Design of Adaptive Communication Channel Buffers for Low-Power Area-Efficient Network-on-Chip Architecture

Avinash Kodi[†], Ashwini Sarathy[*] and Ahmed Louri[*]

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

[*]Department of Electrical and Computer Engineering, University of Arizona, Tucson, AZ 85719

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

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

ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS'07)

Dec 3-4, 2007

Talk Outline

- **Motivation & Introduction**
- **iDEAL – Inter-router Dual-function Energy and Area-efficient Links for NoC architectures**
  - **Link and Router Architecture**
- **Performance Evaluation**
  - **Power & Area estimation for the Links & Routers**
  - **Simulation results for Throughput, Latency & Overall network power**
- **Conclusions**

Motivation

System-on-Chip (SoC) paradigm

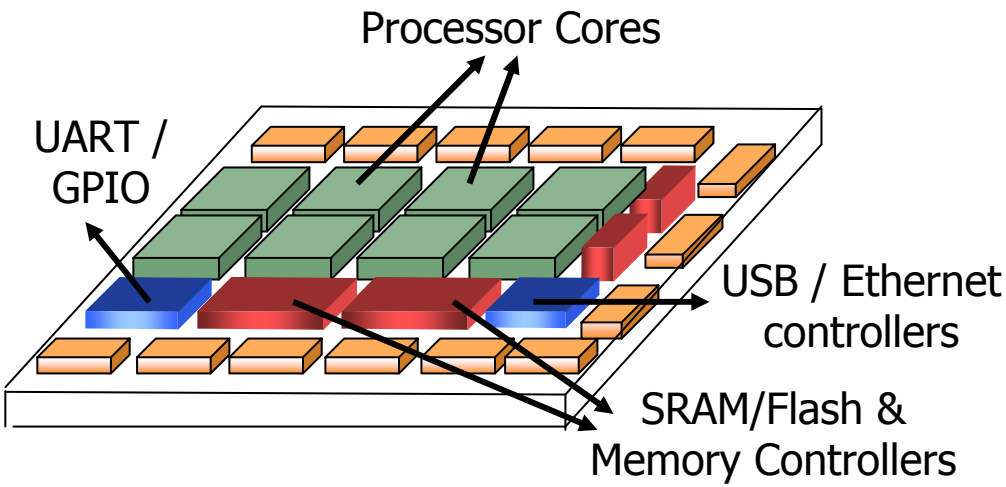

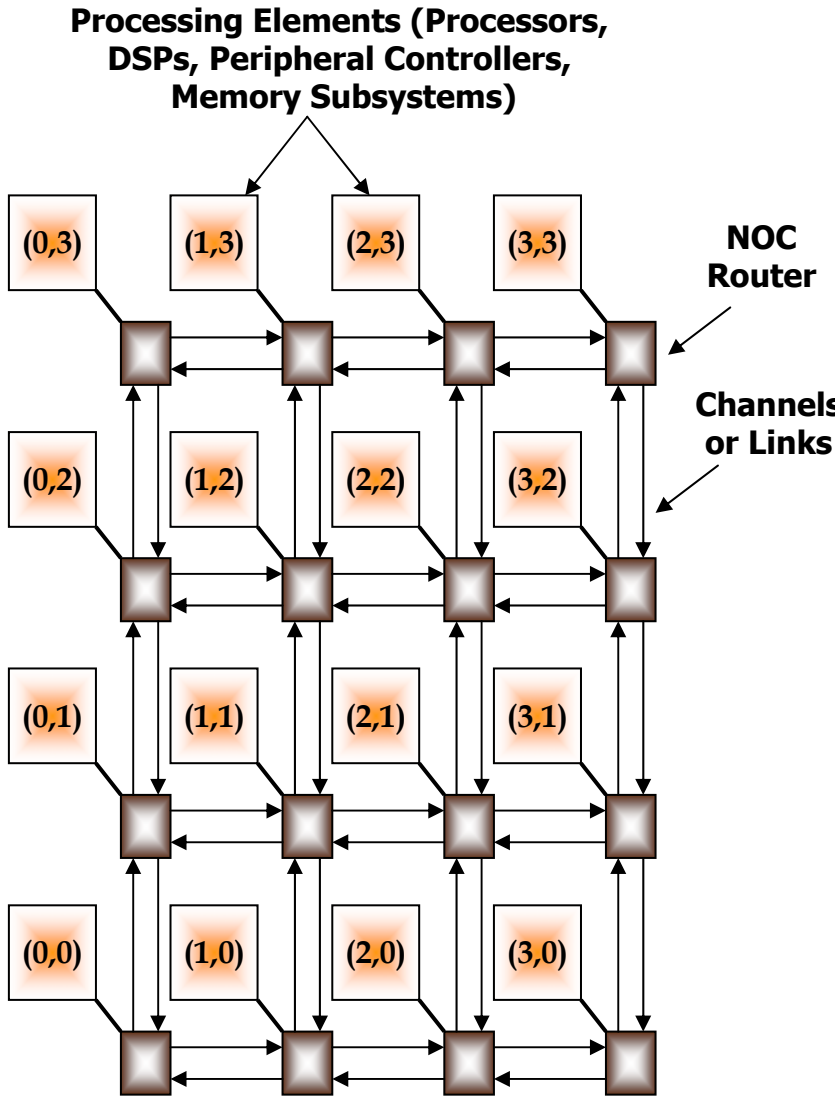

- **Increasing wire delay with decreasing feature size**

- **Scalable, modular interconnect – Network-on-Chip (NoC)**

Motivation

Generic NoC Router

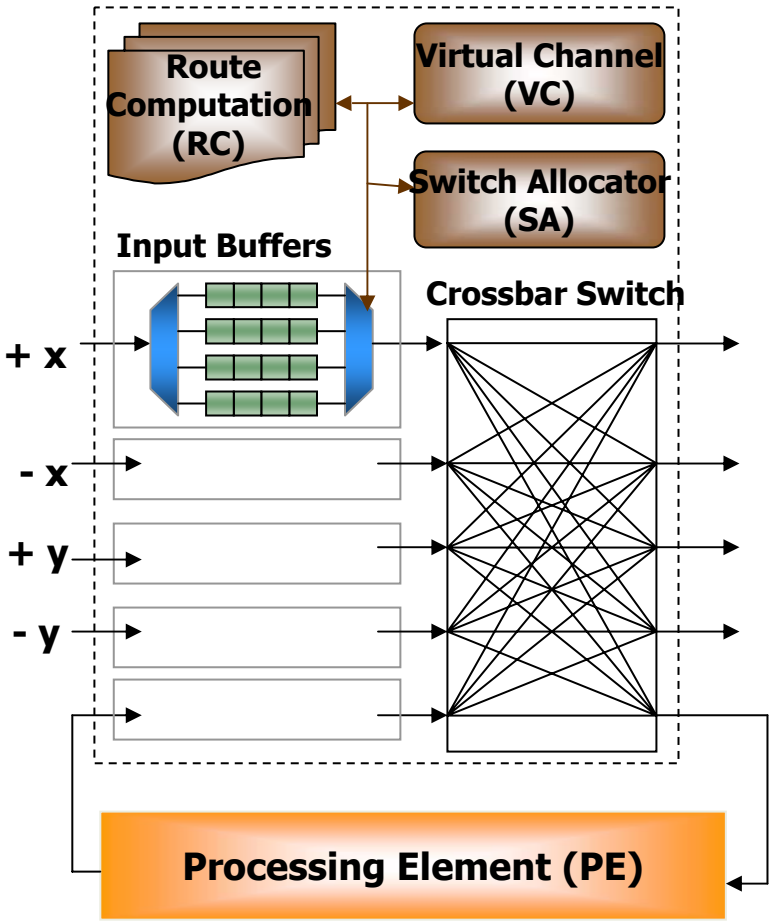

Recent NSF-sponsored workshop on On-Chip Interconnection Networks[1] :

- "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

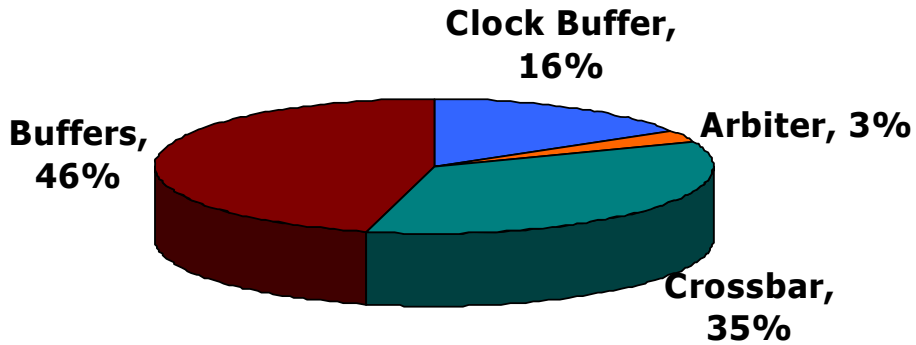

1. Reference : J.D.Owens, W.J.Dally, R.Ho, D.N.Jayasimha, S.W.Keckler and L.S.Peh, "Research Challenges for On-Chip Interconnection Networks", IEEE Micro, vol. 27, no. 5, pp. 96 – 108, September-October 2007.

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

iDEAL Methodology (circuit and architectural techniques)

- Reduce the number of router buffers
- To prevent performance degradation, use adaptive channel buffers to store data along the links when required
- Dynamic buffer allocation within the router buffers

