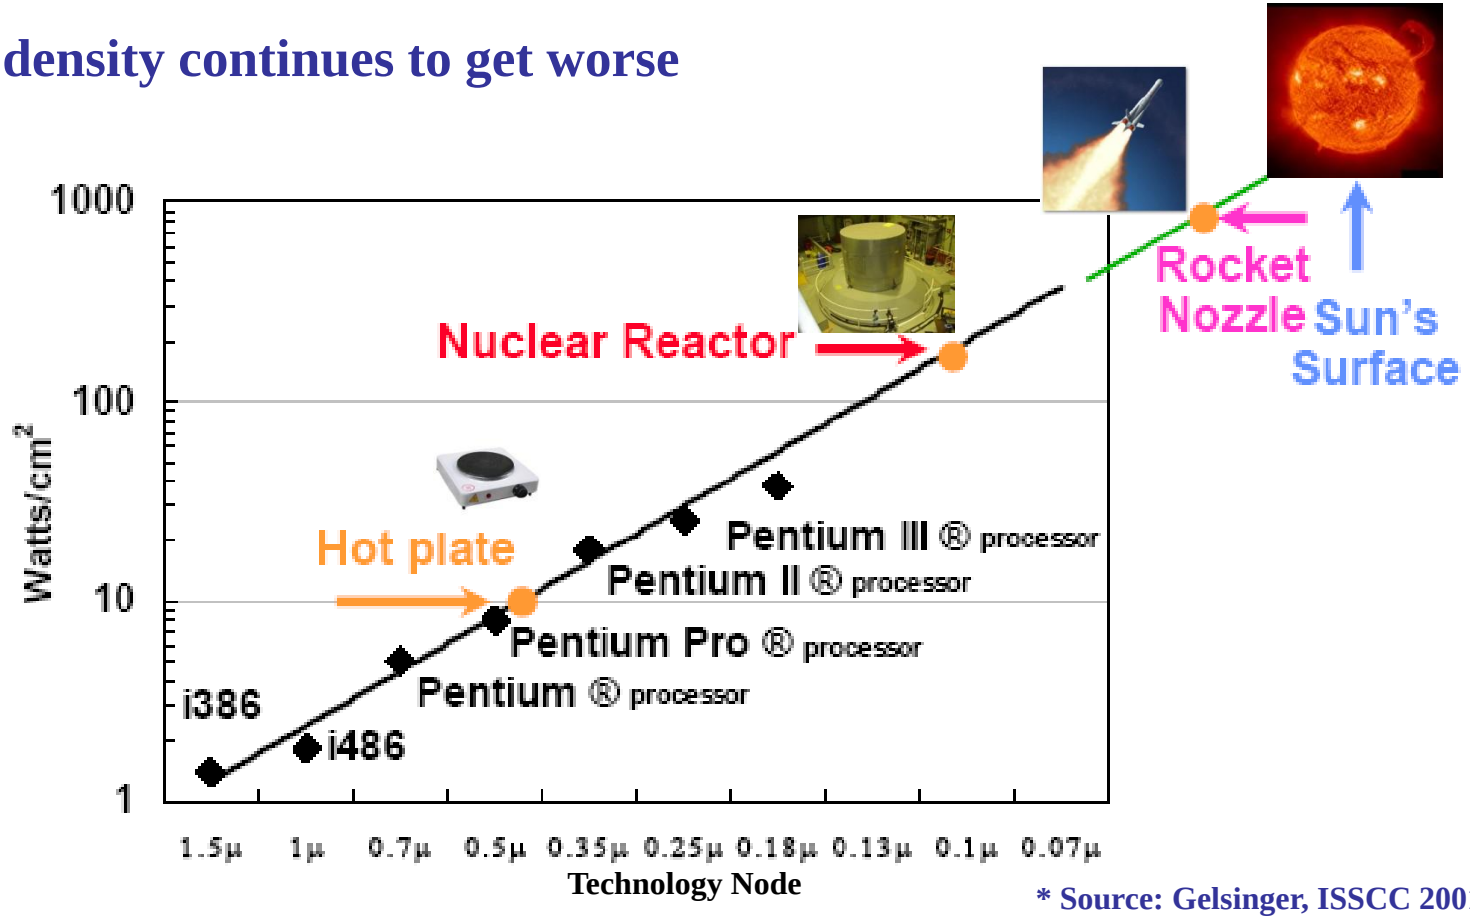

*** Source: Gelsinger, ISSCC 2001**

We need to reduce power consumption.

