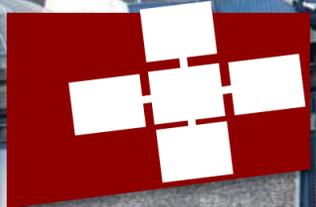


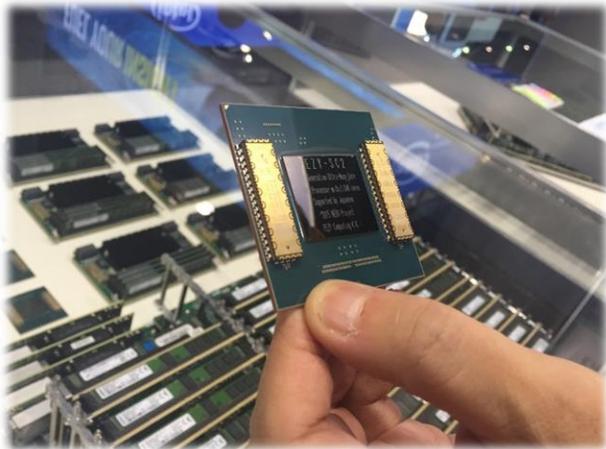
# Slim NoC: A Low-Diameter On-Chip Network Topology for High Energy Efficiency and Scalability

MACIEJ BESTA, SYED MINHAJ HASSAN, SUDHAKAR YALAMANCHILI,  
RACHATA AUSAVARUNGNIRUN, ONUR MUTLU, TORSTEN HOEFLER

The word "SAFARI" is written in a bold, orange, sans-serif font on a white rectangular background.The Georgia Tech logo, consisting of the words "Georgia Tech" in a yellow, sans-serif font on a dark blue background, with a small yellow tower icon to the right.The SPCL logo, featuring the letters "SPCL" in a large, white, bold, sans-serif font, set against a green background with a white outline of a mountain range.

