Adaptive Inter-router Links for Low-Power, Area-Efficient and Reliable Network-on-Chip (NoC) Architectures

Avinash Karanth Kodi†, Ashwini Sarathy*, Ahmed Louri* and Janet Wang*

†Department of Electrical Engineering and Computer Science, Ohio University, Athens, OH 45701

*Department of Electrical and Computer Engineering, University of Arizona, Tucson, AZ 85721

E-mail: kodi@ohio.edu, {sarathya, louri, wml} @ece.arizona.edu

Speaker : Jin Sun

Dept. of Electrical and Computer Engineering, University of Arizona, Tucson, AZ 85721

Sponsored: National Science Foundation (NSF) grants CCR-0538945 and ECCS-0725765
(at the High Performance Computing Architectures and Technologies Lab, University of Arizona, Tucson)

Talk Outline

- **Motivation**
- **iDEAL – inter-router Dual-function Energy and Area-efficient Links for NoC architectures**
  - **Dual-function links**
  - **Design of static & dynamic router buffer**
- **Performance Evaluation**
  - **Power and area estimation**
  - **Simulation results (power, throughput)**
  - **Design headroom for Negative Bias Temperature Instability (NBTI) related reliability issues**
- **Conclusion**

Motivation for Networks-on-Chips (1/2)

- **Technology scaling - Increase in transistor density on a single chip**

- **Chip MultiProcessors (CMPs) and MultiProcessor System-on-Chip (MPSoC) architectures**

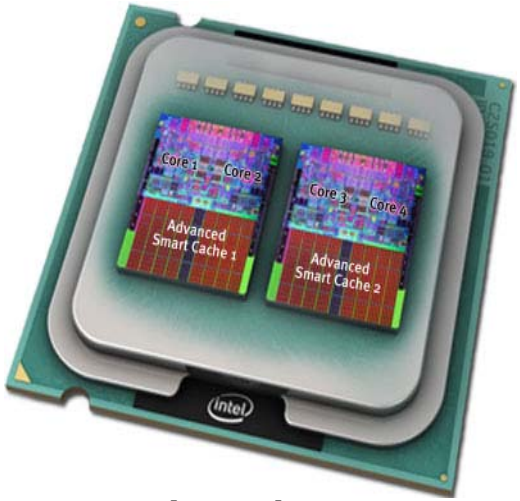

Intel quad core

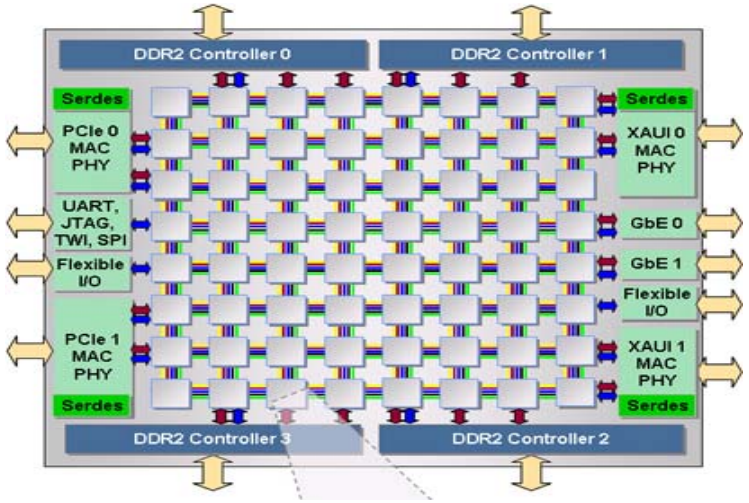

Tilera TILE64

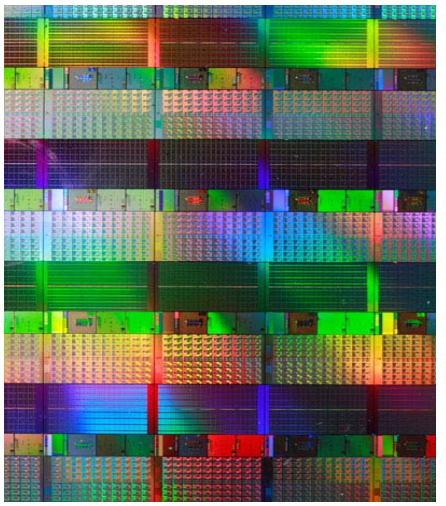

Intel 80-tile Polaris - Wafer

- **Challenges in future scaling of CMPs and MPSoCs :**
  - **Overcoming the increasing wire delays**
  - **Minimizing the power consumption**

Motivation for Networks-on-Chips (2/2)

Wire delay constraint with decreasing feature size

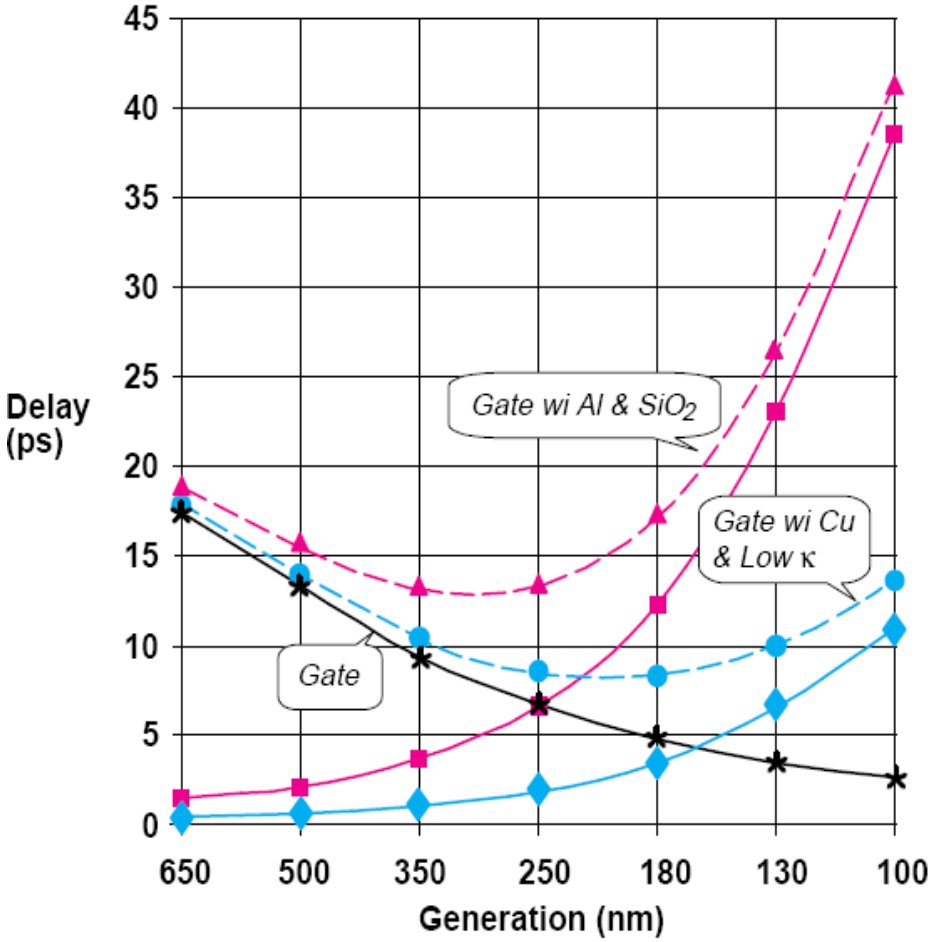

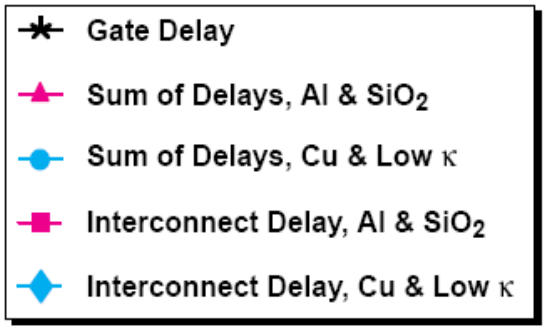

Al	*3.0 μΩ –cm*
Cu	*1.7 μΩ –cm*
SiO_2	*κ = 4.0*
Low κ	*κ = 2.0*
Al & Cu	*.8 μ Thick*
Al & Cu Line	*43 μ Long*

- ***RC*** **delay of local wires increasing by 32 % per year[2]**
- **Cross-chip (global) wire delays increasing by 49 % per year[2]**

1. 1997 Roadmap of the Semiconductor Industry Association
2. W. J. Dally and J. W. Poulton, ***Digital Systems Engineering***, Cambridge University Press, NY, USA, 2001.