Distribution of power consumption

➢**Interconnects consume up to 60%**

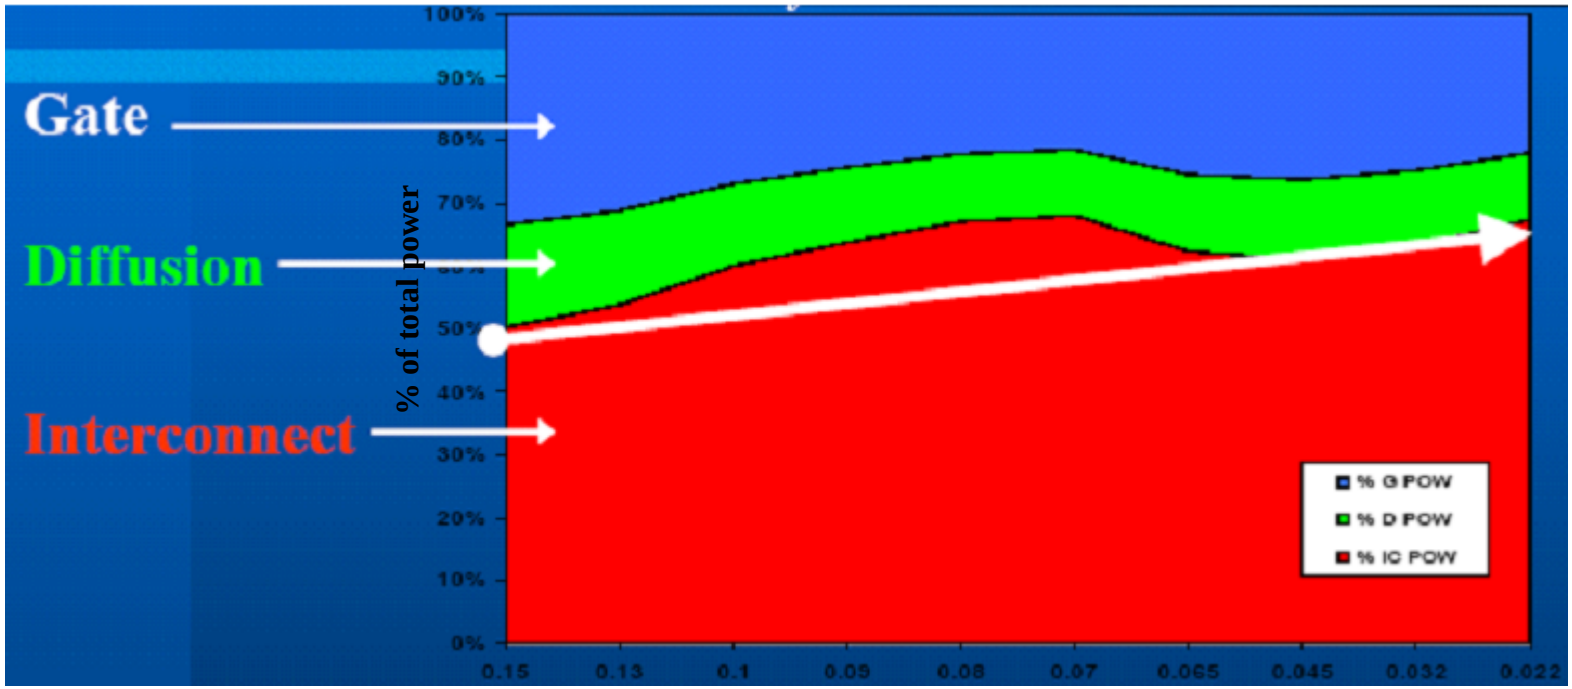

*** Source: ITRS 2001 Edition**

Interconnects hold the lion's share of total power.