# MASSIVELY PARALLEL MANYCORES

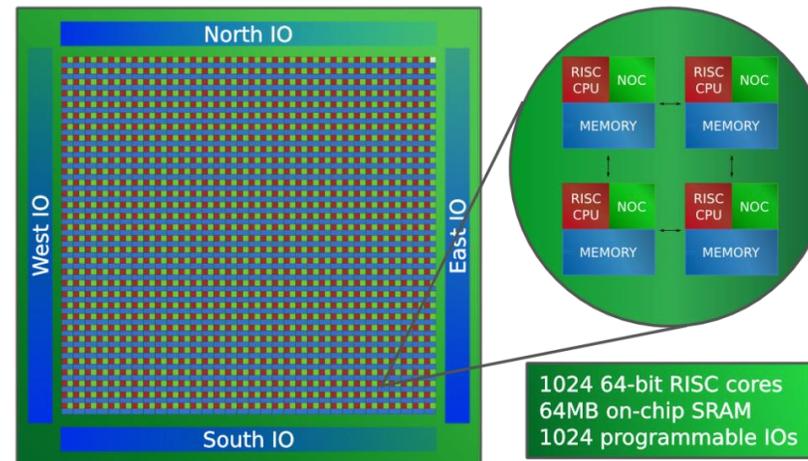
PEZY-SC2: 2048 cores



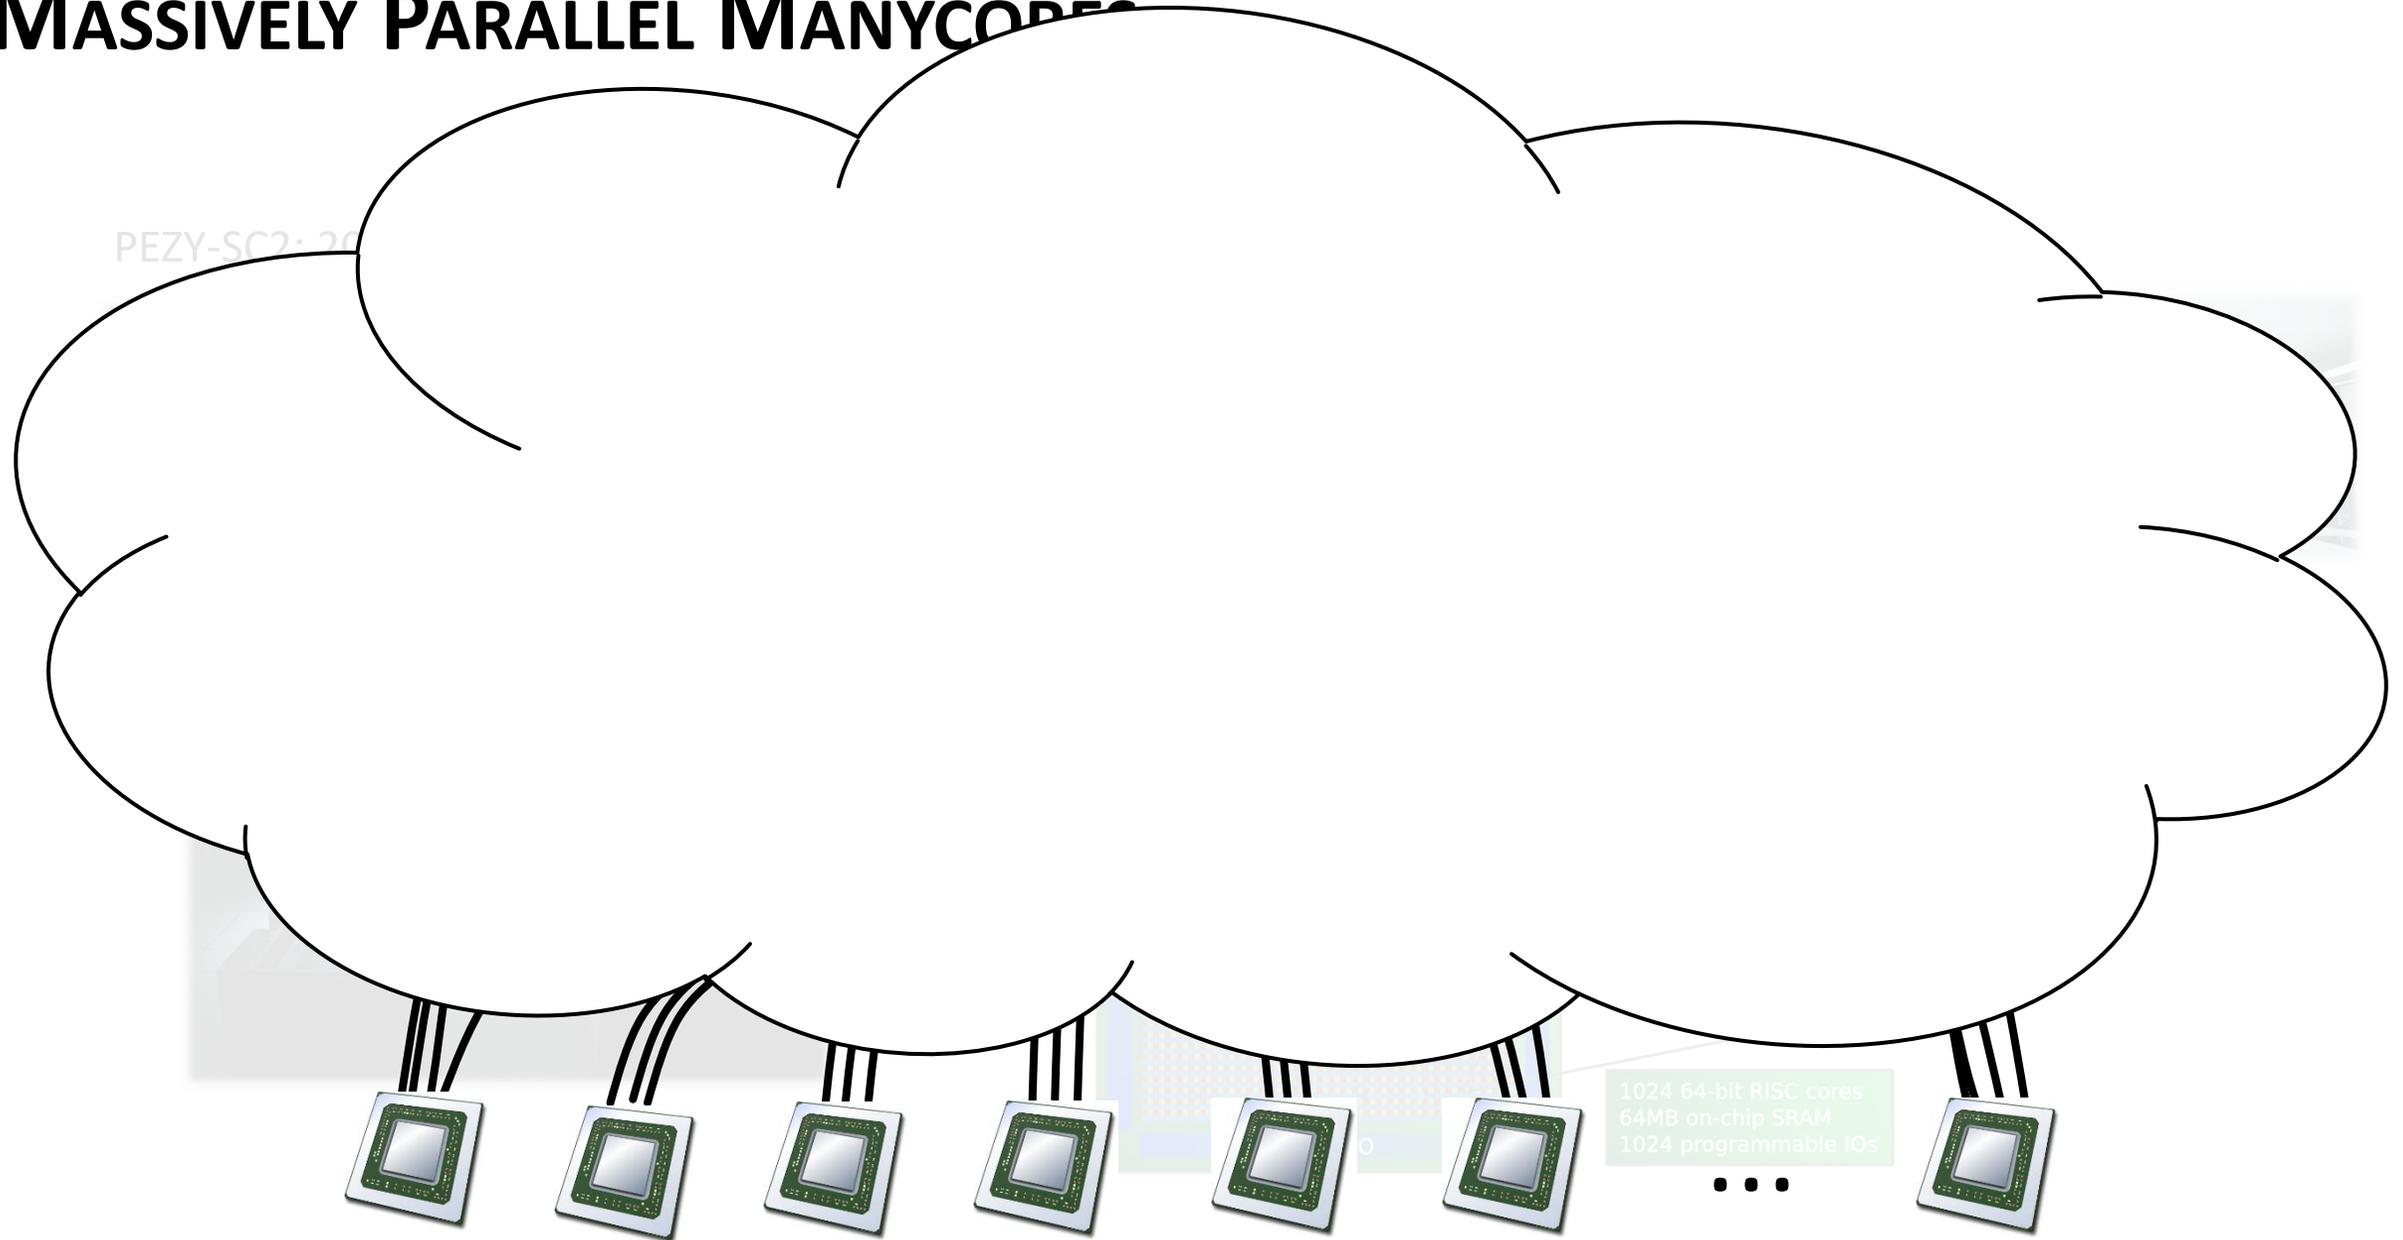
SW26010: 260 cores



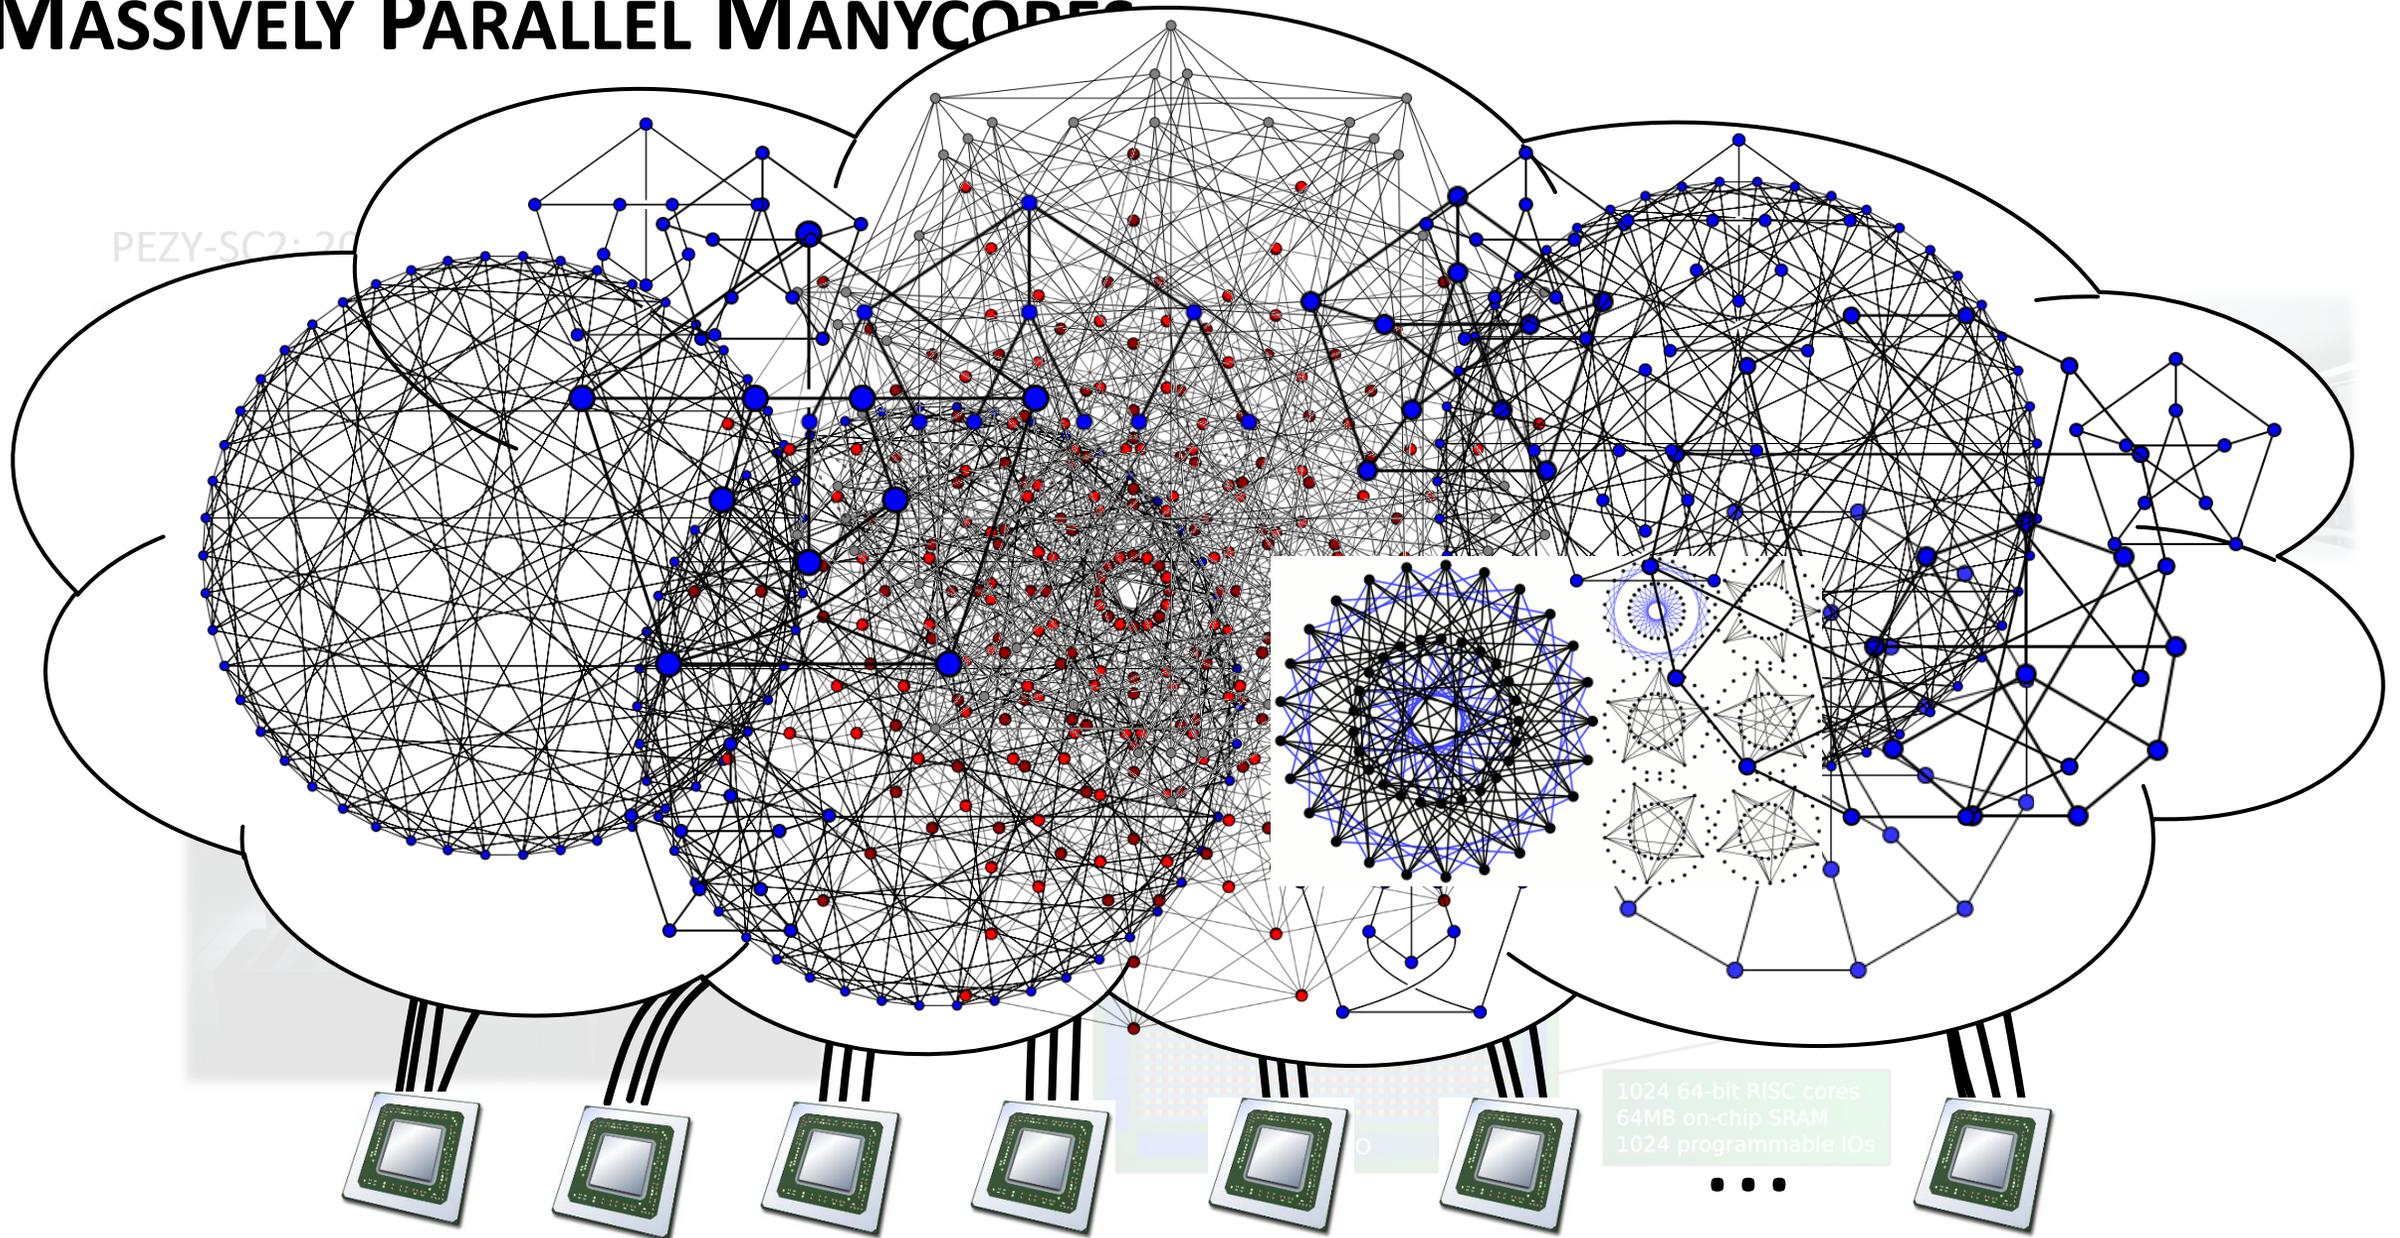
Adapteva Epiphany: 1024 cores



# MASSIVELY PARALLEL MANYCORES

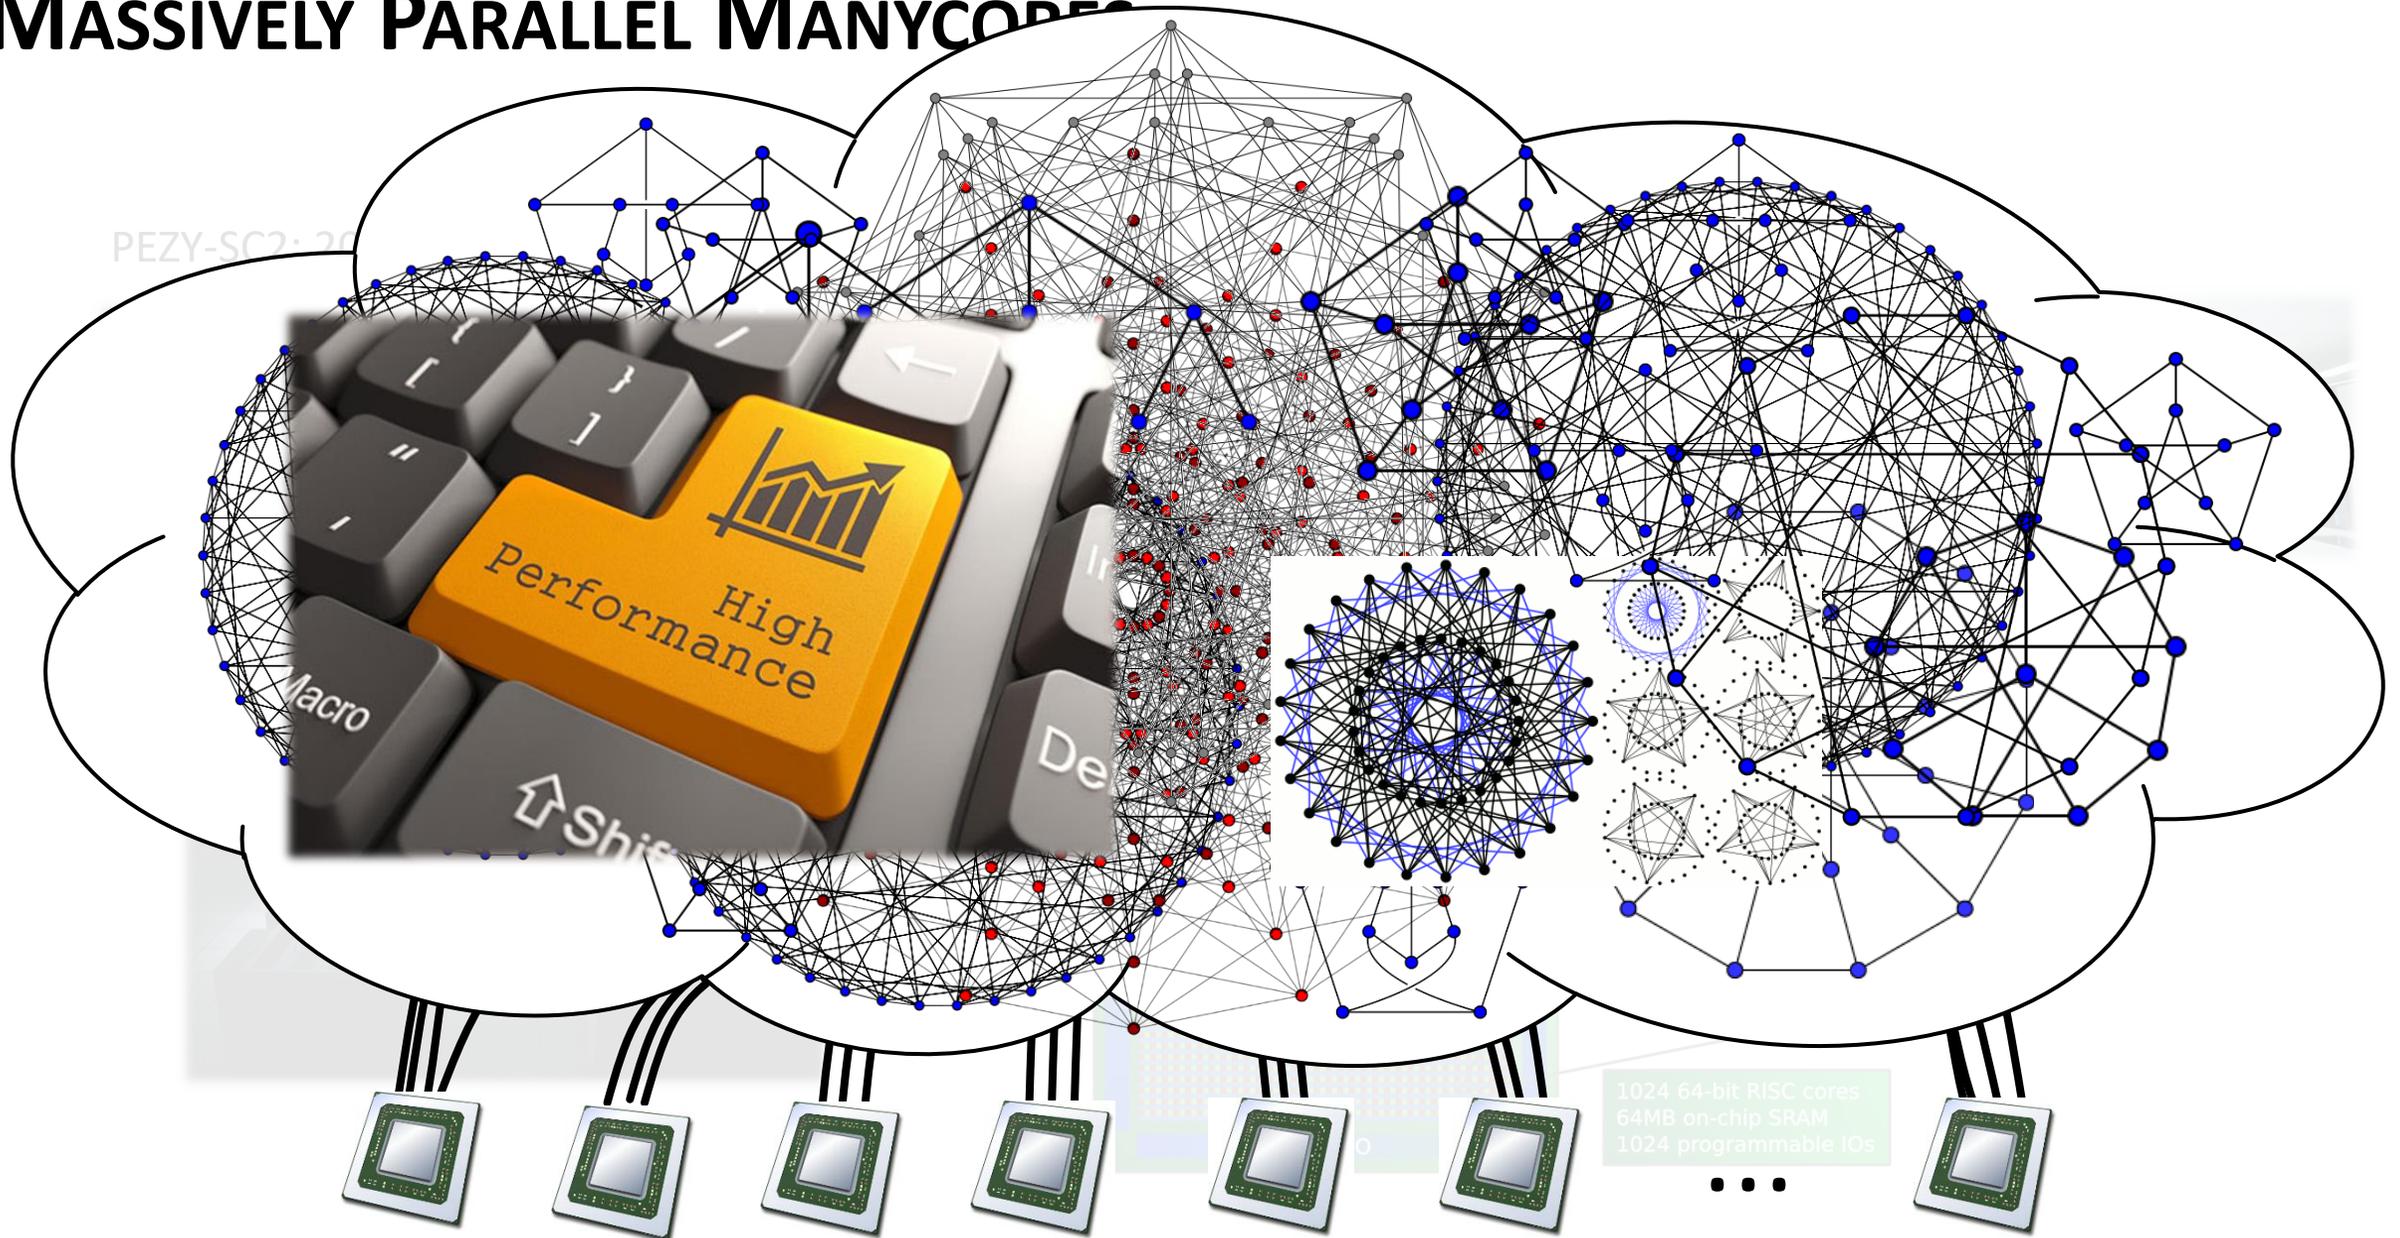


