U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Variable-Width Datapath for On-Chip Network Static Power Reduction

George Michelogiannakis, John Shalf

Postdoctoral Research Fellow

Computer Architecture Laboratory

Lawrence Berkeley National Laboratory

In a Nutshell

- Leakage power is an increasing problem in future or near threshold voltage (NTV) technologies
- Leakage power can be important even at high network loads
- This work proposes variable-width datapaths
  - Parts of channels, buffers, and crossbars can be activated on demand
- We demonstrate an average of 33% total power reduction with PARSEC benchmarks

Today's Menu

- Leakage power / motivation
- Related work
- Variable-width datapaths
- Results
- Conclusions

Leakage Power Contribution

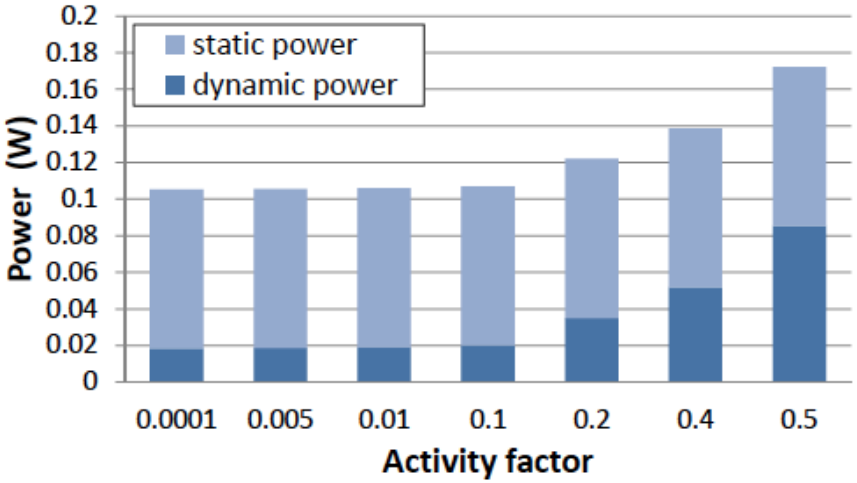

32nm (above). 45nm (below). Orion 2.0.

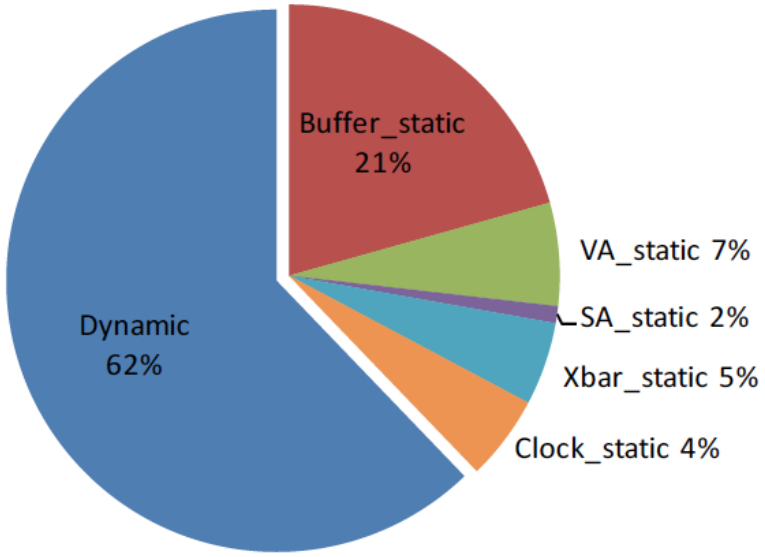

Single router. 45nm. 1.0V Orion 2.0.

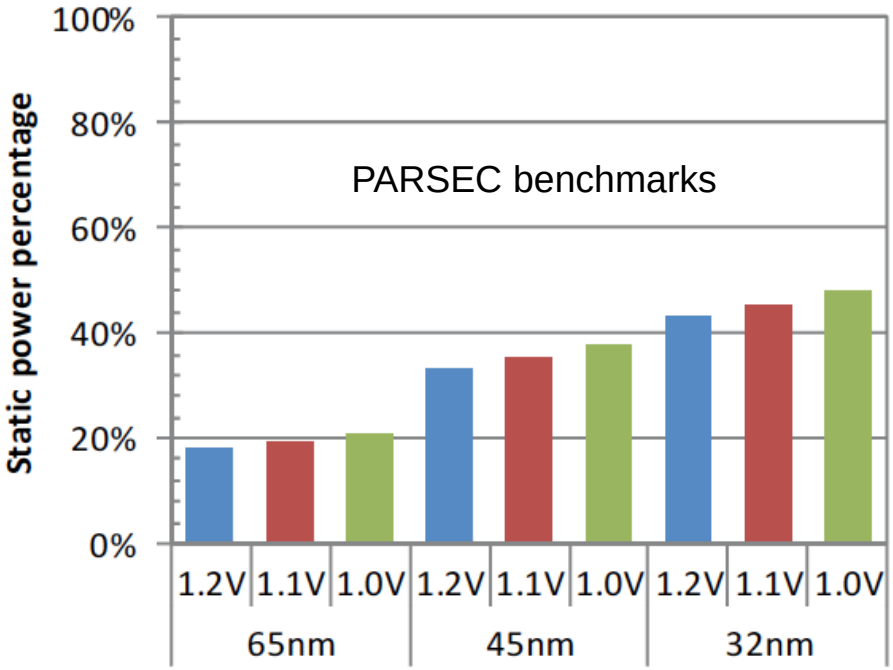

[Top left] “FlexiBuffer: Reducing Leakage Power in On-Chip Network Routers”. DAC 2011

[Rest] “NoRD: Node-Router Decoupling for Effective Power-gating of On-Chip Routers”. MICRO 2012