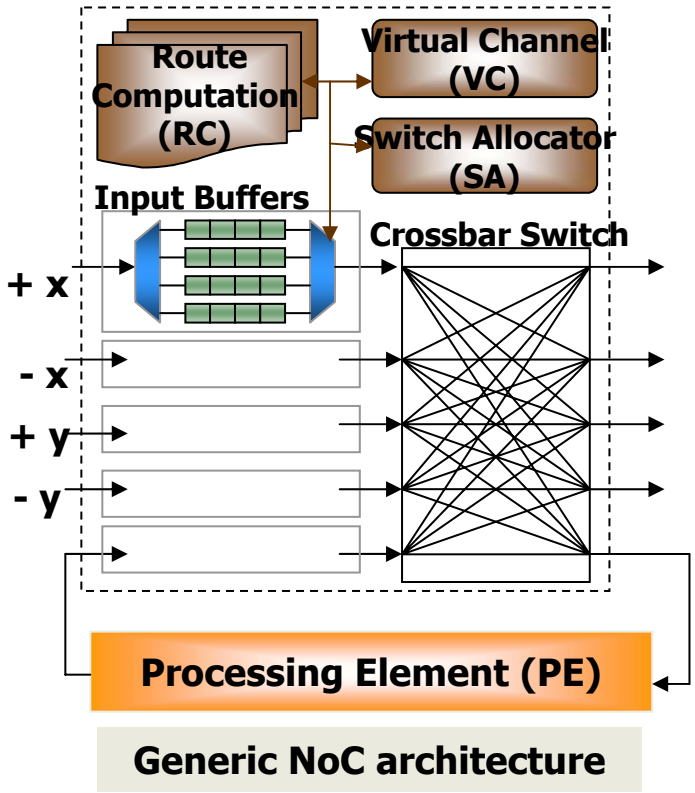

Generic NoC architecture

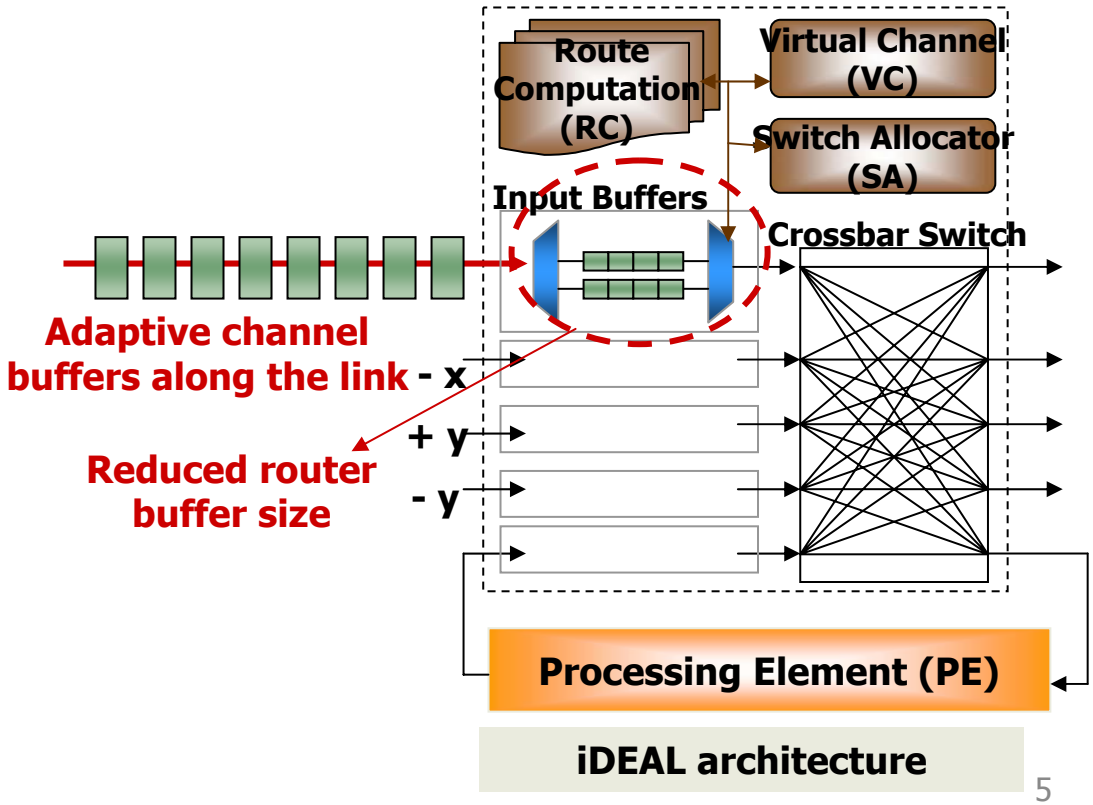

iDEAL architecture

Conventional Links

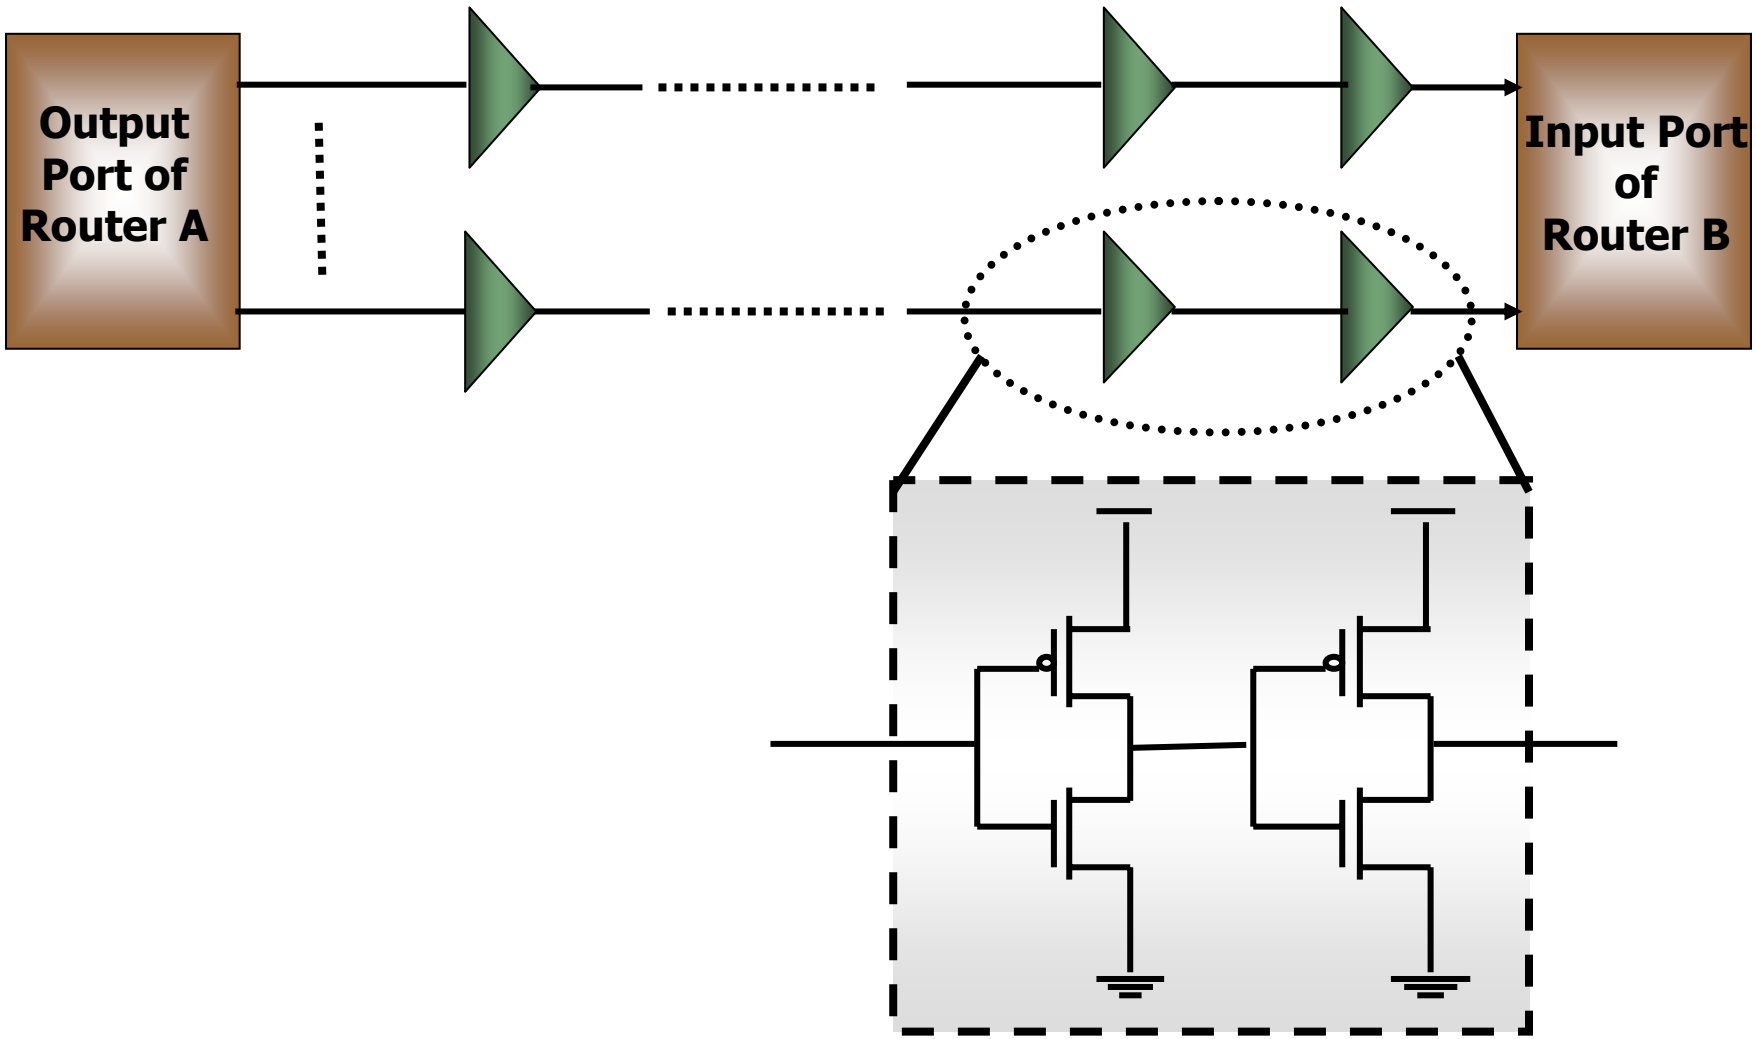

iDEAL – Channel Buffer Design (1/2)

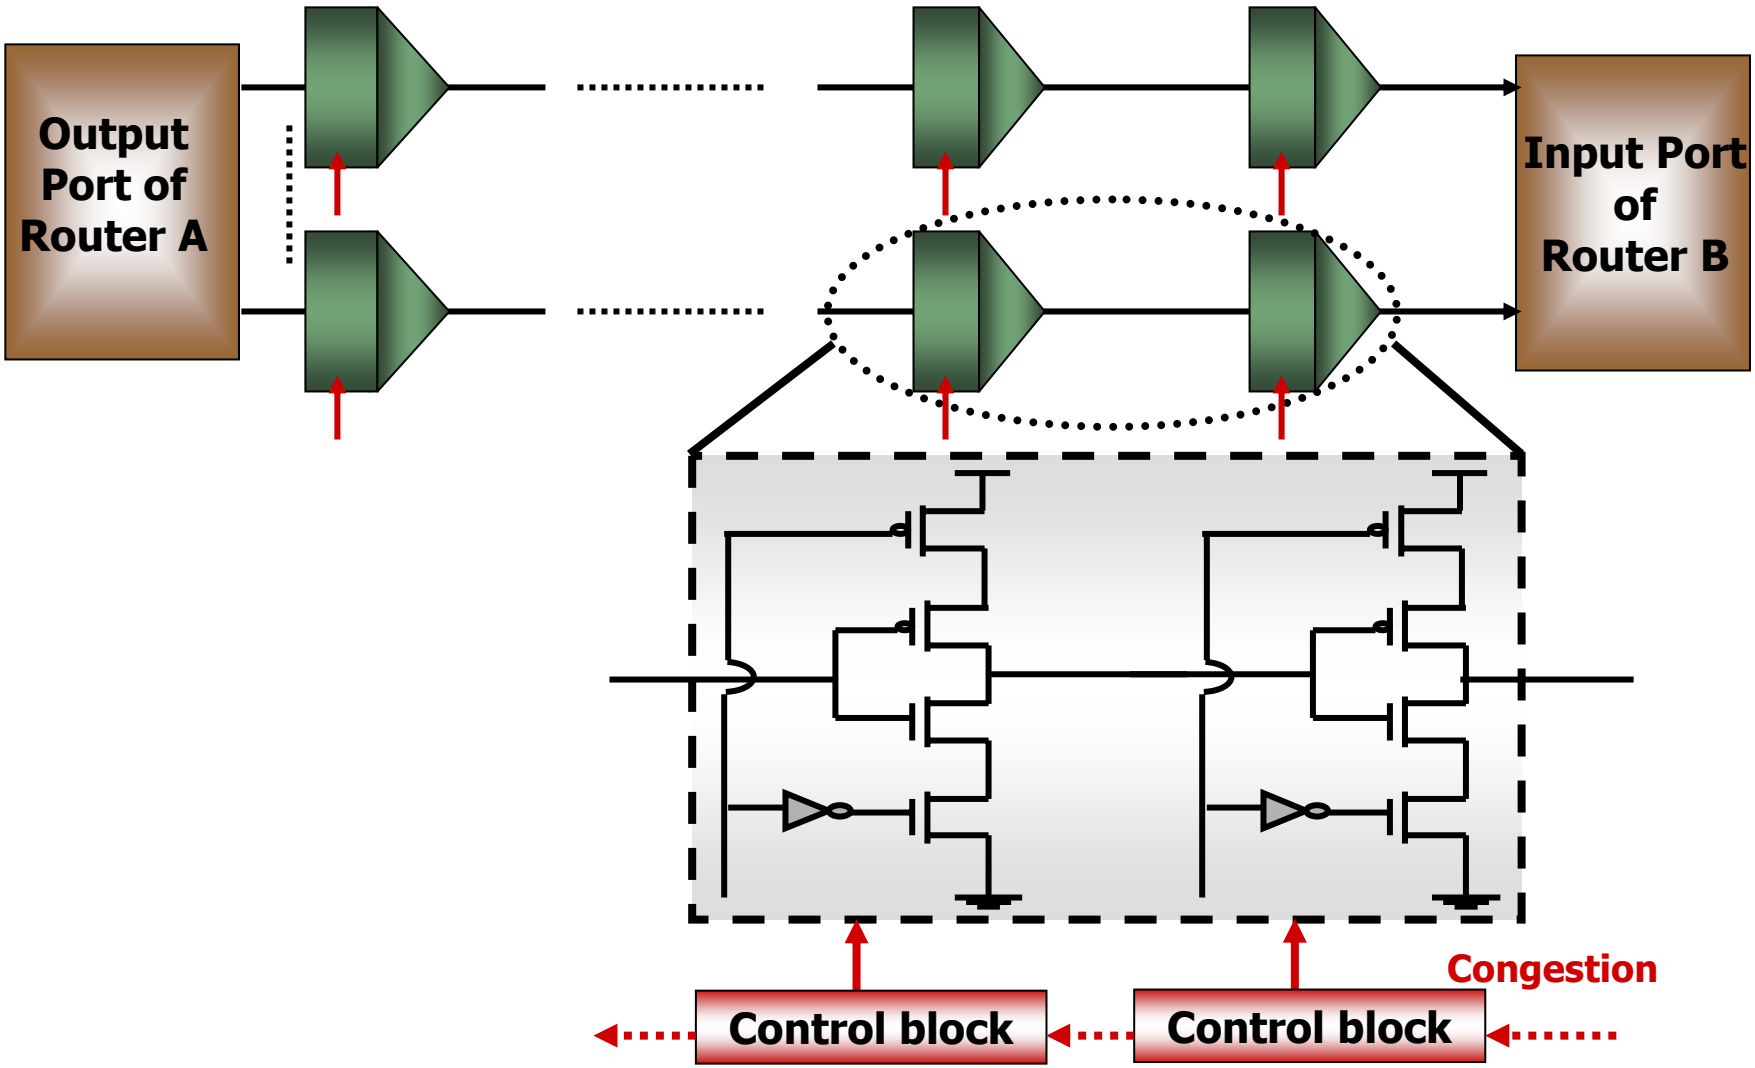

iDEAL – Channel Buffer Design (2/2)

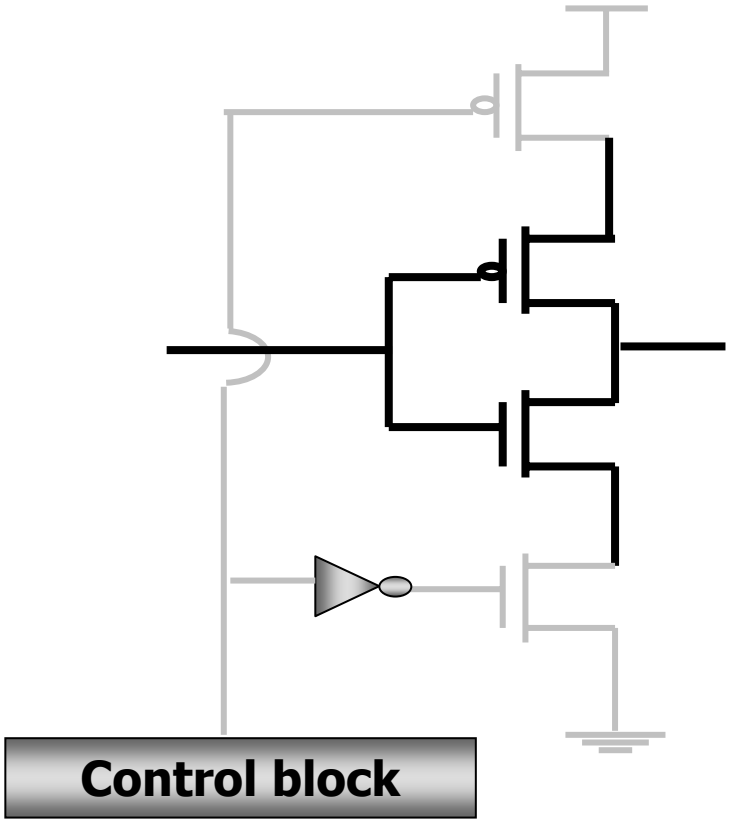

Functions as a conventional repeater when there is no congestion.

Control block is turned 'OFF'.