# MASSIVELY PARALLEL MANYCORES

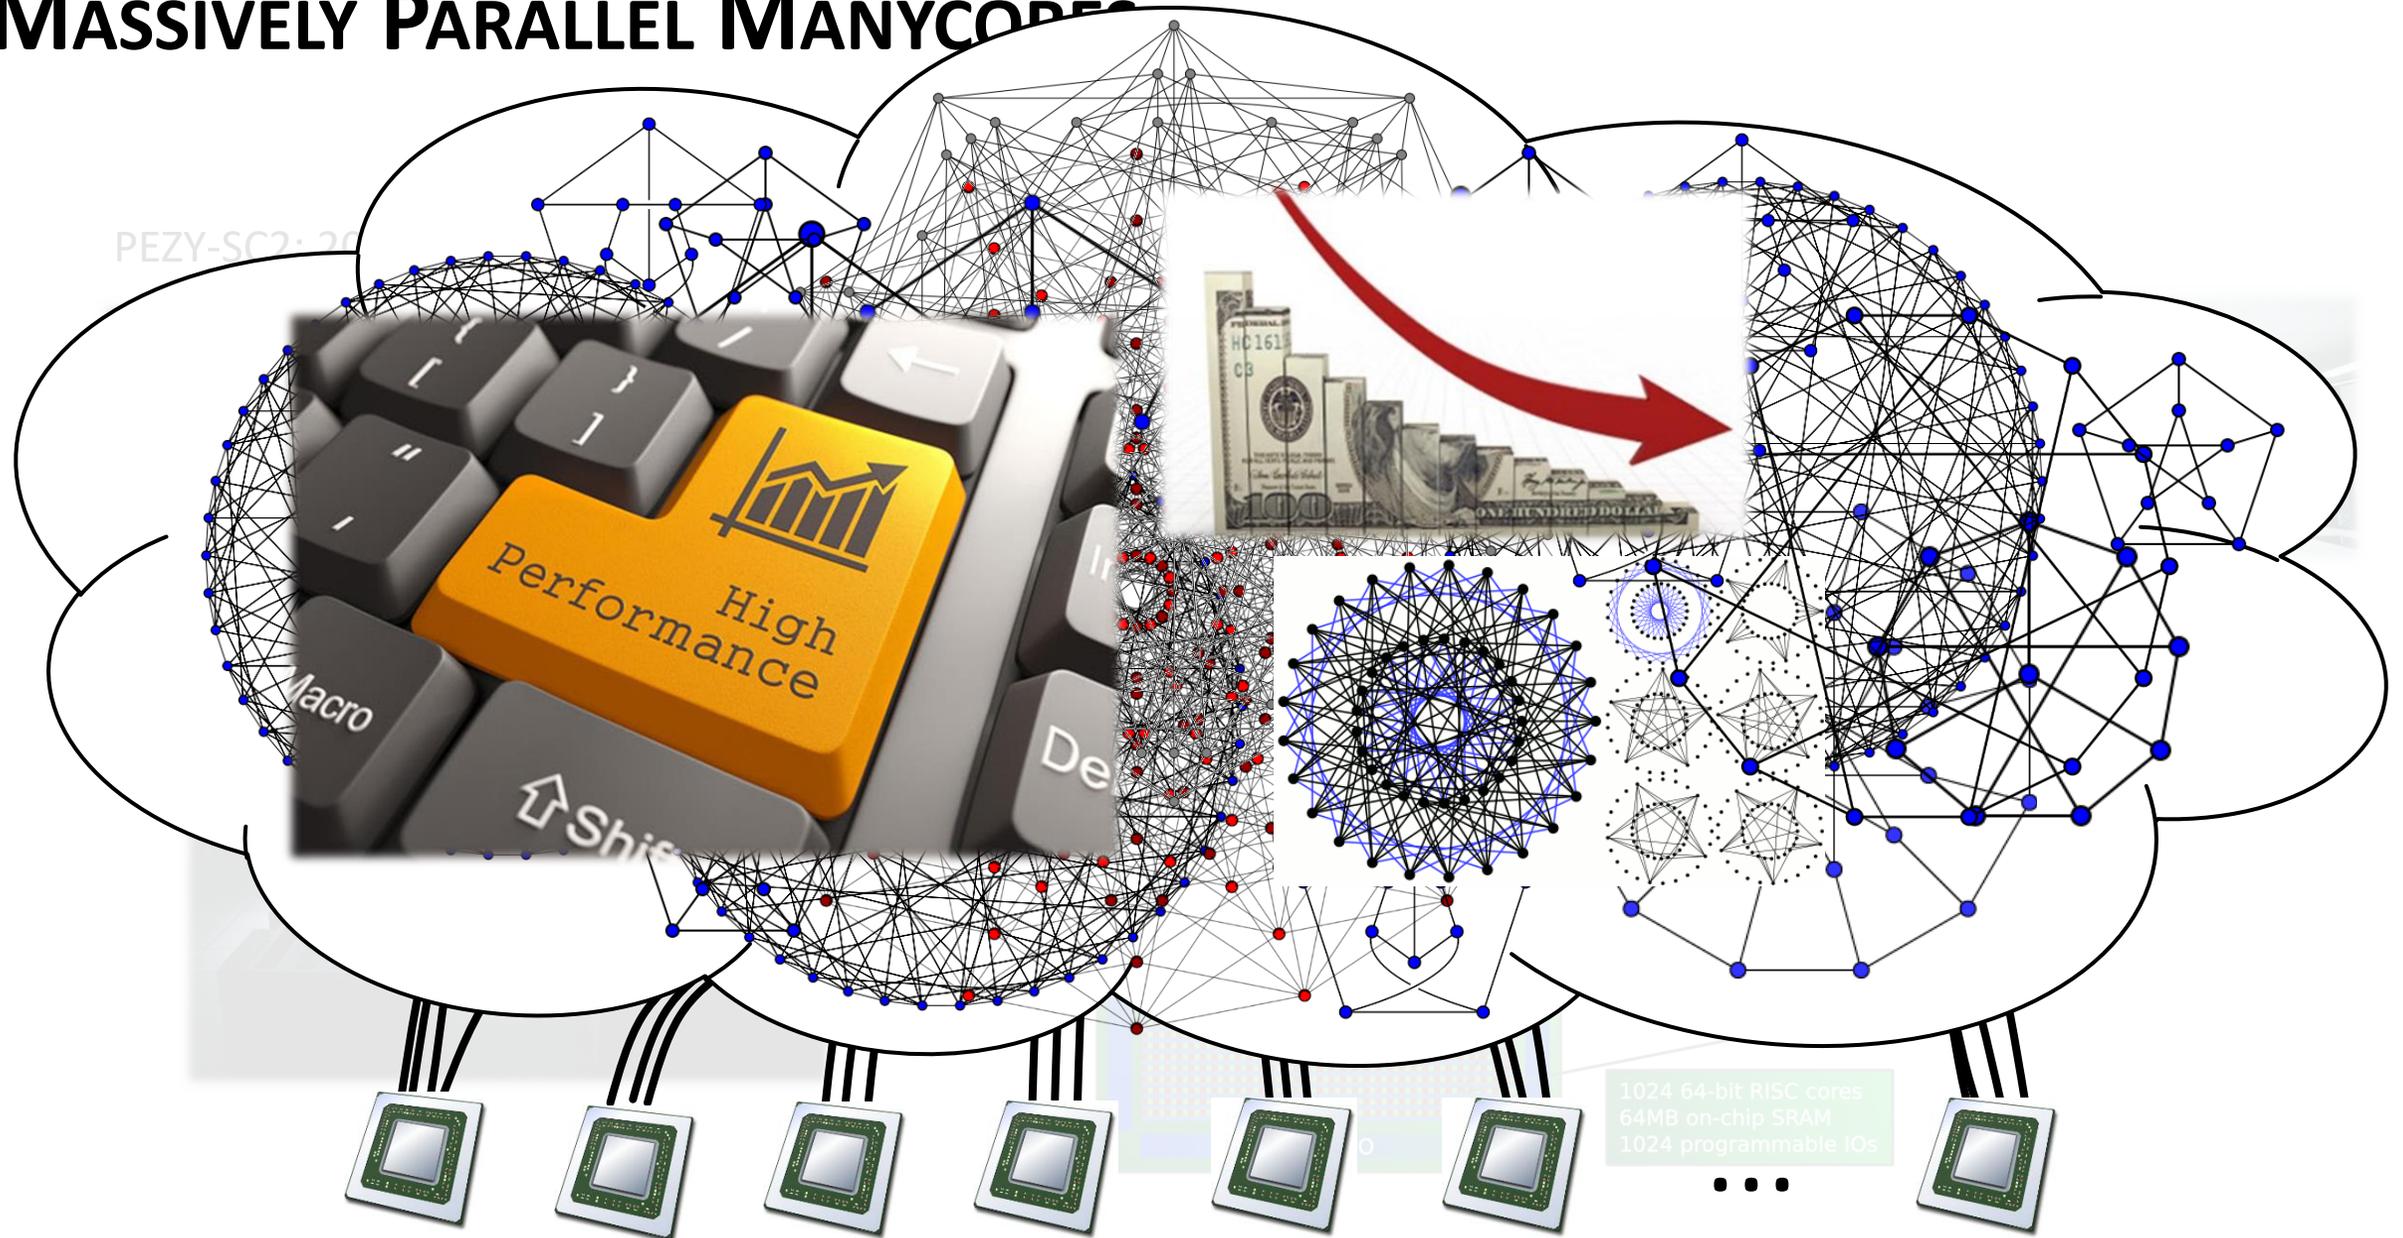


PEZY-SC2: 20

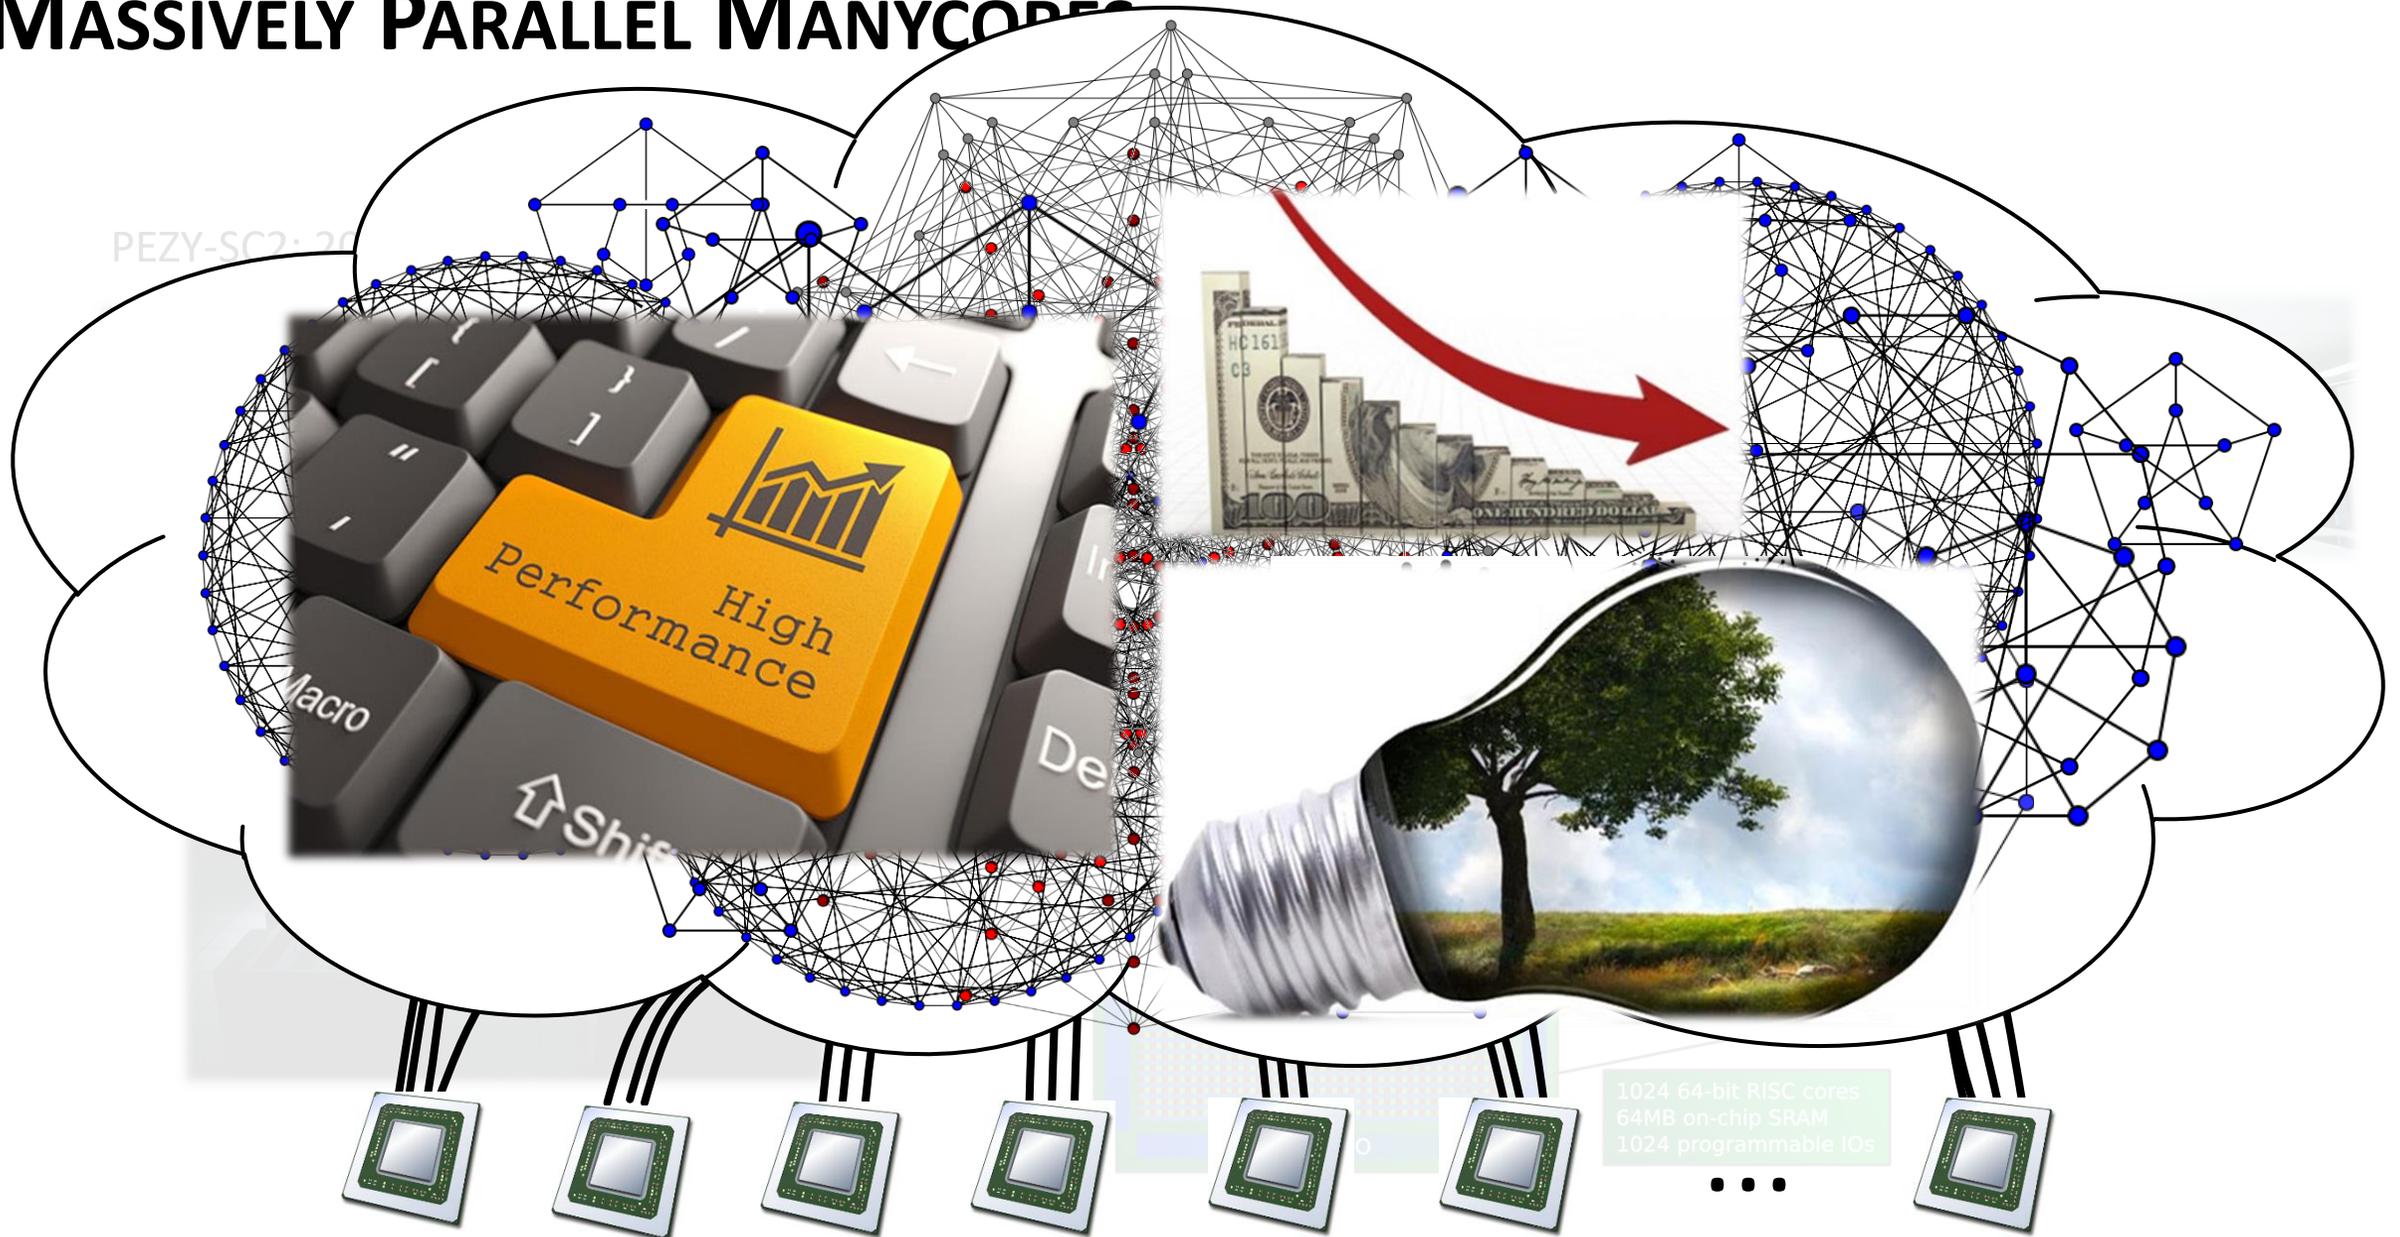
# MASSIVELY PARALLEL MANYCORES



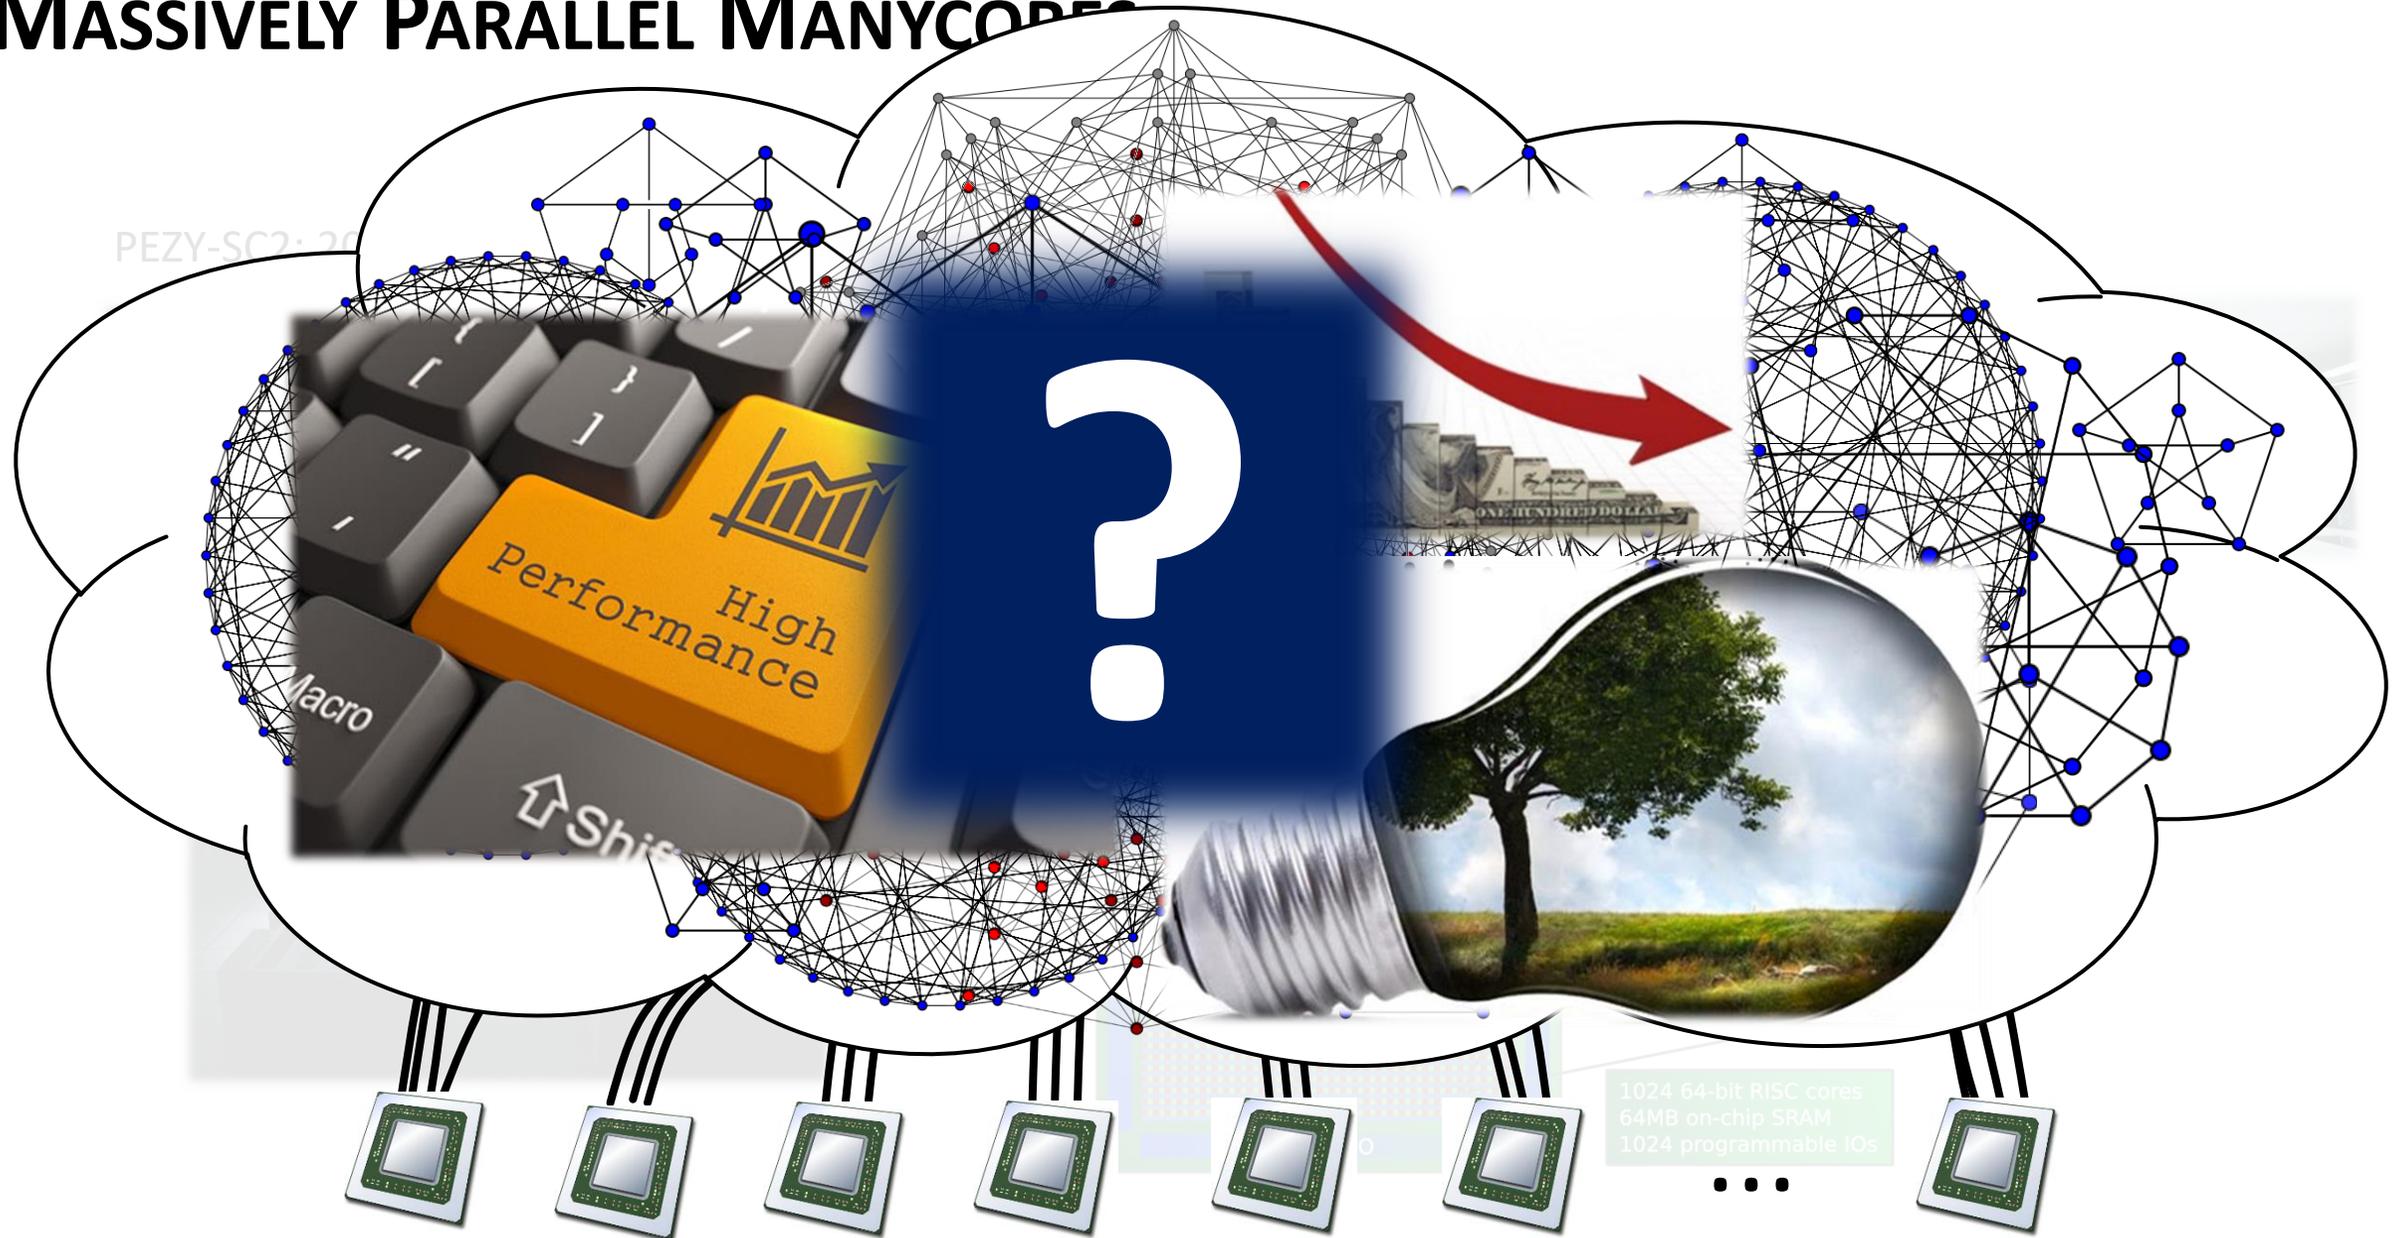
# MASSIVELY PARALLEL MANYCORES



# MASSIVELY PARALLEL MANYCORES

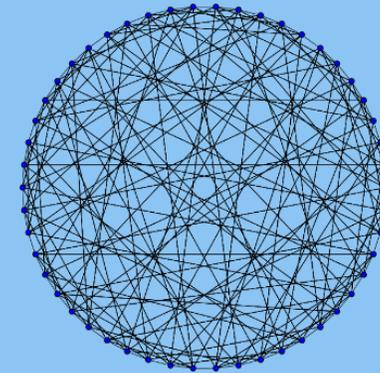


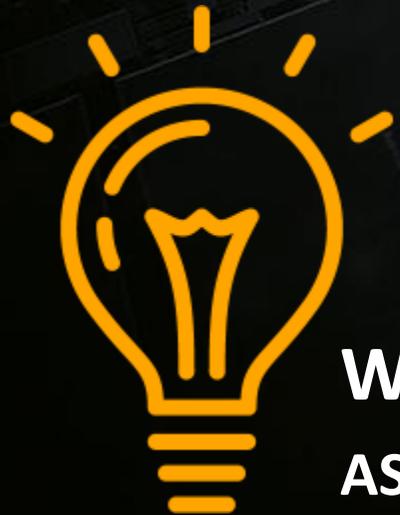