➢NoC context

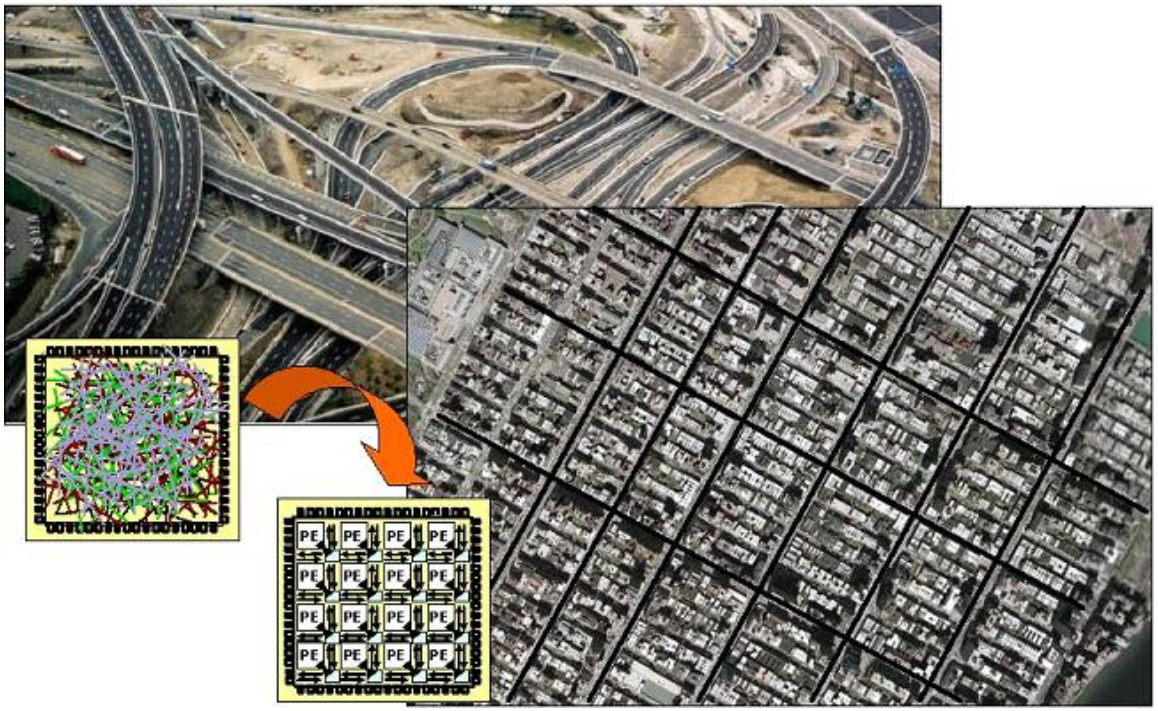

Our goal is to reduce the power consumption on NoC links.

Opportunity: reduce the number of bit transitions on the link

➢**Dynamic power consumption origins form bit transitions**

- **Circuit nodes are being charged**

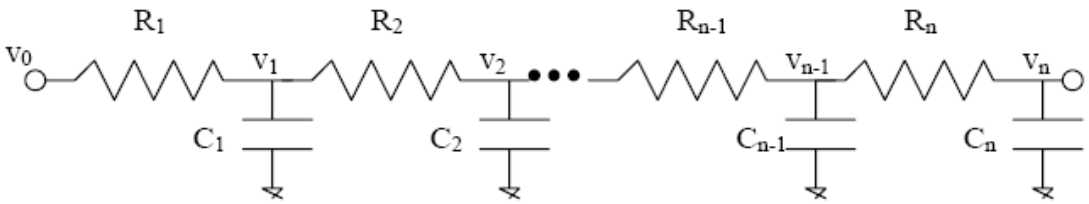

Goal: reduce the number of bit transitions.

Observation

➢There are multiple virtual channels in the router

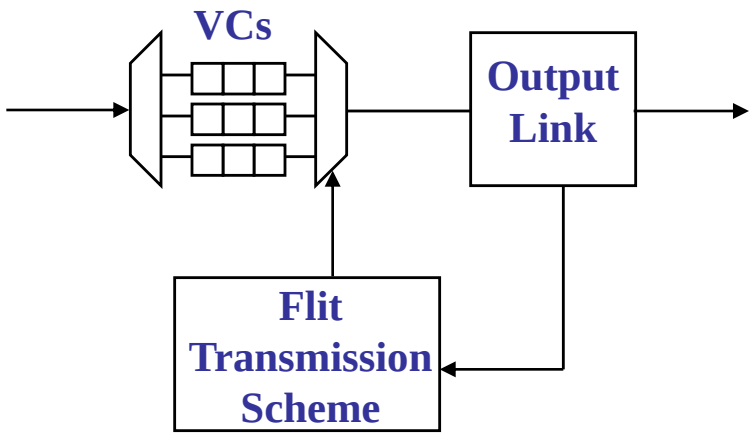

There is flexibility in interleaving.