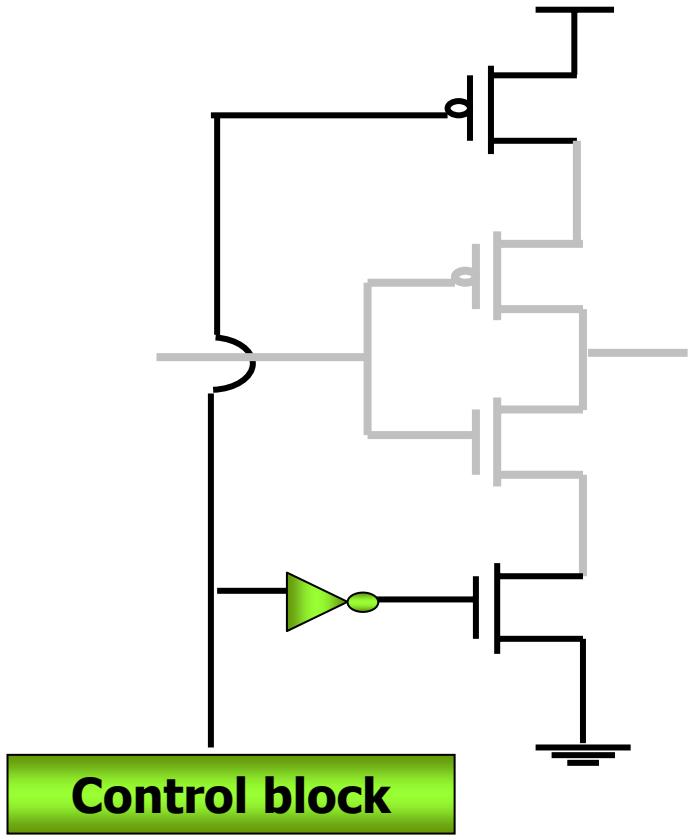

Repeater tri-stated and holds the sampled value, during congestion.

Control block is turned 'ON'.