The Network-on-Chip (NoC) Paradigm

Modular and Scalable NoC architectures

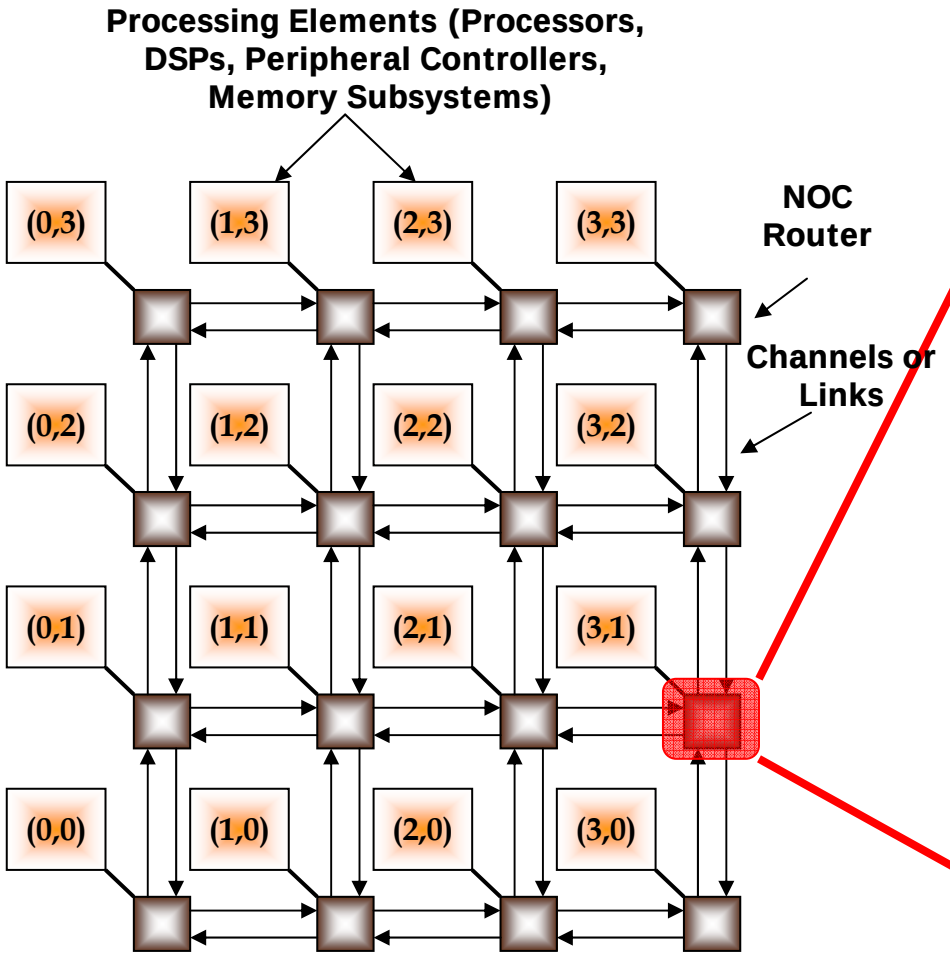

A Generic NoC Router

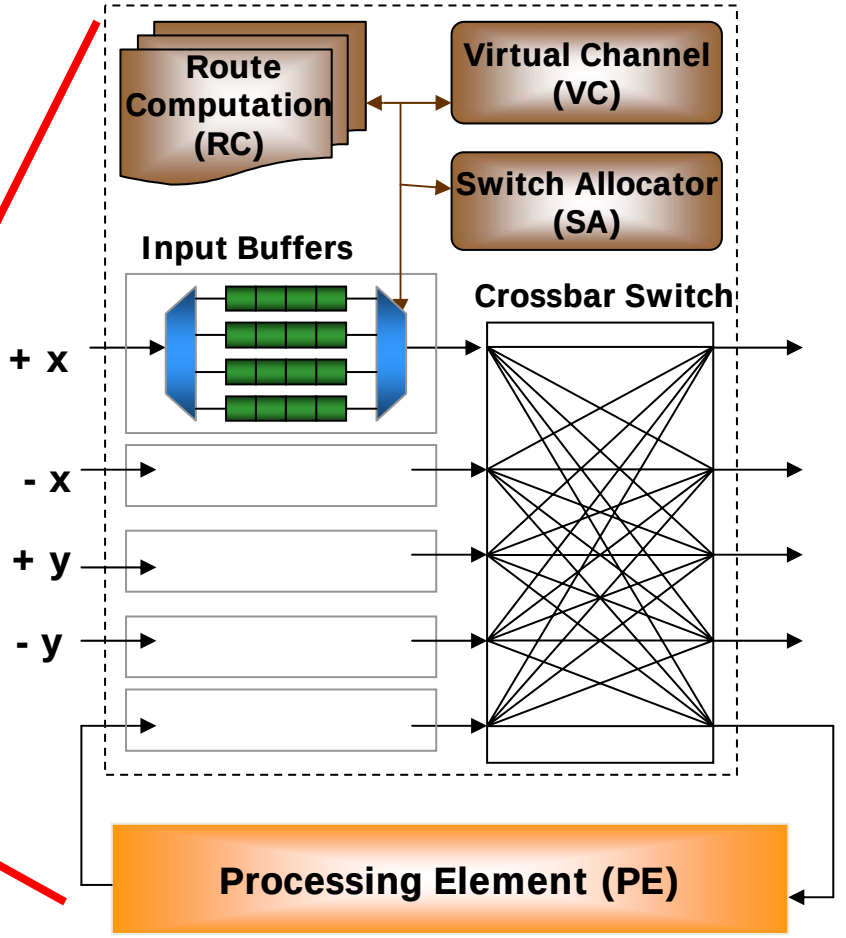

Significance of Router Buffers in NoCs

Recent NSF-sponsored workshop on On-Chip Interconnection Networks[1] :

- **"The most important technology constraint for on-chip networks is power consumption".**
- **Power consumption of OCINs implemented with current techniques exceeds expected needs by a factor of 10.**

Power Break-up in the NoC Router [2]

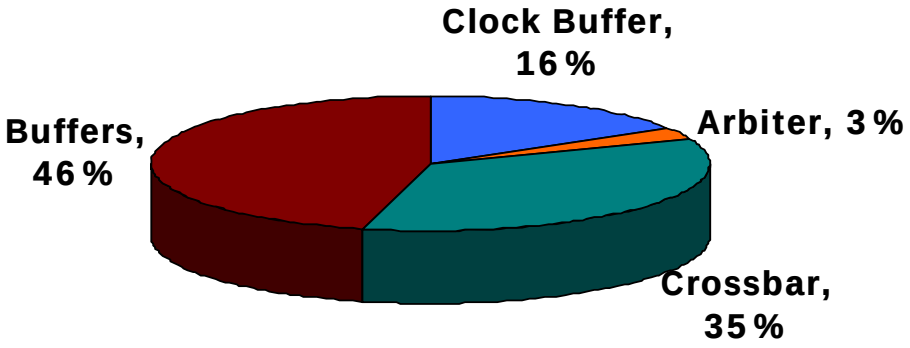

Area profile of a generic NoC Router [2]

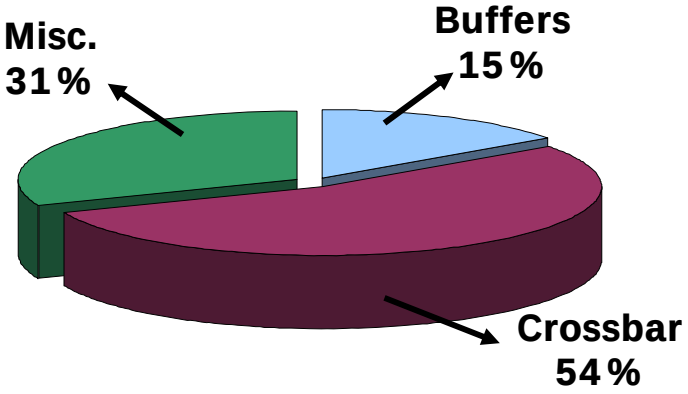

1. J.D.Owens, W.J.Dally, R.Ho, D.N.Jayasimha, S.W.Keckler and L.S.Peh, "Research Challenges for On-Chip Interconnection Networks", IEEE Micro, vol. 27, no. 5, pp. 96 – 108, September-October 2007.

2. S.R.Vangal, *et al.*, "An 80-Tile Sub-100-W TeraFLOPS Processor in 65-nm CMOS", IEEE J. Solid-State Circuits, Vol. 43, no. 1, January 2008.

iDEAL – Inter-router Dual-function Energy and Area-efficient Links for NoC architectures

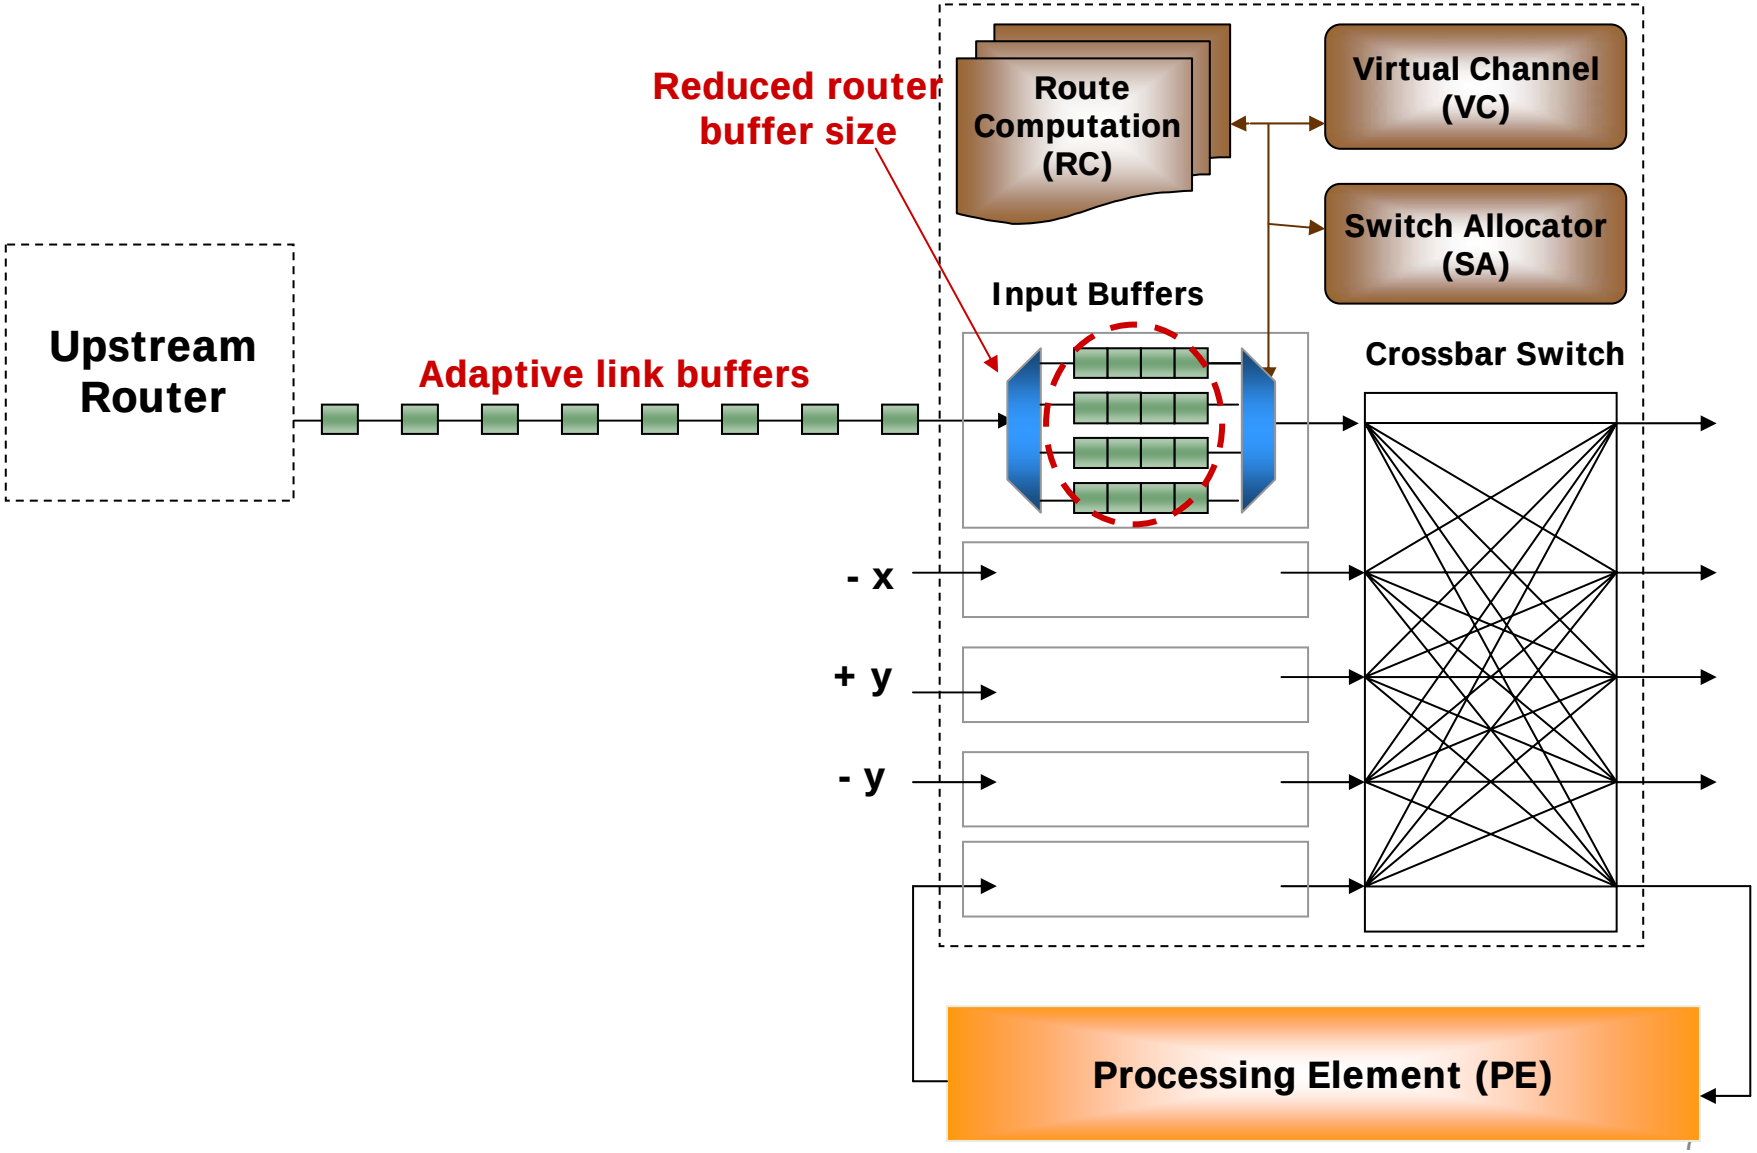

iDEAL – Inter-router Dual-function Energy and Area-efficient Links for NoC architectures

iDEAL Methodology

- Reduce the number of router buffers
(power and area savings)

- Use adaptive link buffers to store data along the links when required
(create more storage)

- Dynamic buffer allocation within the router buffers
(to sustain performance)