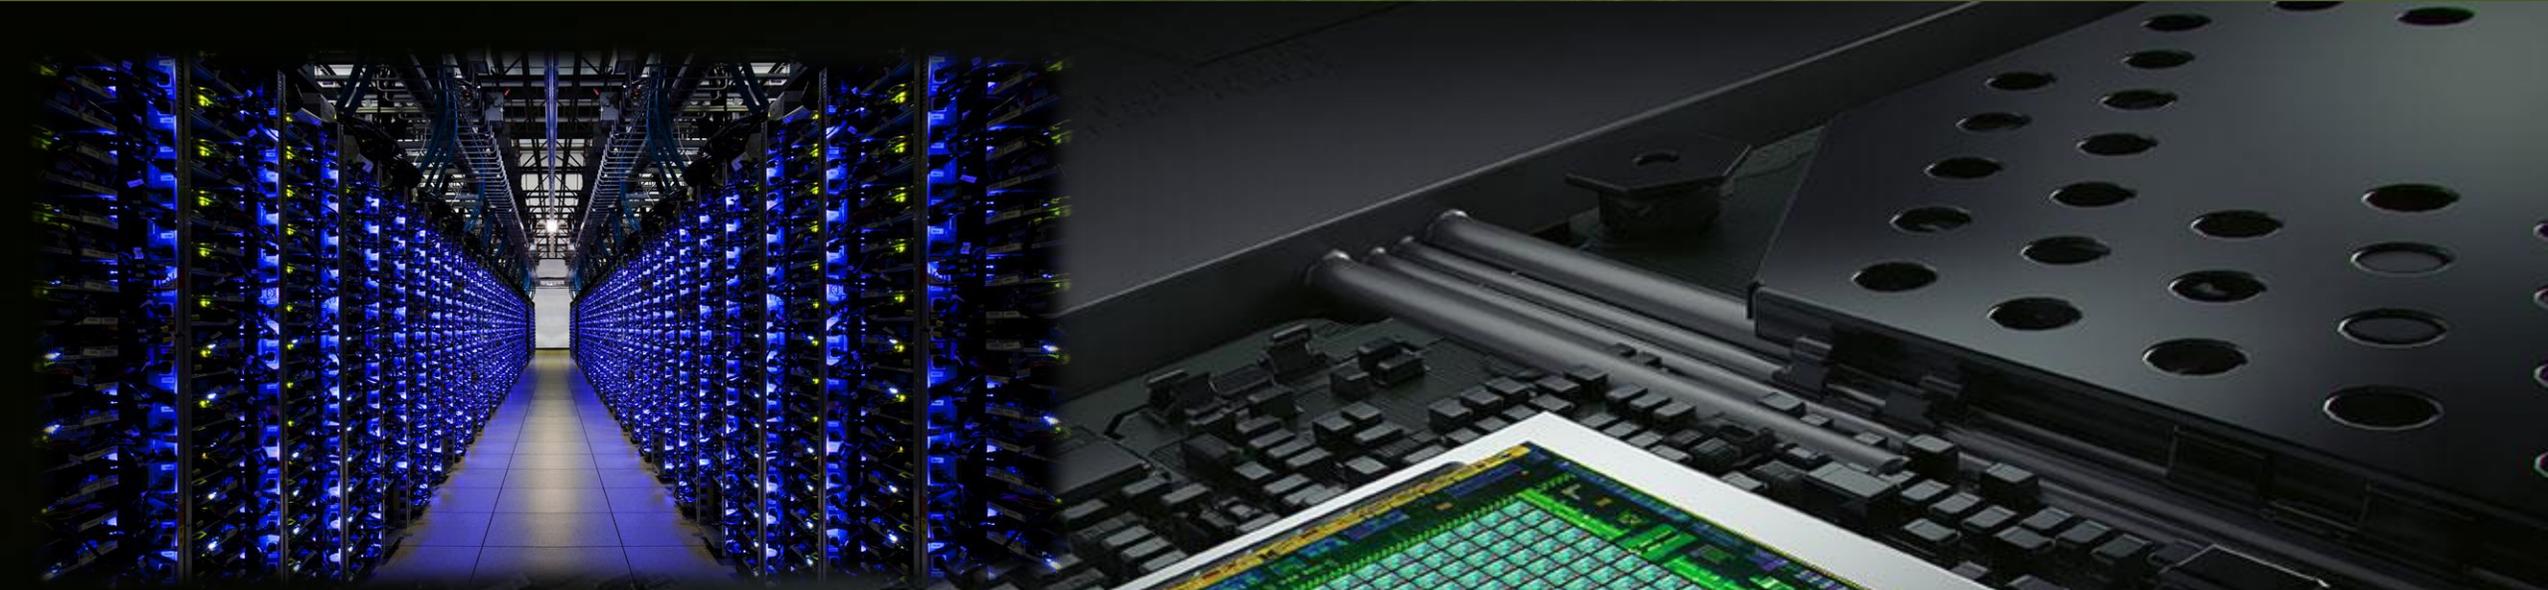
# MASSIVELY PARALLEL MANYCORES



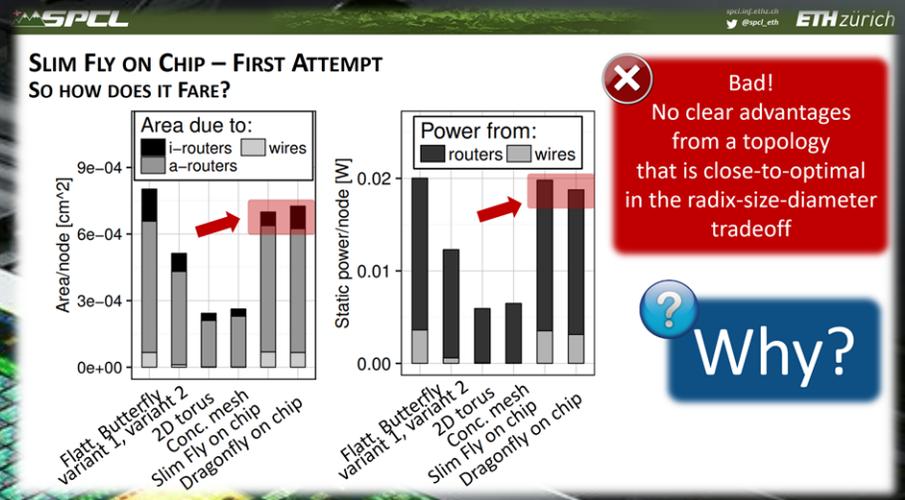
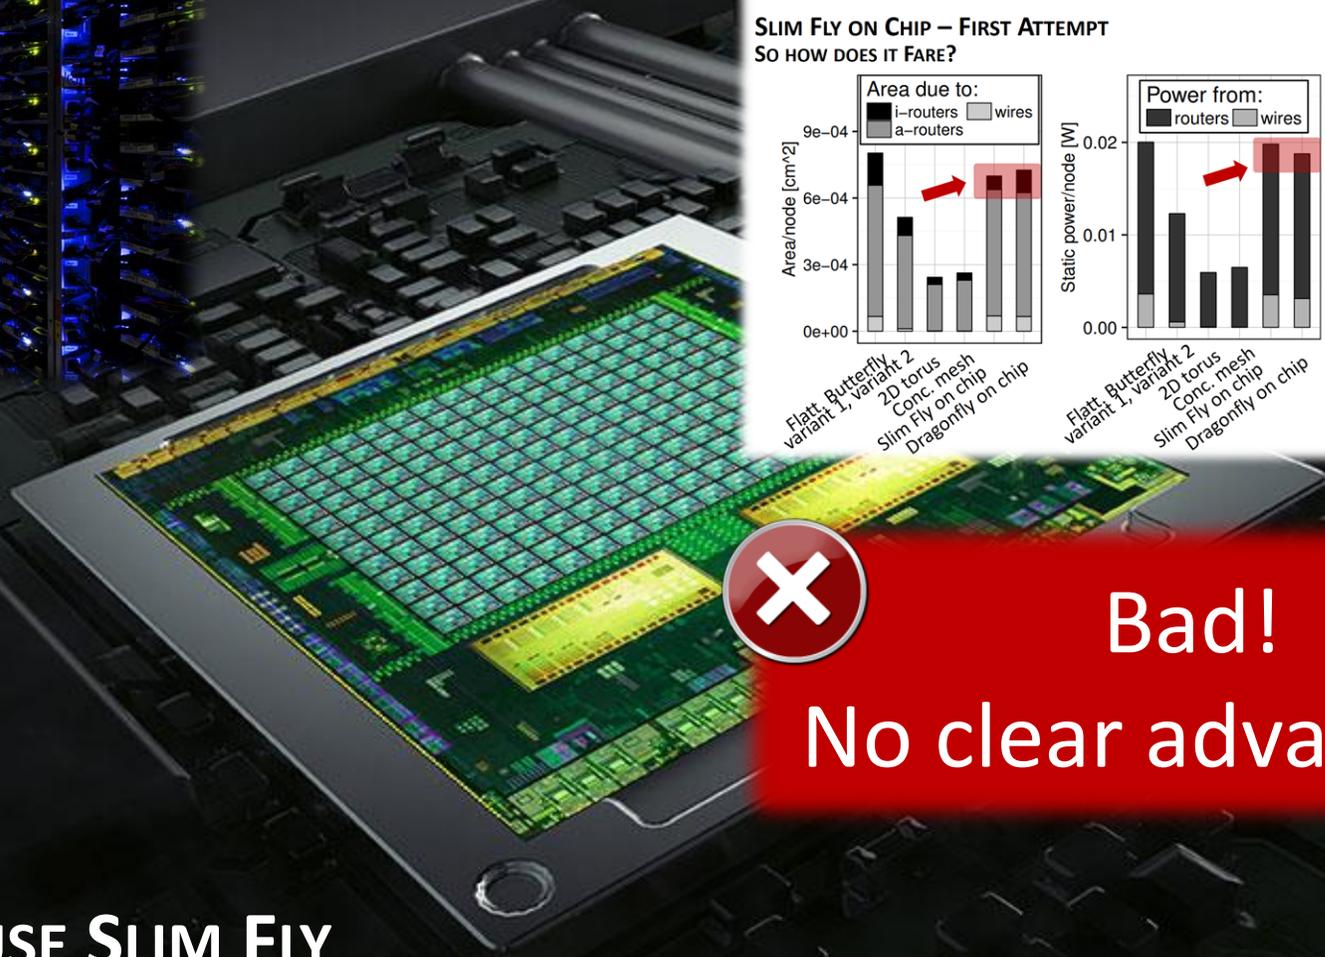


**Slim Fly [1]** ensures the lowest radix (port count) for a given node count and for a fixed diameter (we use two)... Sounds ideal for an on-chip setting?