iDEAL – Control Block

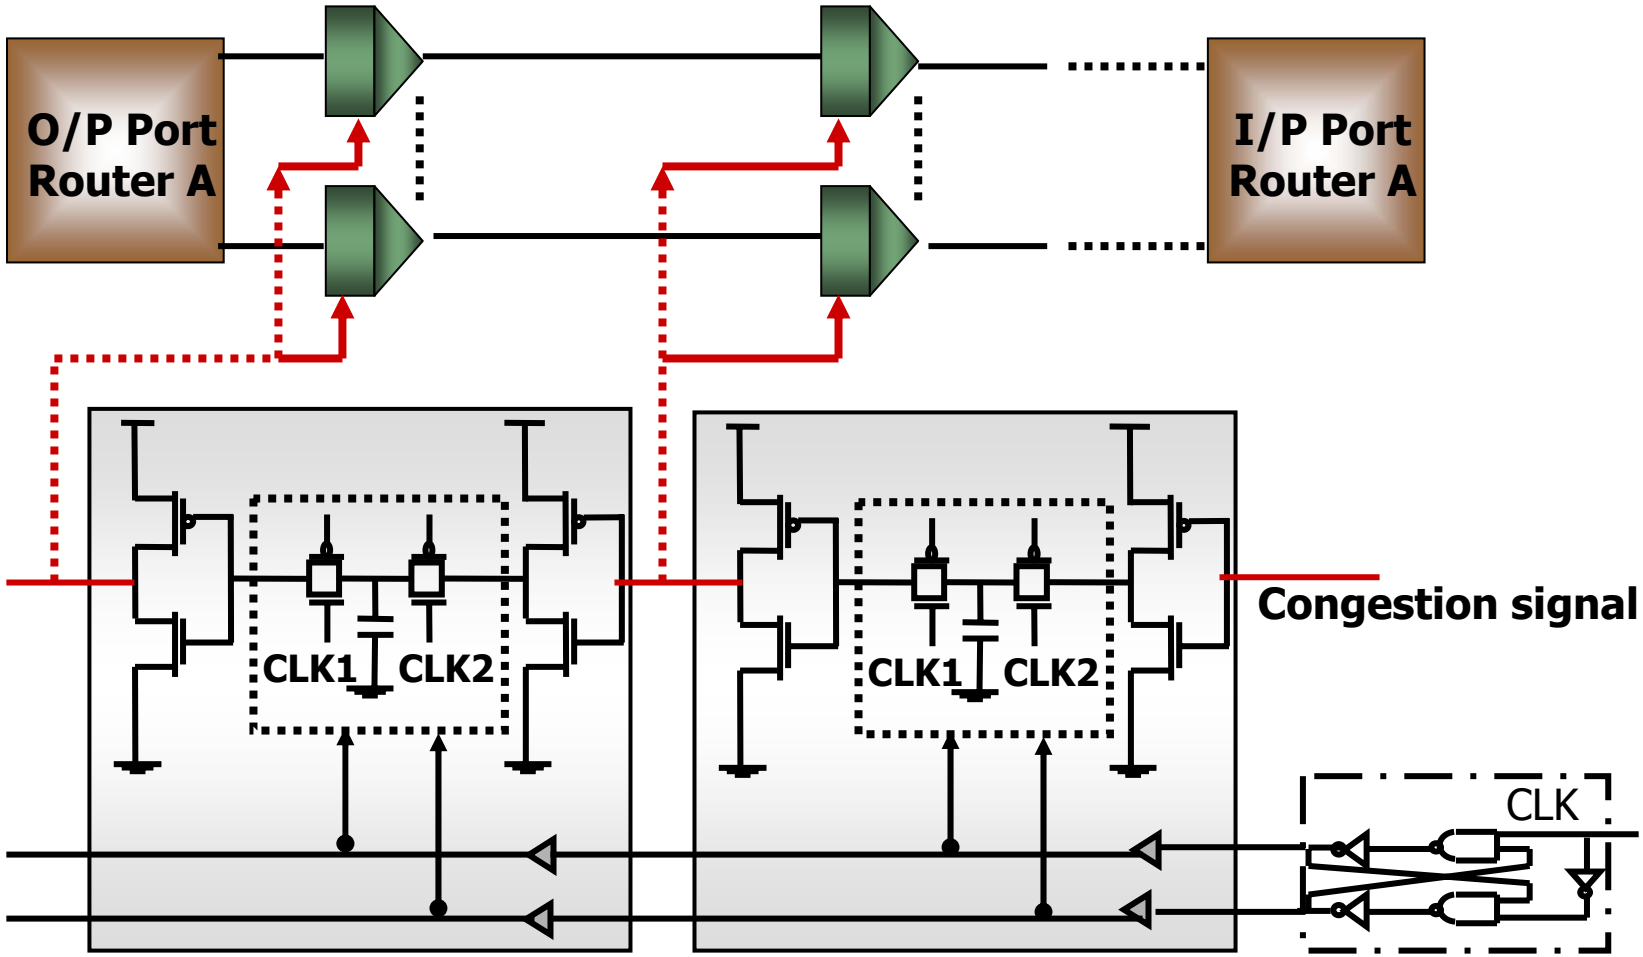

- Power efficient
- Stable at varying frequencies

iDEAL : Dual-function Link

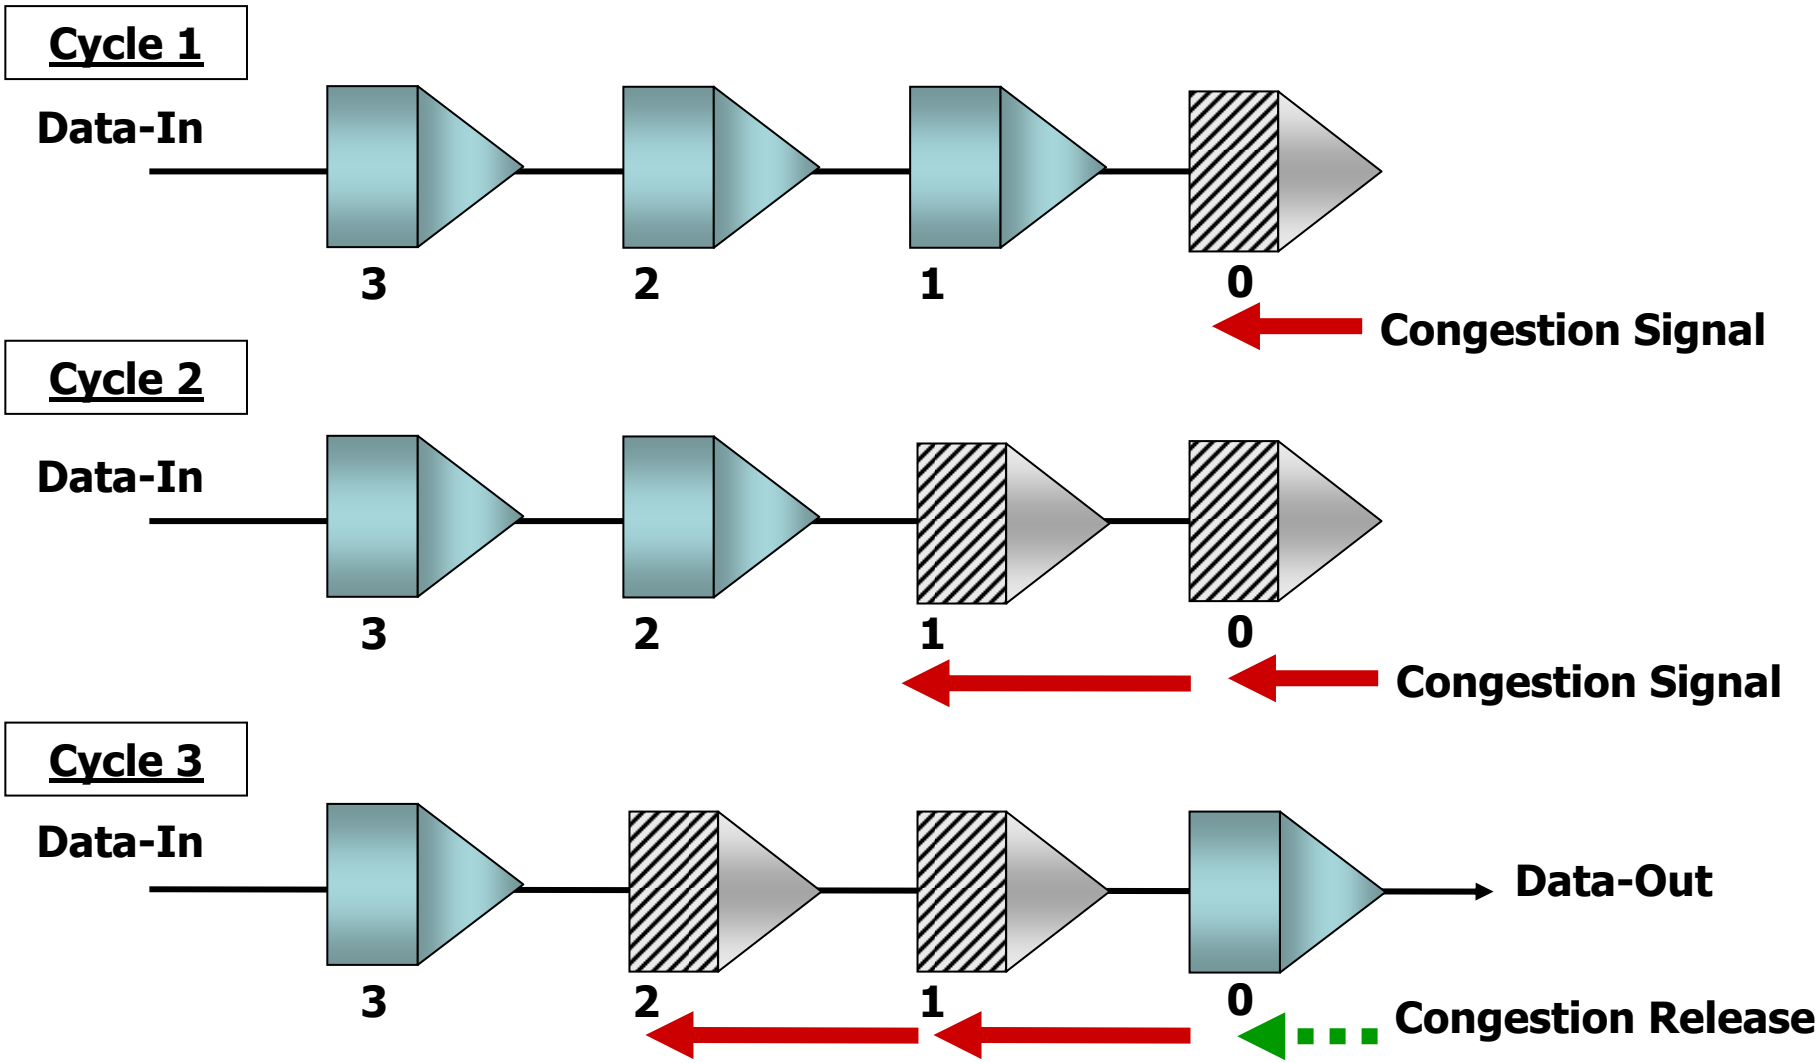

Link - Power & Area Estimation

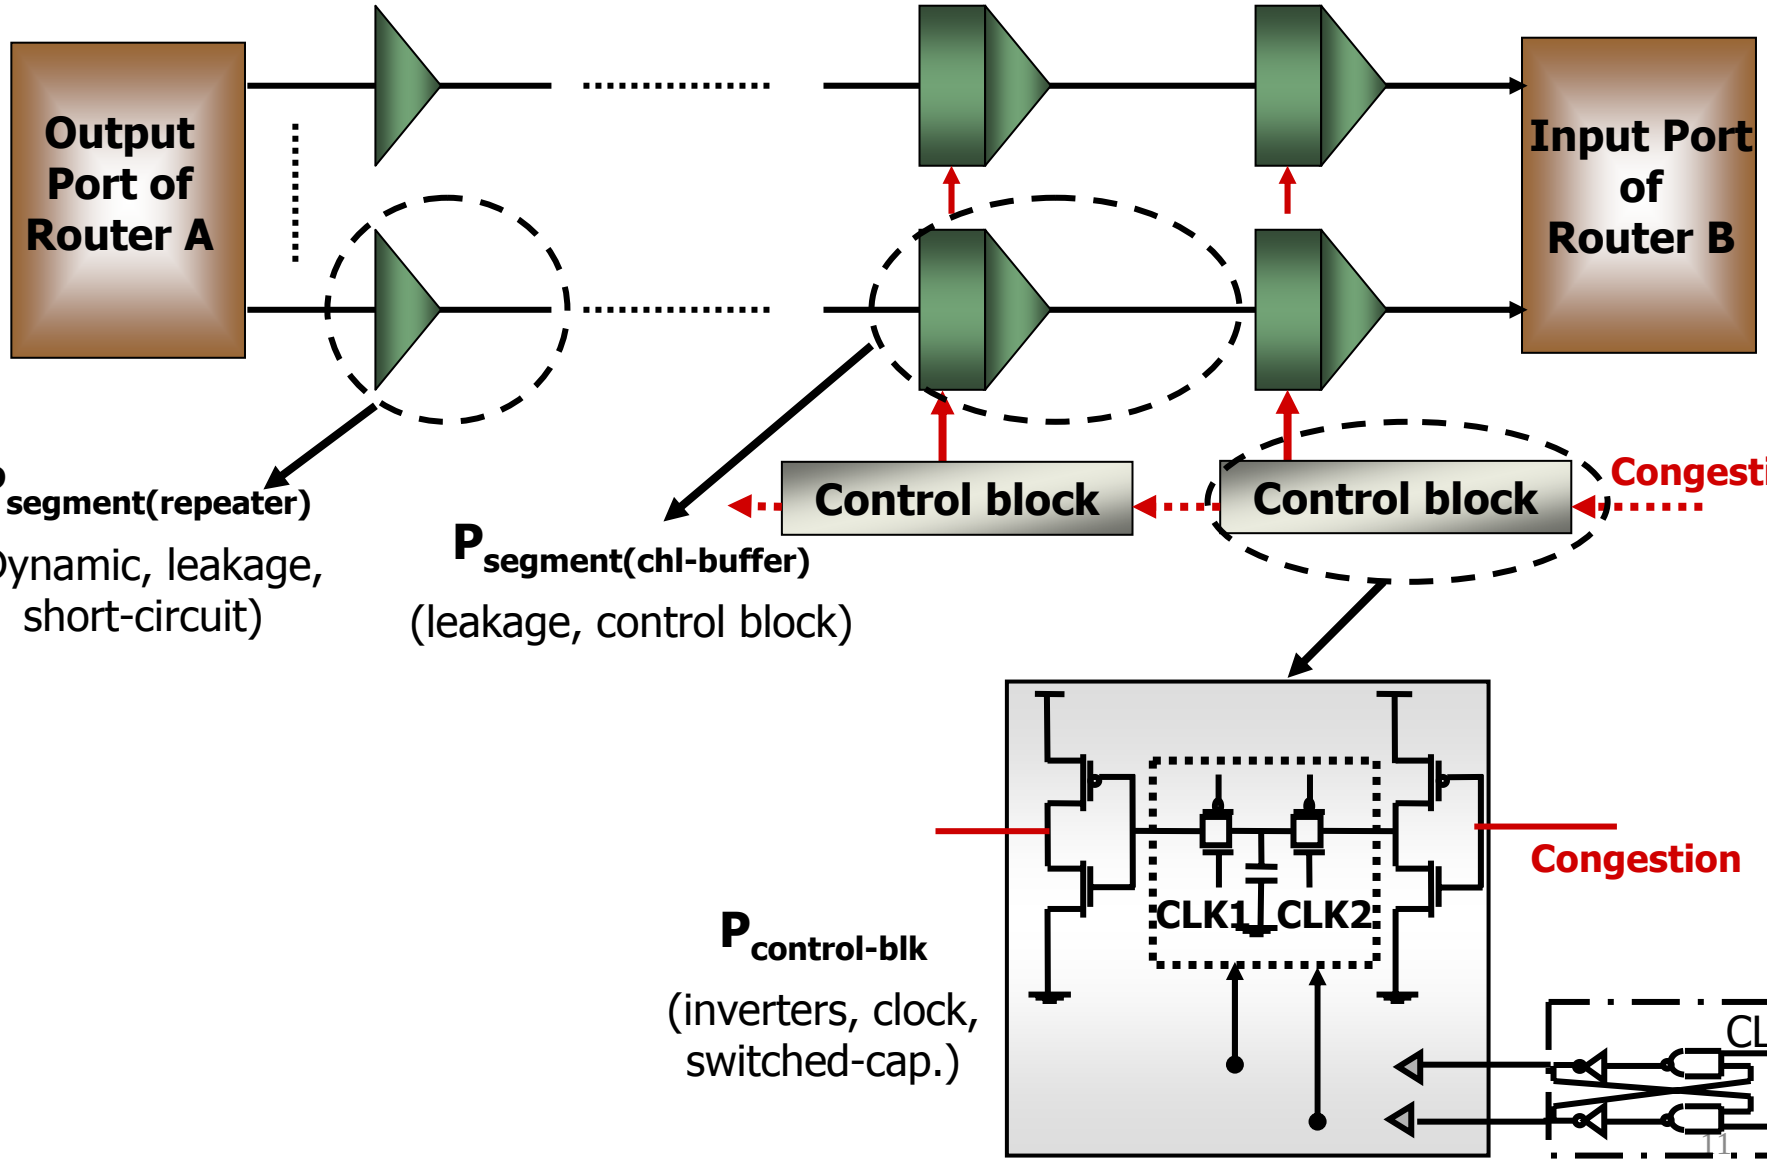

iDEAL – Router Buffer Design

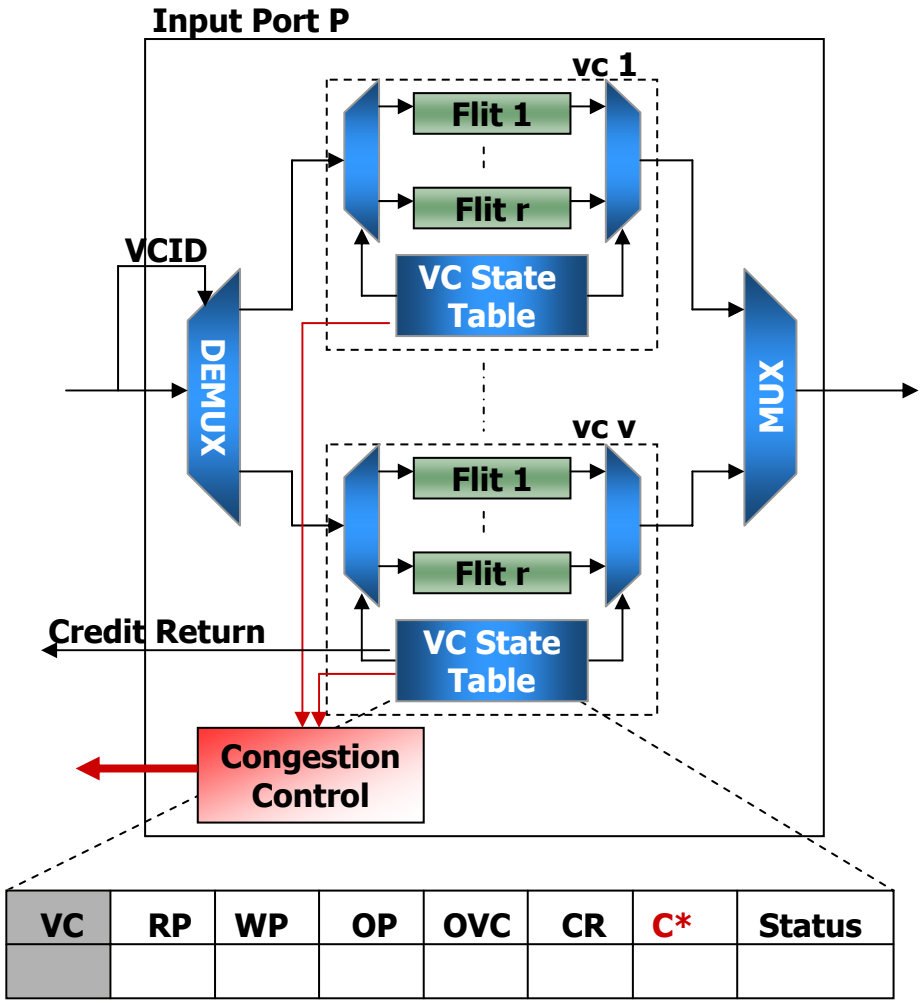

VC	RP	WP	OP	OVC	CR	C*	Status

- 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)

iDEAL – Router Buffer Design

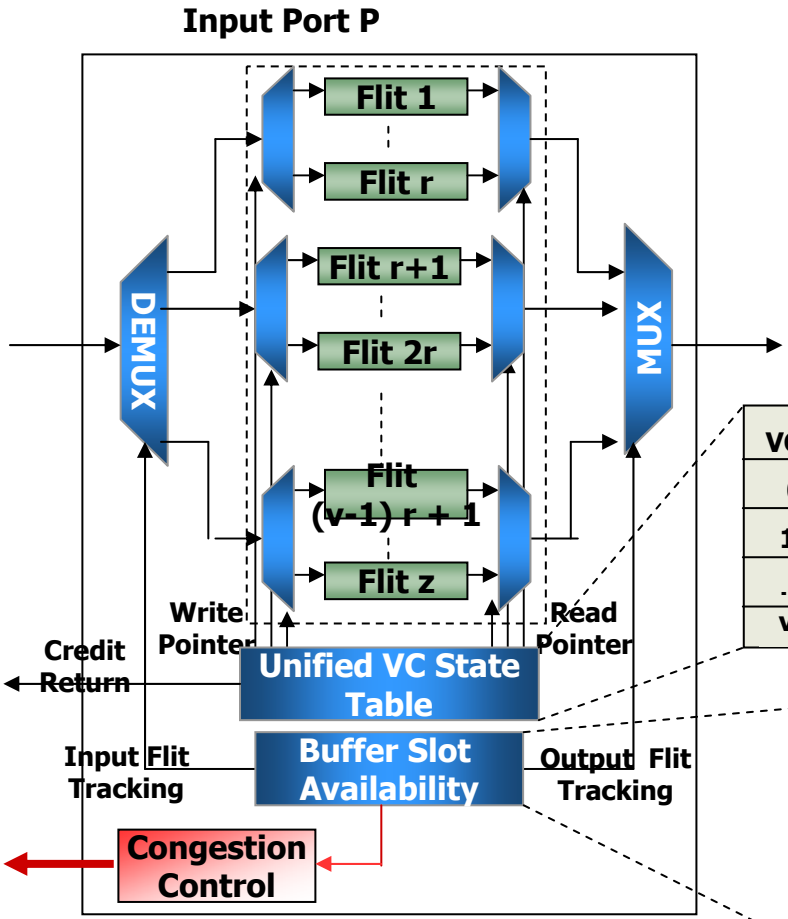

- Dynamic buffer allocation
  - Approximately (z + c)/v buffers per VC (z = router buffers, c = channel buffers, v = # of VCs)

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)

iDEAL – Router Buffer Design

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

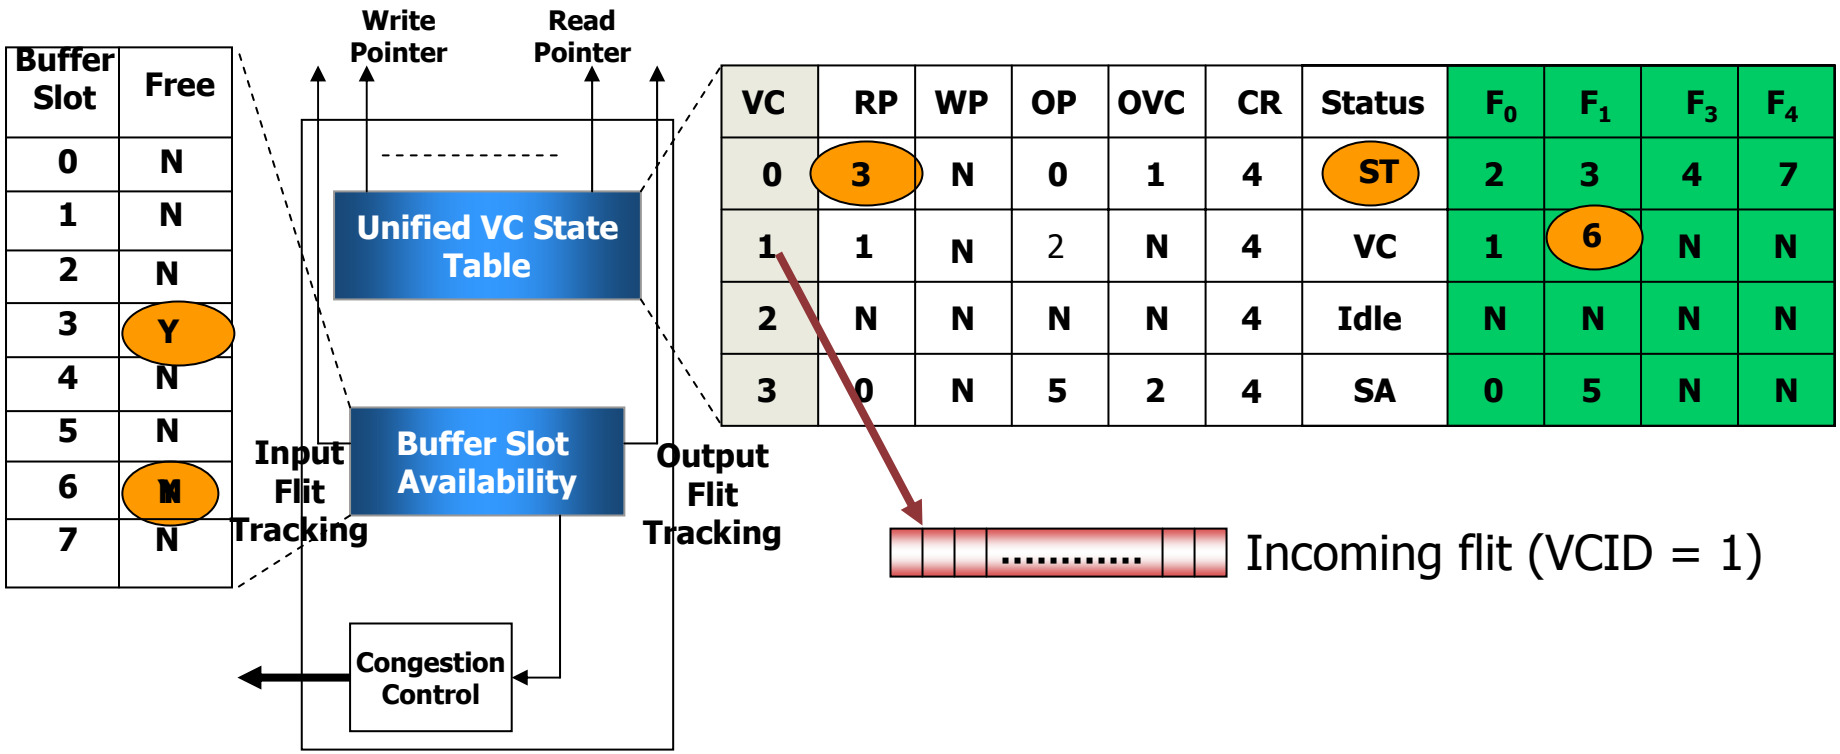

Buffer Slot	Free
0	N
1	N
2	N
3	Y
4	N
5	N
6	N
7	N

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)