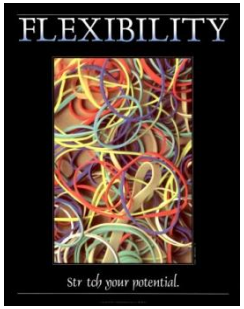

Observation: example

- Two packets in two virtual channels are transmitted
  - A 4-bit flit is sent on the link

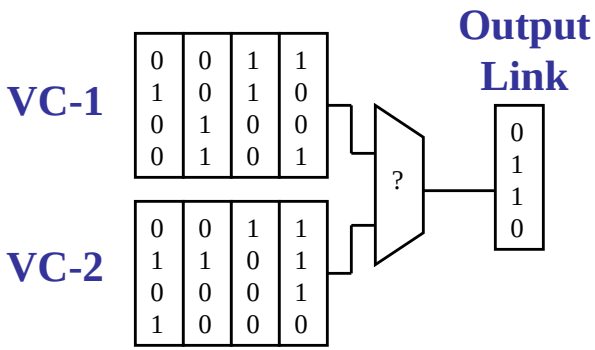

Can we reduce the number of bit transitions?

Selective Packet Interleaving (SPI)

➢**Transmit the flit that entails minimum bit transitions**

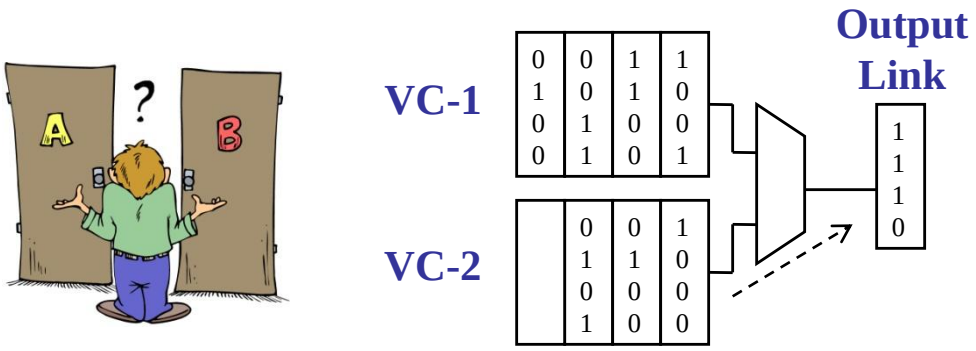

If D(VC1,Link)<= D(VC2,Link) then transmit the flit from VC1
If D(VC2,Link)< D(VC1,Link) then transmit the flit from VC2

SPI multiplexes flits to the router's output link so as to minimize the number of bit transitions from the previously transmitted flit.

SPI scalability

➢If there are more VCs, there is a higher probability for lower bit transitions

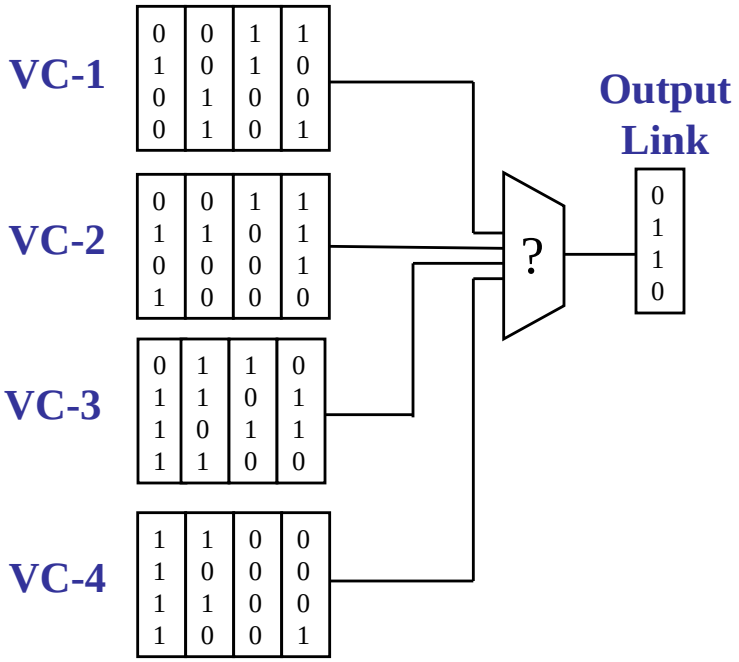

SPI is scalable with the number of VCs.