Conventional (Repeater-Inserted) Links

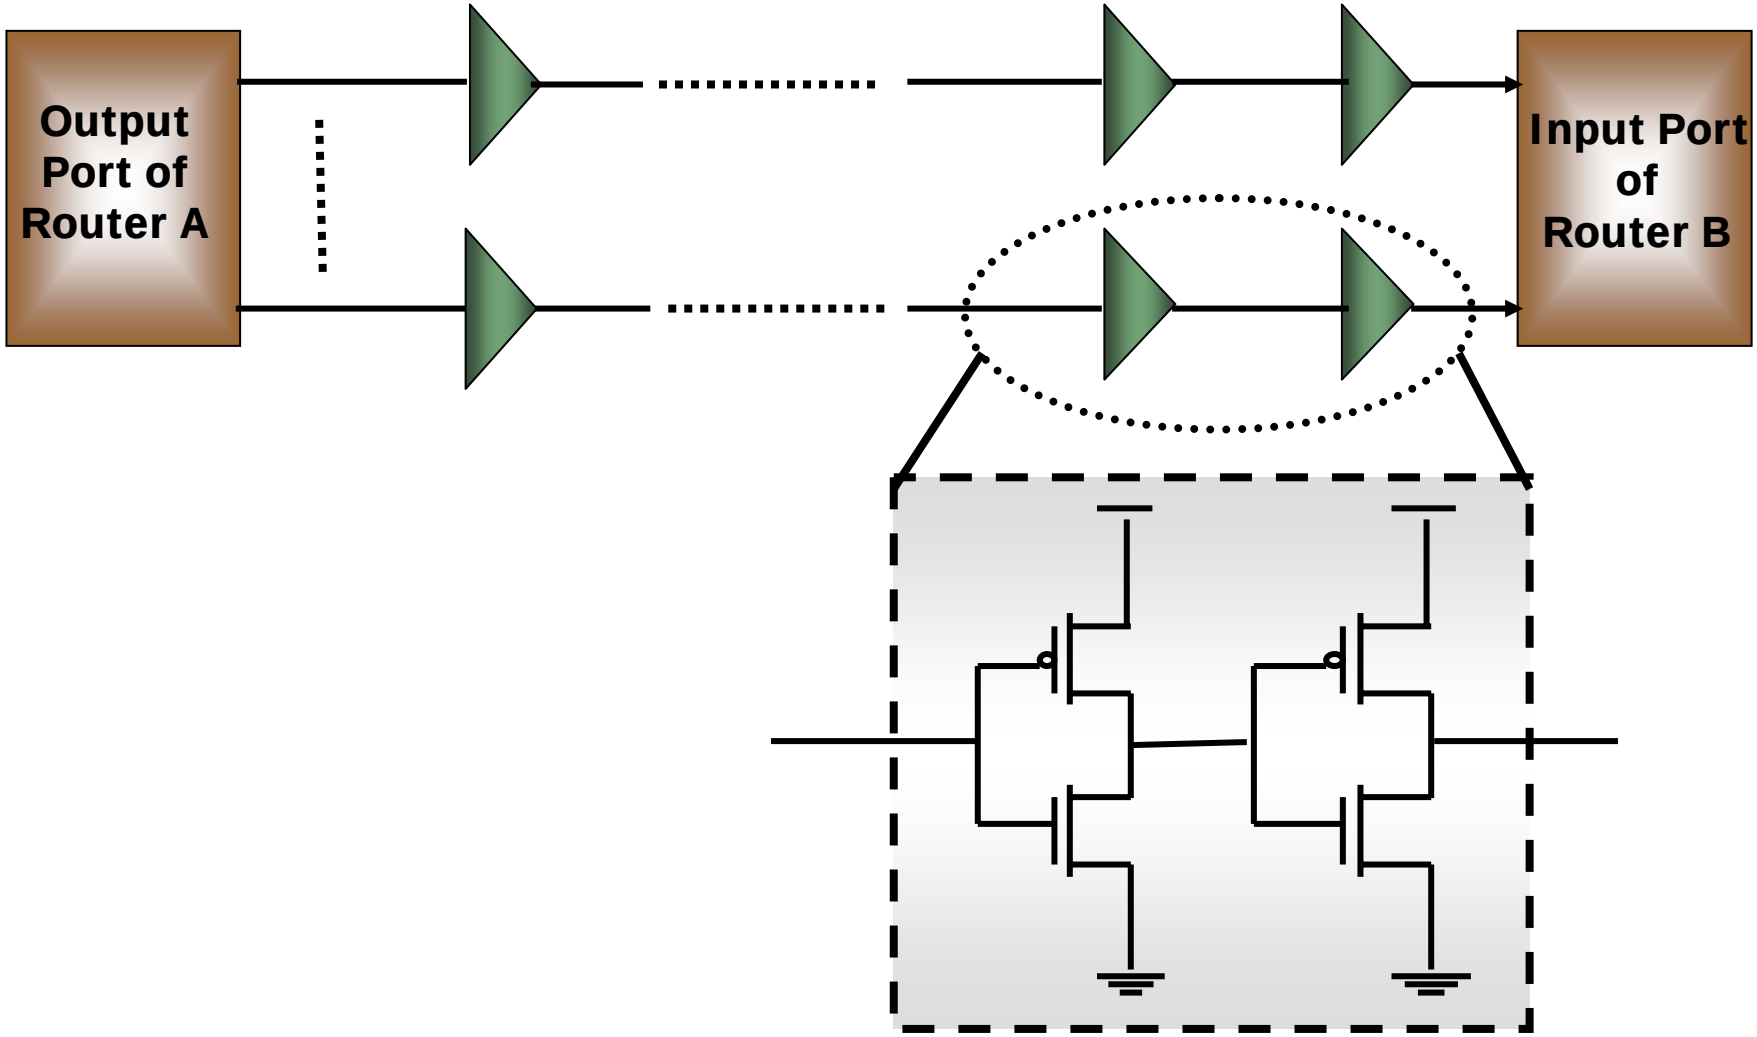

iDEAL – Link Buffer Design (1/2)

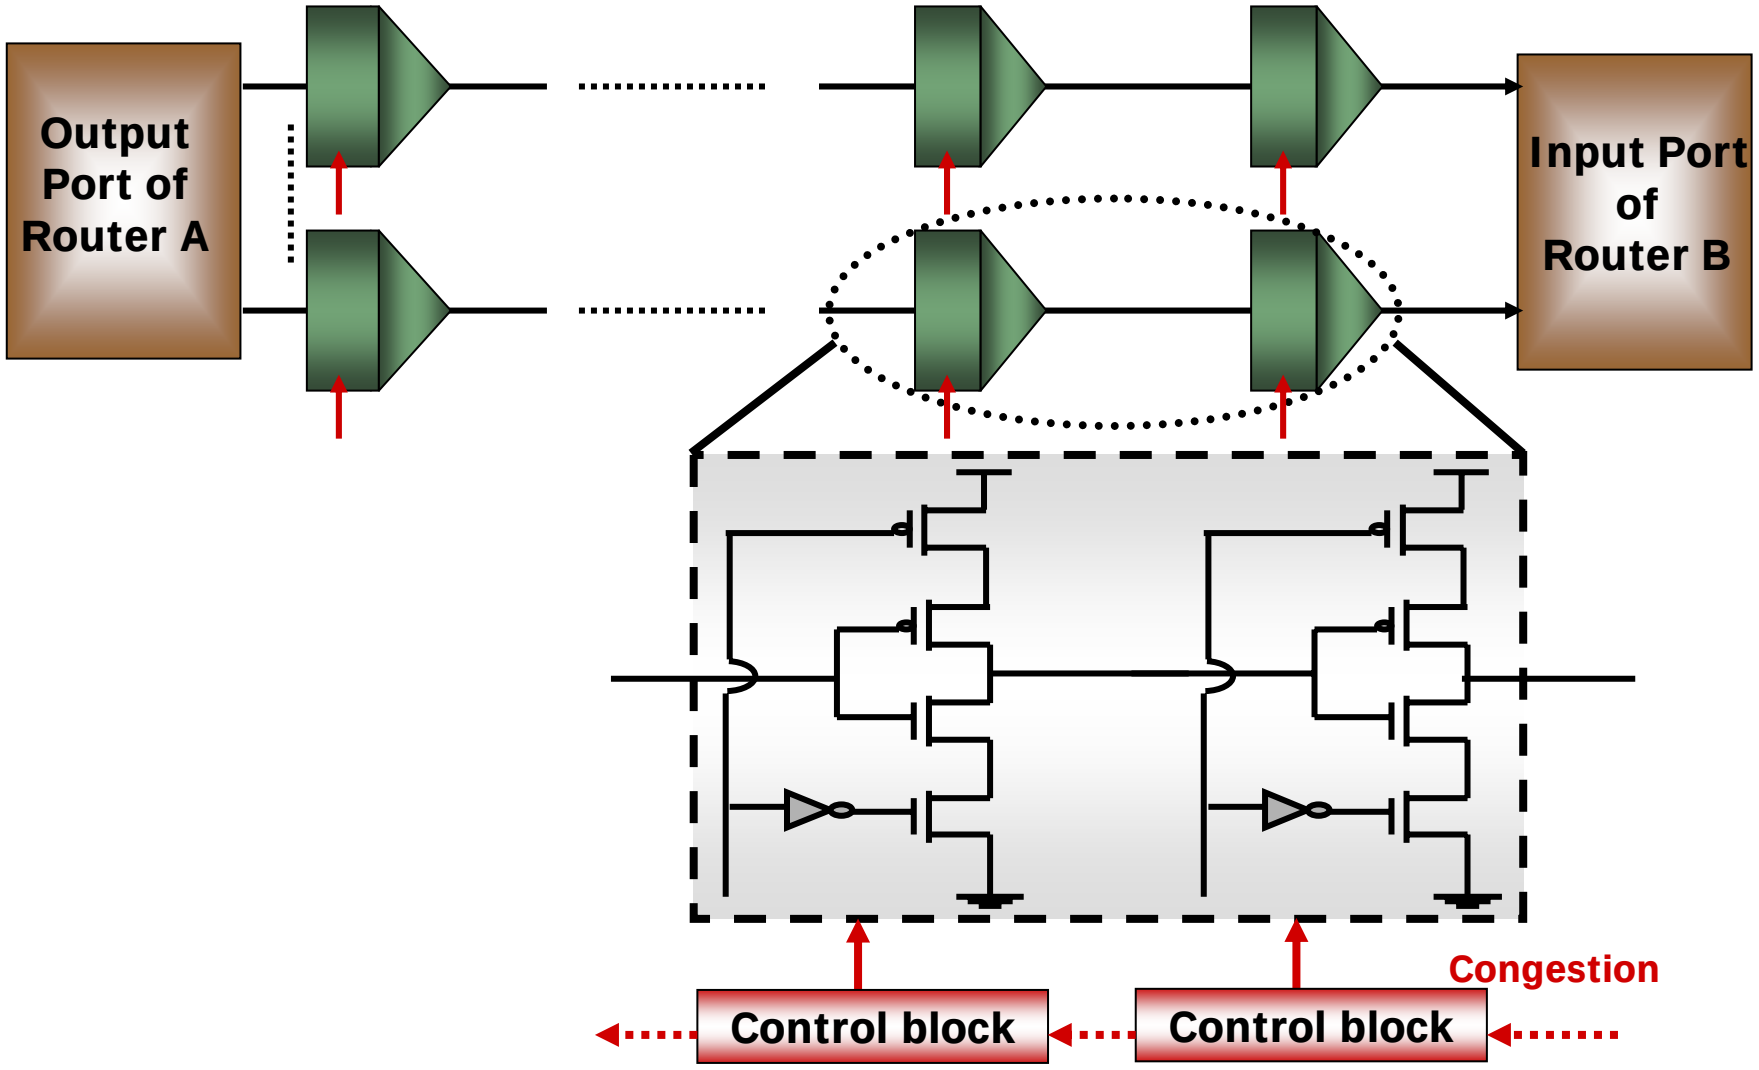

iDEAL – Link Buffer Design (2/2)

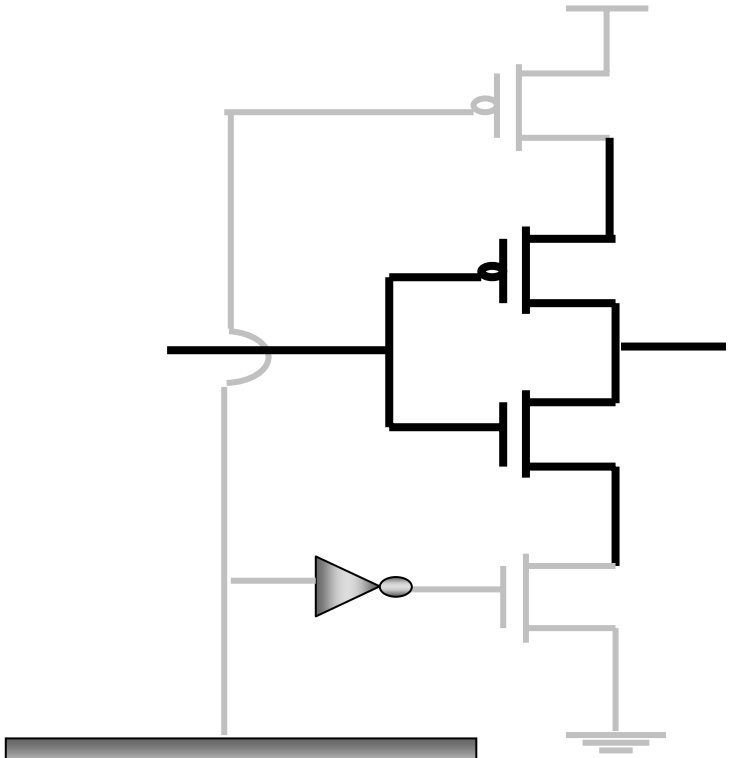

Control block

Functions as a conventional repeater when there is no congestion.