Router - Power & Area Estimation

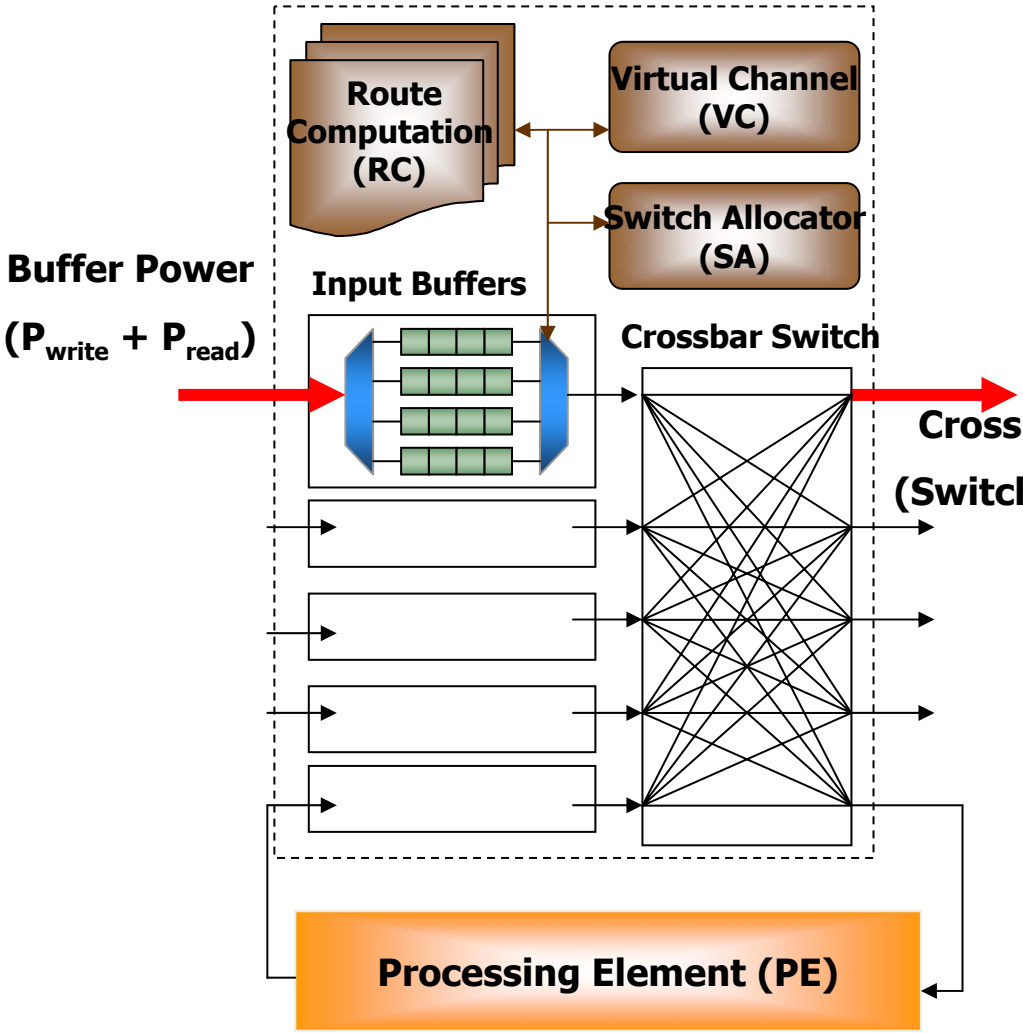

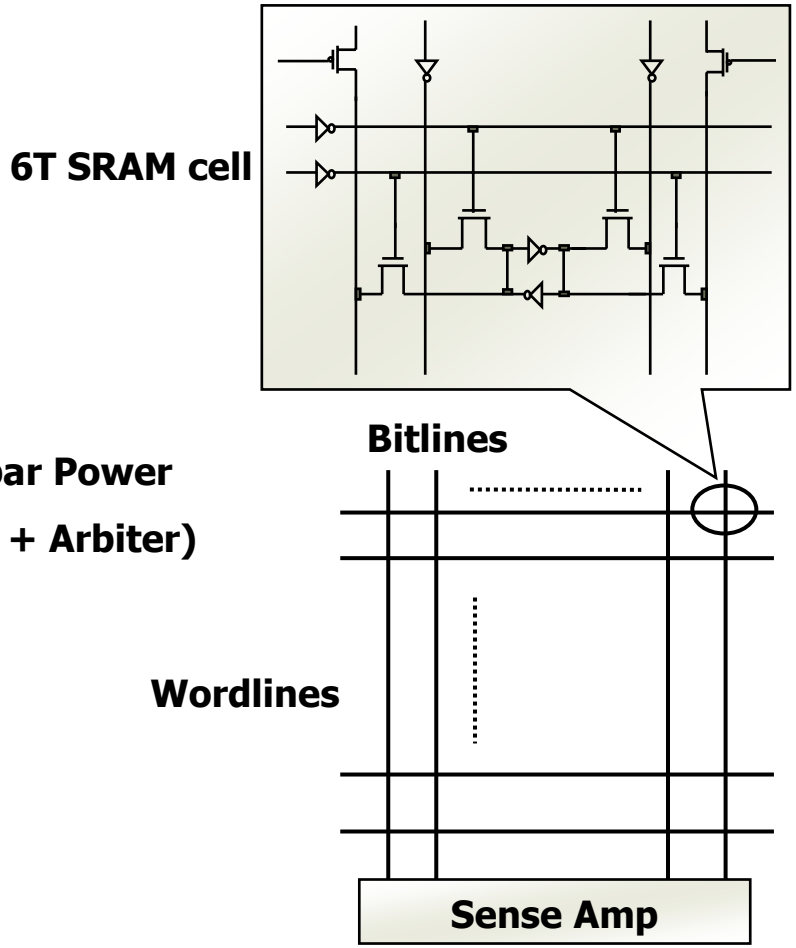

- Power reduces on decreasing the buffer size

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) 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

Power Estimation - Summary

vnV – rnR - cnC	Buffer Power (mW)	% Change	Mesh Link + Control Power (mW)	% Change	Folded Torus Link + Control Power (mW)	% Change
v4-r4-c0	2.020	-	2.032 + 0	-	4.068 + 0	-
v4-r3-c4	1.646	-18.51	2.164 + 0.0122	+ 7.0	4.195 + 0.0122	+ 3.4
v4-r2-c8	**1.272**	**-37.02**	**2.296+0.0205**	**+13.9**	**4.437+0.0205**	**+6.8**
v3-r4-c4	1.646	-18.51	2.164 + 0.0122	+ 7.0	4.195 + 0.0122	+ 3.4
v3-r3-c7	1.365	-32.41	2.263 + 0.0184	+ 12.2	4.294 + 0.0184	+ 6.0
v5-r2-c6	1.459	-27.76	2.230 + 0.0164	+ 10.5	4.261 + 0.0164	+ 5.1
v5-r3-c1	1.926	-4.65	2.065 + 0.0059	+ 1.8	4.096 + 0.0059	+ 0.8

n_V = number of VCs per input port

n_R = number of router buffers per VC

n_C = number of channel buffers

Buffer Power – 8x8 Mesh and Folded Torus

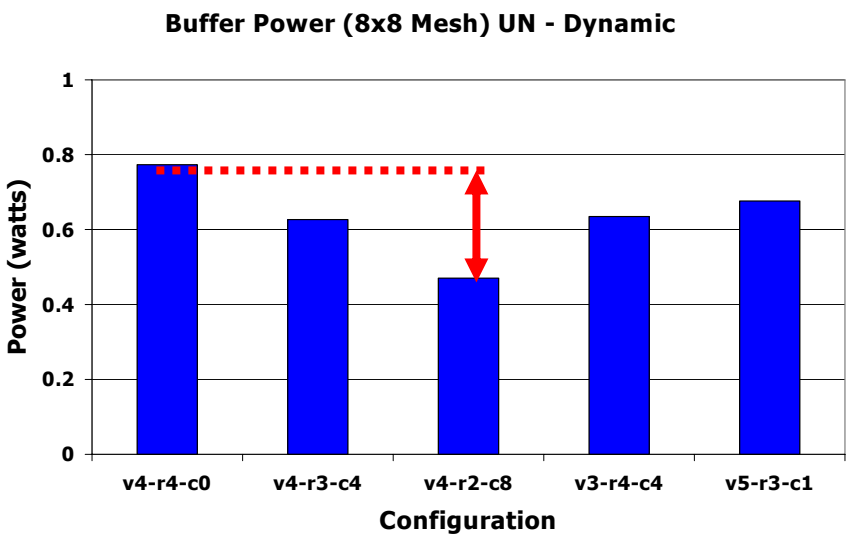

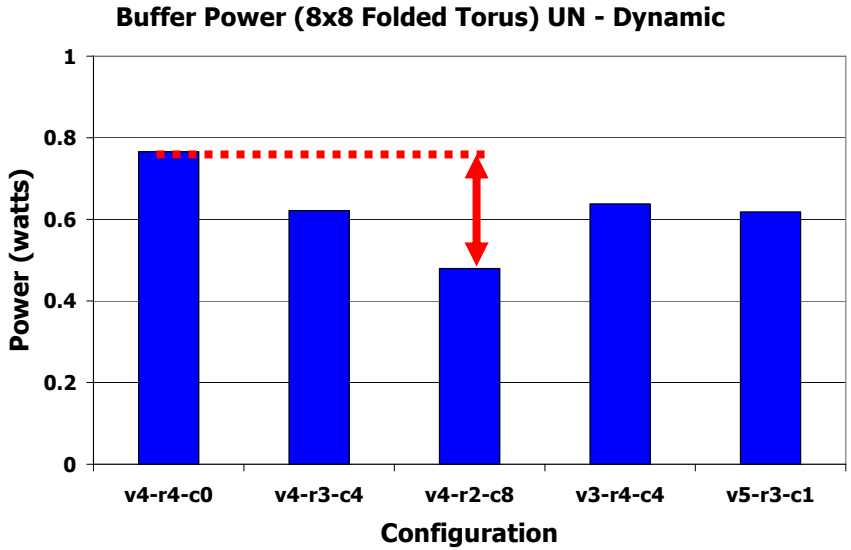

- Uniformly distributed traffic

⇒ **Nearly 40% power savings for 50% buffer size reduction (428), using Dynamic buffer allocation**

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

Throughput – 8x8 Mesh and Folded Torus

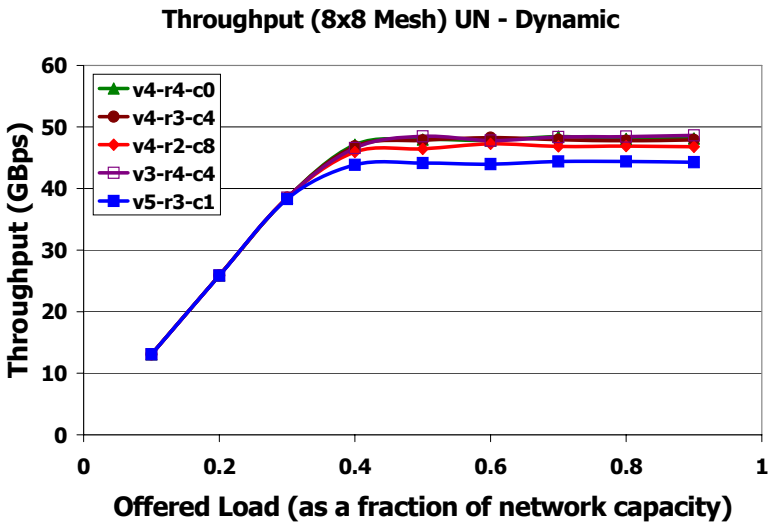

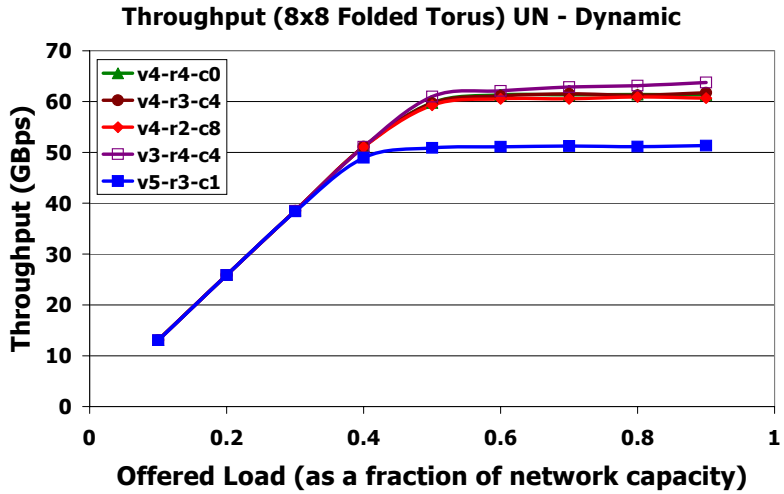

- Uniformly distributed traffic

⇒ **Only about 5% drop in throughput for the 428 case (Dynamic buffer allocation)**