SPI overhead

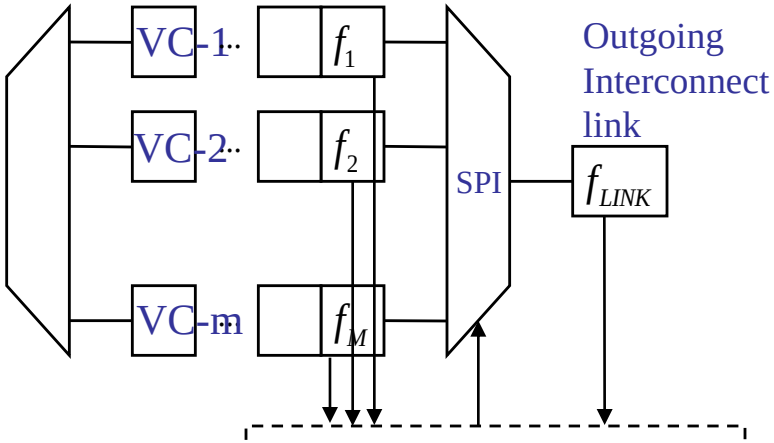

Changes to NoC's router: add SPI modules.

Analysis results

For 8b links, SPI outperforms *Bus Invert* (BI) starting from two VCs.
For 16b and 32b links, SPI outperforms BI starting from three VCs.