Subthreshold Leakage at NTV

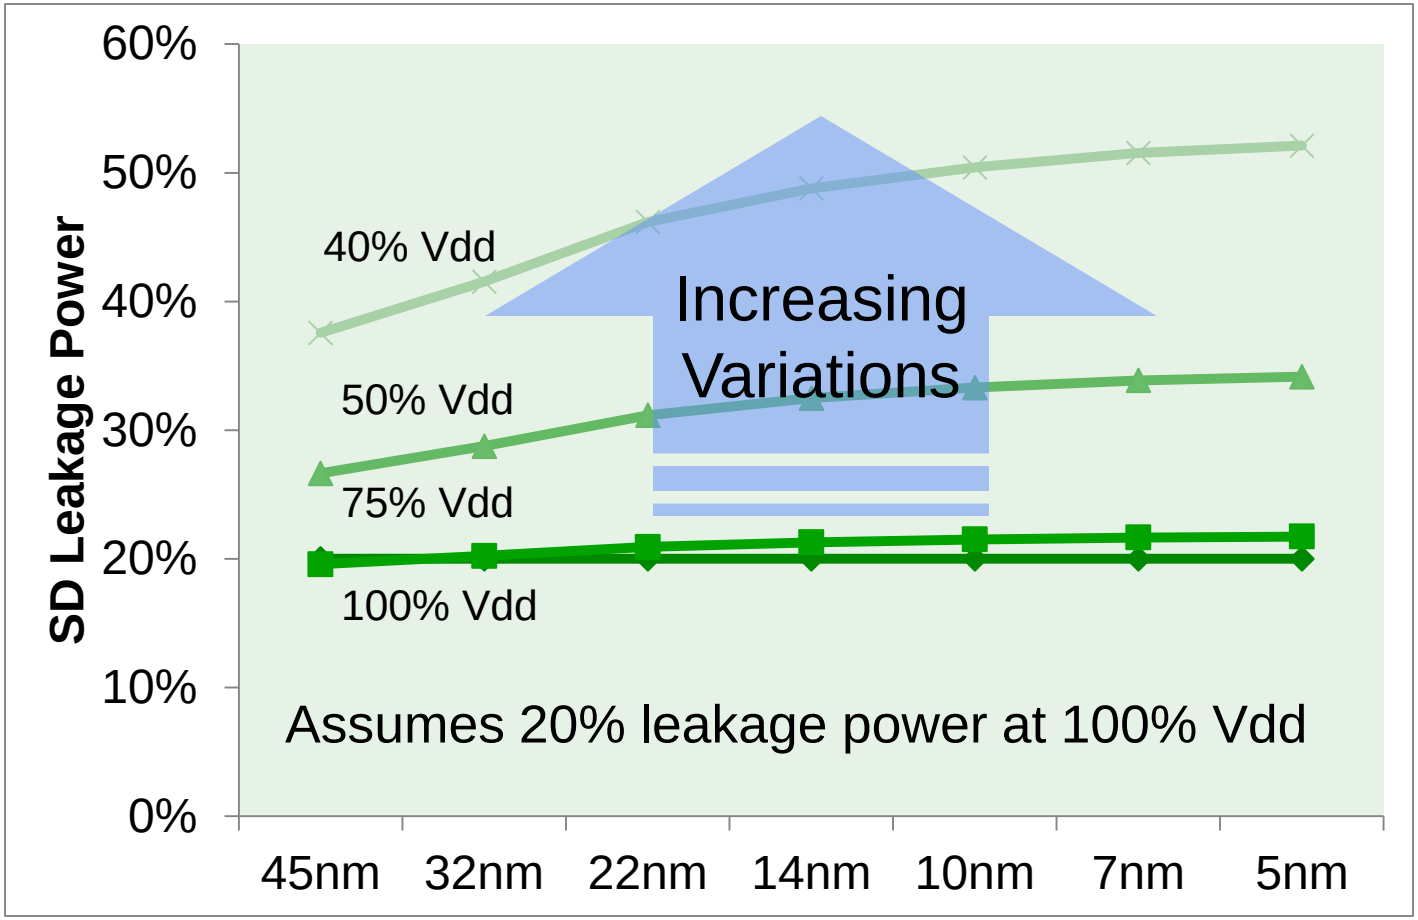

NTV operation reduces total power, improves energy efficiency

Subthreshold leakage power is substantial portion of the total

Near-threshold voltage (NTV) design — Opportunities and challenges. DAC 2012

Applications Load Network Unevenly

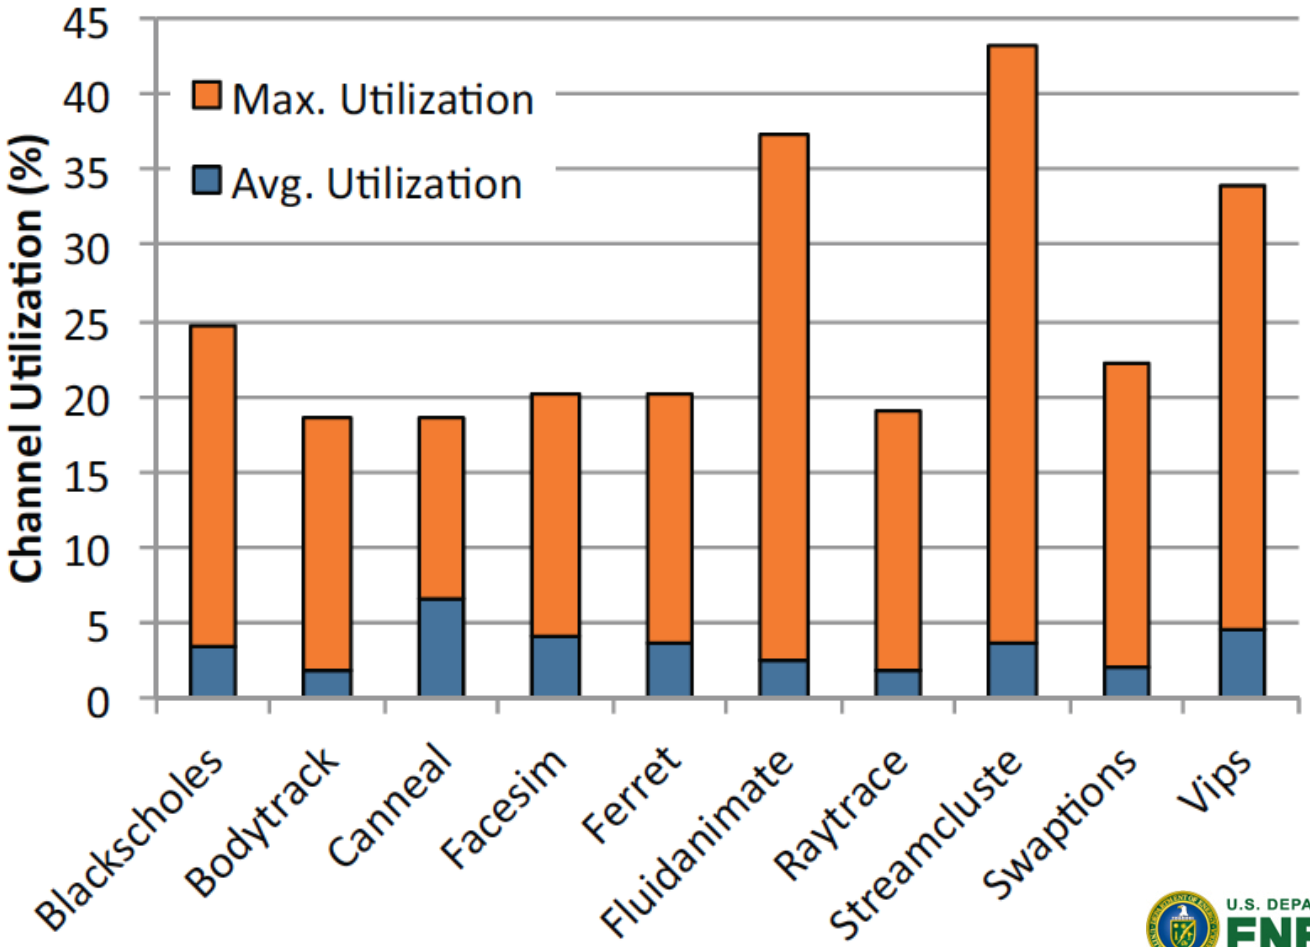

"Fine-grained bandwidth adaptivity in networks-on-chip using bidirectional channels". NOCS 2012

Today's Menu

- Leakage power / motivation
- Related work
- Variable-width datapaths
- Results
- Conclusions

Power Gating

- High threshold voltage “sleep switch” transistor
- Savings when sleep time enough to overcome energy overheads

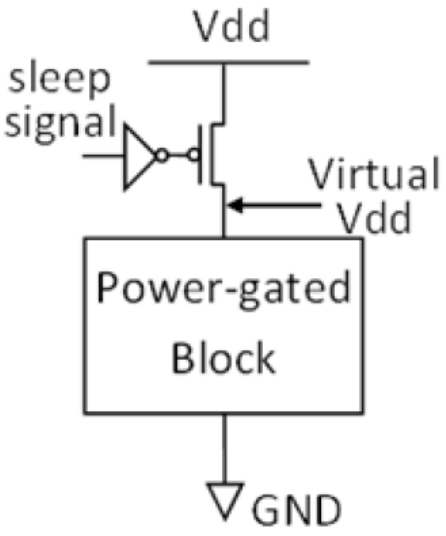

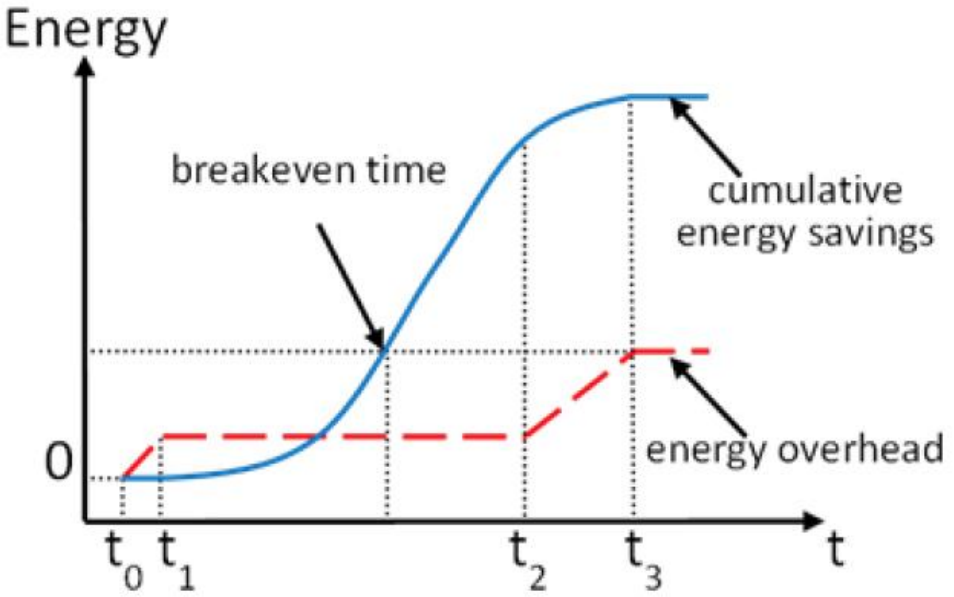

“MP3: Minimizing Performance Penalty for Power-gating of Clos Network-on-Chip”. HPCA 2014

Drowsy SRAMs

- Put SRAM lines into low-power (low voltage) “drowsy” mode
  - Preserves data
- Faster activation than power-gated SRAMs (1-2 cycles)
  - Higher leakage current while drowsy. Higher activation penalty