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

Overall Network Power – 8x8 Mesh and Folded Torus

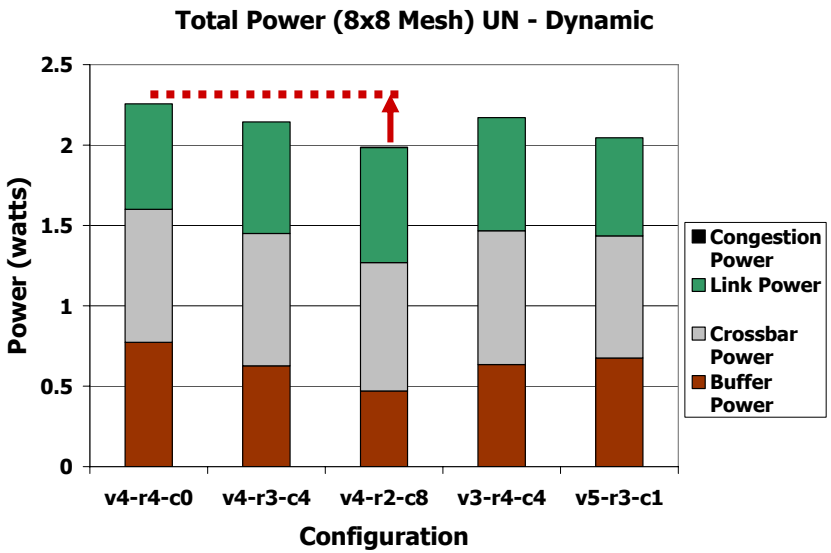

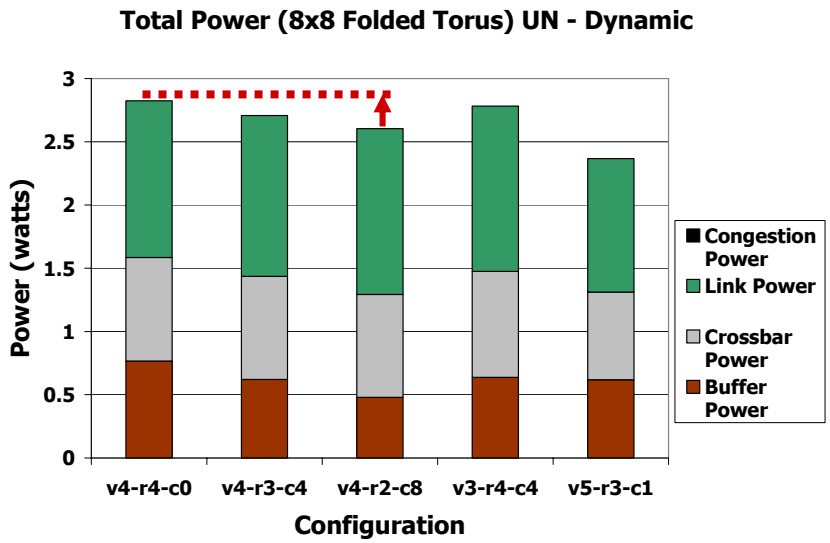

- **Total power consumed for a network load of 0.5**

⇒ Nearly 20% savings for the 428, using Dynamic buffer allocation

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

Buffer Power – 8x8 Mesh – all Traffic Patterns

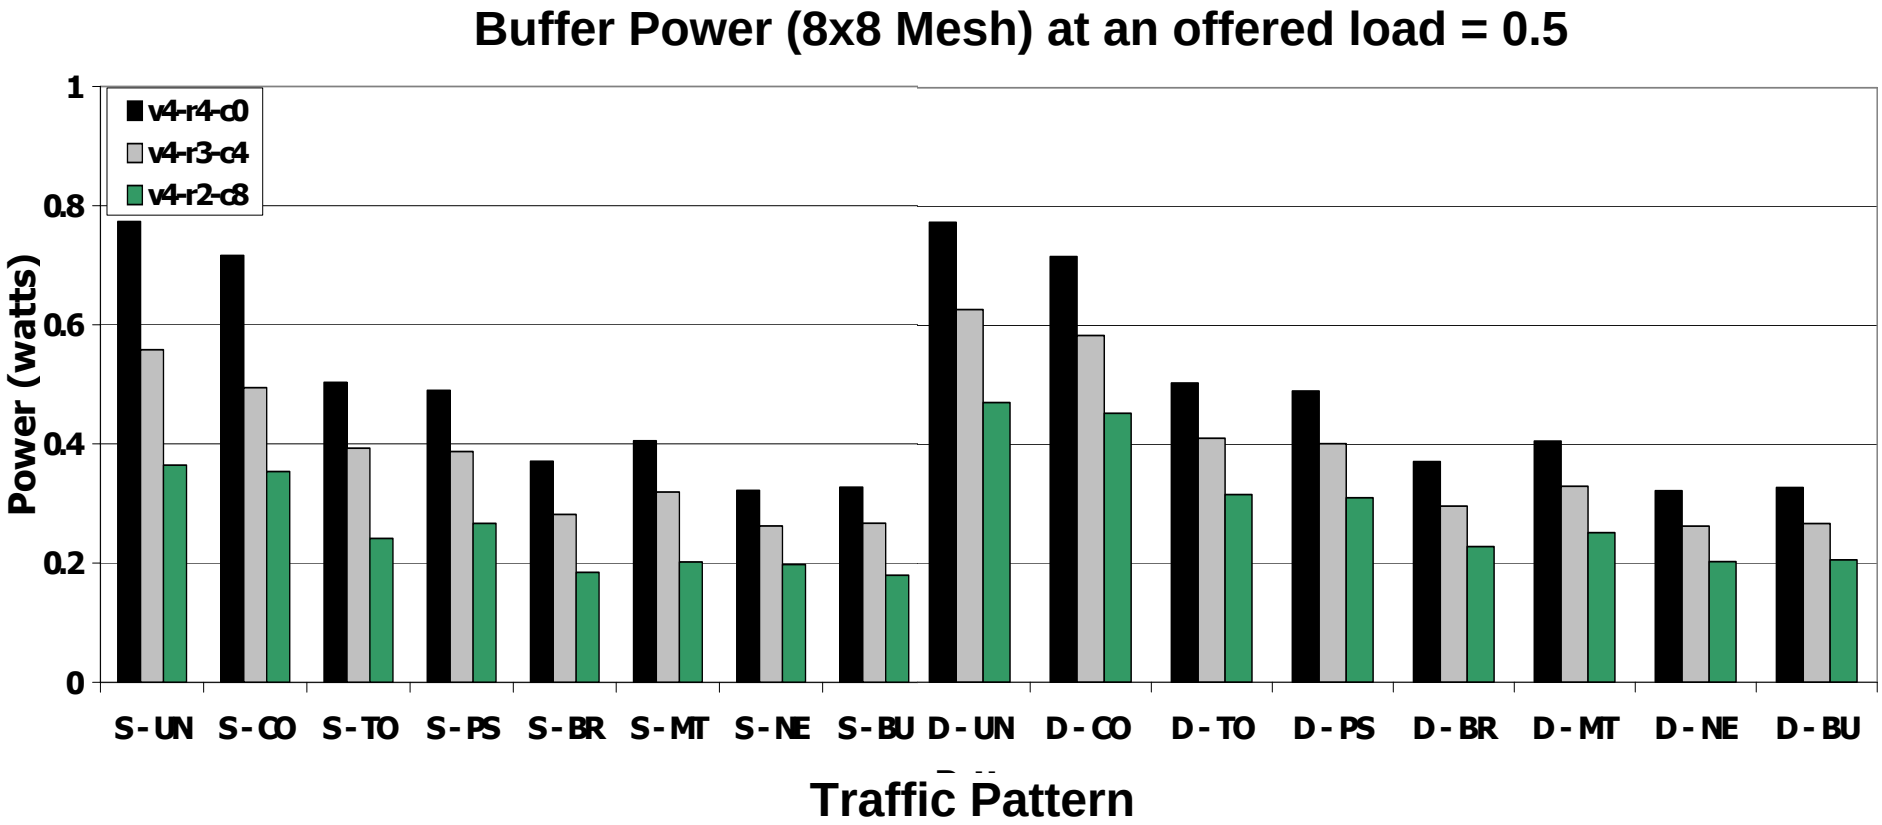

- **Reduction in power for all configurations, under all traffic patterns, compared to the baseline (440)**
- **For example, under Complement traffic the 428 configuration achieves 45% savings under Static allocation and 37.5% savings under Dynamic allocation**

Throughput – 8x8 Mesh – all Traffic Patterns

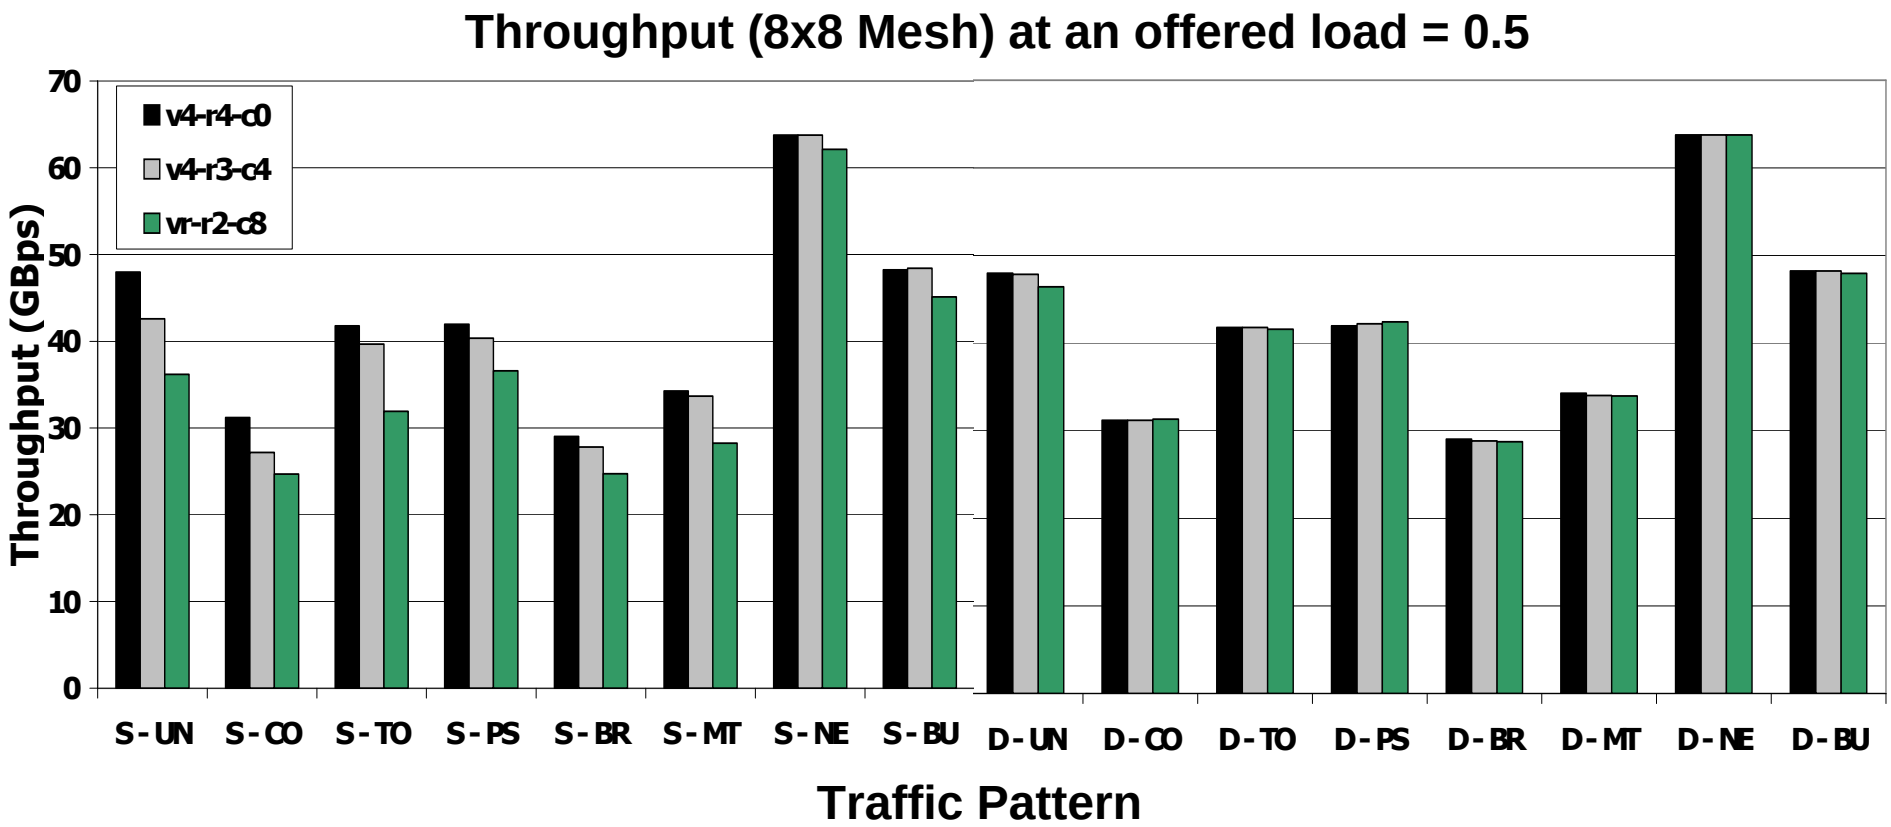

- **No significant decrease in throughput under any traffic pattern, using Dynamic allocation**

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 **40-52% savings in power** is achieved, with a significant reduction in router area. There is only a marginal 1-5% drop in performance, under dynamic buffer allocation.