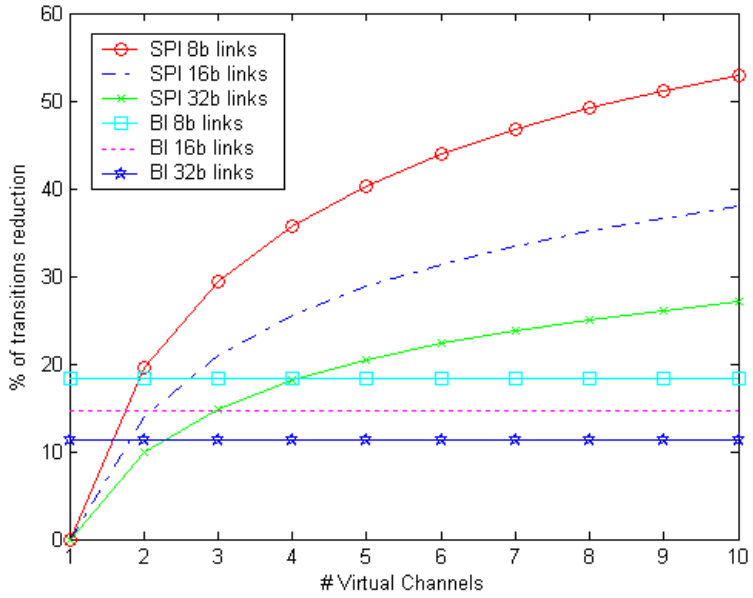

SPI and BI improvement over uncoded

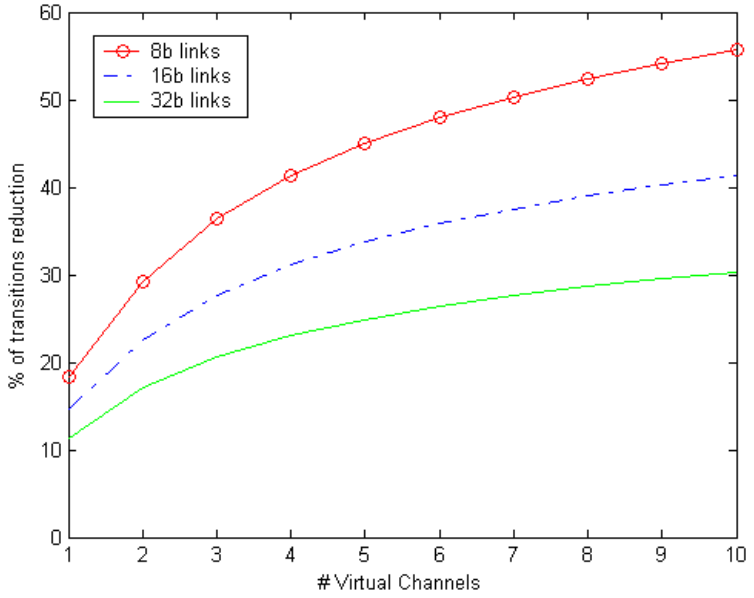

SPI+BI improvement over uncoded

Comparison: SPI vs. SPI+BI

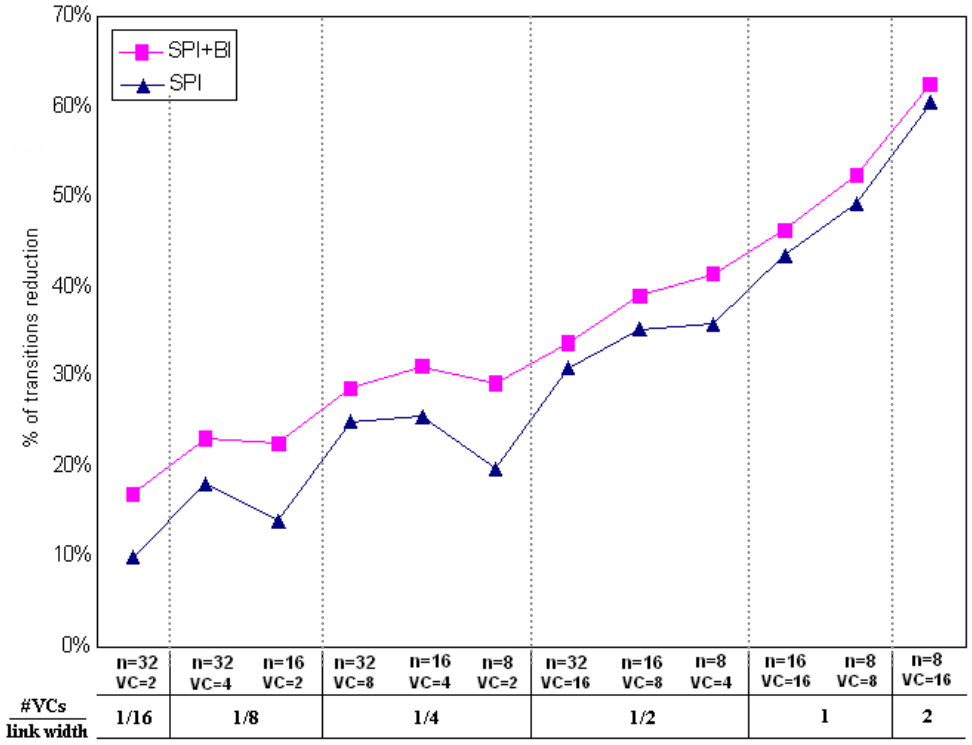

As the ratio #VCs/(link width) increases, the percentage of improvement over uncoded transmissions increases, and the gap between SPI and SPI+BI shrinks.

Benchmark simulations

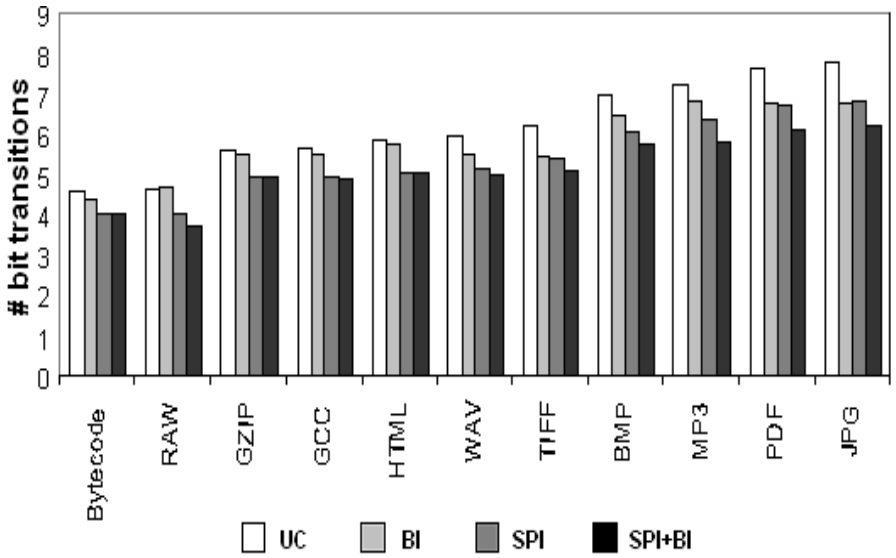

2VCs, 16b links.