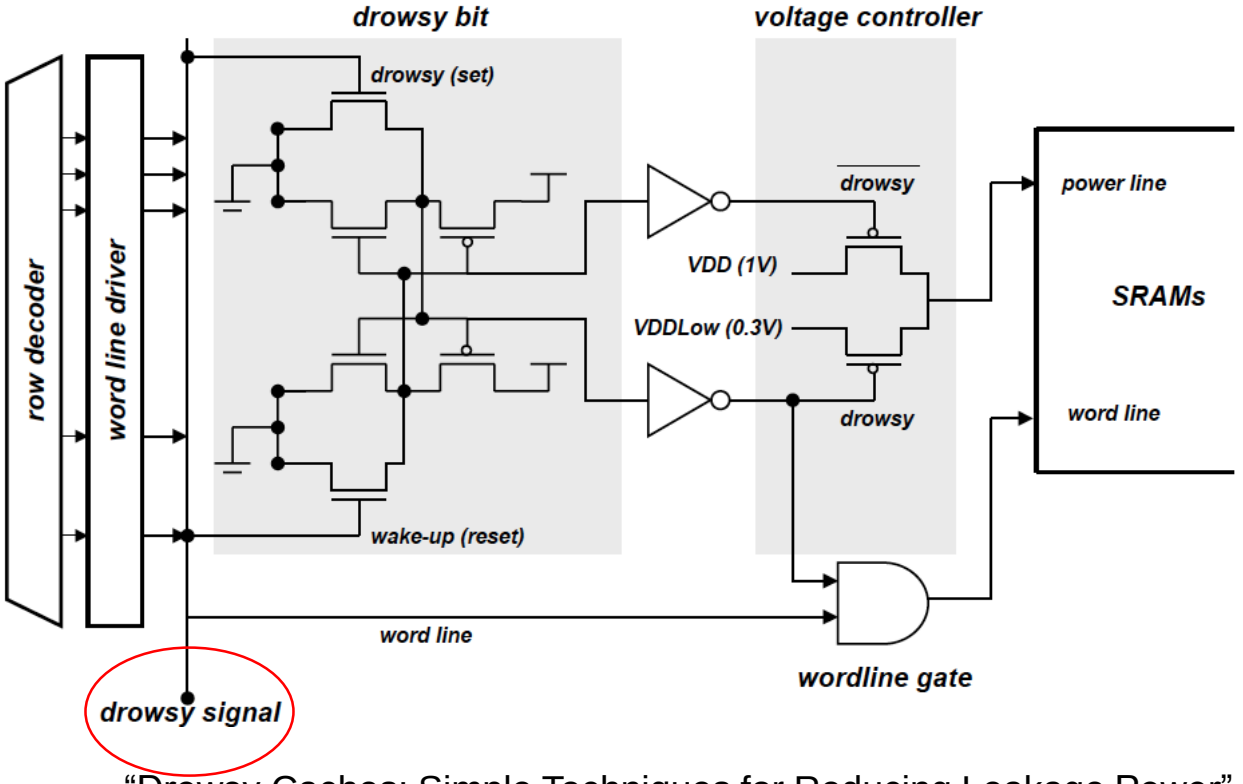

“Drowsy Caches: Simple Techniques for Reducing Leakage Power”. ISCA 2002

CatNap: Multiple Networks

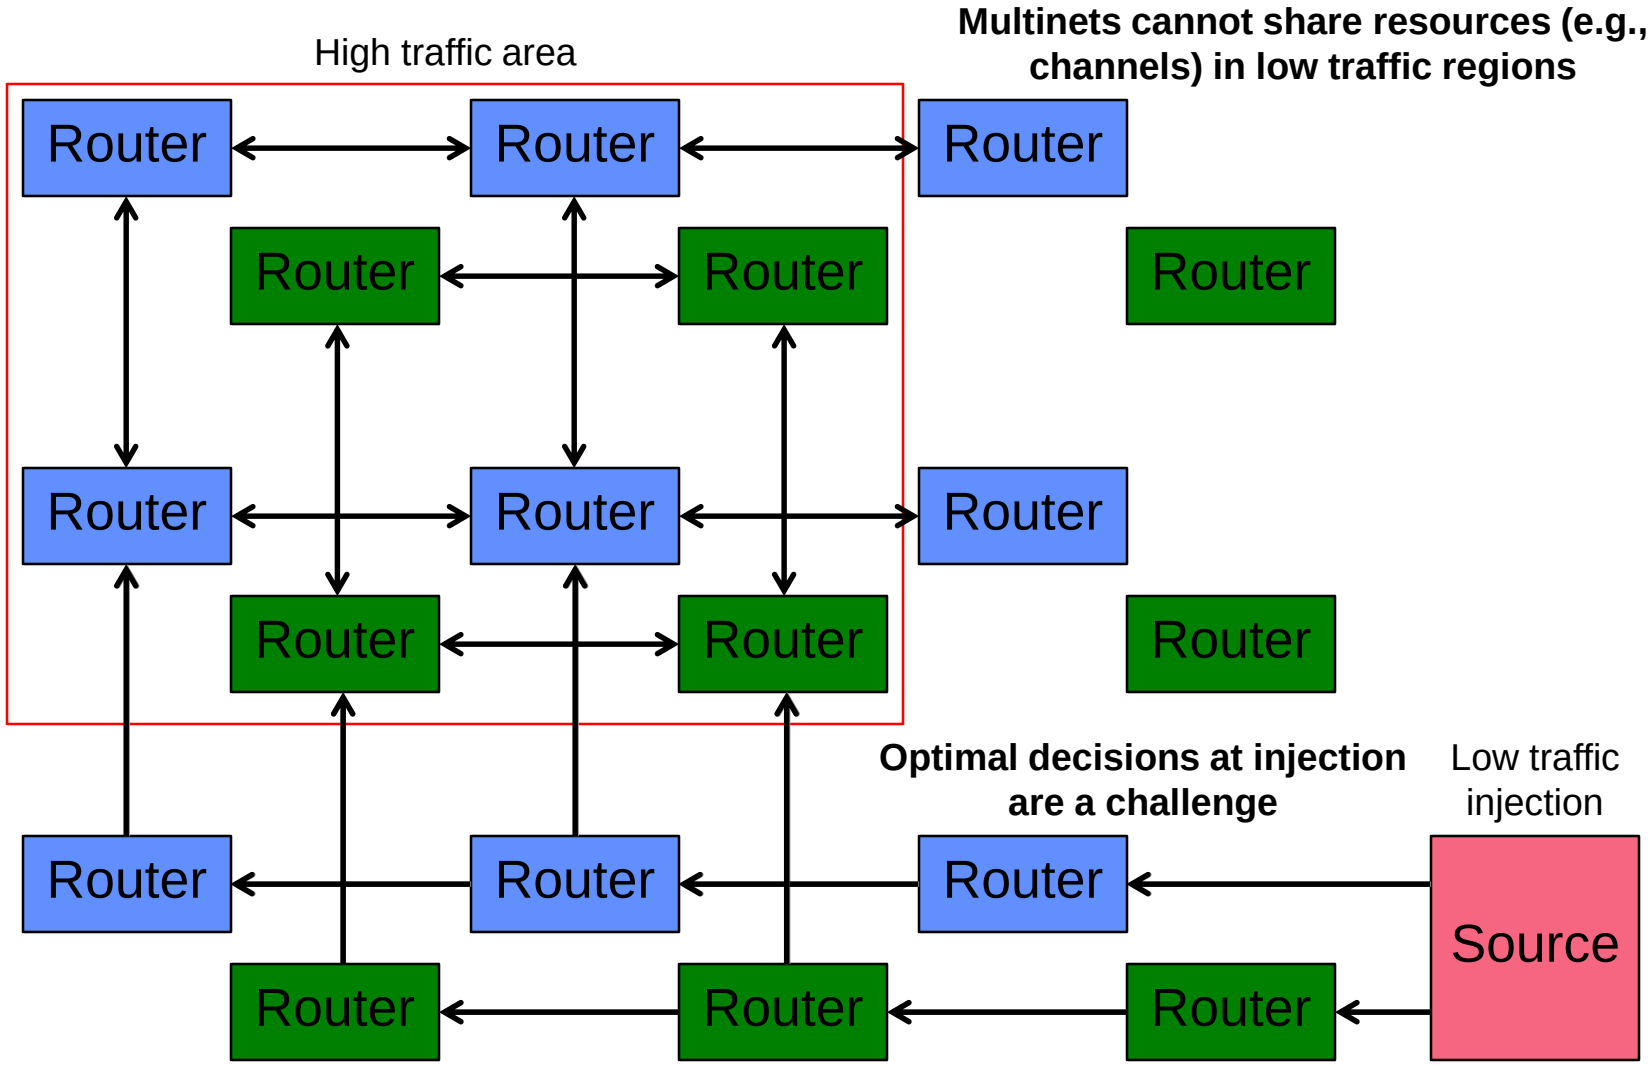

"Catnap: Energy Proportional Multiple Network-on-Chip". ISCA 2013

Today's Menu

- Leakage power / motivation
- Related work
- Variable-width datapaths
- Results
- Conclusions