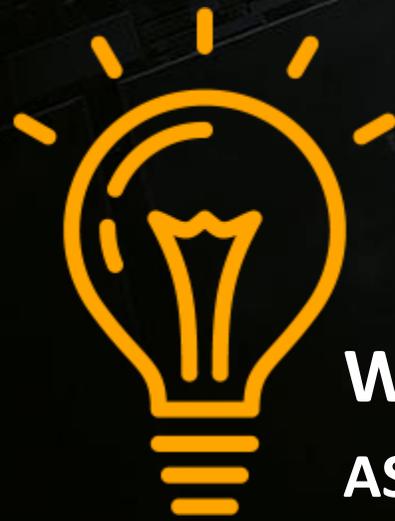
**WHY NOT JUST USE SLIM FLY  
AS AN ON-CHIP NETWORK?**



Bad!  
No clear advantages from a topology that is close-to-optimal in the radix-size-diameter tradeoff

Why?

Bad!  
No clear advantages



WHY NOT JUST USE SLIM FLY AS AN ON-CHIP NETWORK?

# SOLUTION: SLIM NoC

# SOLUTION: SLIM NoC

## NEW COST AND LAYOUT MODELS

! Minimize the average wire length ( $M$ ):

$$M = \frac{\text{Sum of distances}}{\text{Number of links}}$$

# SOLUTION: SLIM NoC

## NEW COST AND LAYOUT MODELS

! Minimize the average wire length ( $M$ ):

$$M = \frac{\text{Sum of distances}}{\text{Number of links}}$$

! Minimize the total buffer area ( $\Delta$ ):

$$\Delta = \sum_{\text{All router pairs } i, j} \begin{array}{l} \text{If } i, j \text{ are connected,} \\ (\varepsilon_{ij} = 1) \text{ add the size} \\ \text{of a buffer from } i \text{ to } j \end{array}$$

# SOLUTION: SLIM NoC

## NEW COST AND LAYOUT MODELS

! Minimize the average wire length ( $M$ ):

$$M = \frac{\text{Sum of distances}}{\text{Number of links}}$$

! Minimize the total buffer area ( $\Delta$ ):

$$\Delta = \sum_{\text{All router pairs } i, j} \begin{cases} \text{If } i, j \text{ are connected,} \\ (\varepsilon_{ij} = 1) \text{ add the size} \\ \text{of a buffer from } i \text{ to } j \end{cases}$$

## NON-PRIME FINITE FIELDS

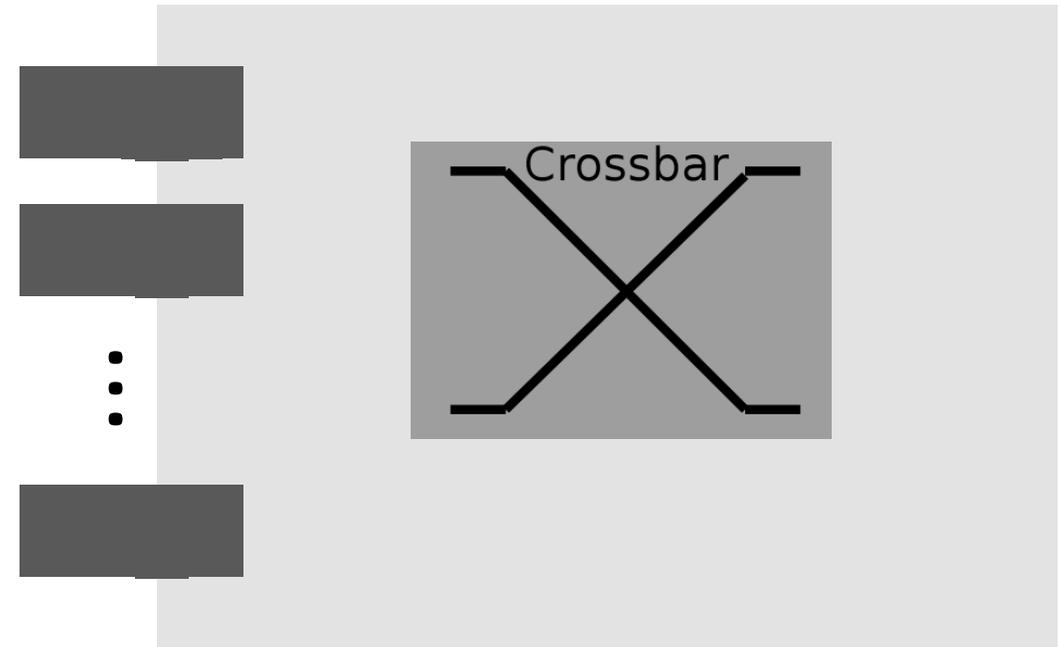
$$\mathcal{F}_q = \mathbb{Z}/q\mathbb{Z} = \{x_0, x_1, \dots, x_{q-1}\}$$

+	0 1 2 u v w x y z	×	0 1 2 u v w x y z	elem	-elem
0	0 1 2 u v w x y z	0	0 0 0 0 0 0 0 0 0	0	0
1	1 2 0 v w u y z x	1	0 1 2 u v w x y z	1	2
2	2 0 1 w u v z x y	2	0 2 1 x z y u w v	2	0
u	u v w x y z 0 1 2	u	0 u x 2 w z 1 v y	u	x
v	v w u y z x 1 2 0	v	0 v z w x 1 y 2 u	v	z
w	w u v z x y 2 0 1	w	0 w y z 1 u v x 2	w	y
x	x y z 0 1 2 u v w	x	0 x u 1 y v 2 z w	x	u
y	y z x 1 2 0 v w u	y	0 y w v 2 x z u 1	y	w
z	z x y 2 0 1 w u v	z	0 z v y u 2 w 1 x	z	v

**Addition**

**Multiplication Inverse**

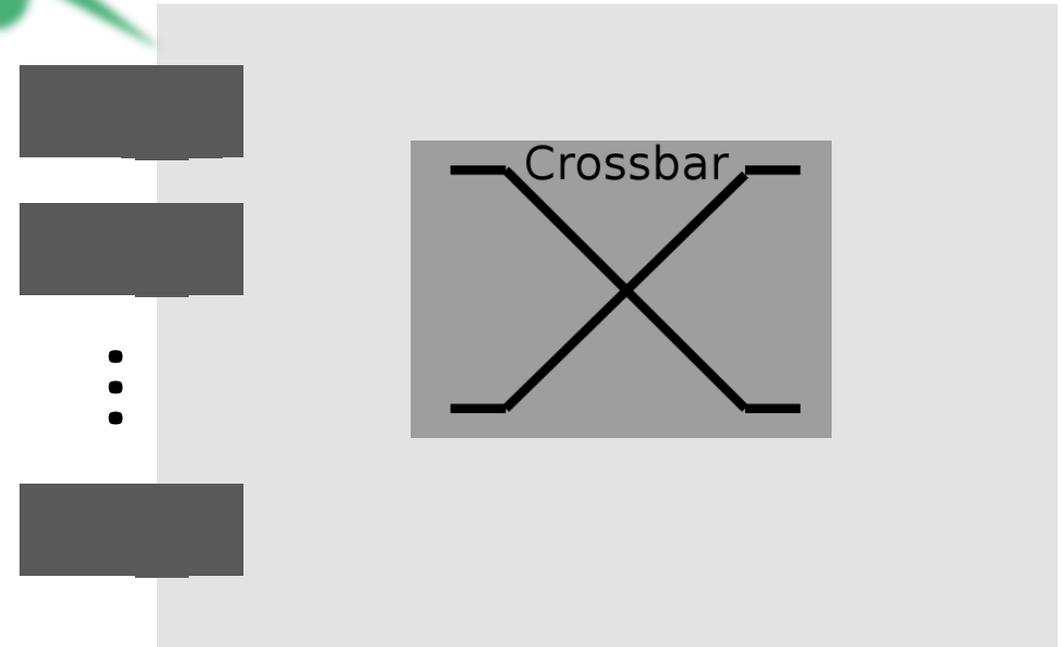
# SLIM NoC ROUTER MICROARCHITECTURE



# SLIM NoC ROUTER MICROARCHITECTURE

**ENHANCEMENT 1:**  
ELASTIC BUFFER LINKS [1]  
+ ELASTISTORE [2]

No deadlock



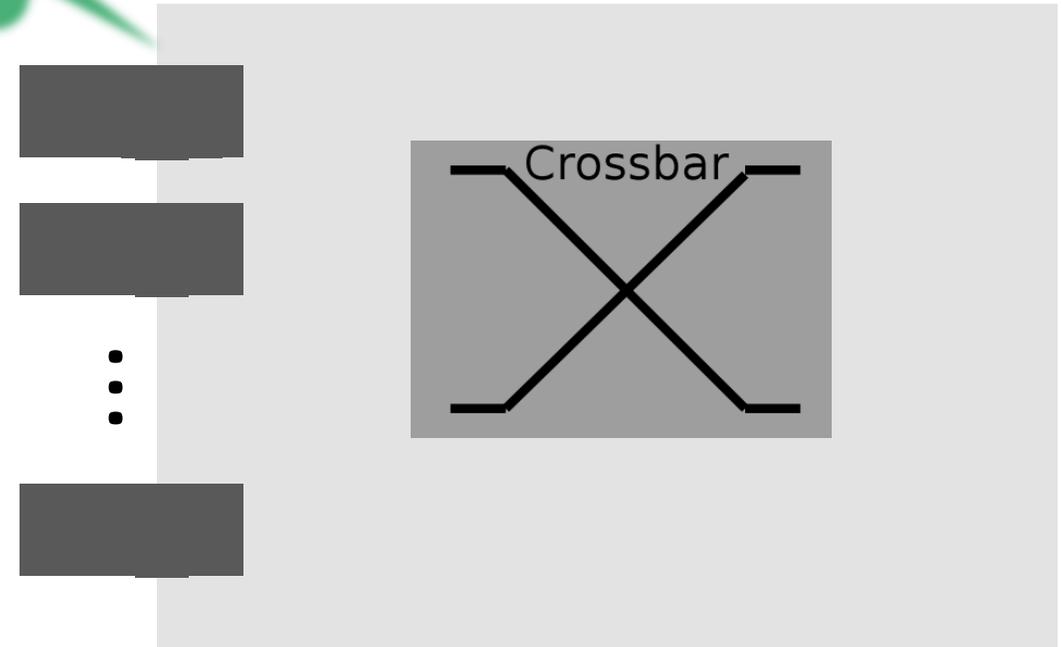
[1] G. Michelogiannakis et al. Elastic-Buffer Flow Control for On-Chip Networks. HPCA'09.

[2] I. Seitanidis et al. ElastiStore: An Elastic Buffer Architecture for Network-on-Chip Routers. DATE'14.

# SLIM NoC ROUTER MICROARCHITECTURE

**ENHANCEMENT 1:**  
ELASTIC BUFFER LINKS [1]  
+ ELASTISTORE [2]

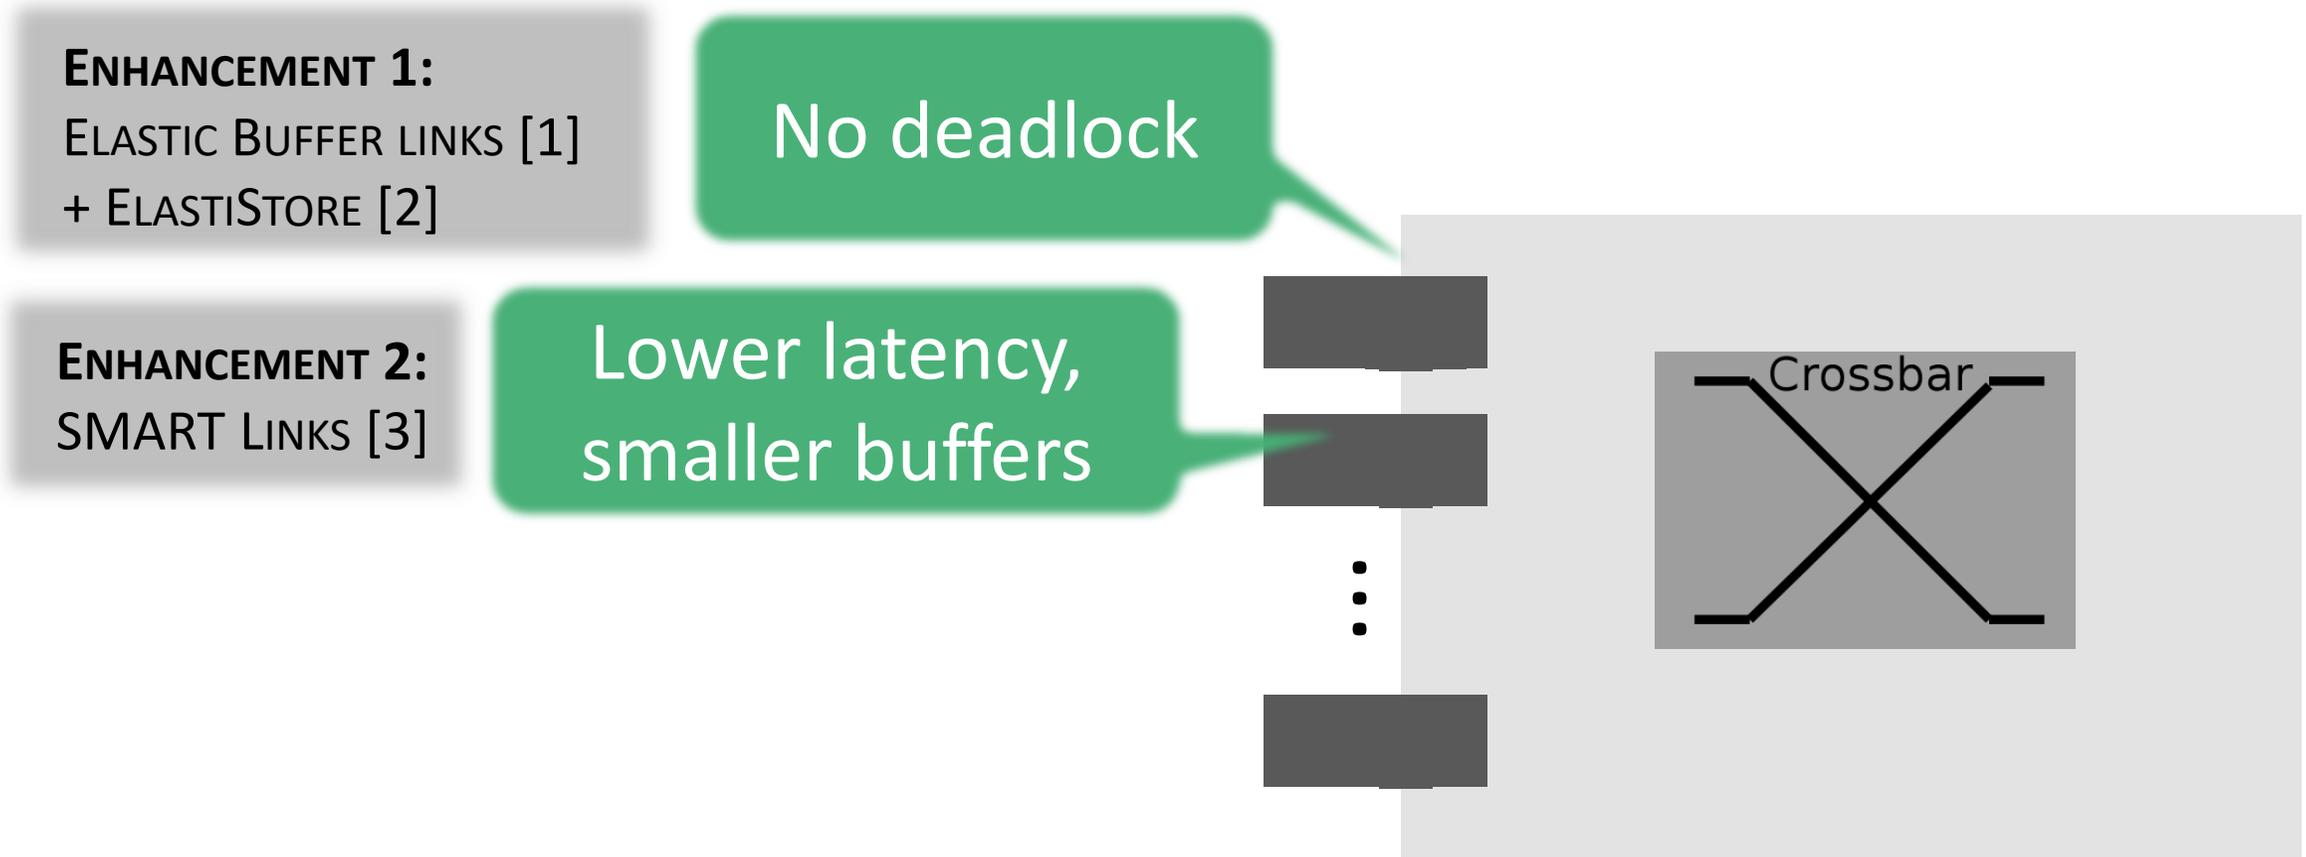
No deadlock



[1] G. Michelogiannakis et al. Elastic-Buffer Flow Control for On-Chip Networks. HPCA'09.