Control block is turned ‘OFF’.

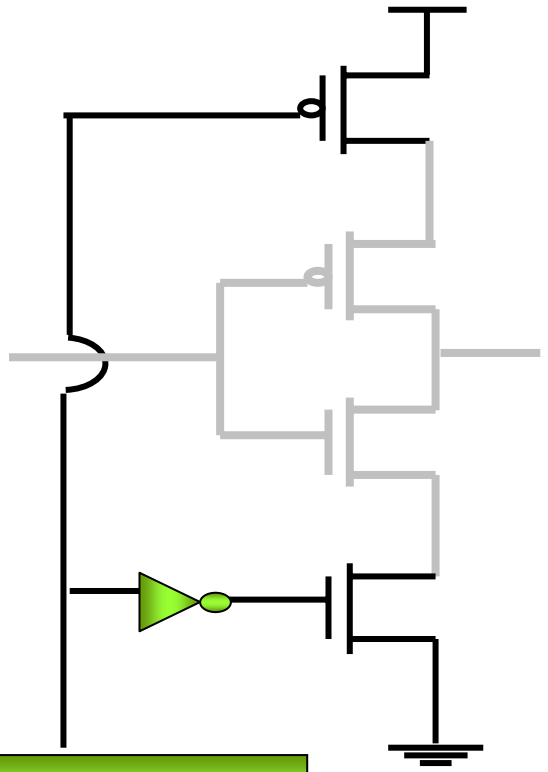

Control block

Repeater tri-stated and holds the sampled value, during congestion.

Control block is turned ‘ON’.

Control Block

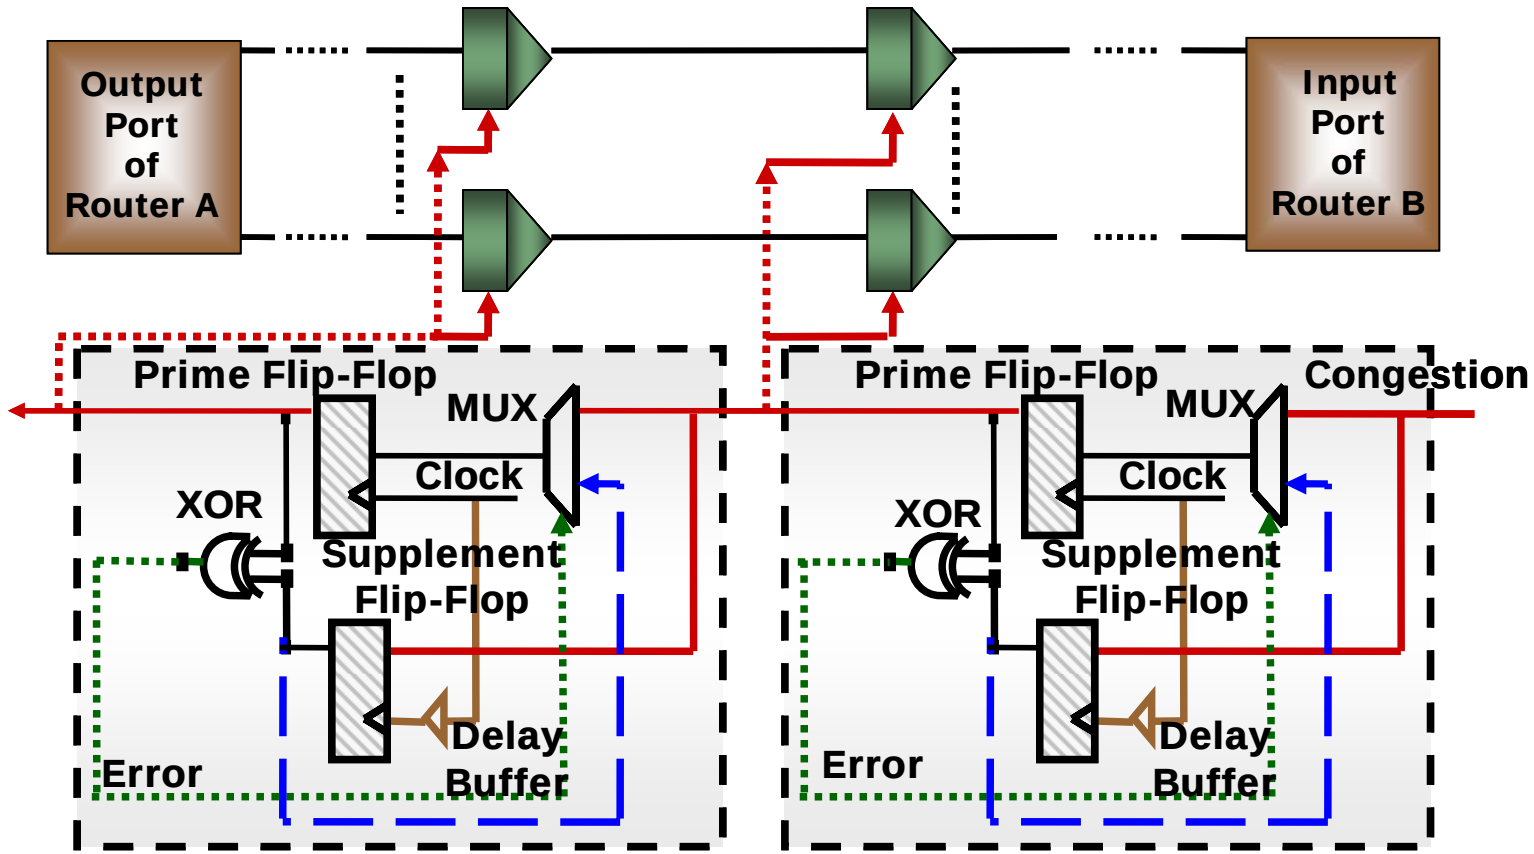

- Uses double-sampling technique[2] for stable error-free operation under varying frequencies
- Power consumption (about 6 µW when enabled)

2. M.Nicolaidis, "Time redundancy based soft-error tolerance to rescue nanometer technologies", Proc. 17th IEEE VLSI Test Symposium, San Diego, USA, April 25-30, 1999, pp. 86 – 94.