Variable-Width Channels

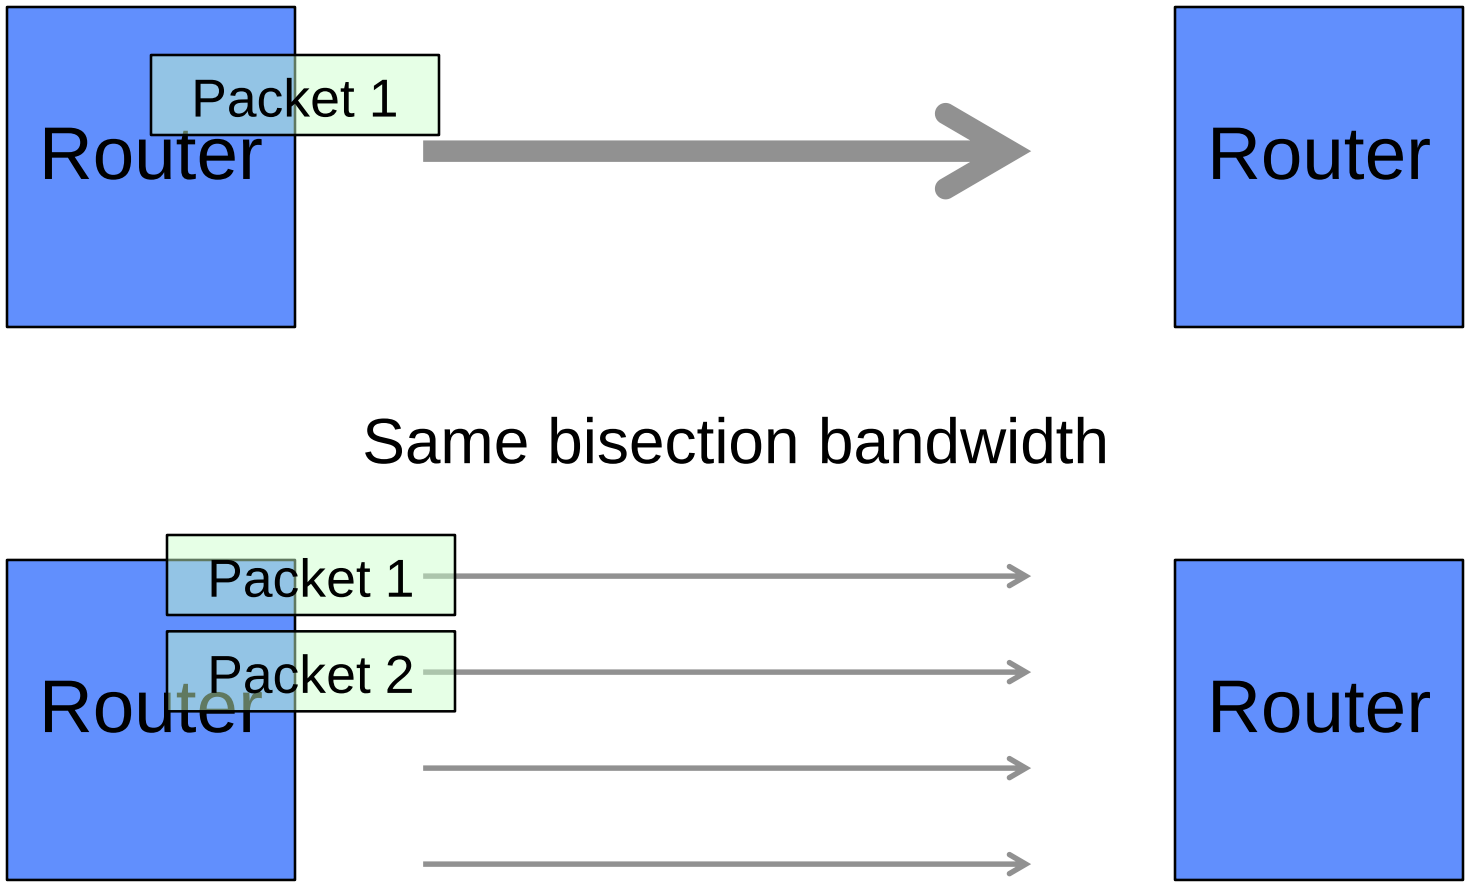

Input Buffer Gating

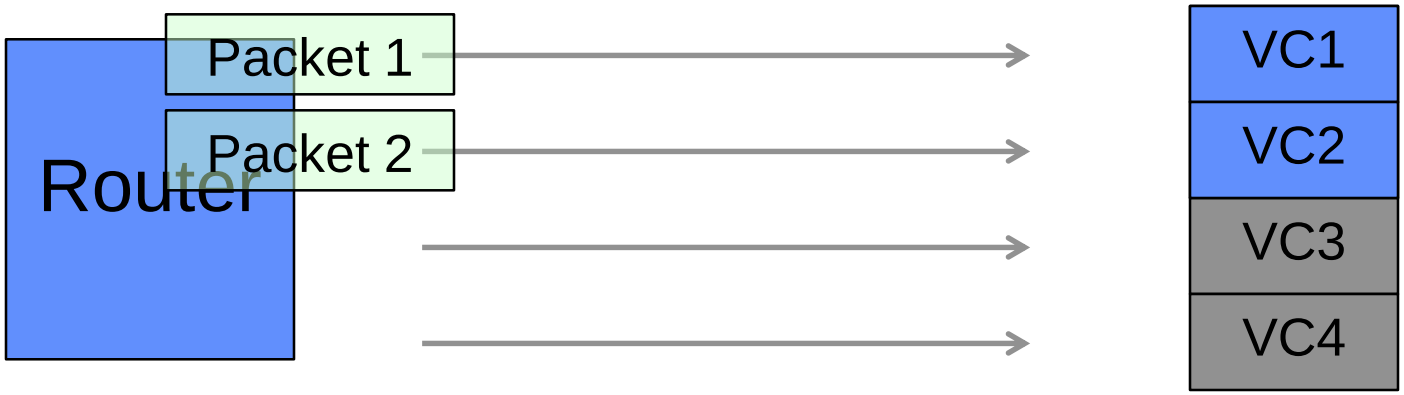

- If flits from any channel lane can choose any VC, that necessasitates multiplexers

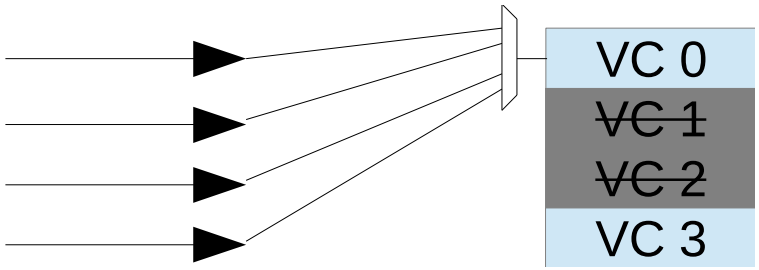

- We map channel lanes to use only a subset of VCs (1-1 with equal VCs and lanes)