[2] I. Seitanidis et al. ElastiStore: An Elastic Buffer Architecture for Network-on-Chip Routers. DATE'14.

# SLIM NoC ROUTER MICROARCHITECTURE



[1] G. Michelogiannakis et al. Elastic-Buffer Flow Control for On-Chip Networks. HPCA'09.

[2] I. Seitanidis et al. ElastiStore: An Elastic Buffer Architecture for Network-on-Chip Routers. DATE'14.

[3] C.-H. O. Chen et al. SMART: A Single-Cycle Reconfigurable NoC for SoC Applications. DATE'13.

# SLIM NoC ROUTER MICROARCHITECTURE

## ENHANCEMENT 1:

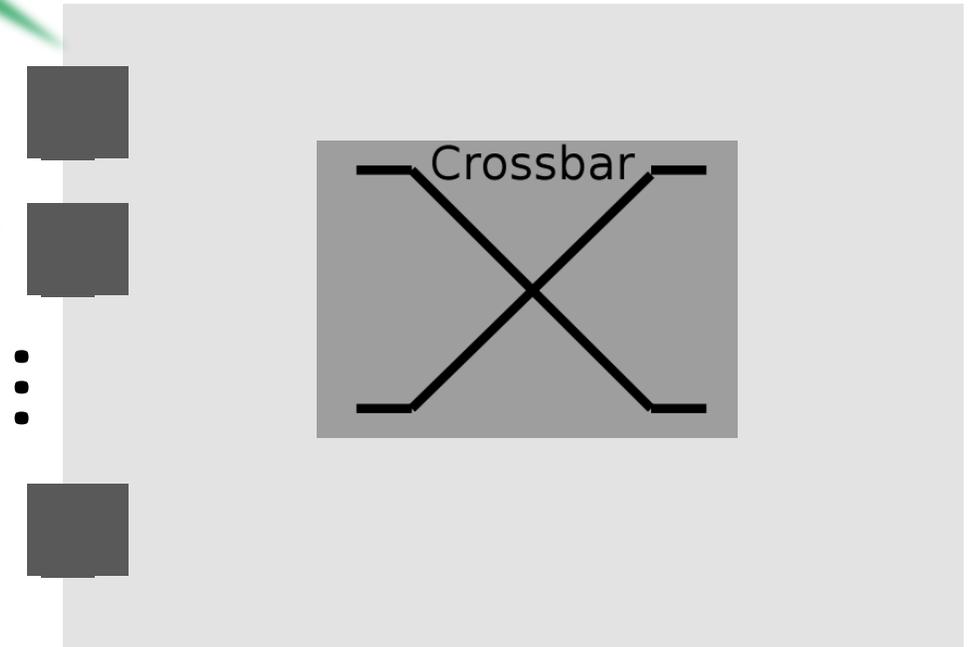
ELASTIC BUFFER LINKS [1]  
+ ELASTISTORE [2]

No deadlock

## ENHANCEMENT 2:

SMART LINKS [3]

Lower latency,  
smaller buffers

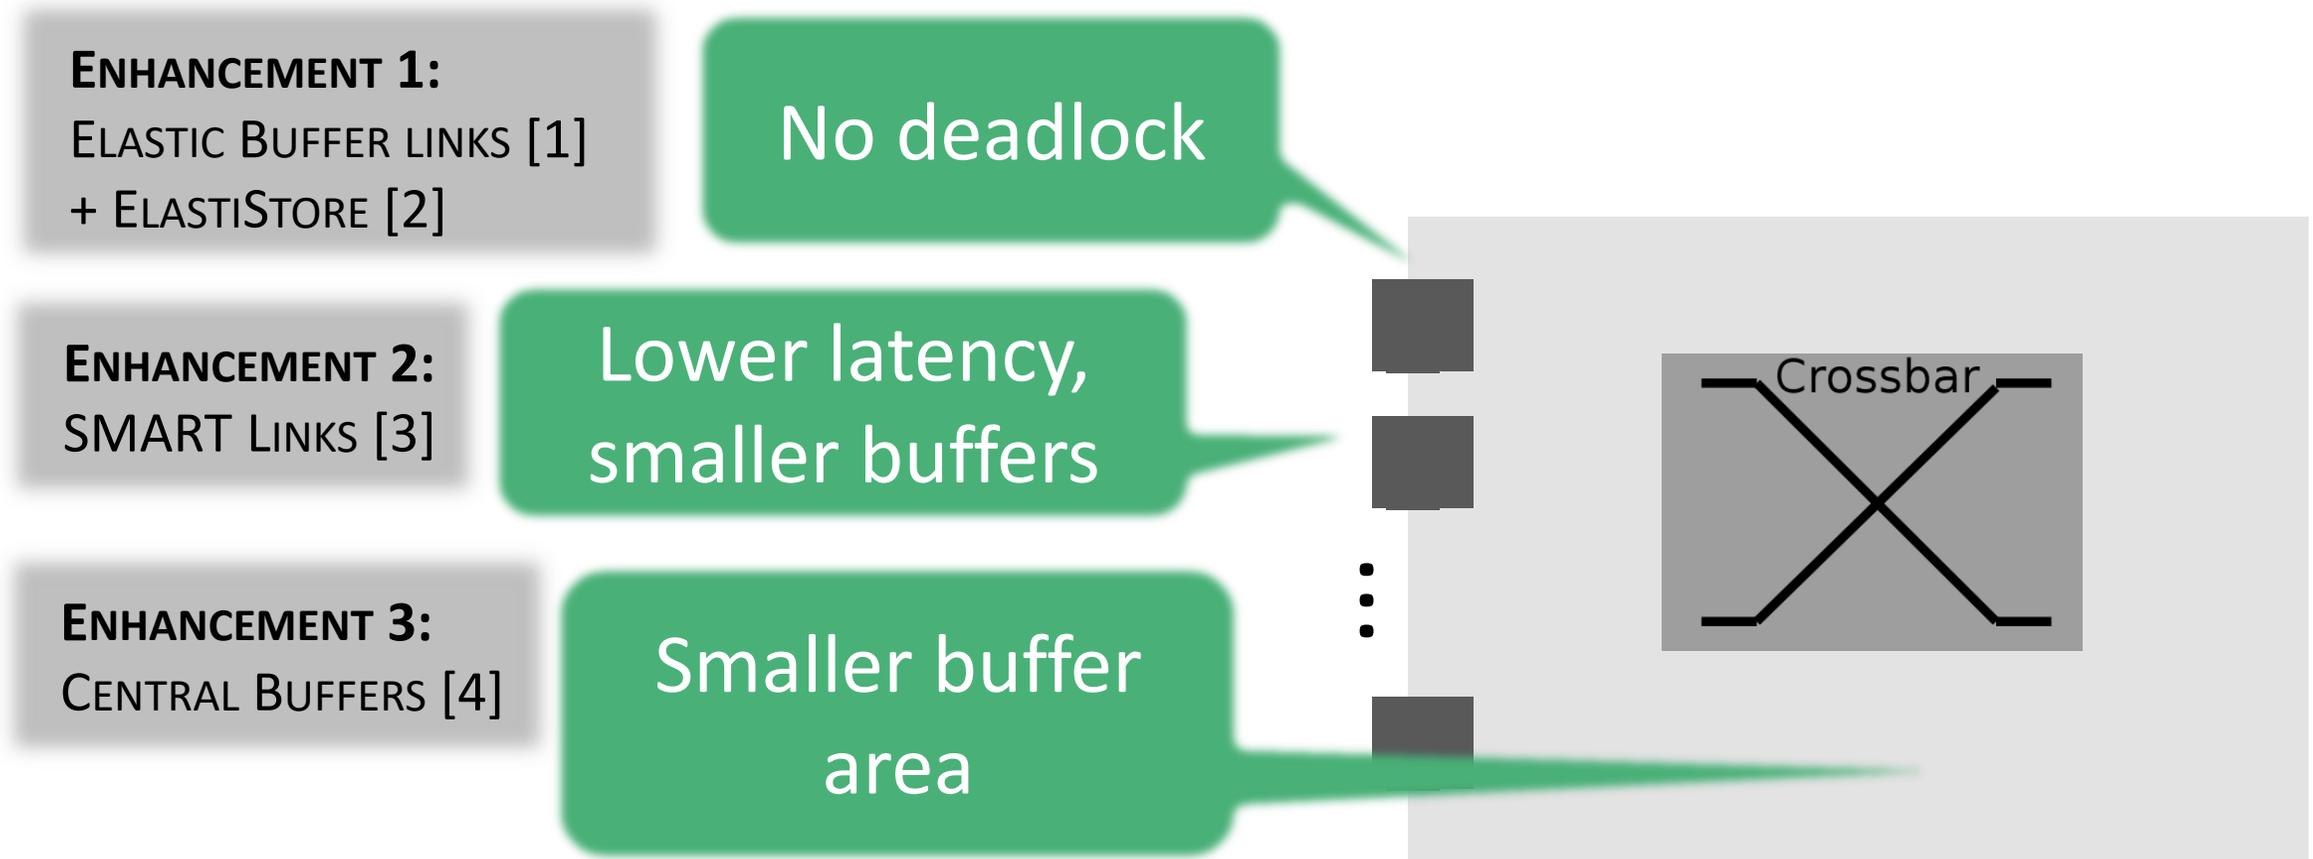


[1] G. Michelogiannakis et al. Elastic-Buffer Flow Control for On-Chip Networks. HPCA'09.

[2] I. Seitanidis et al. ElastiStore: An Elastic Buffer Architecture for Network-on-Chip Routers. DATE'14.

[3] C.-H. O. Chen et al. SMART: A Single-Cycle Reconfigurable NoC for SoC Applications. DATE'13.

# SLIM NoC ROUTER MICROARCHITECTURE



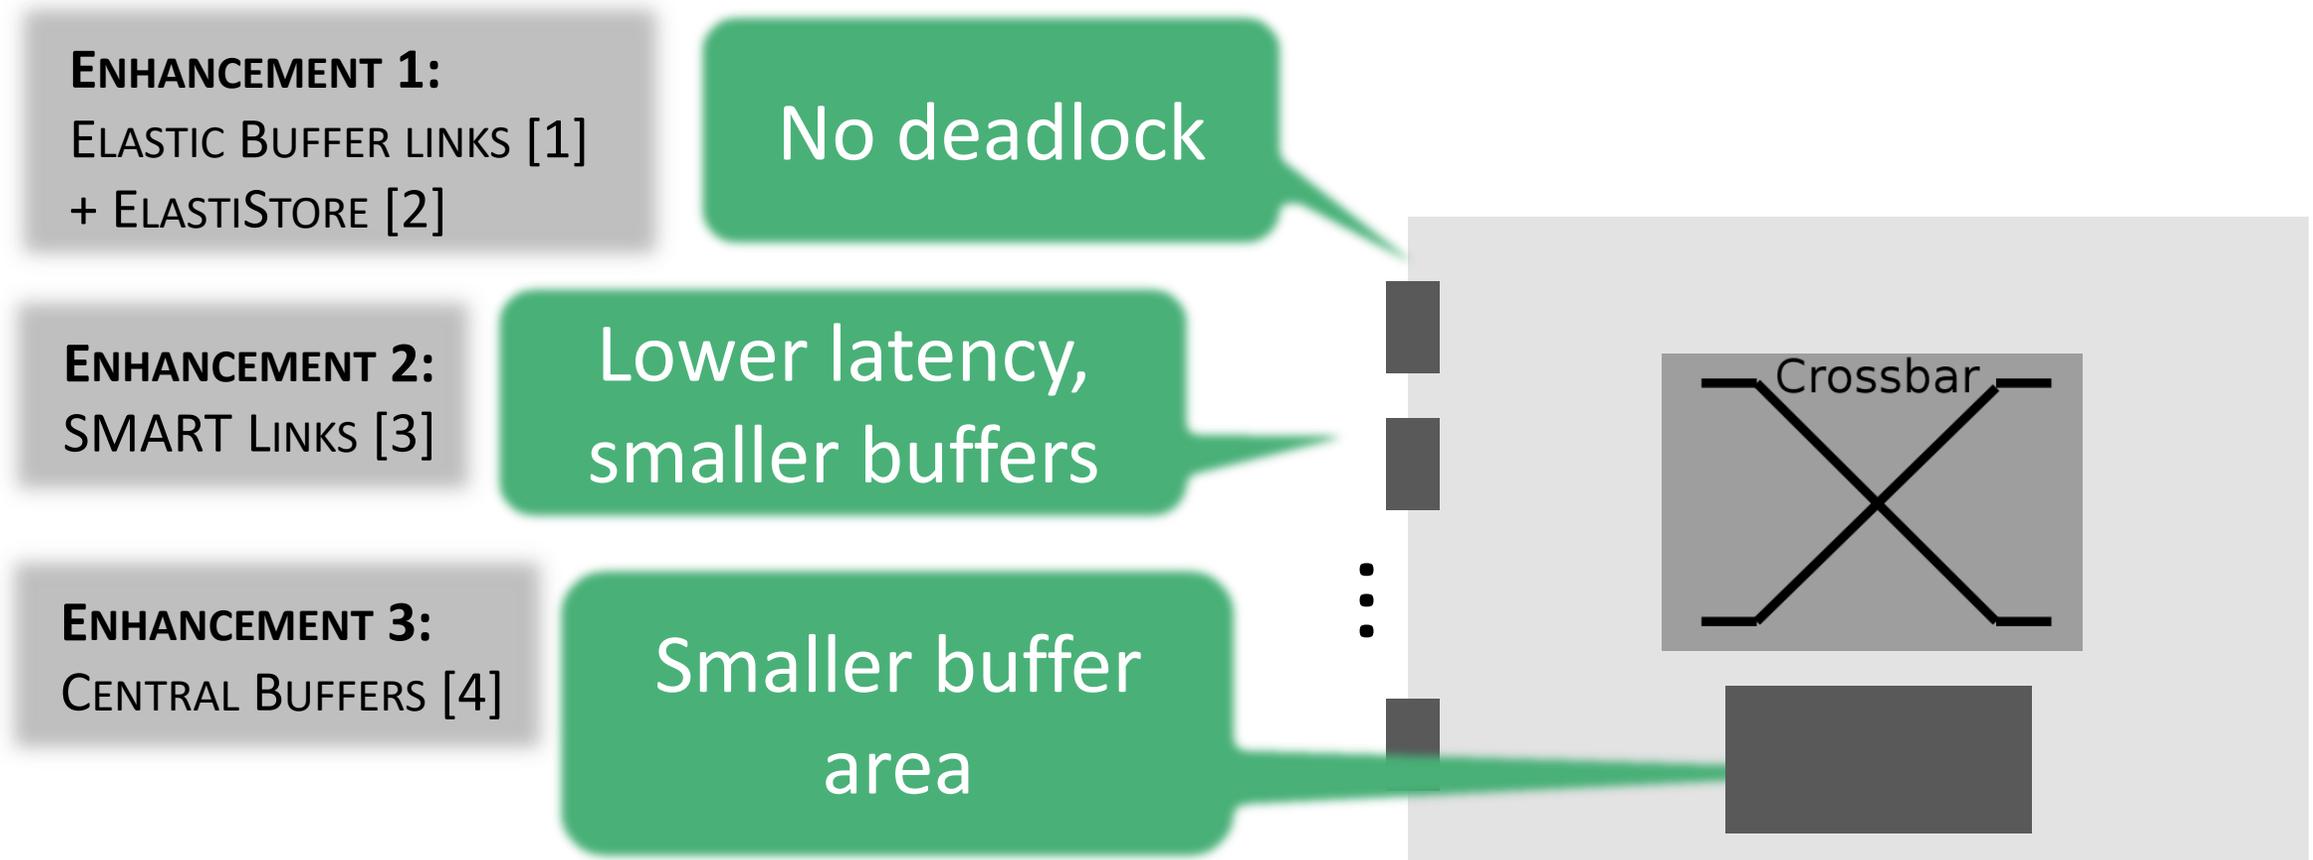
[1] G. Michelogiannakis et al. Elastic-Buffer Flow Control for On-Chip Networks. HPCA'09.