Dual-function Link

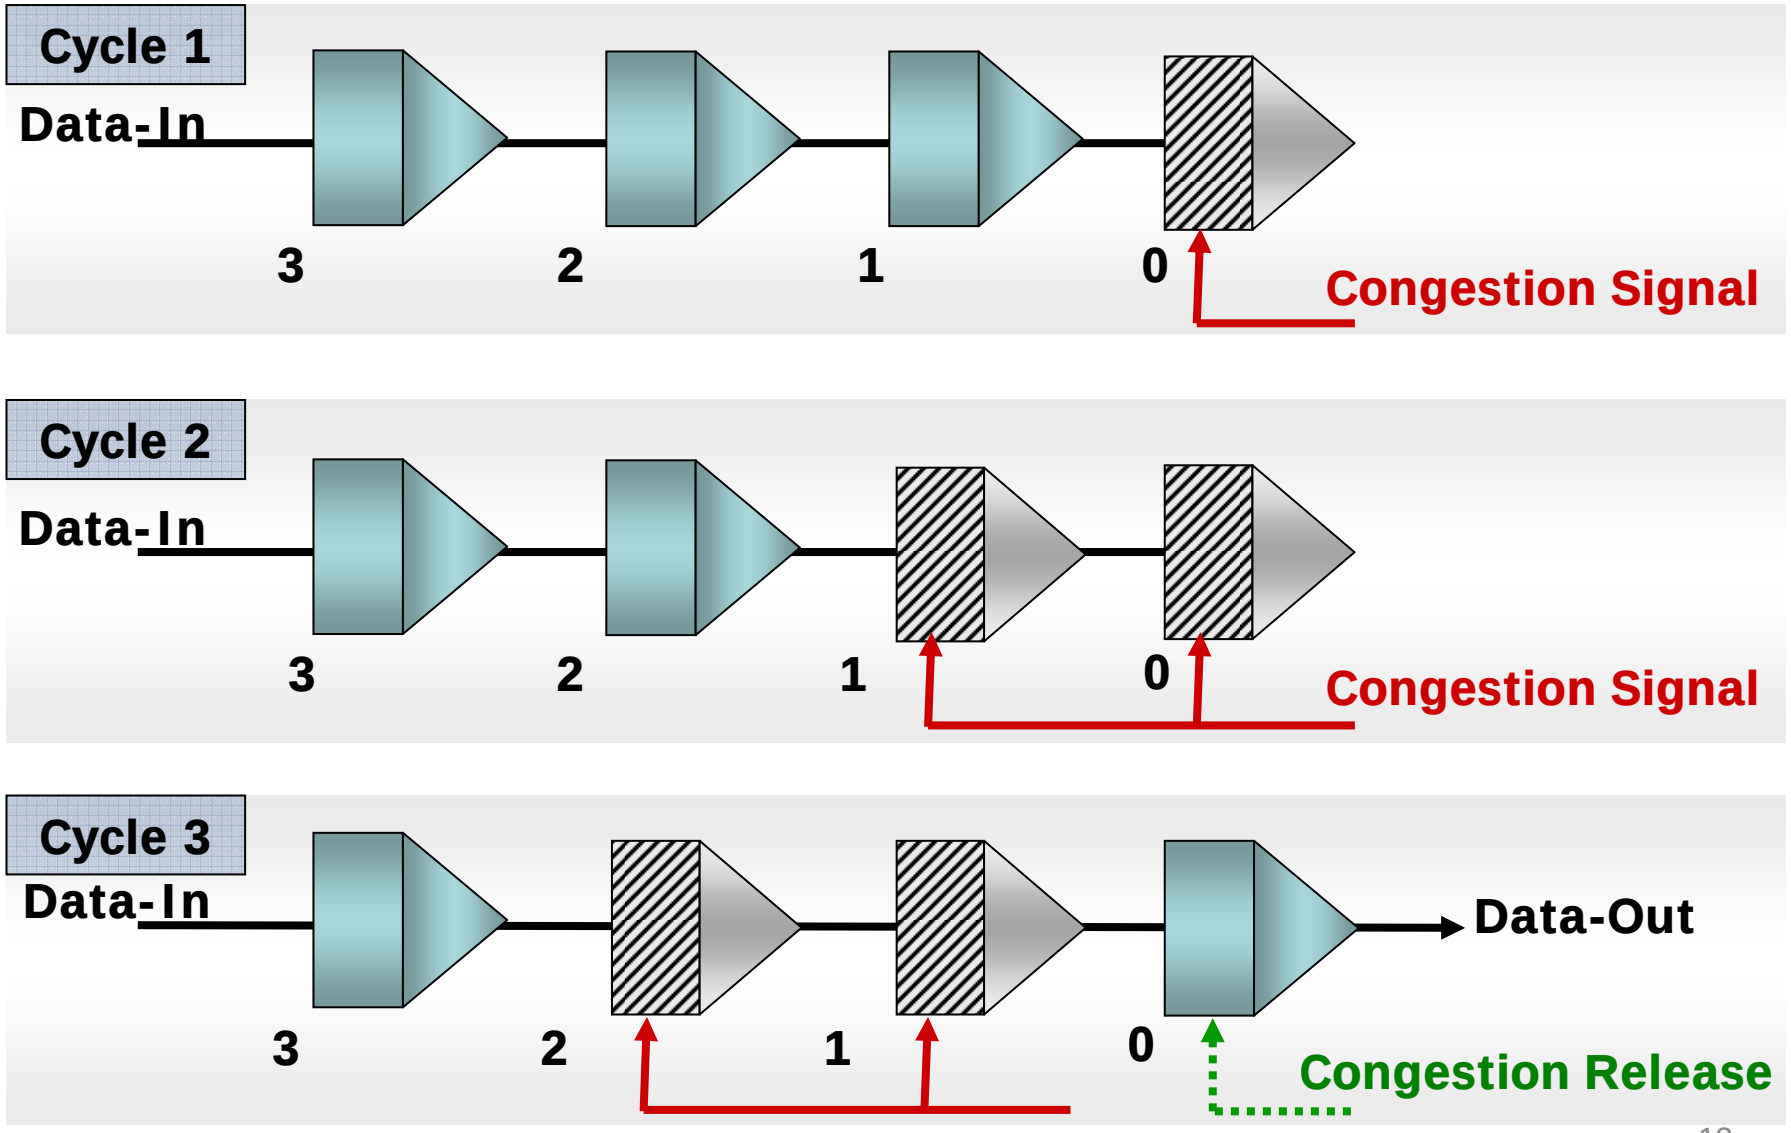

Statically Allocated Router Buffer

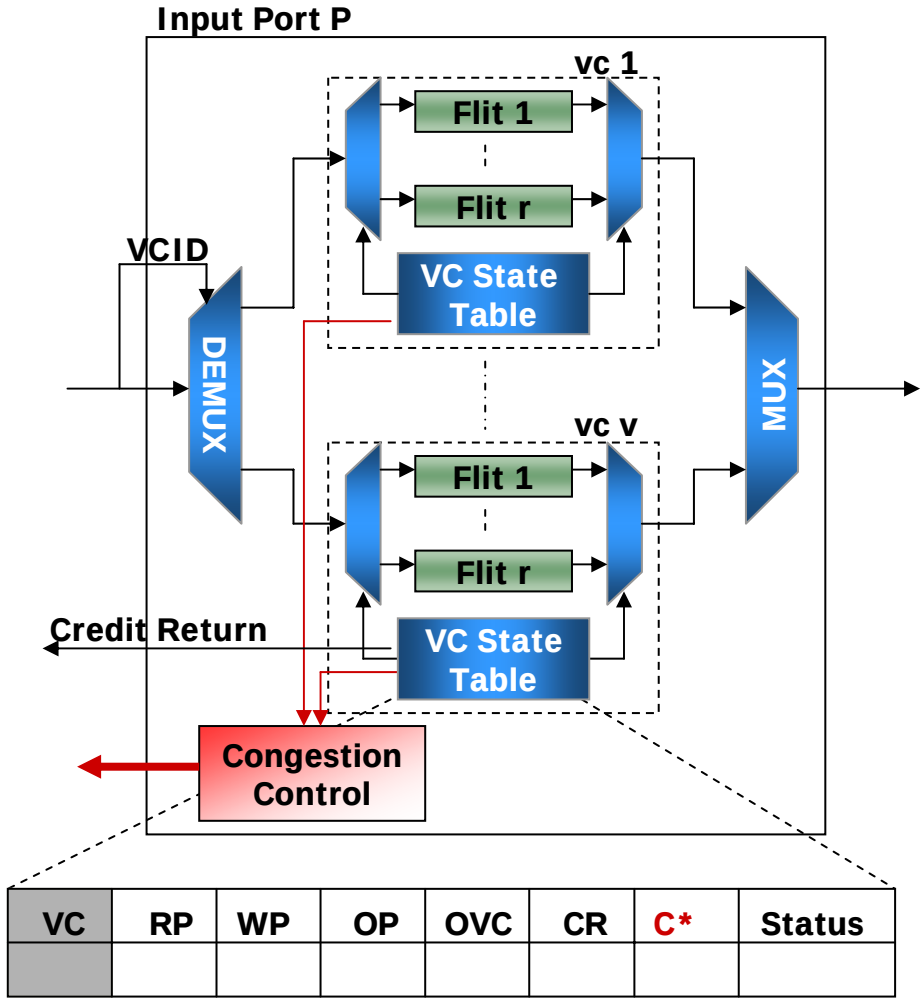

VC	RP	WP	OP	OVC	CR	C*	Status

- Static buffer allocation
  - Fixed number of buffers per VC
  - HoL blocking

RP = read pointer, WP = write pointer, OP = output port, OVC = output VC, CR = credits, C* = congestion
Status = status of the VC (idle, waiting, RC, VA, SA, ST)

Dynamically Allocated Router Buffer (1/2)

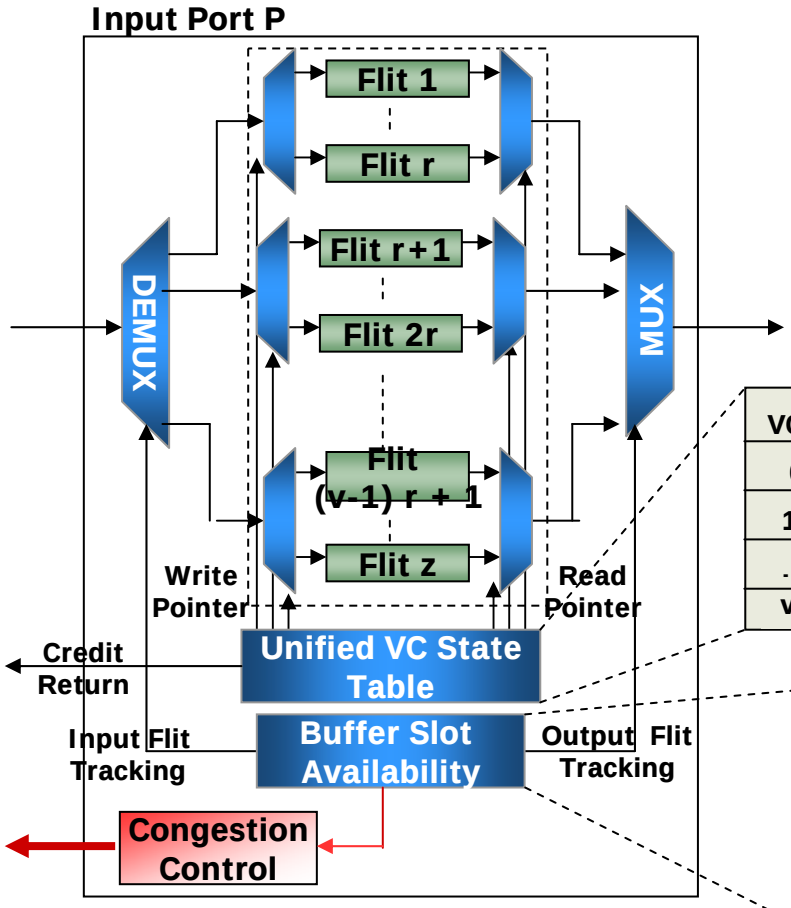

- Dynamic buffer allocation
  - Approximately (z + c)/v buffers per VC (z = router buffers, c = channel buffers, v = # of VCs)

VC	RP	WP	OP	OVC	CR	Status	F_0	F_1	...	$F_{(z+c)/v}$
0	3	N					N	3	...	N
1	6	N					N	6	...	N
...	...	...					...	...	...	...
v	5	N					N	5	...	N

Buffer Slot	Free
1	Y
2	N
⋮	⋮
z	N

RP = read pointer, WP = write pointer, OP = output port, OVC = output VC, CR = credits, C* = congestion
Status = status of the VC (idle, waiting, RC, VA, SA, ST)

Dynamically Allocated Router Buffer (2/2)

- **Example illustrating Dynamic buffer allocation in iDEAL**

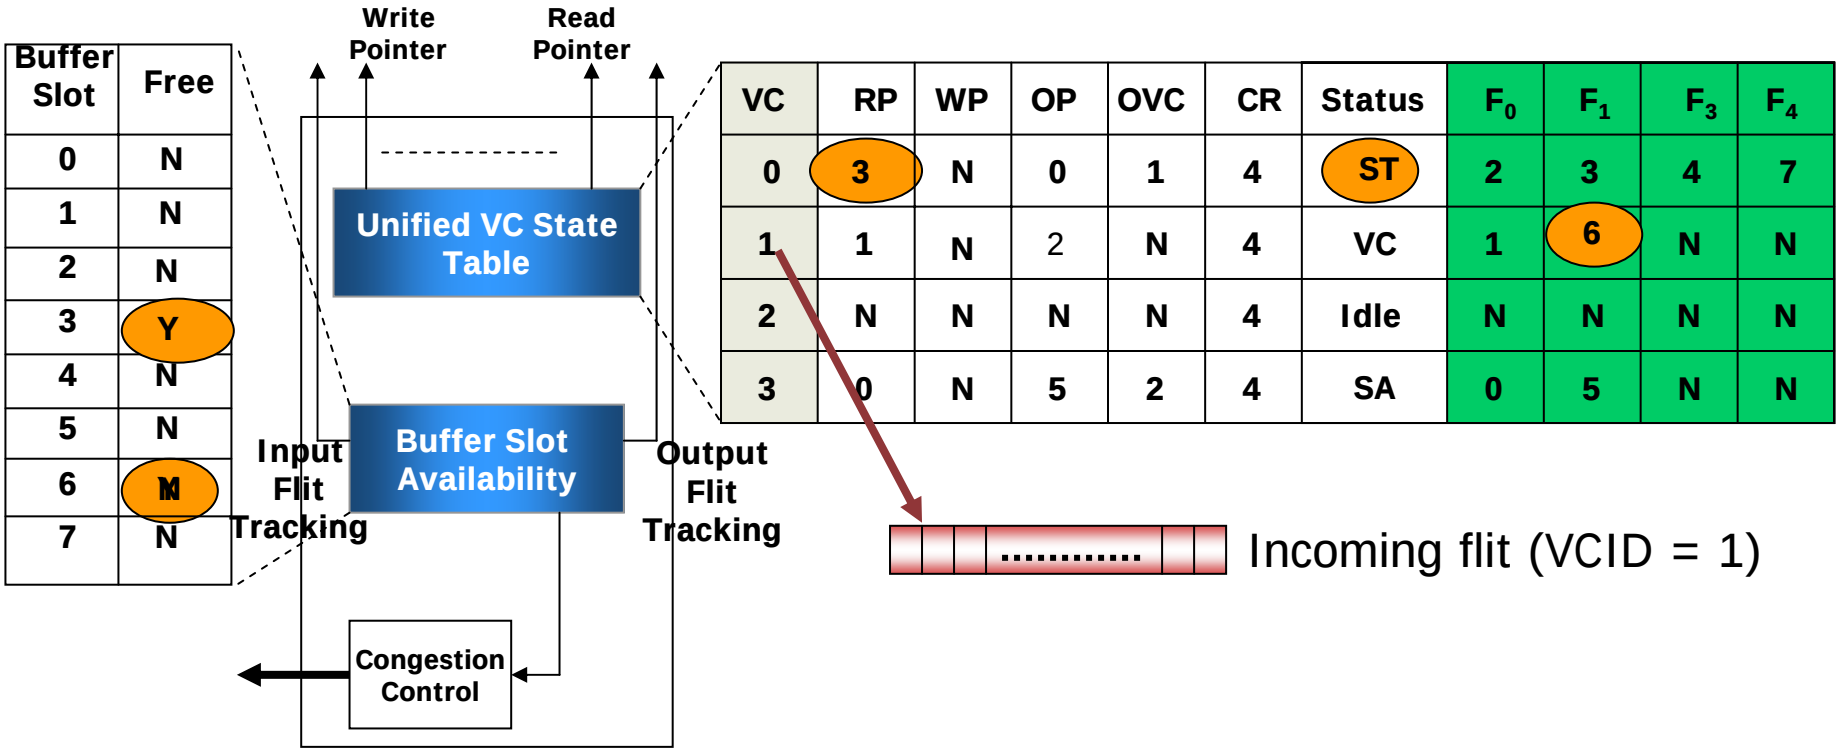

Buffer Slot	Free
0	N
1	N
2	N
3	Y
4	N
5	N
6	N
7	N

VC	RP	WP	OP	OVC	CR	Status	F_0	F_1	F_3	F_4
0	3	N	0	1	4	ST	2	3	4	7
1	1	N	2	N	4	VC	1	6	N	N
2	N	N	N	N	4	Idle	N	N	N	N
3	0	N	5	2	4	SA	0	5	N	N