"Adding slow-silent virtual channels for low-power on-chip networks". NOCS 2008

Crossbar Gating

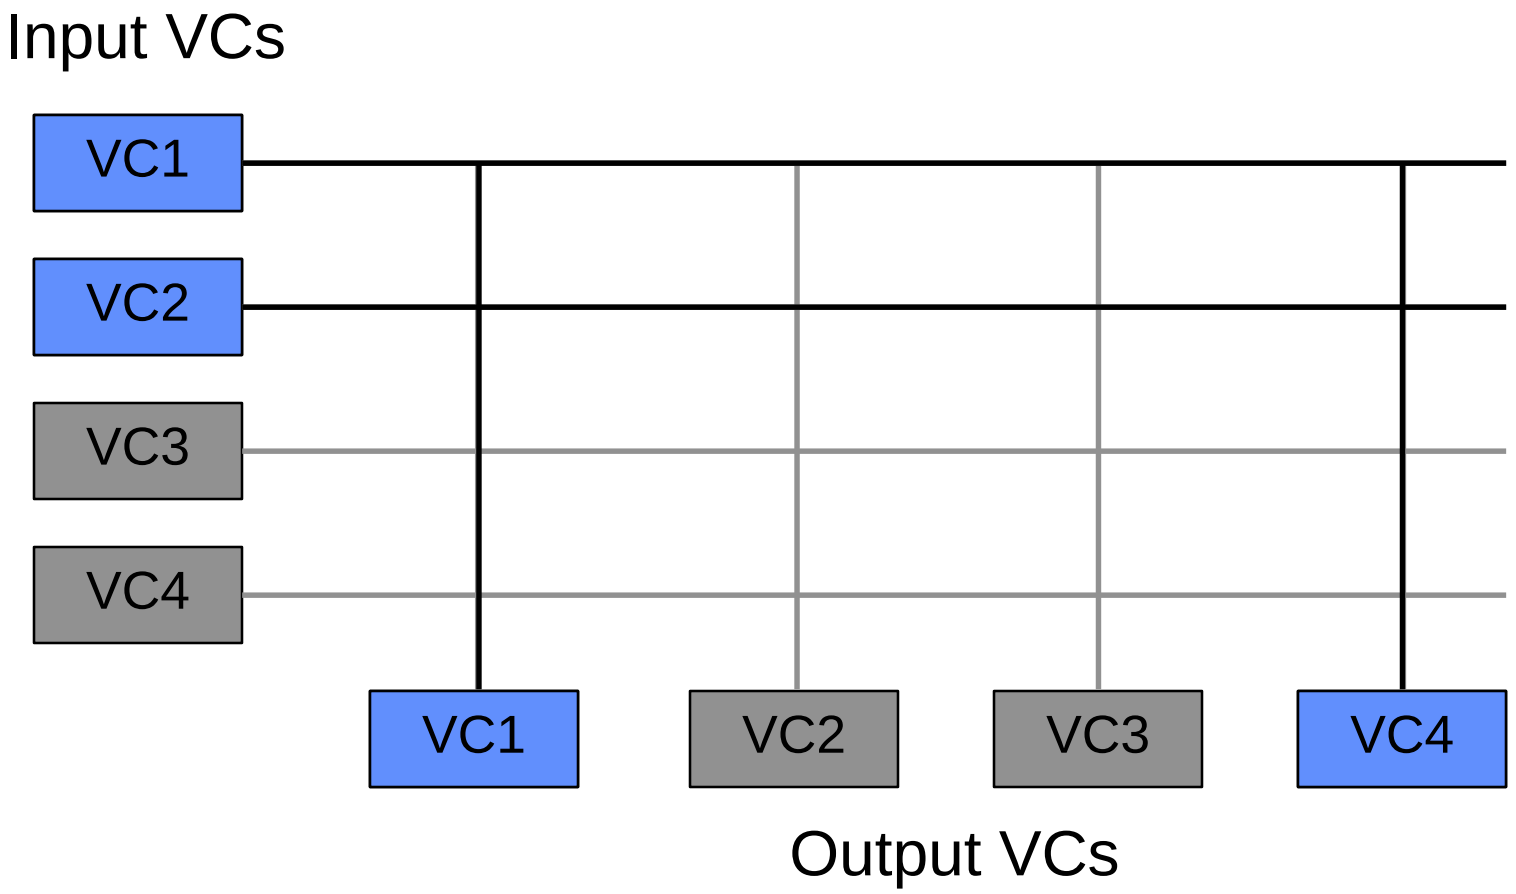

"Segment gating for static energy reduction in networks-on-chip". NoCArc 2009

Activation Mechanism

- Flits winning switch allocation (SA) activate in the next router:
  - Output channels and switch lanes (3 cycles)
  - Input buffers (1 cycle with drowsy SRAMs)
    - No false activations
  - With the below 4-stage router pipeline, no activation stalls

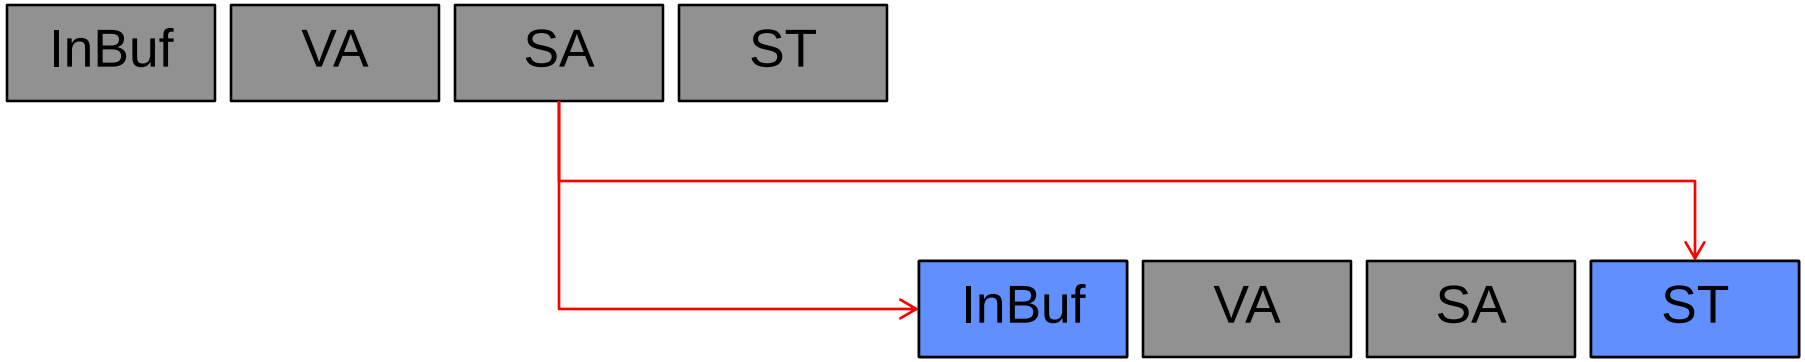

Router Impact

- ABN switch allcocators: (Inputs x VCs) x (Outputs x ChanLanes)
  - As long as ChanLanes no greater than VCs, switch allocator no more complex than VC allocator
  - If VC and switch allocators in different pipeline stages, router cycle time does not extend

- VC allocators consume 2-10mW and occupy 5000um
  - Both very small percentages of the router
  - Therefore increase of switch allocator's cost insigificant

"Allocator implementations for network- on-chip routers". SC 2009

Today's Menu

- Leakage power / motivation
- Related work
- Variable-width datapaths
- Results
- Conclusions

Methodology

- Booksim network simulator
- 8x8 Mesh. DOR
- We compare:
  - Single-lane: Single-lane power-gated network
  - ABN: Flits choose a lane based on their output VC
  - Multinets: Multiple power-gated subnetworks
- Router pipeline previously presented
- 2 VCs as baseline. Normalize for buffer size by adjusting VCs
- Activation and deactivation delays (65nm at 1GHz):
  - Channel and crossbar activation delay: 3 cycles
  - Channel and crossbar activation wait: 15 cycles
  - Channel and crossbar deactivation wait: 6 cycles
  - Buffer (VC) deactivation wait: 3 cycles