[2] I. Seitanidis et al. ElastiStore: An Elastic Buffer Architecture for Network-on-Chip Routers. DATE'14.

[3] C.-H. O. Chen et al. SMART: A Single-Cycle Reconfigurable NoC for SoC Applications. DATE'13.

[4] S. Hassan and S. Yalamanchili. Centralized Buffer Router: A Low Latency, Low Power Router for High Radix NoCs. NOCS'13.

# SLIM NoC ROUTER MICROARCHITECTURE



[1] G. Michelogiannakis et al. Elastic-Buffer Flow Control for On-Chip Networks. HPCA'09.

[2] I. Seitanidis et al. ElastiStore: An Elastic Buffer Architecture for Network-on-Chip Routers. DATE'14.

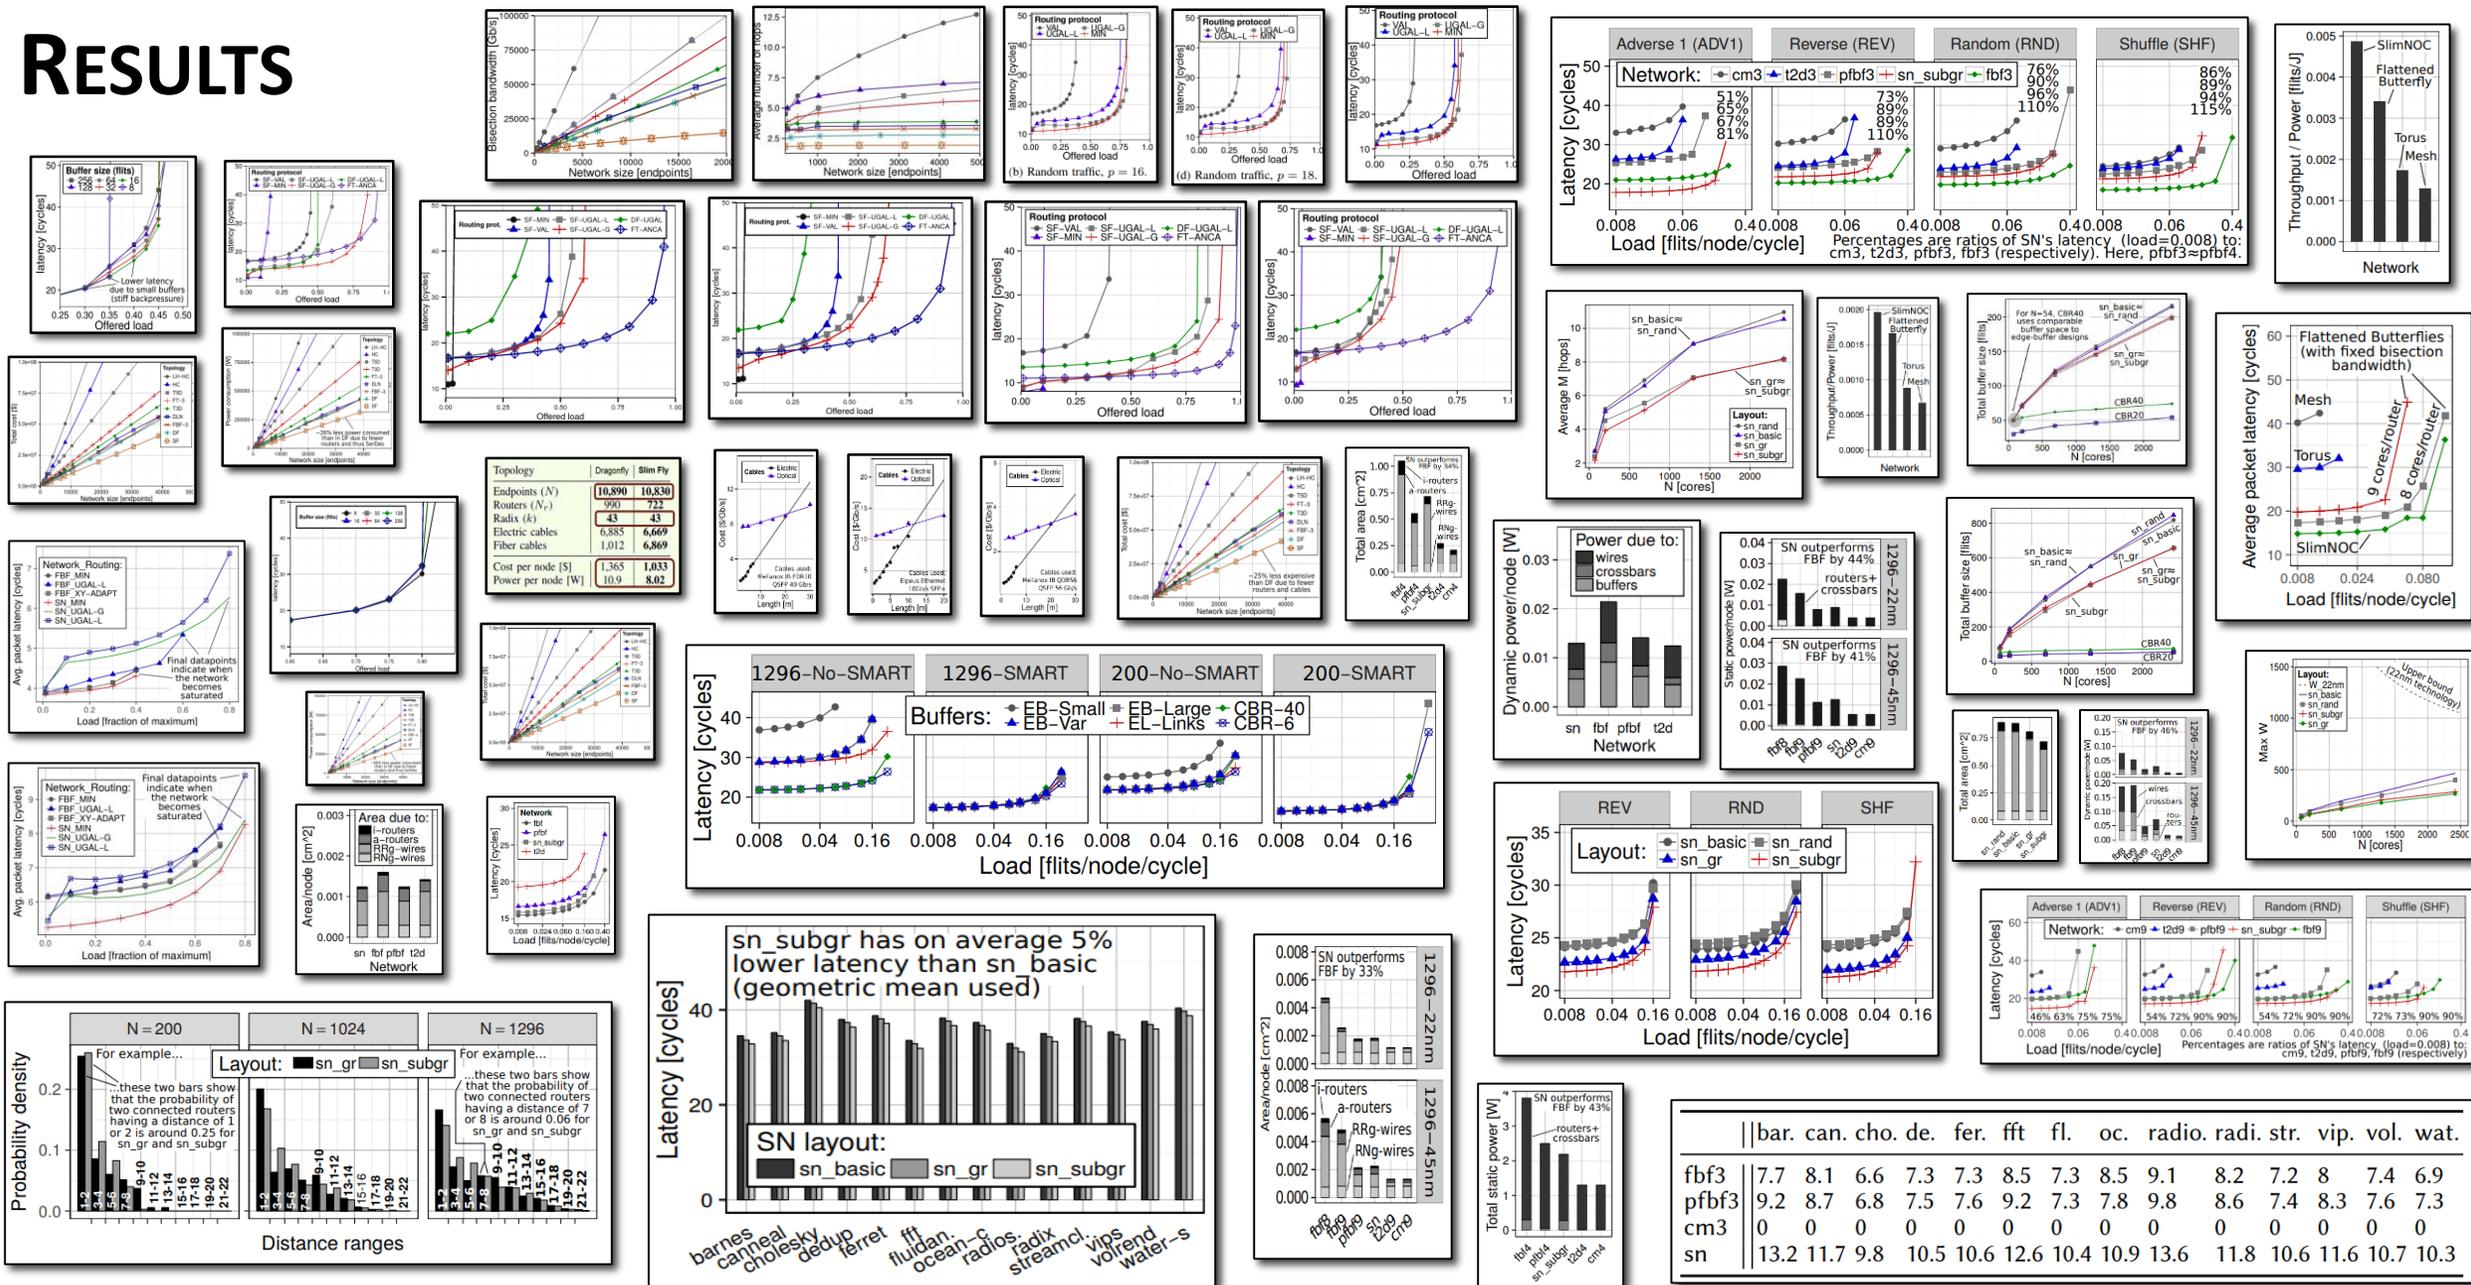
[3] C.-H. O. Chen et al. SMART: A Single-Cycle Reconfigurable NoC for SoC Applications. DATE'13.

[4] S. Hassan and S. Yalamanchili. Centralized Buffer Router: A Low Latency, Low Power Router for High Radix NoCs. NOCS'13.



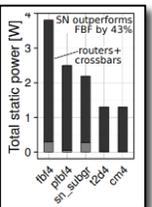
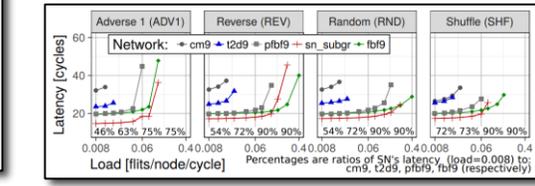
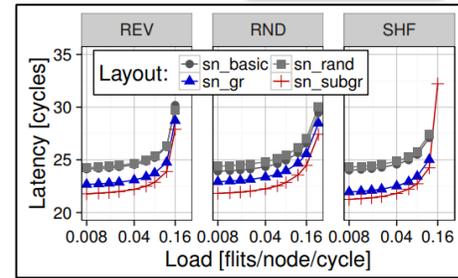
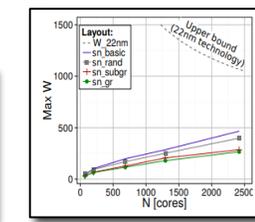
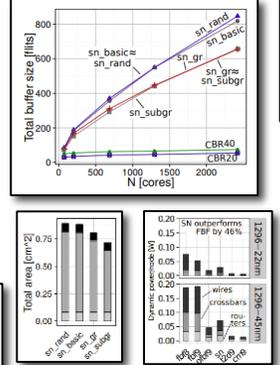
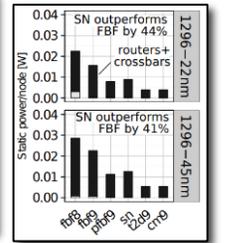
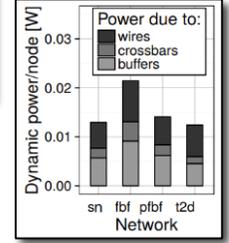
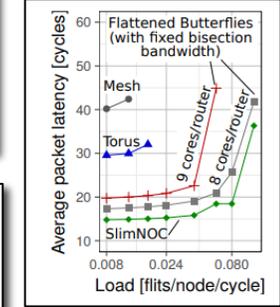
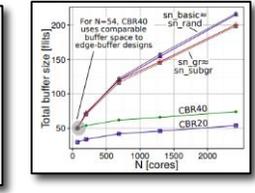
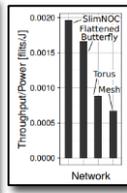
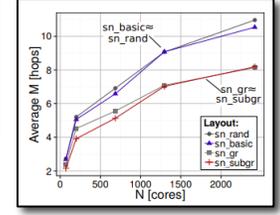
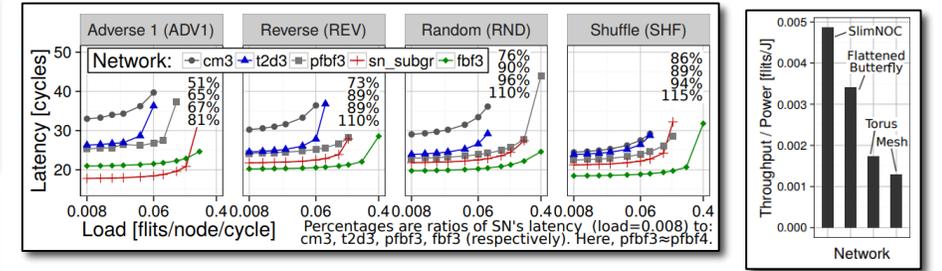
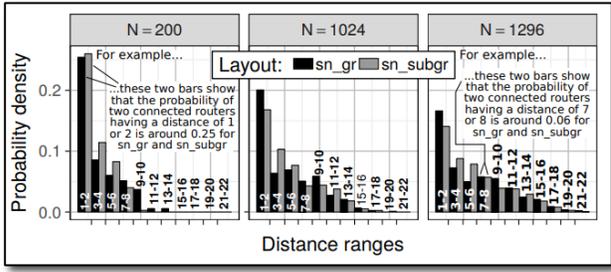
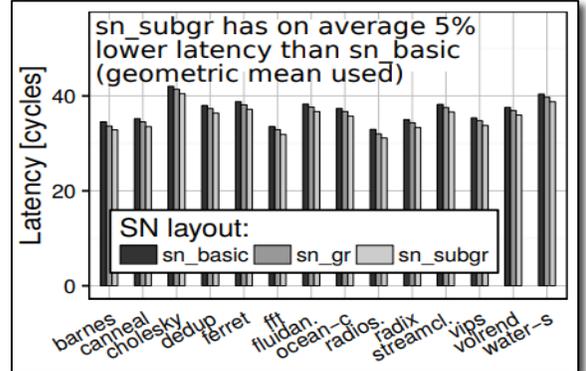
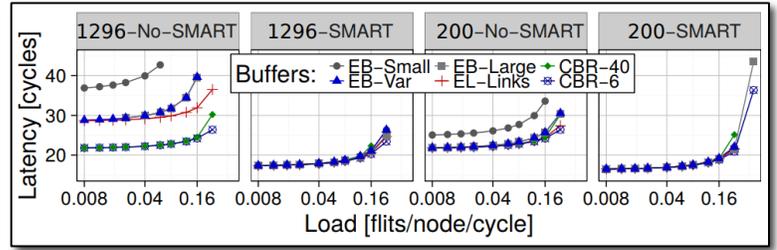
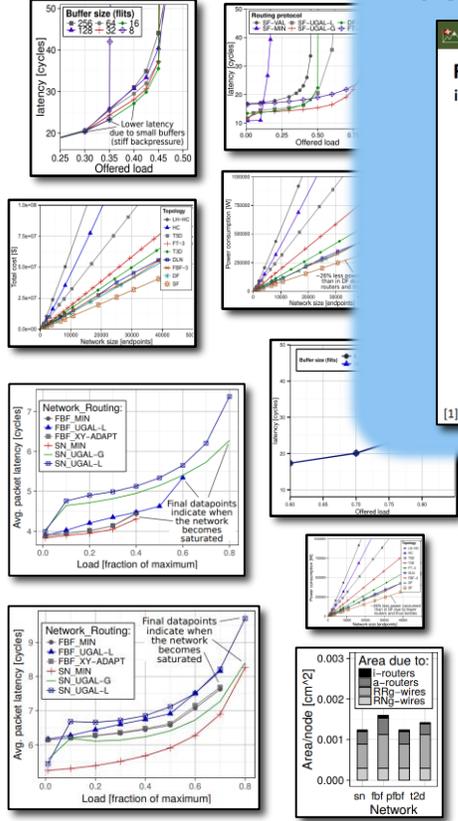
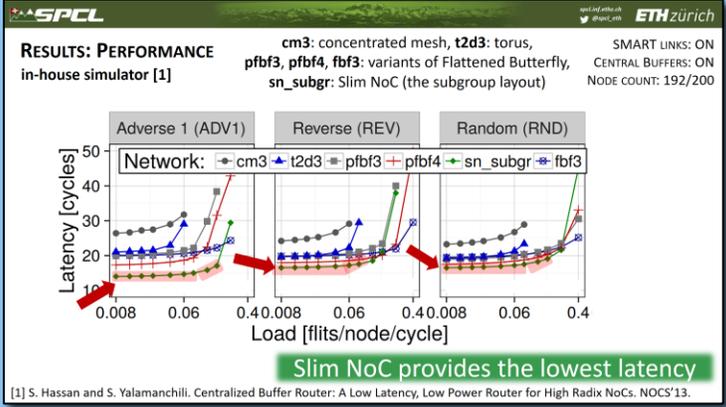