RP = read pointer, WP = write pointer, OP = output port, OVC = output VC, CR = credits, C* = congestion
Status = status of the VC (idle, waiting, RC, VA, SA, ST)

Performance Evaluation

- Evaluated on a cycle-accurate on-chip network simulator
- Simulated 8 x 8 Mesh and 8 x 8 Folded Torus topologies
- Synthetic benchmarks such as uniform, and non-uniform workloads (Butterfly, Complement, Perfect Shuffle, Matrix Transpose, Bit Reversal) as well as SPLASH-2 suite benchmarks (FFT, LU, MP3D, WATER, RADIX) were evaluated
- Parameters evaluated include throughput, latency and overall network power
- Considered 5 different configurations – (vn_V – rn_R – cn_C)

  (n_V = No. of VCs per input port, n_R = No. of router buffers per VC, n_C = number of channel buffers)
  - Baseline = 440
  - 434, 428, 344, 531

Different Test Cases Considered

Baseline NoC Router with 4 VCs per input port, 4 128-bit router buffers per VC and no adaptive link buffers

v4-r4-c0

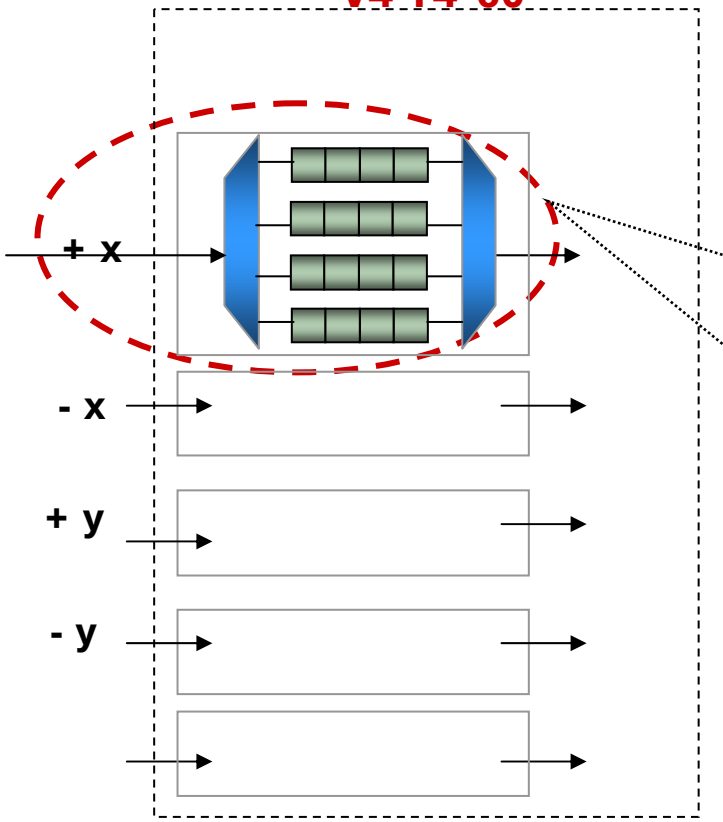

Case with 4 VCs per input port, 2 128-bit router buffers per VC and 8 adaptive link buffers - v4-r2-c8

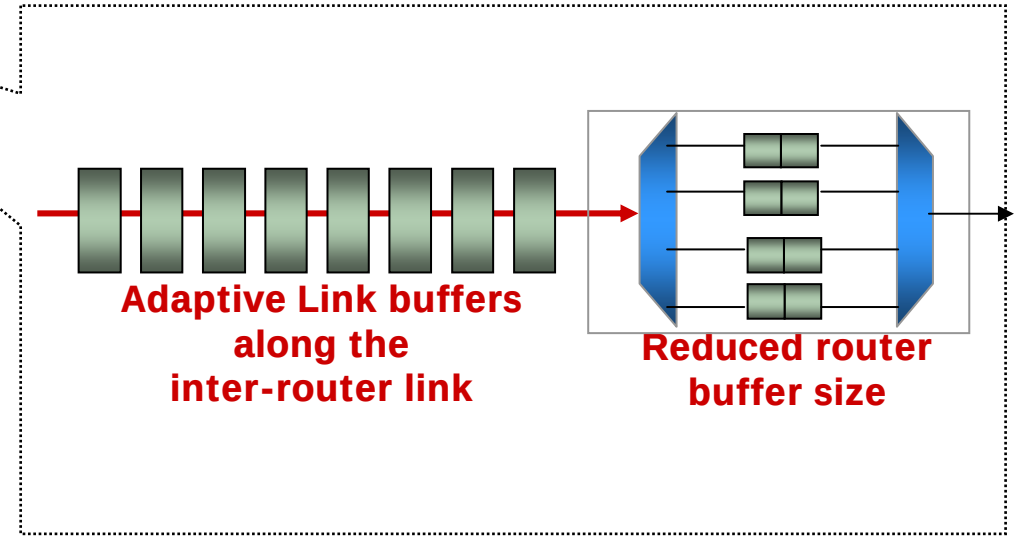

Power Estimation – Summary with values from Synopsys Design Compiler

vn_V – rn_R – cn_C	Buffer Power (mW)	Mesh Link + Control Power (mW)	Folded Torus Link + Control Power (mW)	Mesh Total Power (Buffer + Link) (mW)	% Change	Folded Torus Total Power (Buffer + Link) (mW)	% Change
v4-r4-c0	19.54	2.45 + 0	3.94 + 0	21.99	-	23.48	-
v4-r3-c4	14.51	2.90 + 0.0012	4.39 + 0.0012	17.42	-20.78	18.91	-19.46
v4-r2-c8	**11.57**	**3.55 + 0.02**	**5.04 + 0.02**	**15.14**	**-31.15**	**16.63**	**-29.17**
v3-r4-c4	15.09	2.90 + 0.0012	4.39 + 0.0012	18.00	-18.14	19.49	-16.99
v3-r3-c7	12.56	3.49 + 0.0180	4.98 + 0.0180	16.06	-26.96	17.55	-25.25

n_V = number of VCs per input port, n_R = number of router buffers per VC, n_C = number of link buffers

Area Estimation – Summary with values from Synopsys Design Compiler

vn_V – rn_R - cn_C	Buffer Area (μm^2)	Link Repeater Area (μm^2)	Total Buffer + Link Area (μm^2)	% Change
v4-r4-c0	81,407	8,960	90,367	-
v4-r3-c4	63,991	6,656	70,647	-21.8
v4-r2-c8	**48,066**	**10,240**	**58,306**	**-35.5**
v3-r4-c4	63,250	6,656	69,906	-22.6
v3-r3-c7	50,373	9,344	59,717	-34.0

n_V = number of VCs per input port, n_R = number of router buffers per VC, n_C = number of link buffers

Network Simulation Results (1/6)

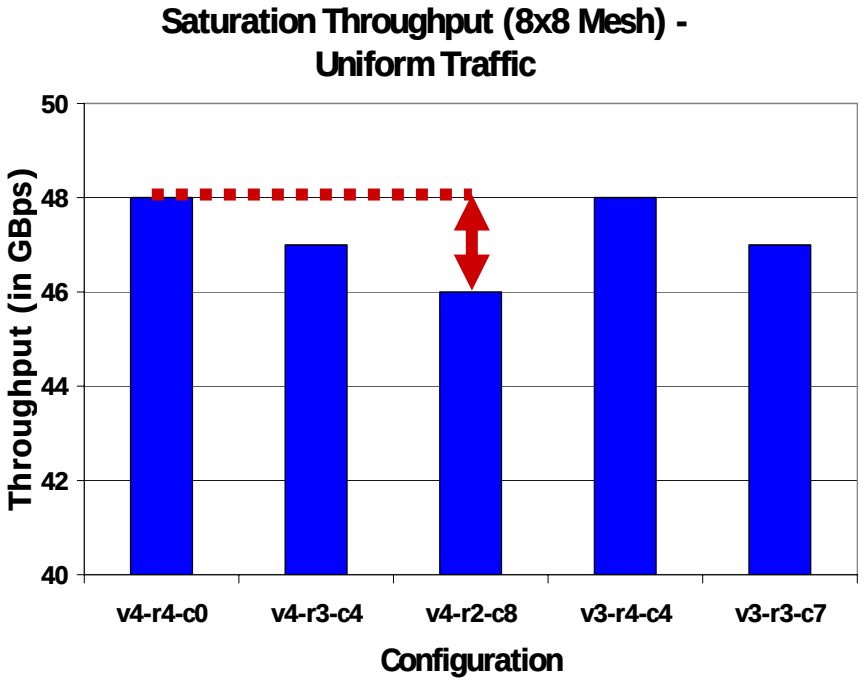

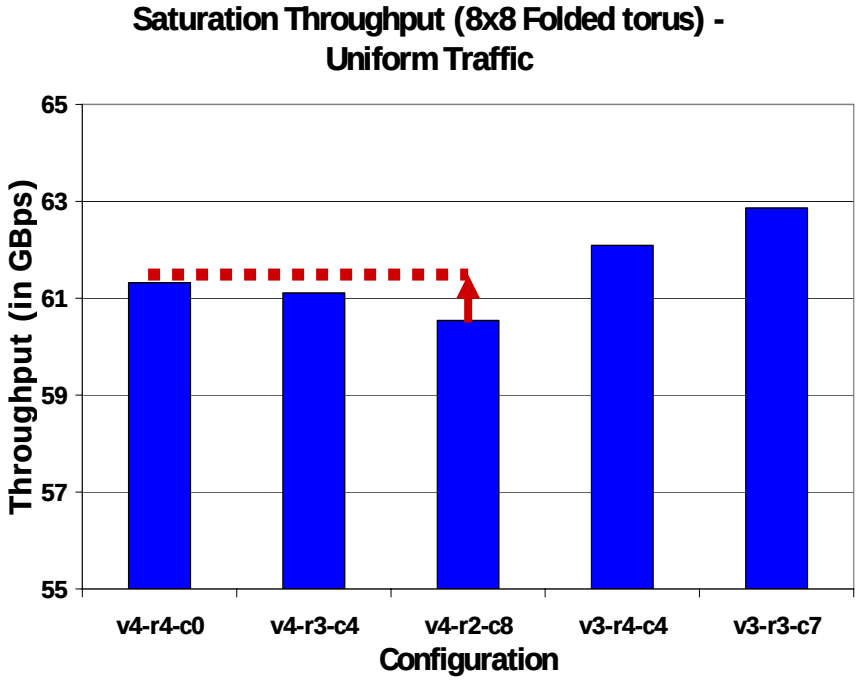

- Uniformly distributed traffic

⇒ **Only about 3 % drop in throughput for the v4-r2-c8 case**

(428 = 4 VCs per port, 2 router buffers per VC, 8 link buffers)

Network Simulation Results (2/6)