Two Subnetworks/Lanes. Static Power

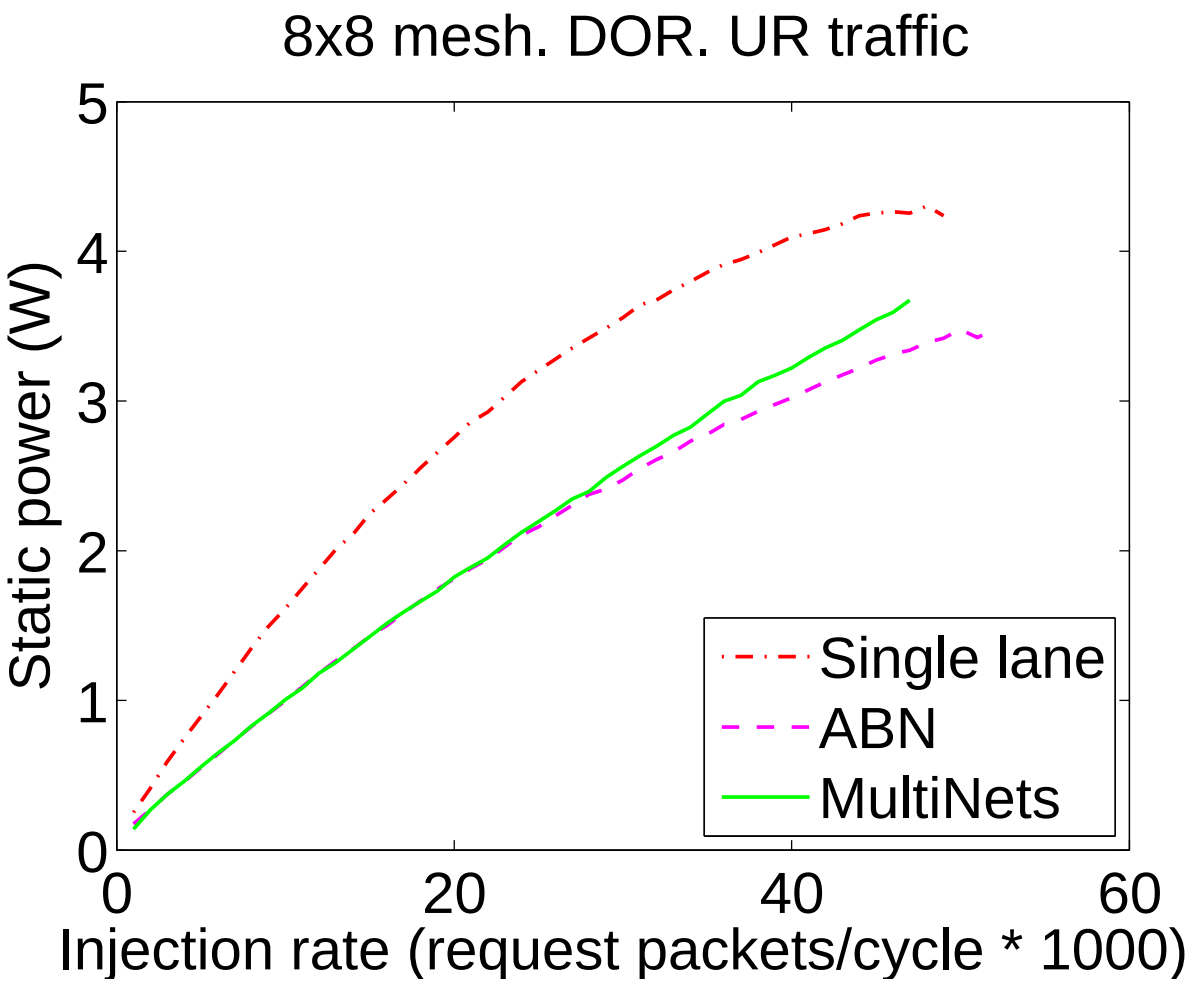

UR worst case for ABNs

ABNs better powers down resources at high loads

Two Subnetworks/Lanes. Dynamic Power

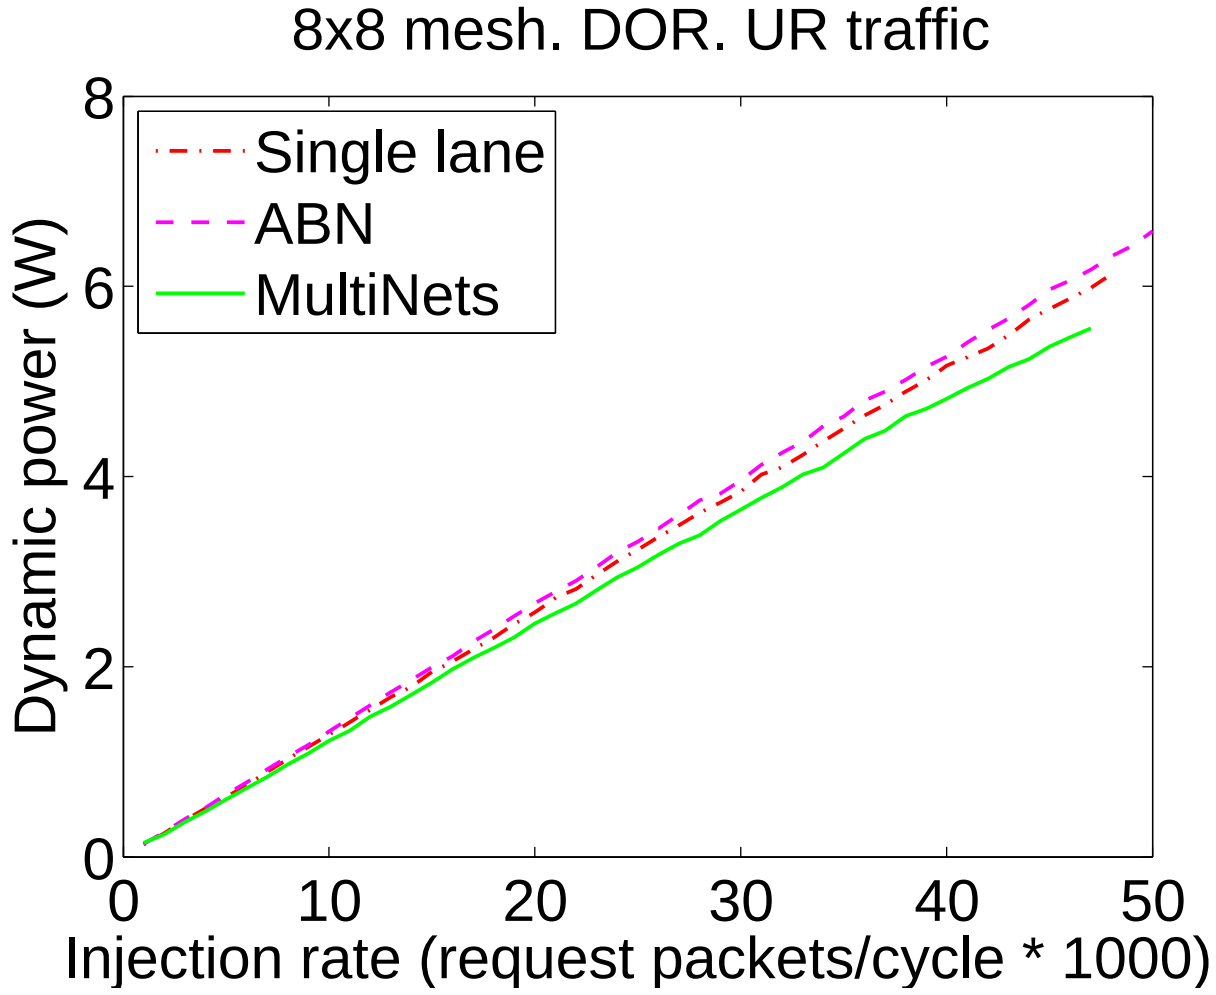

UR worst case for ABNs

Multinets takes advantage of lower-radix switches

Two Subnetworks. Latency

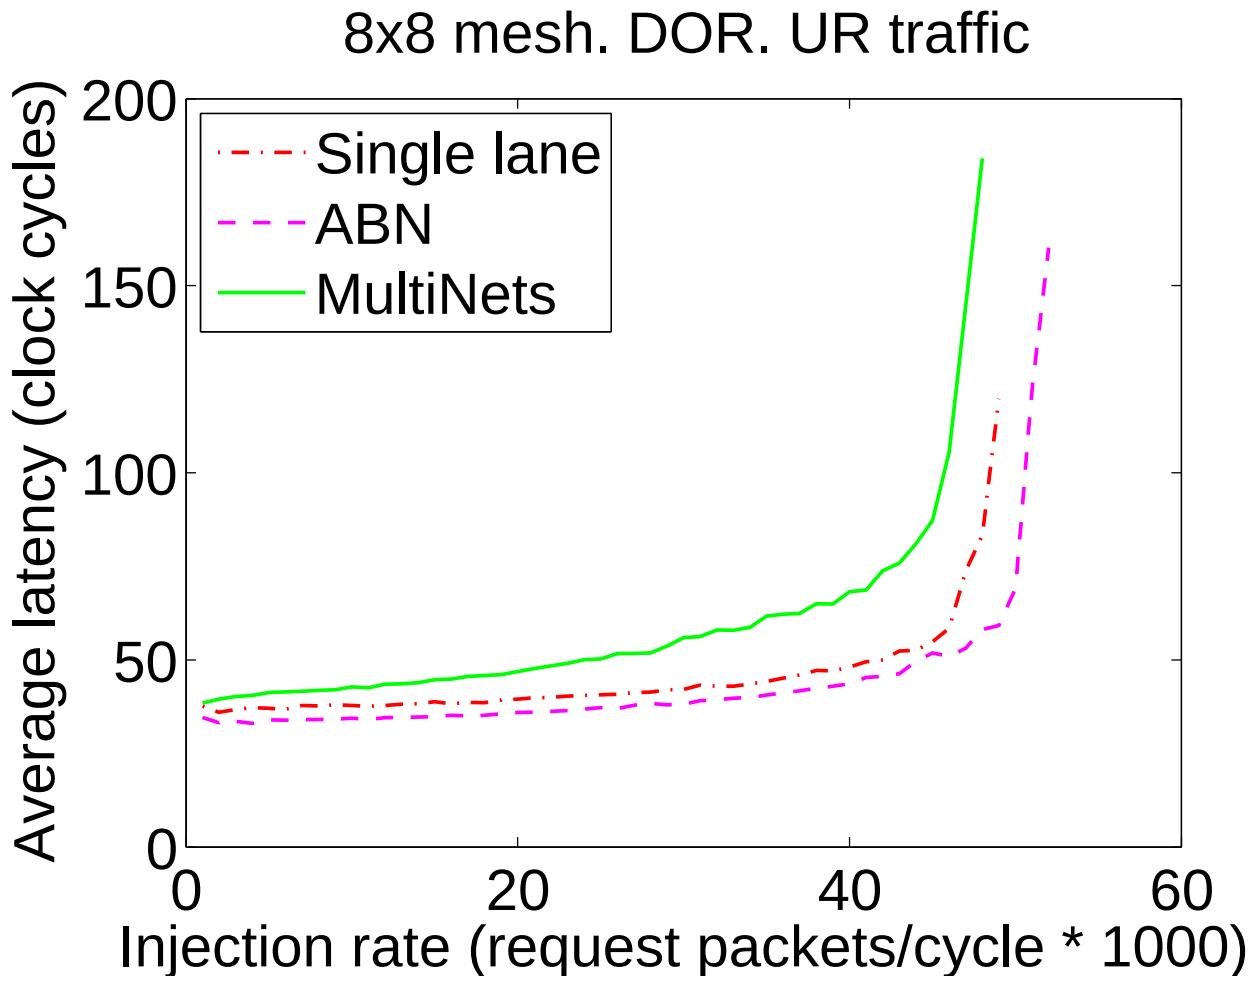

Multinets cannot make perfect injection decisions or use resources in another subnetwork after injection to combat transient imbalance