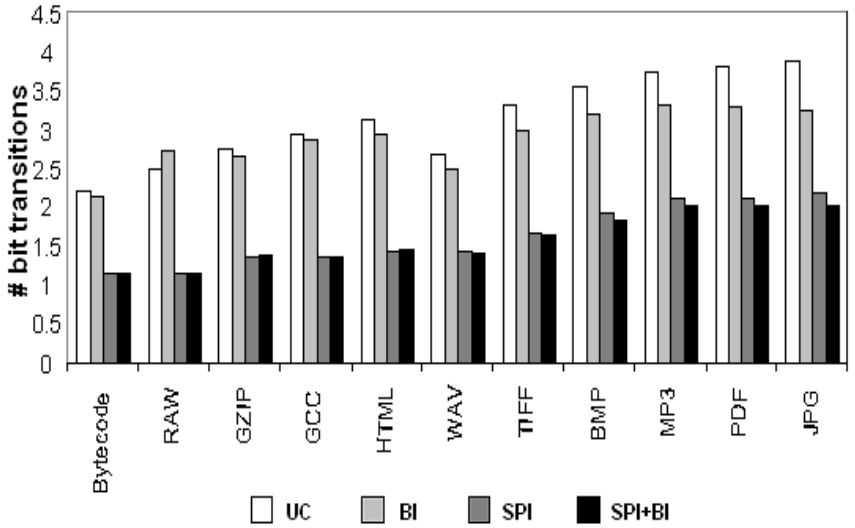

8VCs, 8b links.

The percentage of improvement is similar for all benchmarks.
With two VCs and 16b links, the improvements are 10% to 13%.
With eight VCs and 8b links, the improvements are 45%-55%.

Power reduction

SPI was implemented in VLSI design tools to test for power consumption

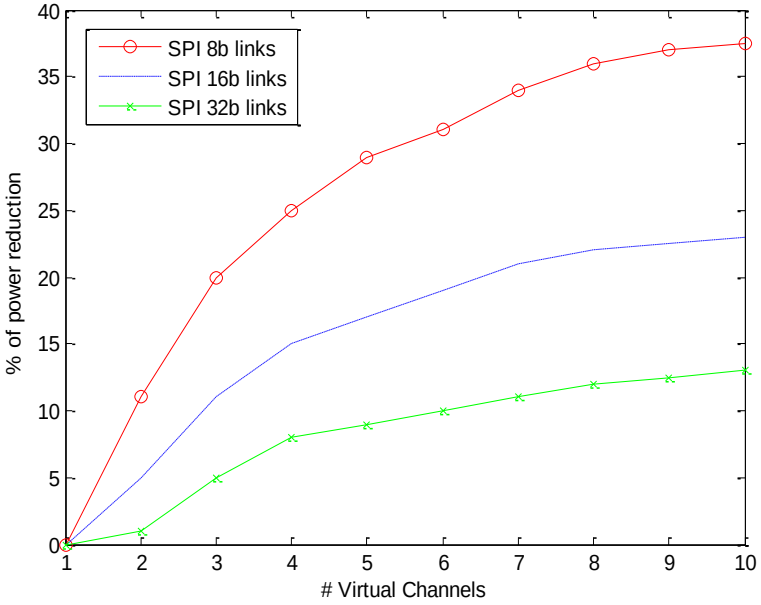

For 8b width links and four VCs, we observe more than 25% power reduction.
For 16b width links and four VCs, we observe 15% power reduction.
For 32b width links and four VCs, the reduction is about 10%.

Summary of selective packet interleaving

- ***Selective packet interleaving*** **statistically reduces the number of bit transitions**
- ***Selective packet interleaving*** **can reduce the power consumption of the NoC**
  - **Less bit transitions**
- ***Selective packet interleaving*** **is scalable with the number of VCs**

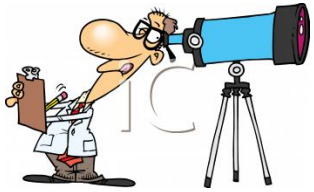

Thank you for your time.

Questions session.