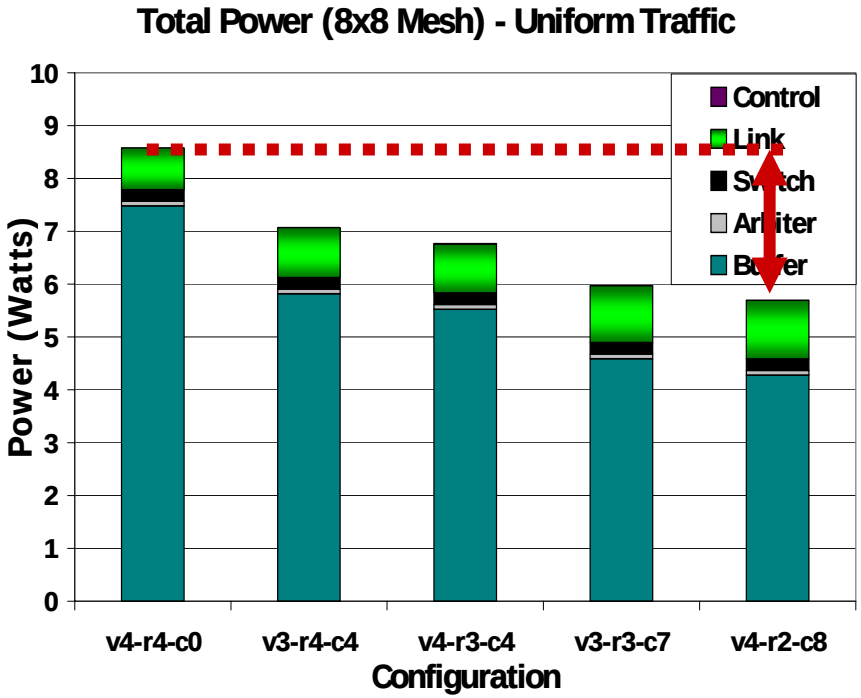

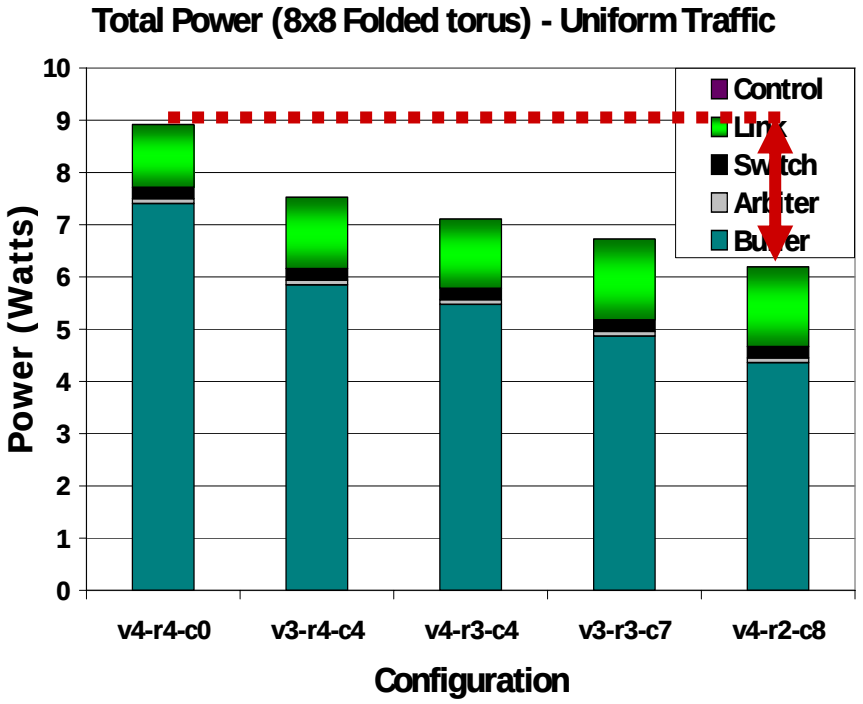

- **Total power consumed for a network load of 0.5**

⇒ **Nearly 30 % savings in overall network power for the v4-r2-c8 case**

(428 = 4 VCs per port, 2 router buffers per VC, 8 link buffers)

Buffer Power – Synthetic Traffic (3/6)

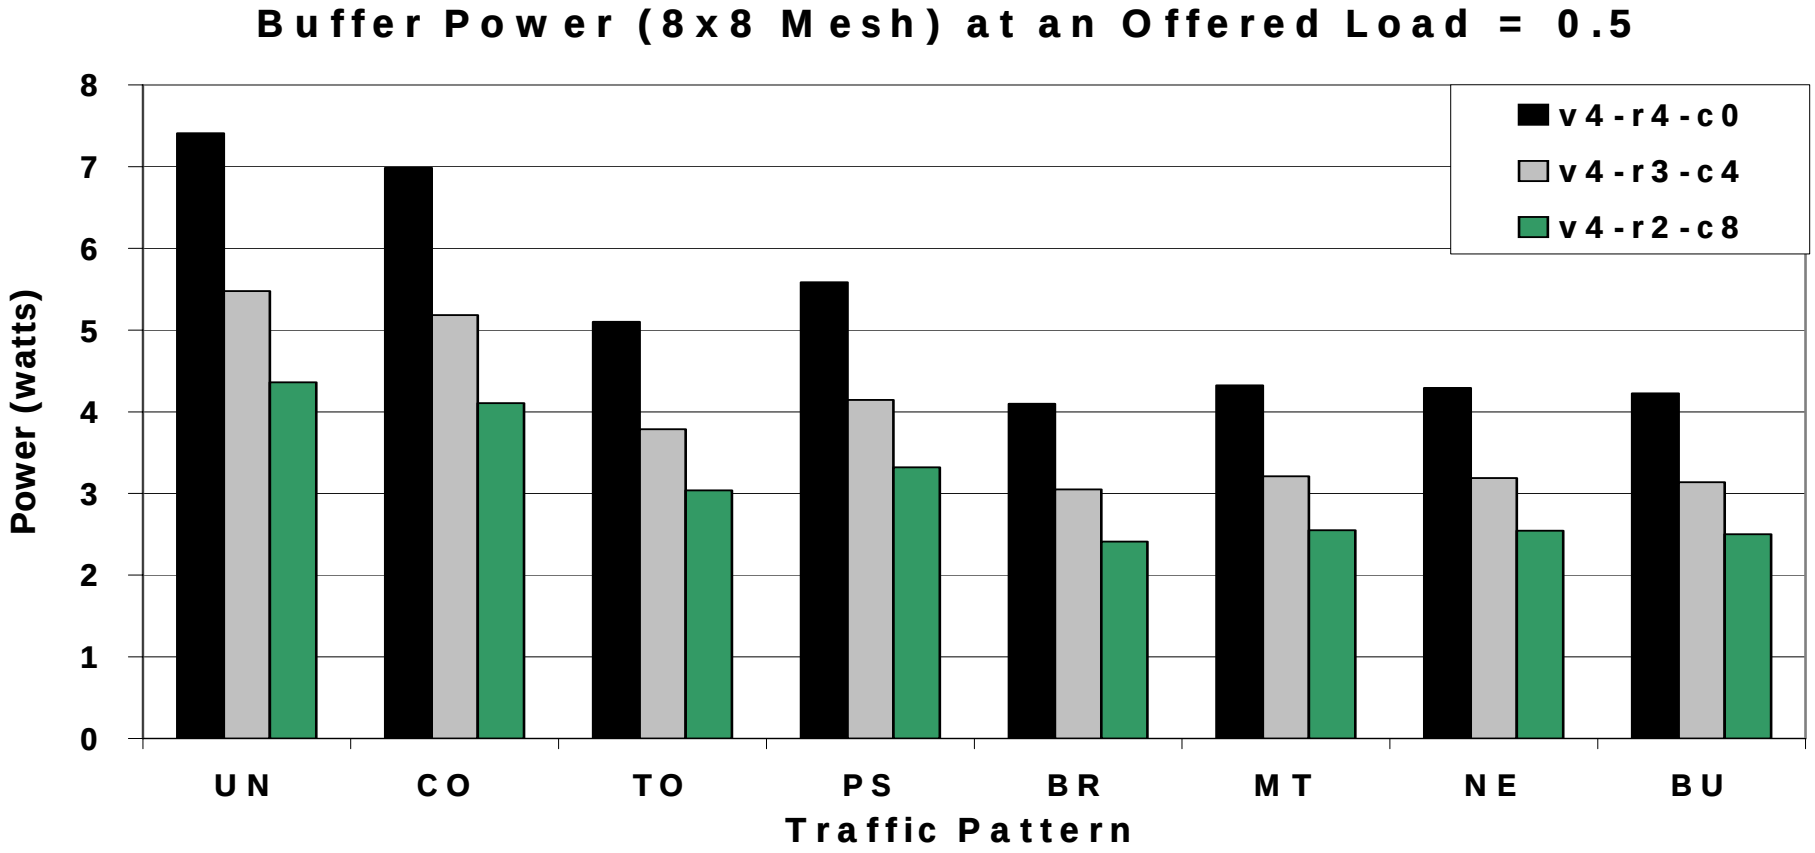

- Reduction in power for all configurations, under all traffic patterns, compared to the baseline (440)
- For example, under Complement traffic the 4-2-8 configuration achieves 40% savings

Throughput – Synthetic Traffic (4/6)

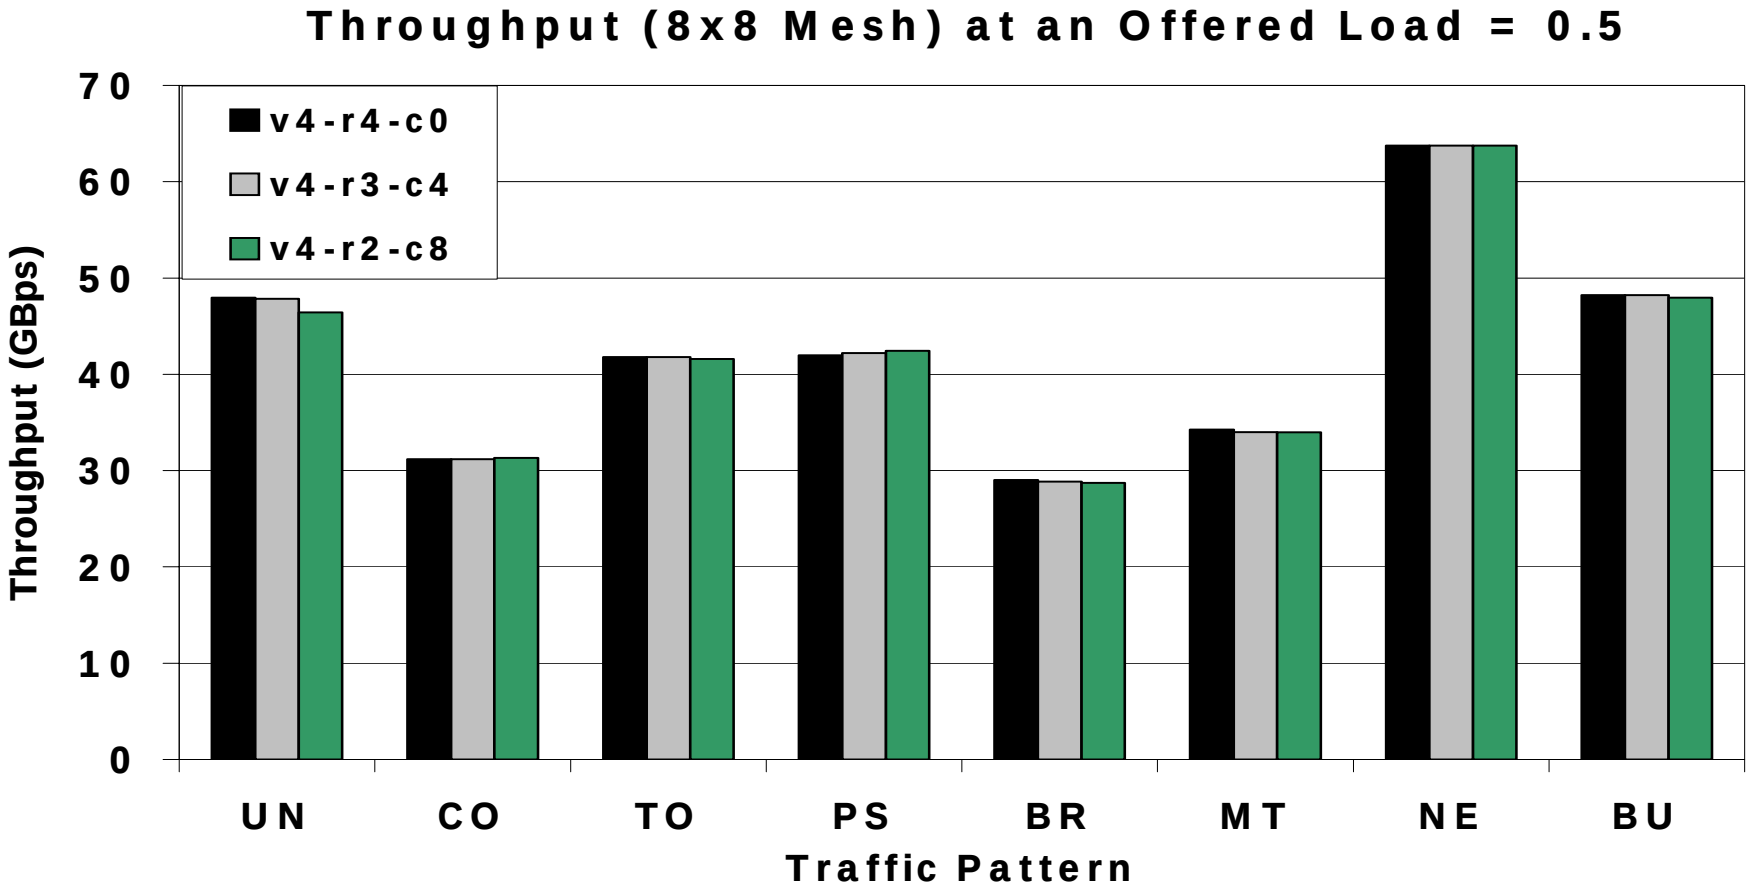

- No significant decrease in throughput under any traffic pattern, using Dynamic allocation