PARSEC Results

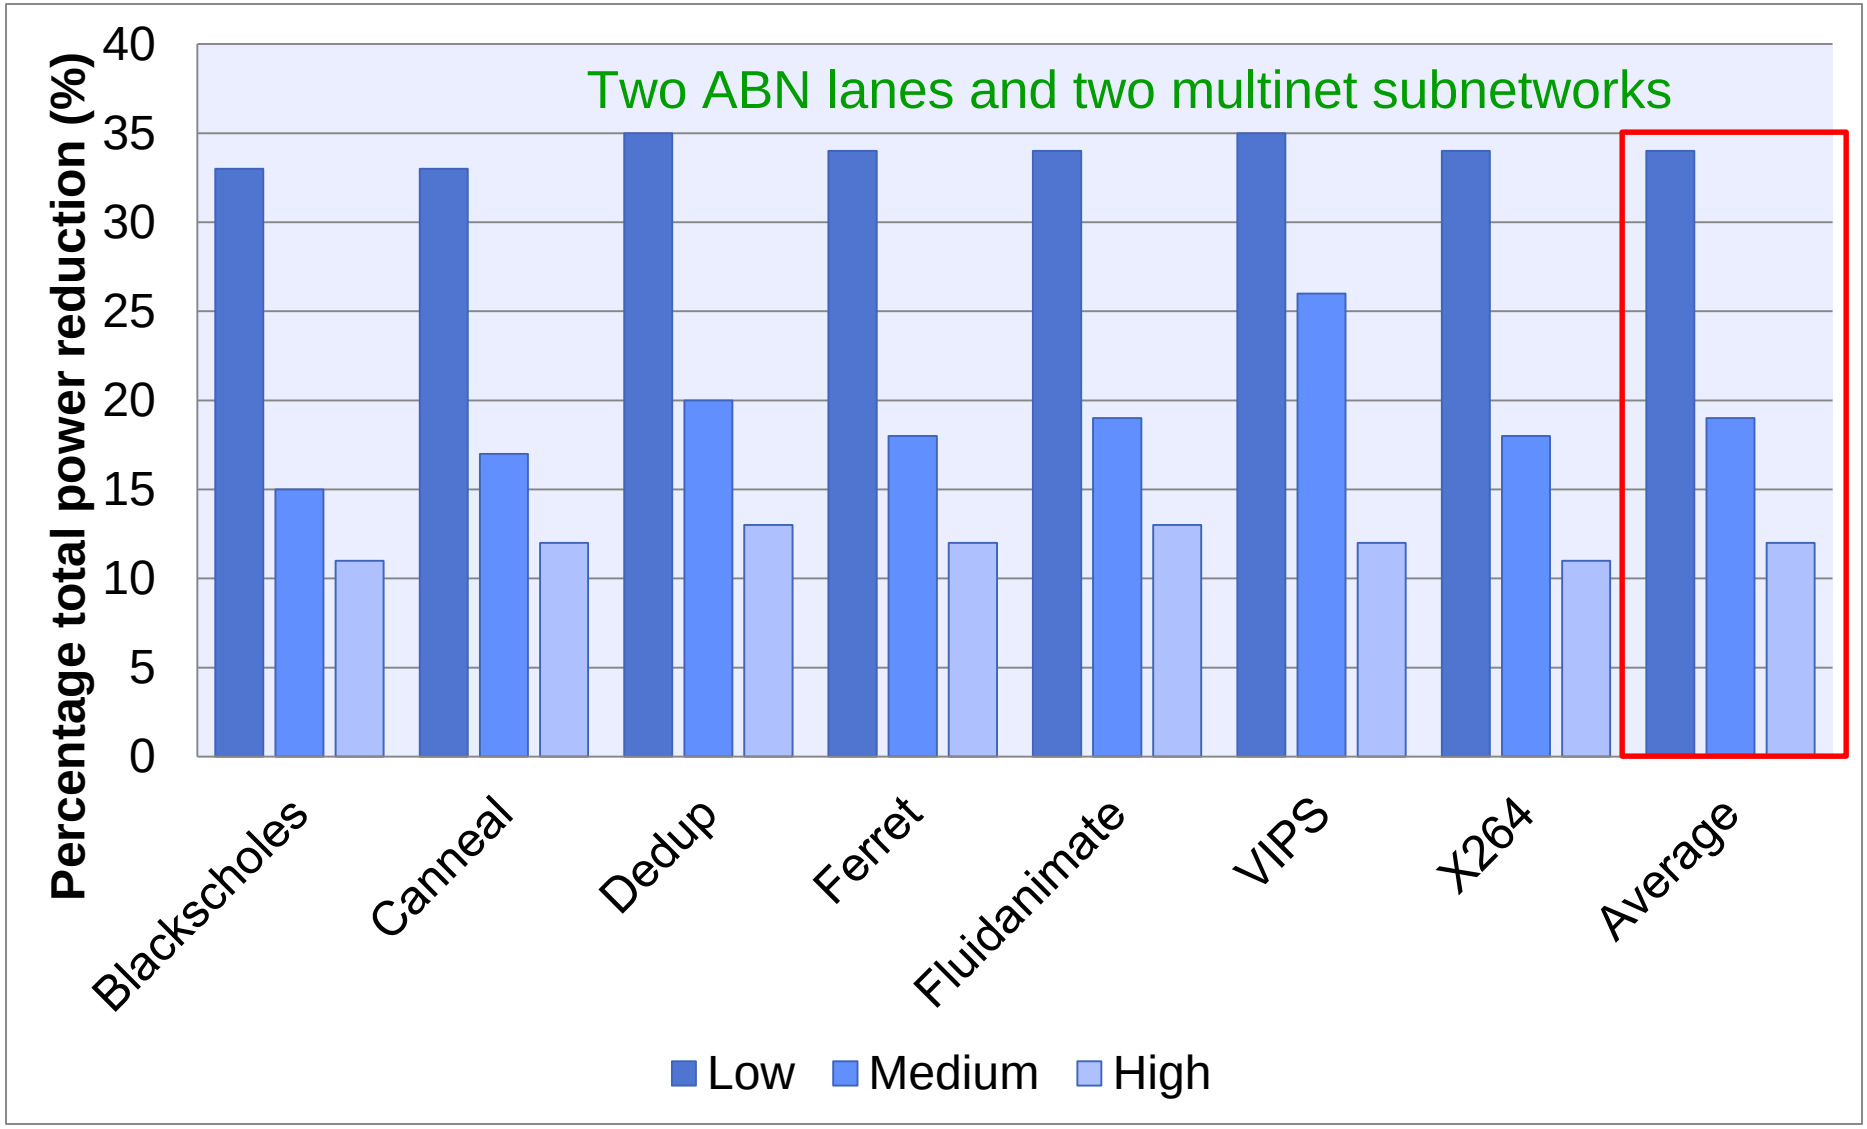

Scaling Up

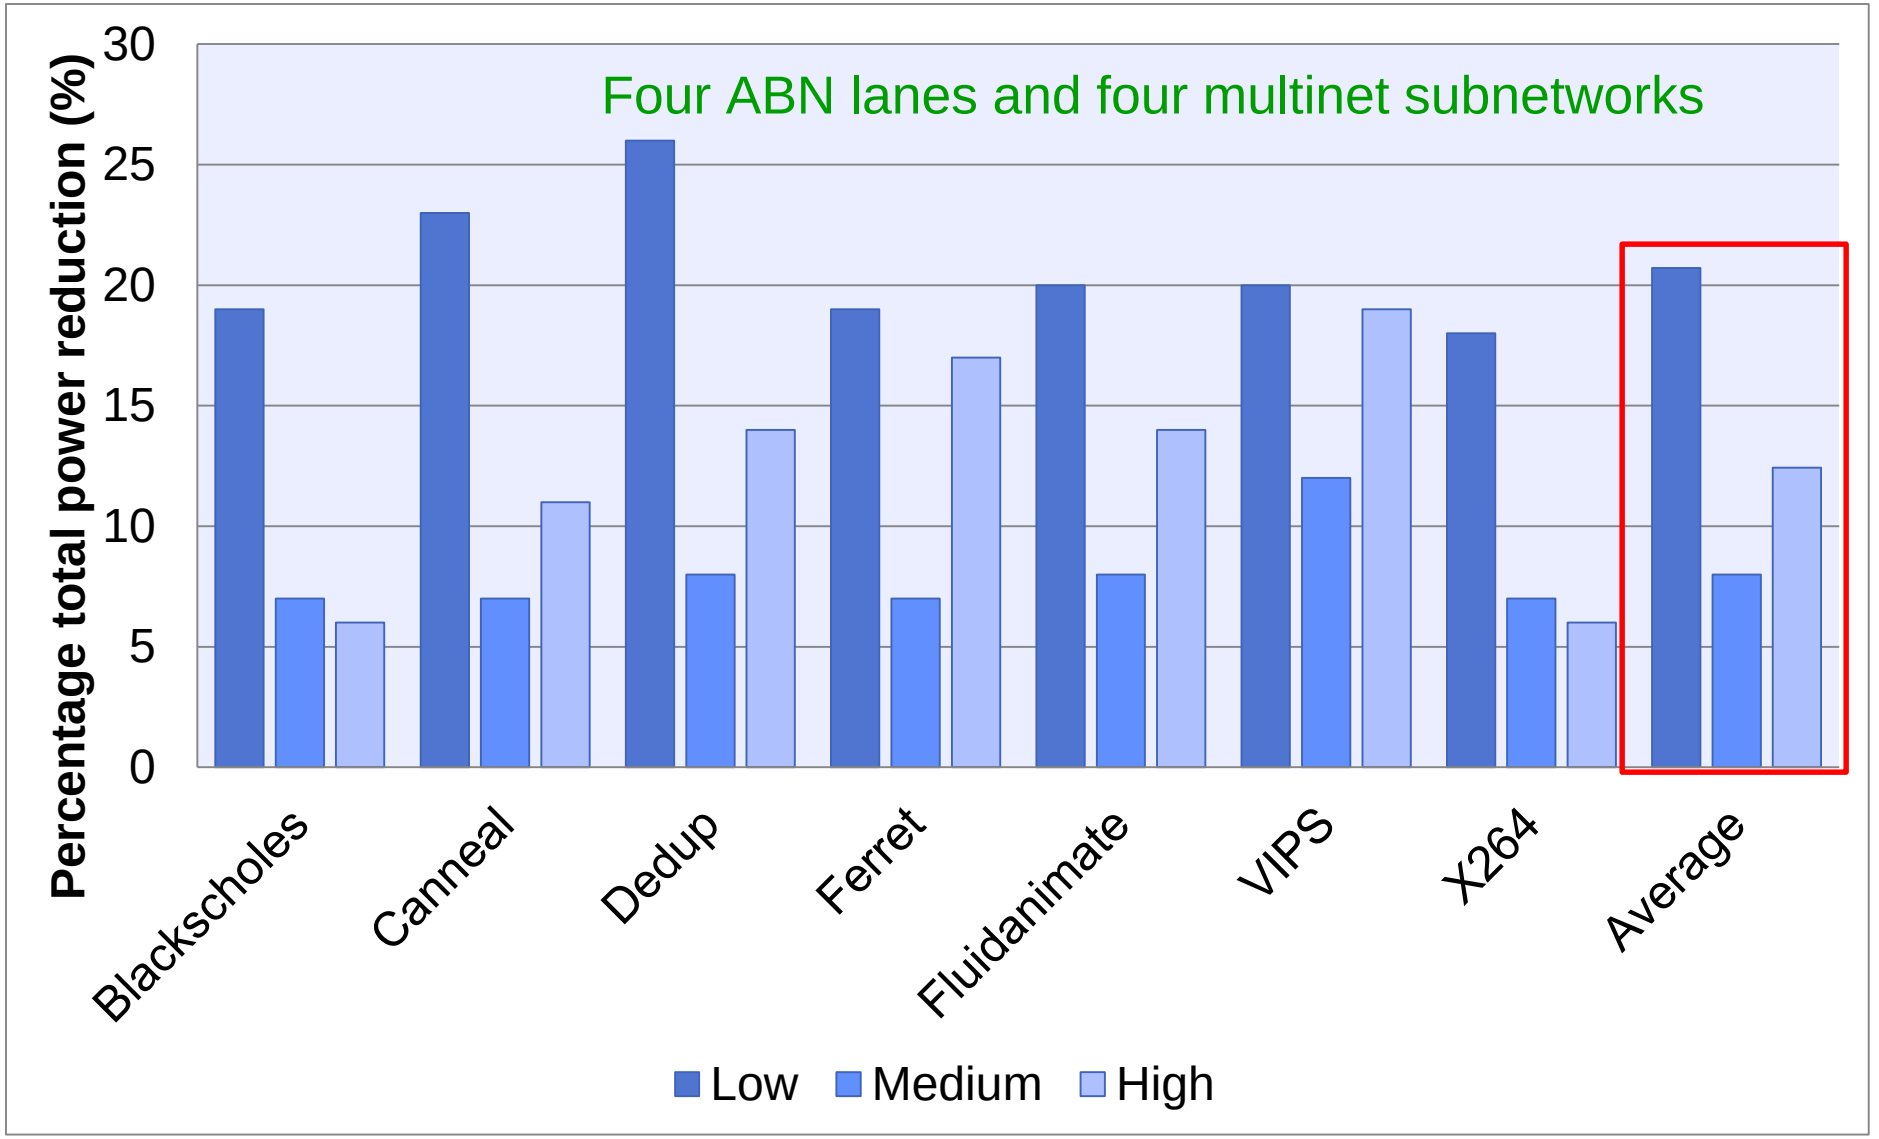
Four ABN lanes and four multinet subnetworks
Percentage total power reduction (%)
30
25
20
15
10
5
0
Blackscholes
Canneal
Dedup
Ferret
Fluidanimate
VIPS
X264
Average
Low
Medium
High

Scaling Up

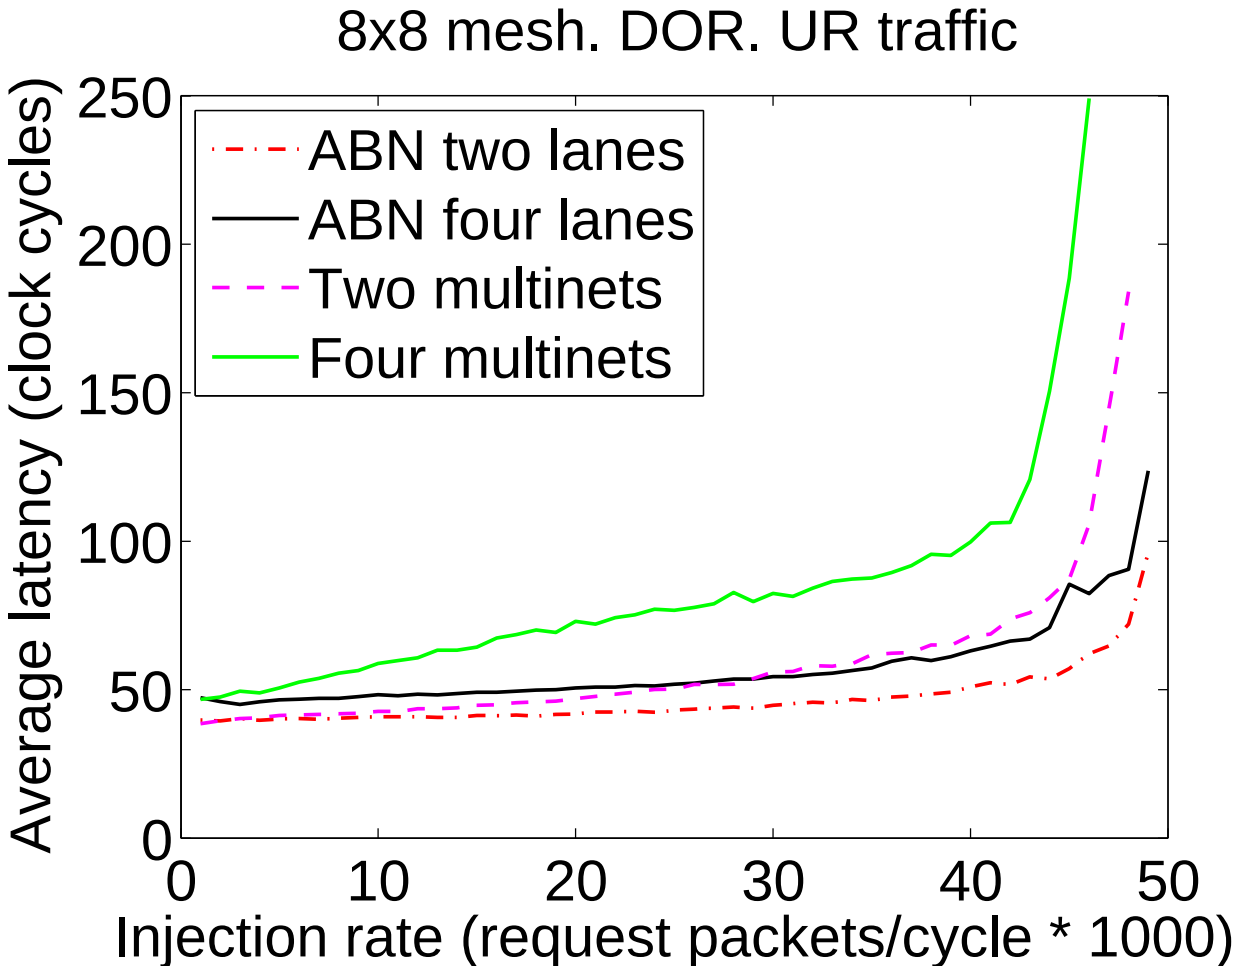

Effects of transient imbalance in multinets are intensified

Today's Menu

- Leakage power / motivation
- Related work
- Variable-width datapaths
- Results
- Conclusions

Conclusions

- Leakage power is a growing concern in future technologies
- Dividing datapaths in lanes provides more flexibility than multi-network approaches
  - But there are tradeoffs
- Using drowsy SRAMs allows hiding the activation delay without false activations
  - Can change with shallow router pipelines
- We demonstrate an average of 33% total power reduction with PARSEC benchmarks

Questions?

Acknowledgment: D.O.E. office of science