- Future work will involve (a) Simulation using real-application traces (b) Exploring architectural improvements such as aggressive speculation in the credit loop

Backup Slides

Area Estimation – Summary with values from Synopsys Design Compiler

vnV – rnR - cnC	Buffer Area (μm^2)	Link Repeater Area (μm^2)	Total Buffer + Link Area (μm^2)	% Change
v4-r4-c0	81,407	32	81,439	-
v4-r3-c4	63,991	52	64,011	-21.40
v4-r2-c8	**48,066**	**80**	**48,146**	**-40.88**
v3-r4-c4	63,250	52	63,302	-22.27
v3-r3-c7	50,373	73	50,446	-38.05
v5-r2-c6	53,712	66	53,778	-33.96
v5-r3-c1	73,797	38	73,803	-9.37

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

Latency – 8x8 Mesh and Folded Torus

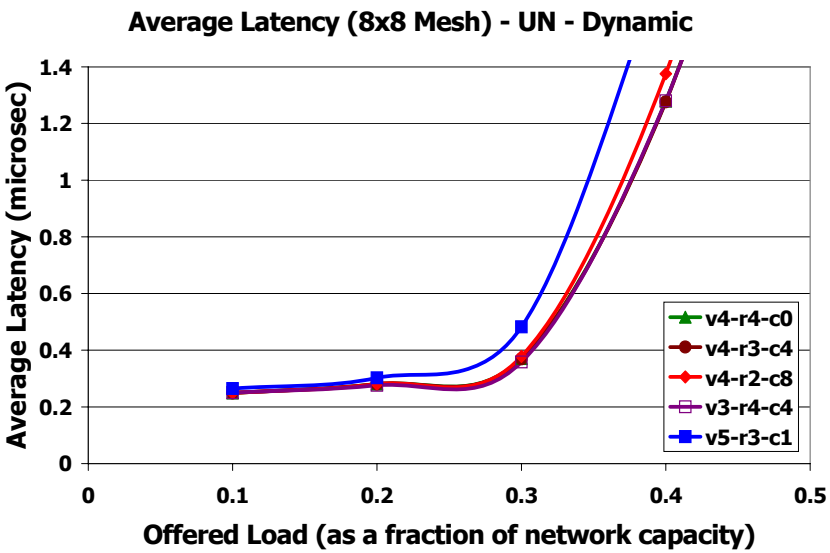

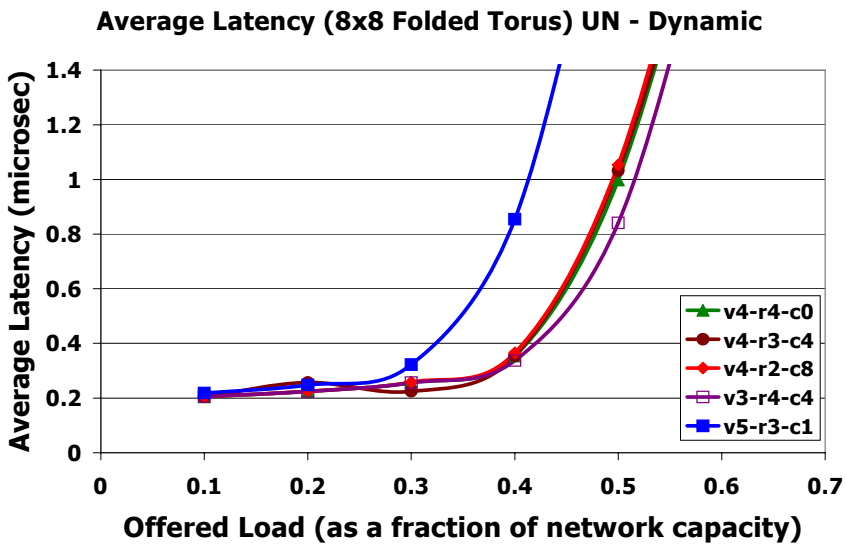

- Uniformly distributed traffic

⇒ For all cases (except 531), saturation for a network load of about 0.3 in case of Mesh and about 0.4 in case of Folded torus

Comparison with FC-CB and DAMQ

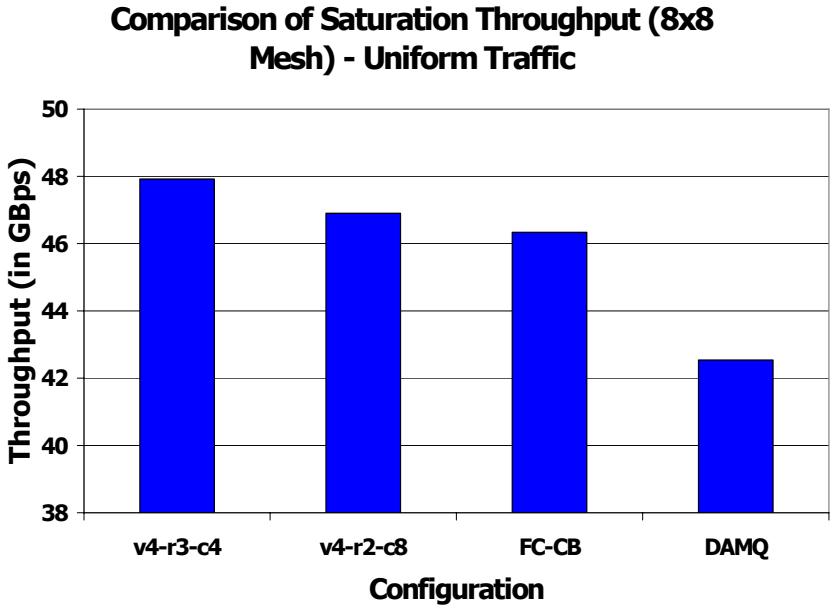

- FC-CB shows similar performance as the dynamically allocated 440 case
- 434 and 428 achieve nearly 4% increase in saturation throughput compared to FC-CB
- 428 achieves nearly 12.5% improvement in saturation throughput compared to DAMQ

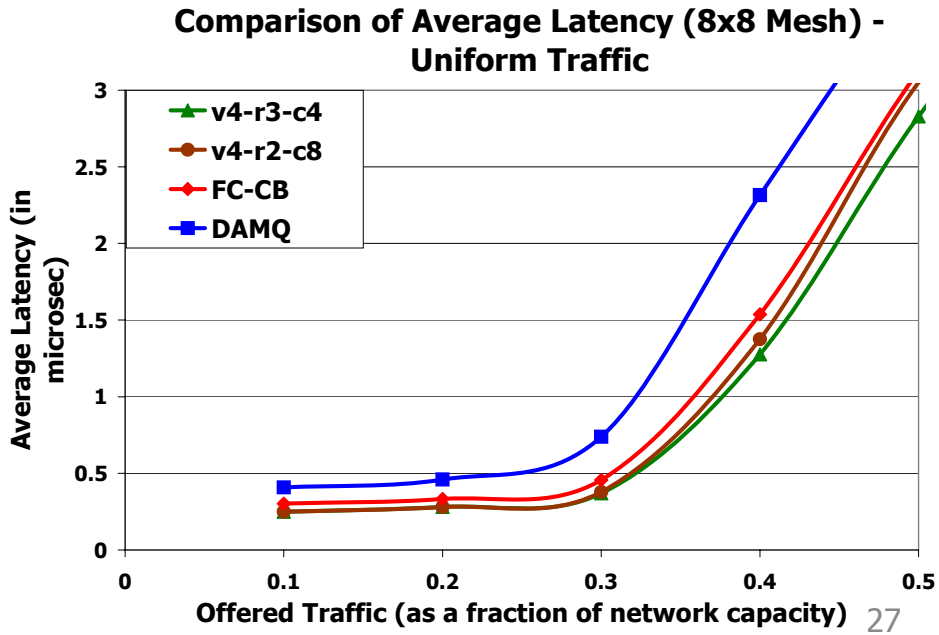

Power calculations using Synopsys Power Compiler

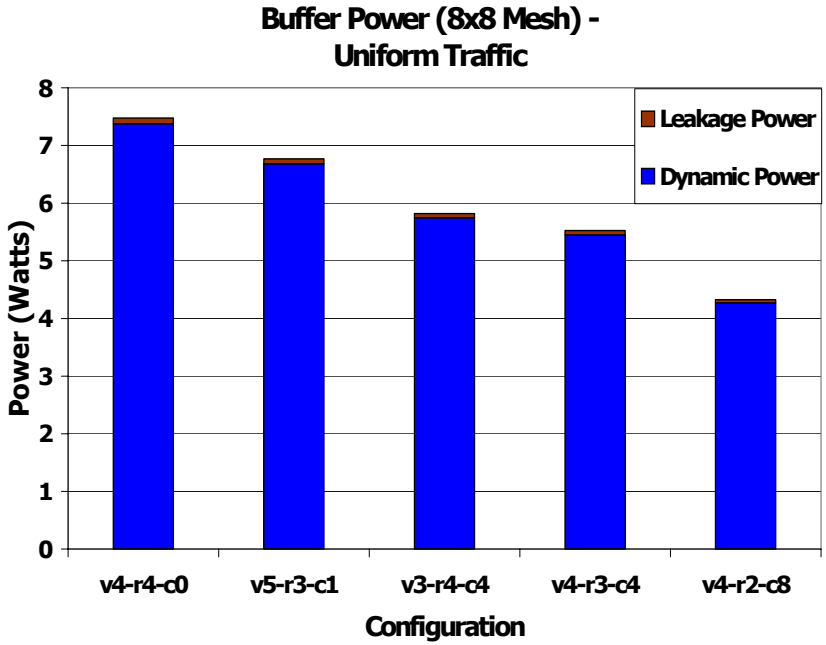

- 428 case shows nearly 40% reduction in buffer power alone
- Nearly 30% decrease in overall network power for the 428 case