Simulation results for the SPLASH-2 suite (5/6)

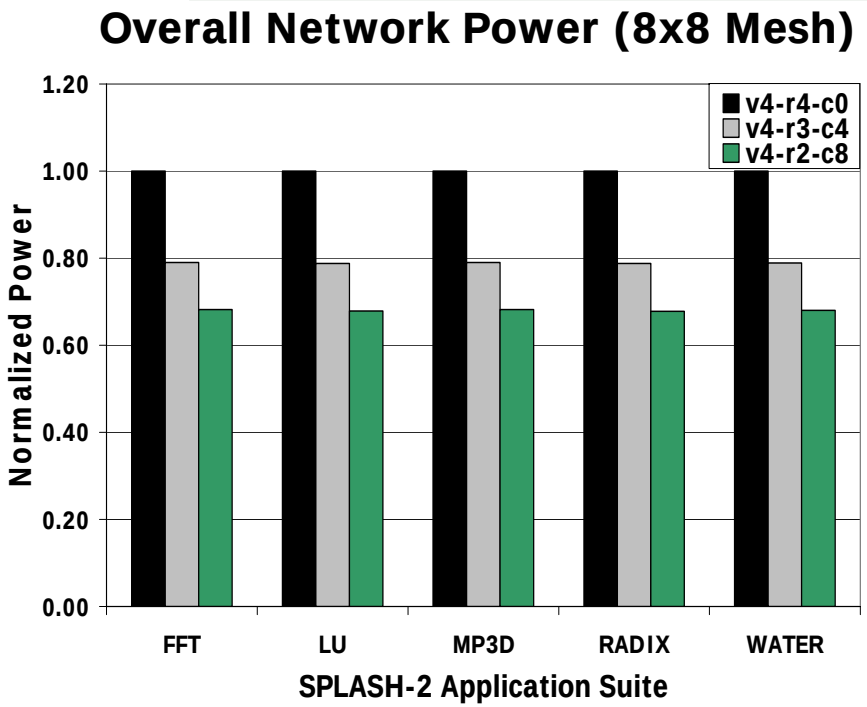

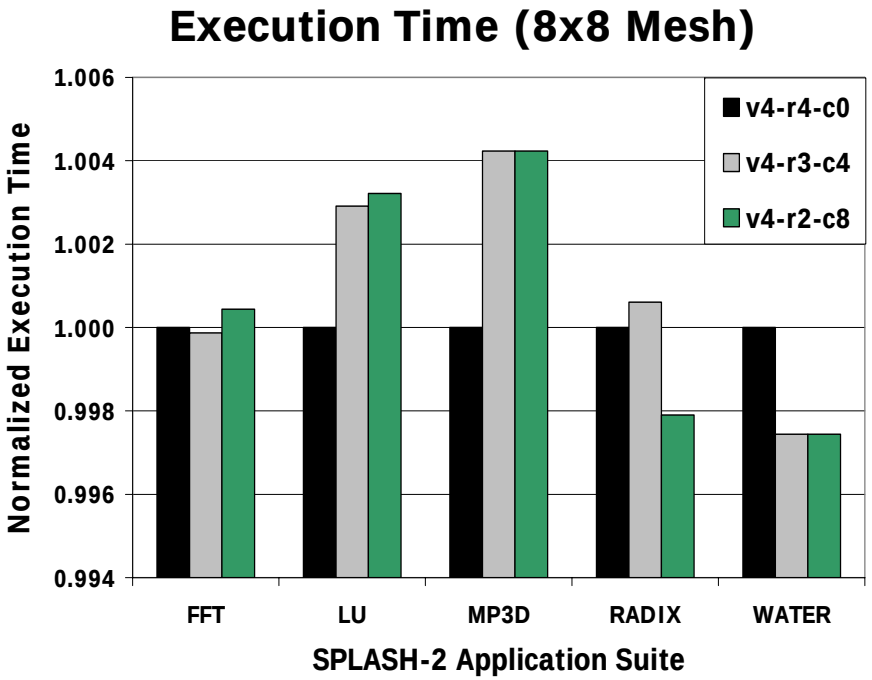

- **Nearly 30% savings in overall network power for the v4-r2-c8 case with only about 1% drop in performance**

(v4-r2-c8 = 4 VCs per port, 2 router buffers per VC, 8 link buffers)

Aggressive Speculation (6/6)

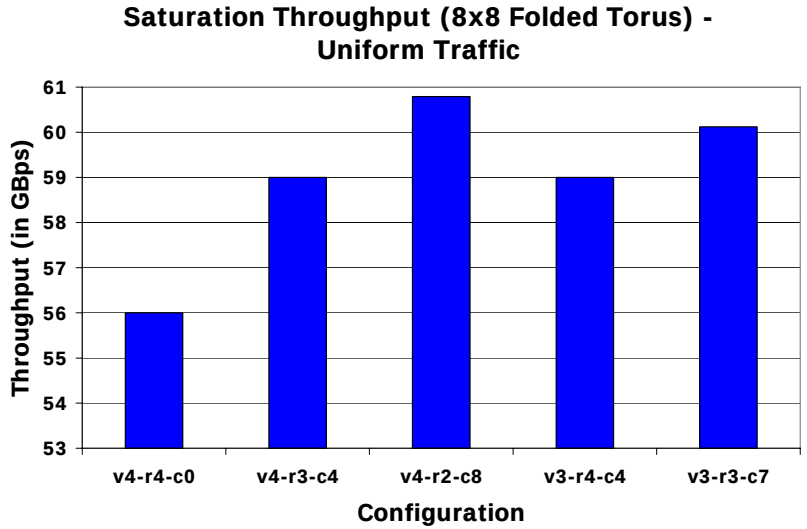

- Aggressive speculation by increasing the number of credits available to 8
- Additional credits are accounted for by the link buffers

⇒ Saturation throughput improves by 10% for the 428 case

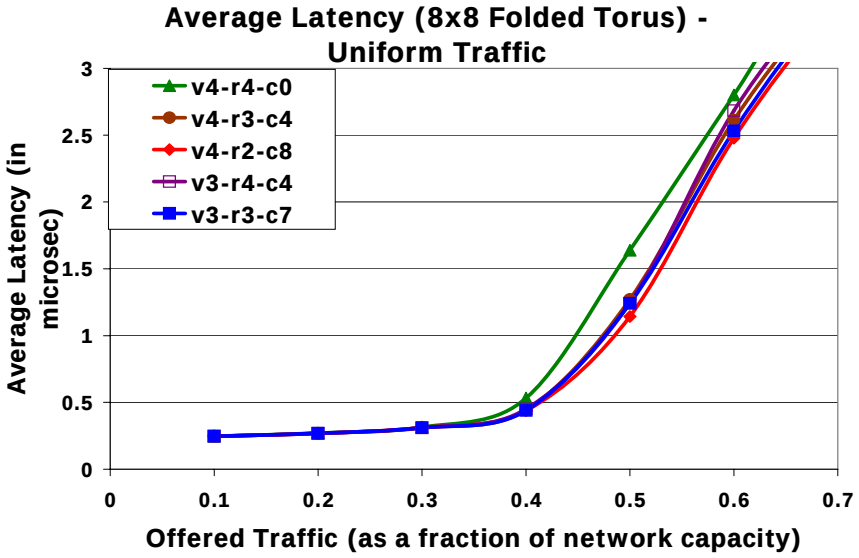

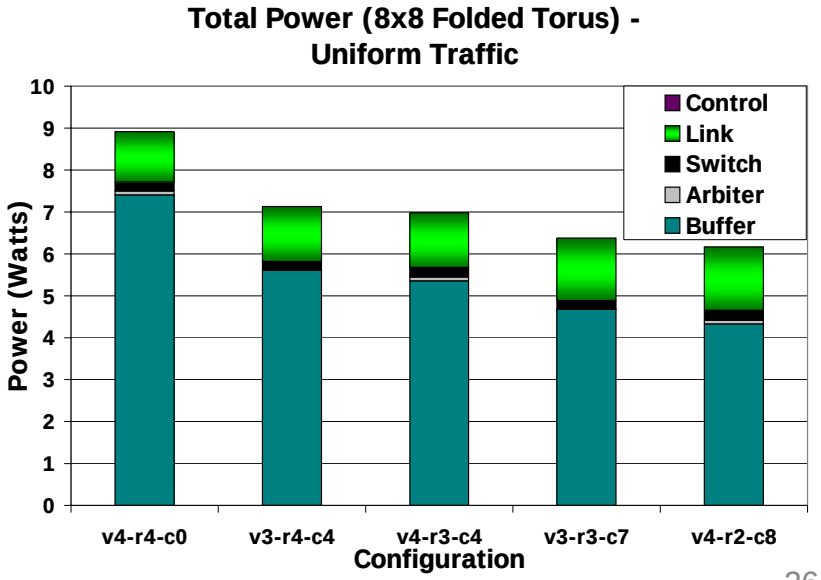

Design Headroom for NBTI-related issues

- Negative Bias Temperature Instability (NBTI) affects PMOS gates with a '0' input. Additional circuits may be employed to change the input of 'Idle' PMOS gates to '1' [3].

- iDEAL architecture 'moves some of the buffers from the router to the links' and provides a low-power area-efficient solution with
  - Reduced power density and temperature in the routers, thereby alleviating NBTI-related degradation
  - Sufficient design headroom for the Thermal Design Power (TDP) of the auxiliary circuits

$$TDP_{NoC} \leqslant P_{saved}$$

$$TDP_{NoC} = K \times S_r$$

(P_{saved} = power reduction achieved by iDEAL, K = impact of the area on the TDP, S_r = area reduction achieved by iDEAL)

3. J.Abella, X.Vera and A.Gonzalez, "Penelope: The NBTI-aware processor", Proc. 40th Annual ACM/IEEE Intl. Symp. Micro Arch. (MICRO-40), Chicago, IL, USA, December 1-5, 2007, pp. 85 – 96.

Conclusion

- **iDEAL** architecture provides a low-power area-efficient solution for NoCs, by reducing power consumption through circuit-level and architecture-level techniques.

- Simulation results show that by reducing the buffer size in half, a **30% savings in overall network power** and **35% savings in overall area** is achieved.
- There is only a marginal 1-3% drop in performance, under dynamic buffer allocation. (note: Performance degradation has been solved already)
- Further, the significant reduction in power and area provide sufficient headroom for monitoring NBTI effects.

Questions?

kodi@ohio.edu