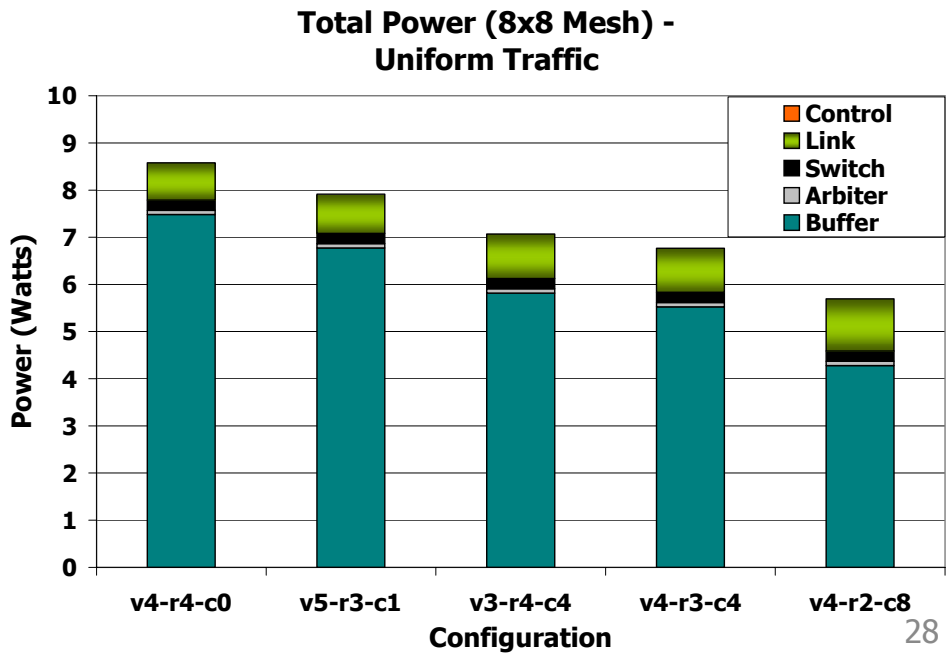

Data flow Control Simulated with Synopsys VCS

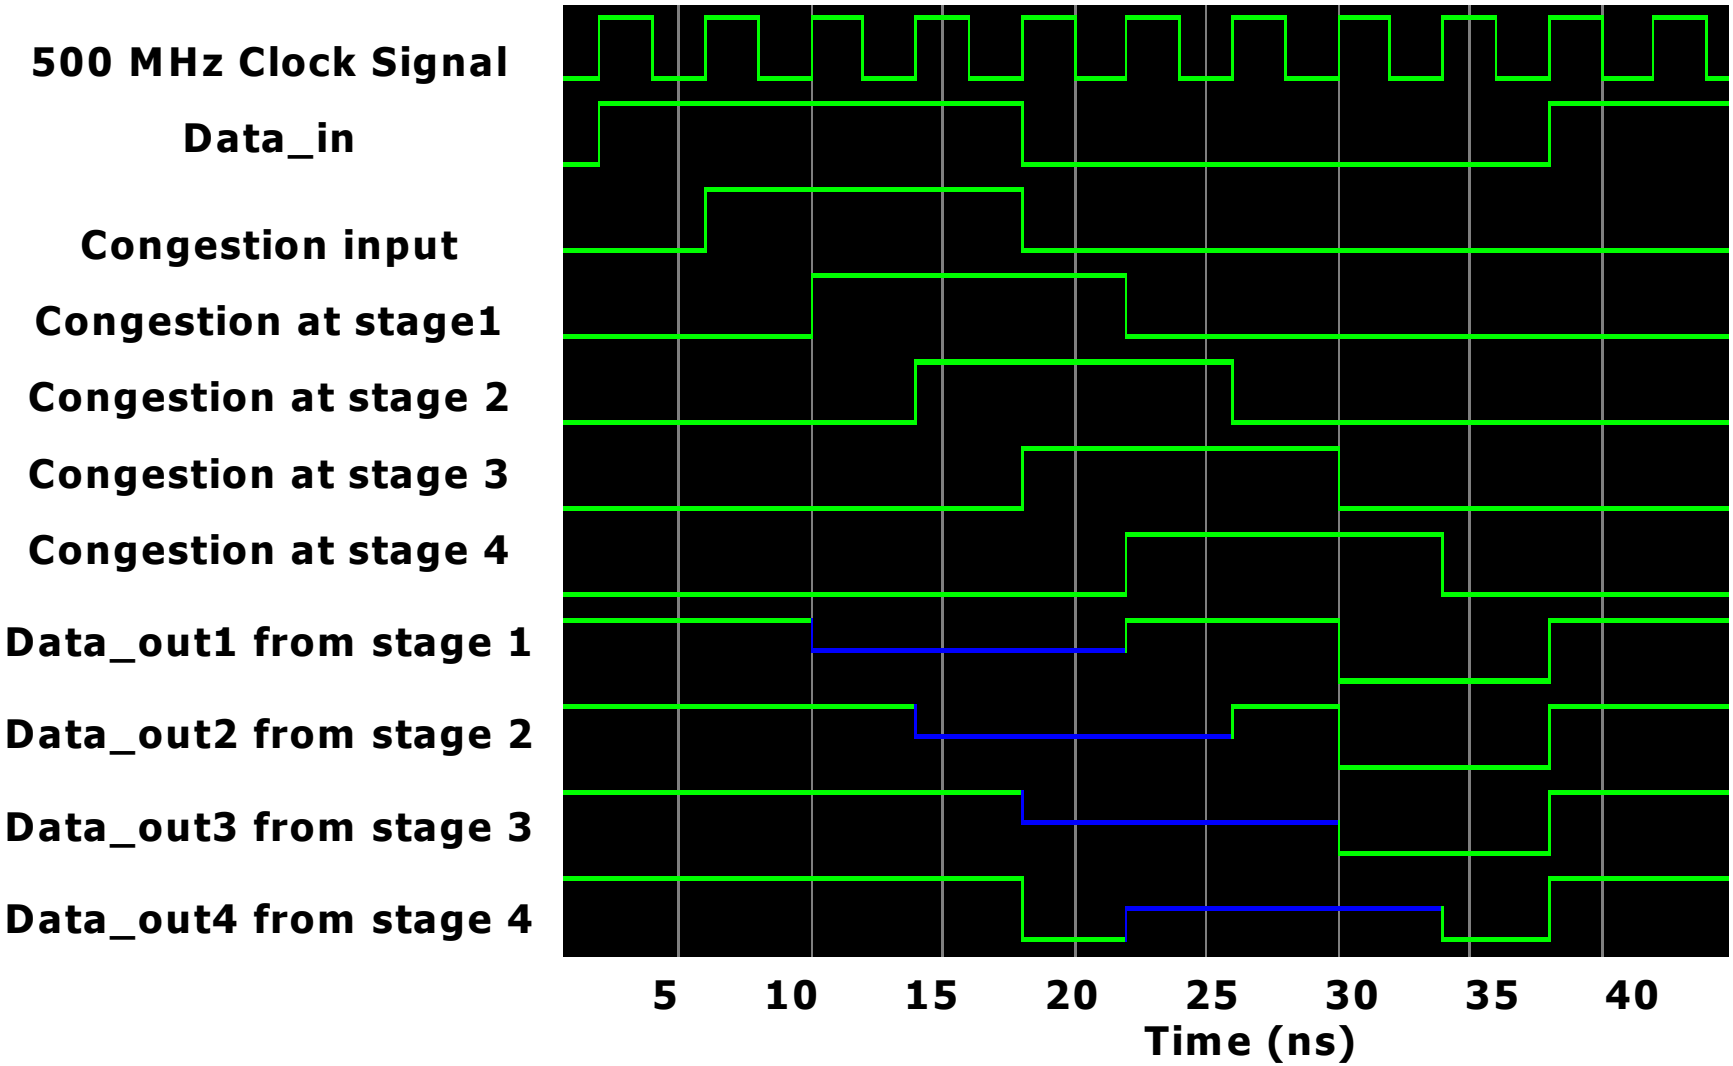

Router - Power Estimation

Component	Power / Area Calculation	Explanation
C_{buf}	$(1/2 \times W \times L \times C_{ox}) + (W \times L_{ov} \times C_{ox})$	C_{buf} = additional capacitance due to three-state repeater along the links W, L = Width & Length of min. sized inverter C_{ox} = oxide capacitance L_{ov} = gate-drain/source overlap length
$\acute{P}_{dynamic}$	$a \times [k(C_o + C_p + C_{buf}) + \ell C_w] \times V_{DD}^2 \times freq$	a = activity factor, k = repeater sizing, ℓ = repeater spacing C_o = diffusion capacitance C_p = gate capacitance C_w = wire capacitance V_{DD} = supply voltage freq = operating frequency
$\acute{P}_{leakage}$	$2 \times [1/2 \times V_{DD} \times (I_{off}(W_n + W_p)k)]$	I_{off} = subthreshold leakage current W_n (W_p) = width of the NMOS (PMOS) in the repeater
$\acute{P}_{short\text{-}ckt}$	$a \times t_{rise} \times W_n \times k \times V_{DD} \times I_{sc} \times freq$	t_{rise} = rise time of the short-ckt current I_{sc}

iDEAL – Control Block

- Self-checking Double-sampling technique for the Control block

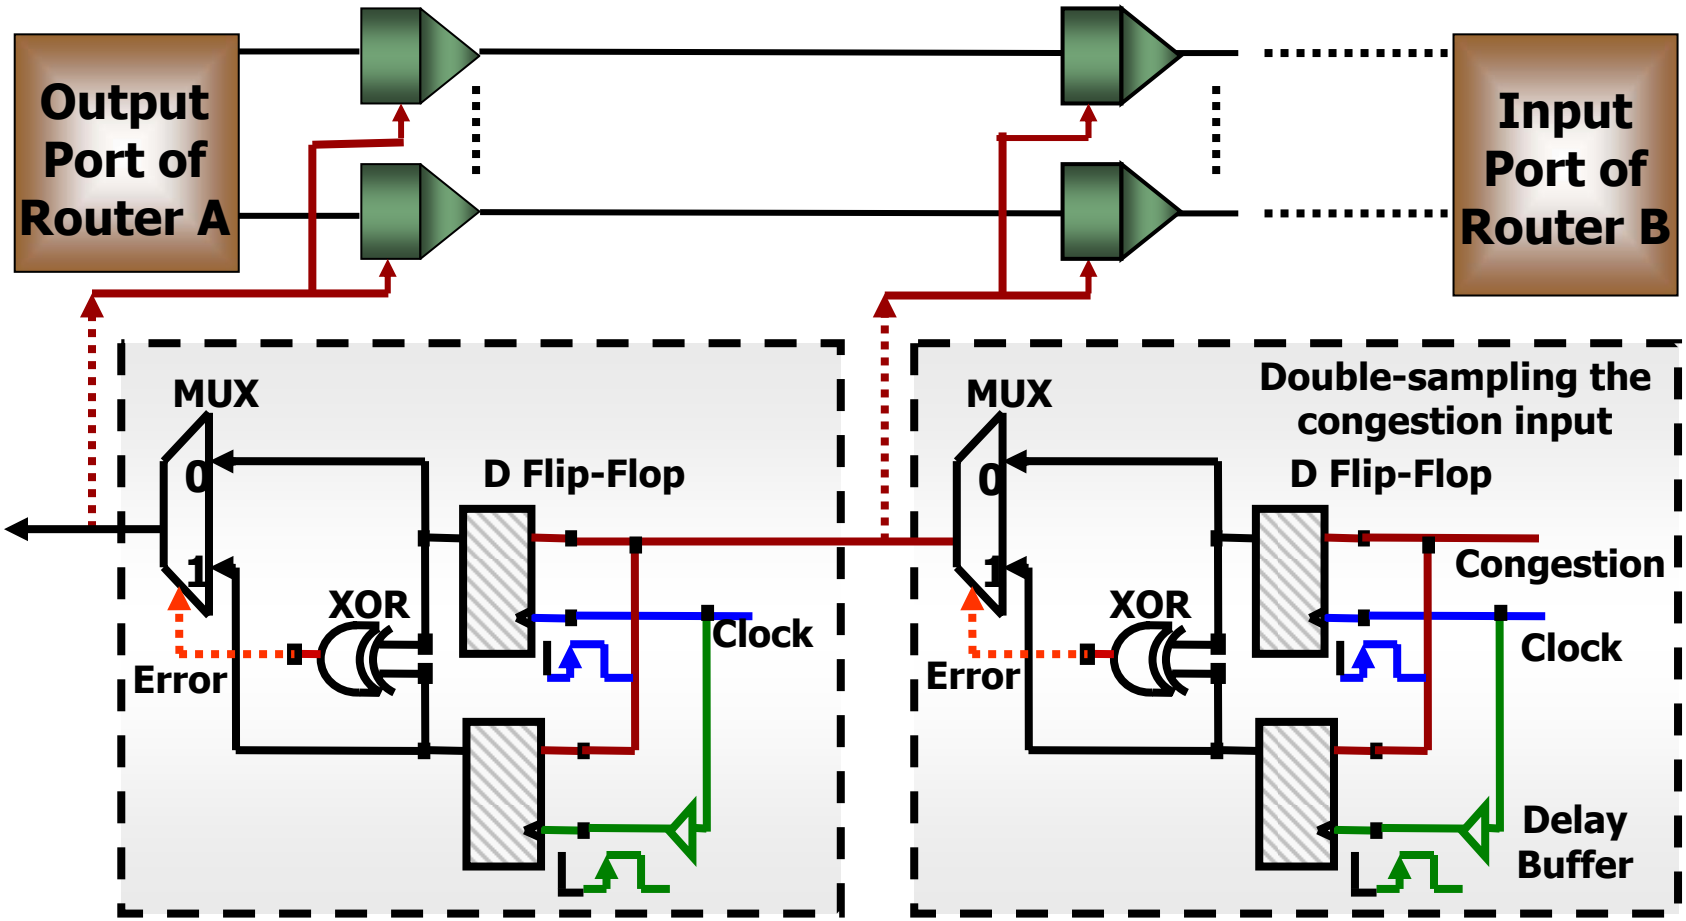

- Slightly more power (0.02 uW v/s 0.06 uW) and area, but more reliable

Aggressive Speculation

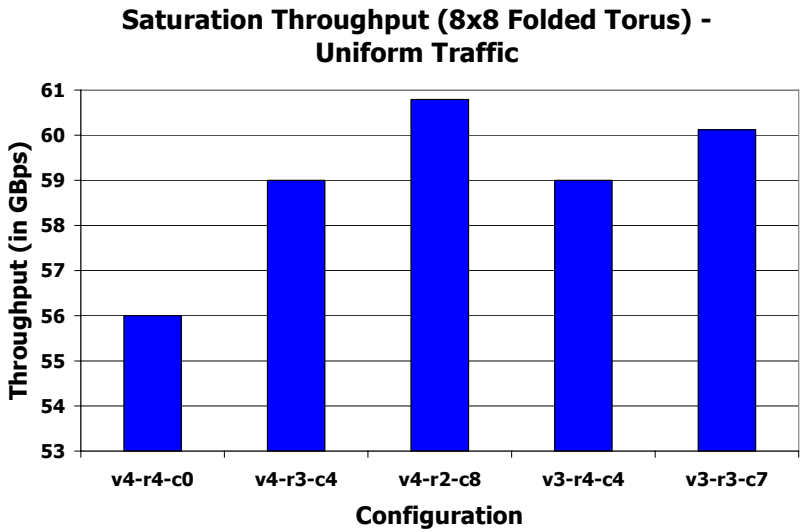

- Aggressive speculation by increasing the number of credits available to 8
- Additional credits are accounted for by the channel buffers

⇒ Saturation throughput improves by 10% for the 428 case

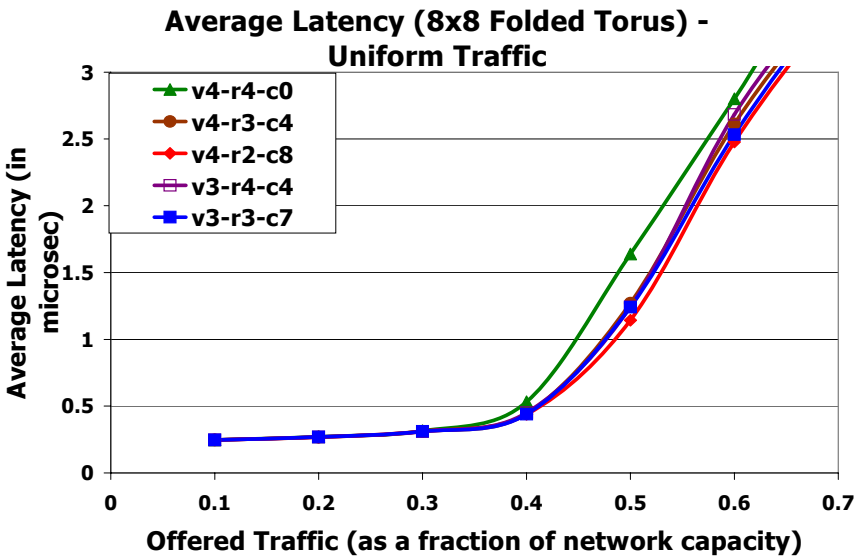

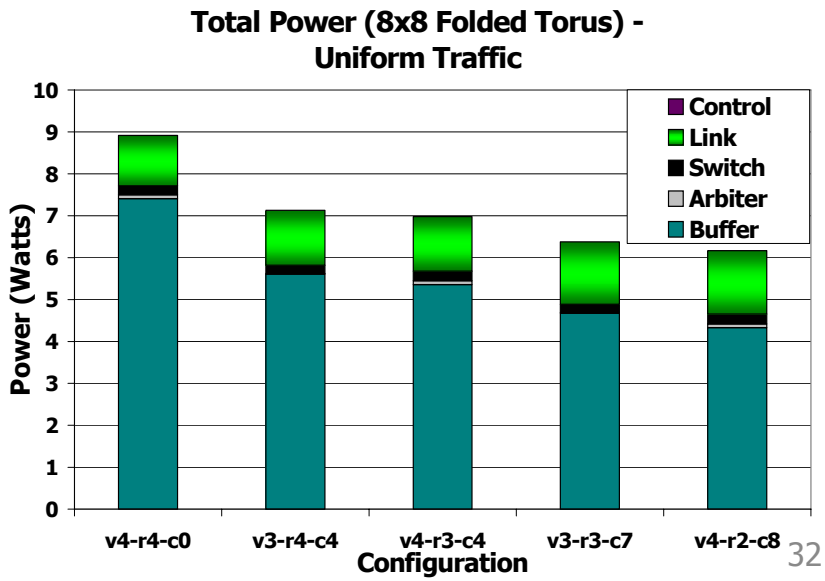

Power Estimation – Summary with values from Synopsys Power Compiler

vnV – rnR - cnC	Buffer Power (mW)	Mesh Link + Control Power (mW)	Total Power (Buffer + Link) (mW)	% Change
v4-r4-c0	19.54	2.45	21.99	-
v4-r3-c4	14.51	2.91	17.42	-20.78
v4-r2-c8	**11.57**	**3.57**	**15.14**	**-31.15**
v3-r4-c4	15.09	2.91	18.00	-18.14
v3-r3-c7	12.56	3.50	16.06	-26.96
v5-r2-c6	14.41	3.31	17.72	-19.41
v5-r3-c1	19.29	2.81	22.10	+ 0.50

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