# RESULTS



# RESULTS

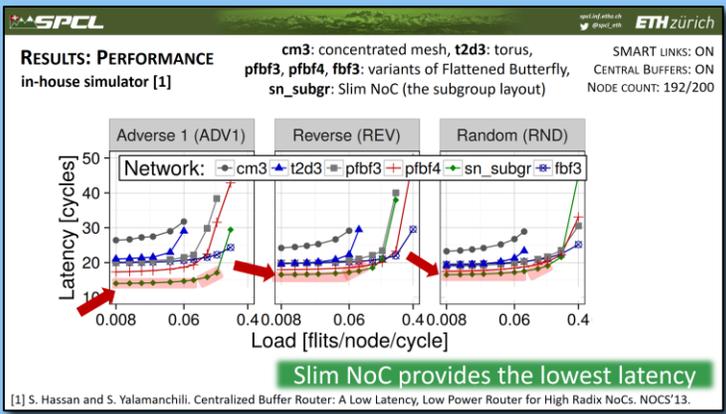
## A LOW-LATENCY TOPOLOGY



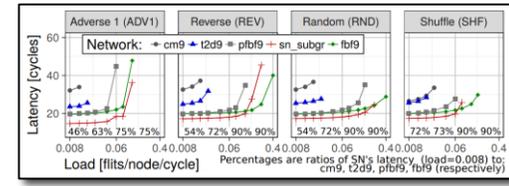
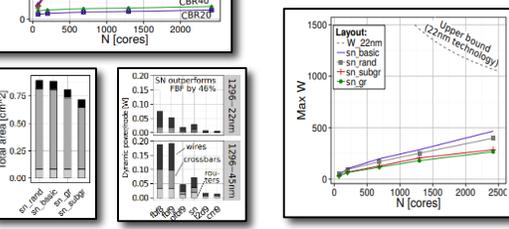
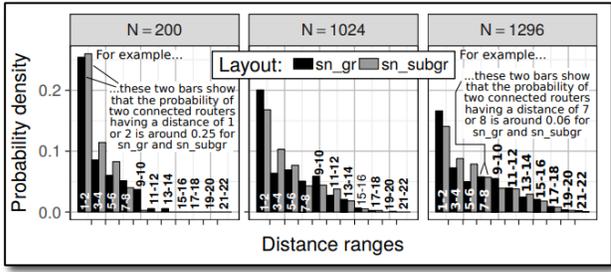
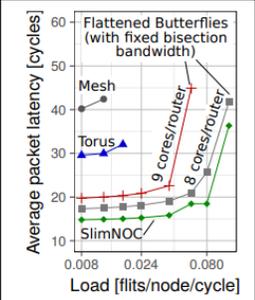
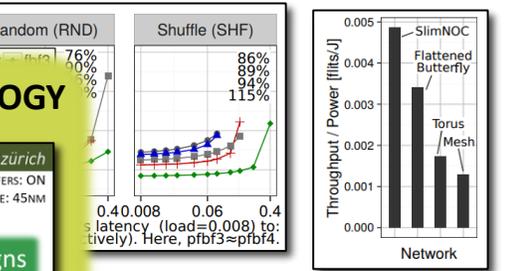
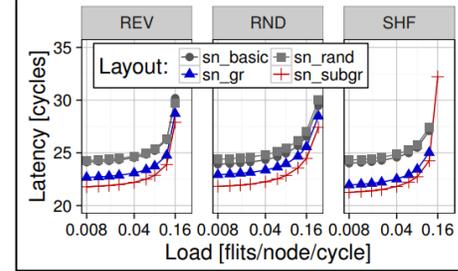
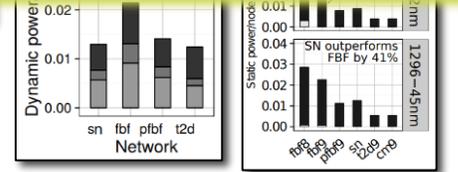
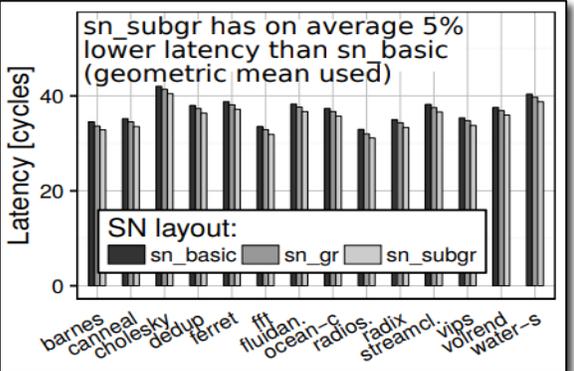
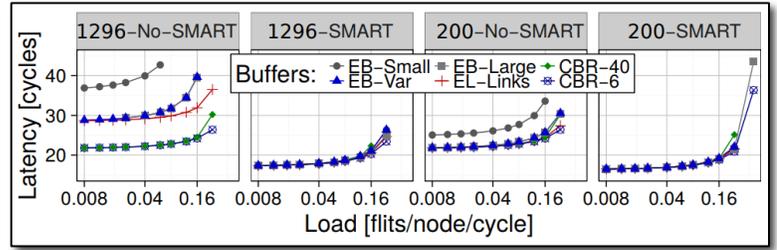
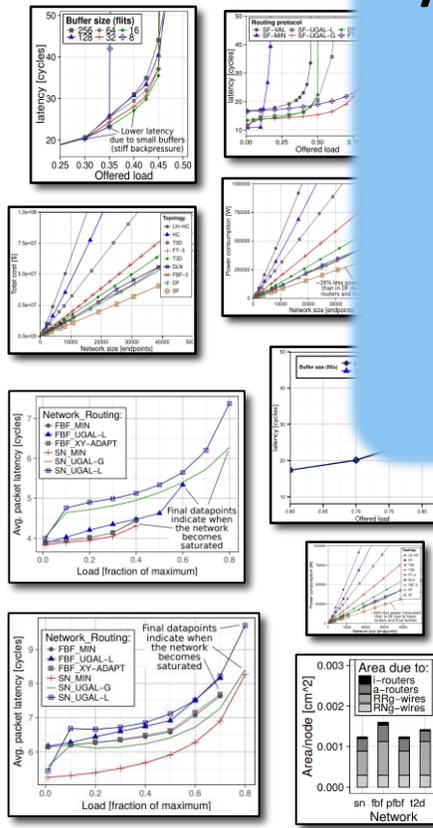
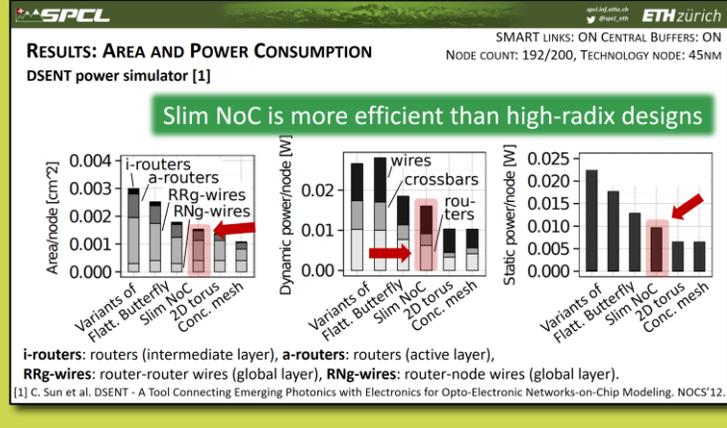
	bar.	can.	cho.	de.	fer.	fft	fl.	oc.	radio.	radi.	str.	vip.	vol.	wat.
fbf3	7.7	8.1	6.6	7.3	7.3	8.5	7.3	8.5	9.1	8.2	7.2	8	7.4	6.9
pfbf3	9.2	8.7	6.8	7.5	7.6	9.2	7.3	7.8	9.8	8.6	7.4	8.3	7.6	7.3
cm3	0	0	0	0	0	0	0	0	0	0	0	0	0	0
sn	13.2	11.7	9.8	10.5	10.6	12.6	10.4	10.9	13.6	11.8	10.6	11.6	10.7	10.3

# RESULTS

## A LOW-LATENCY TOPOLOGY



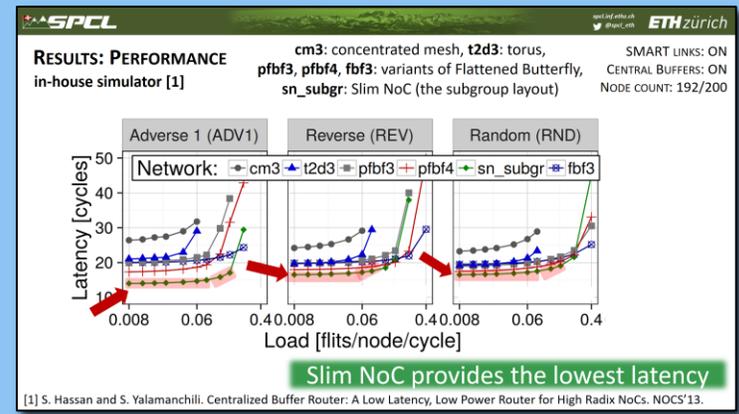
## AN AREA- AND ENERGY-EFFICIENT TOPOLOGY



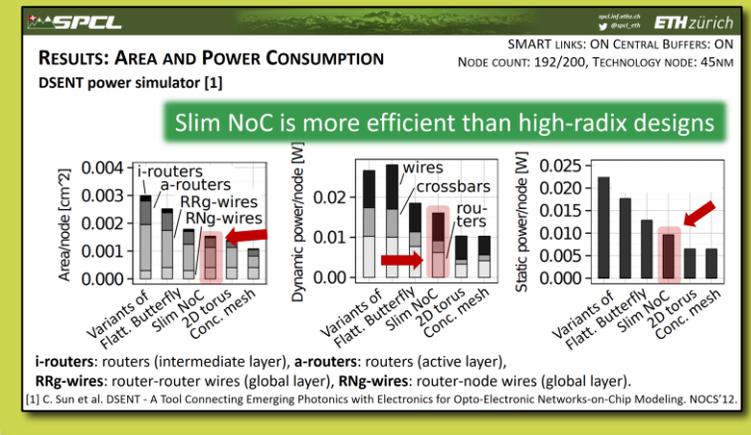
	bar.	can.	cho.	de.	fer.	fft	fl.	oc.	radio.	radi.	str.	vip.	vol.	wat.
fbf3	7.7	8.1	6.6	7.3	7.3	8.5	7.3	8.5	9.1	8.2	7.2	8	7.4	6.9
pfbf3	9.2	8.7	6.8	7.5	7.6	9.2	7.3	7.8	9.8	8.6	7.4	8.3	7.6	7.3
cm3	0	0	0	0	0	0	0	0	0	0	0	0	0	0
sn	13.2	11.7	9.8	10.5	10.6	12.6	10.4	10.9	13.6	11.8	10.6	11.6	10.7	10.3

# RESULTS

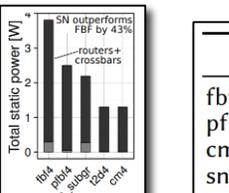
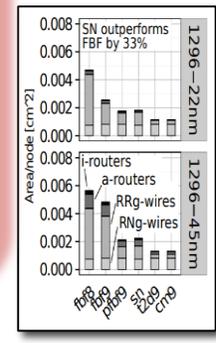
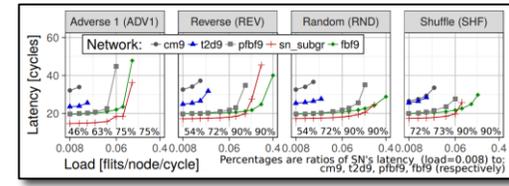
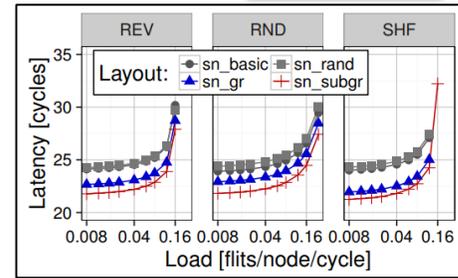
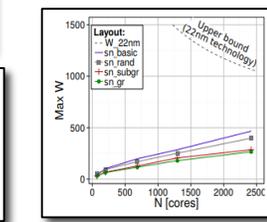
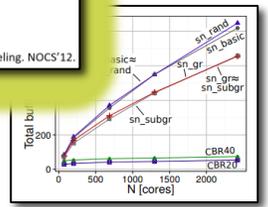
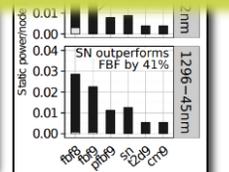
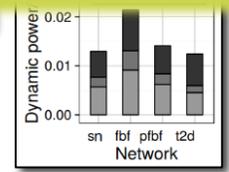
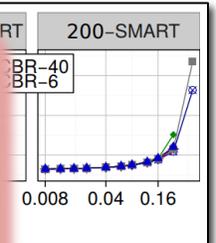
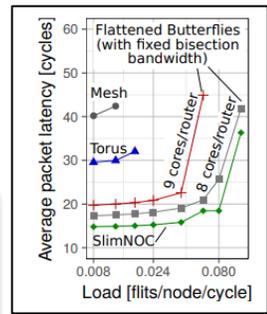
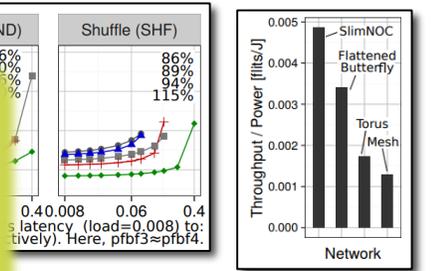
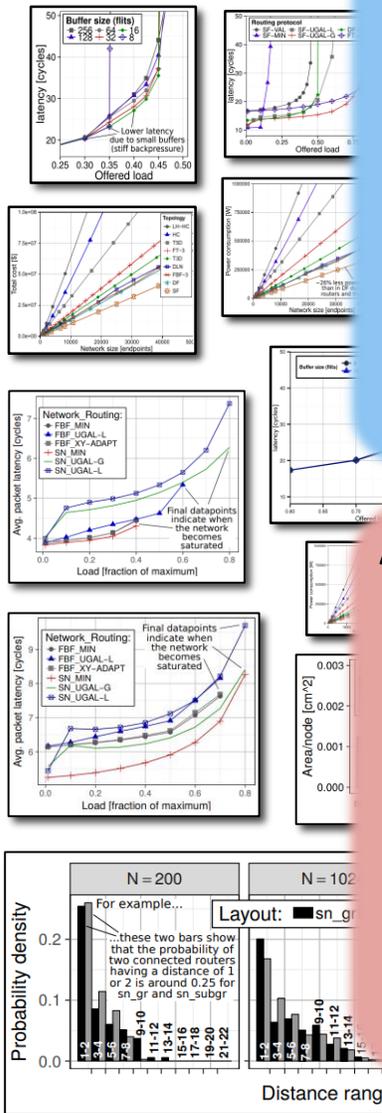
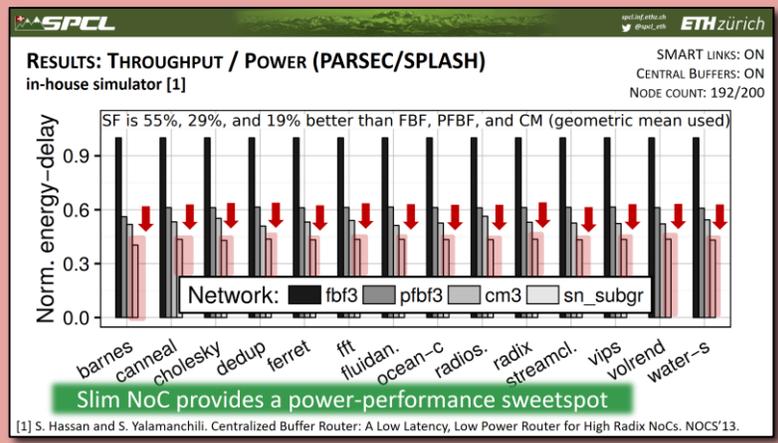
## A LOW-LATENCY TOPOLOGY



## AN AREA- AND ENERGY-EFFICIENT TOPOLOGY



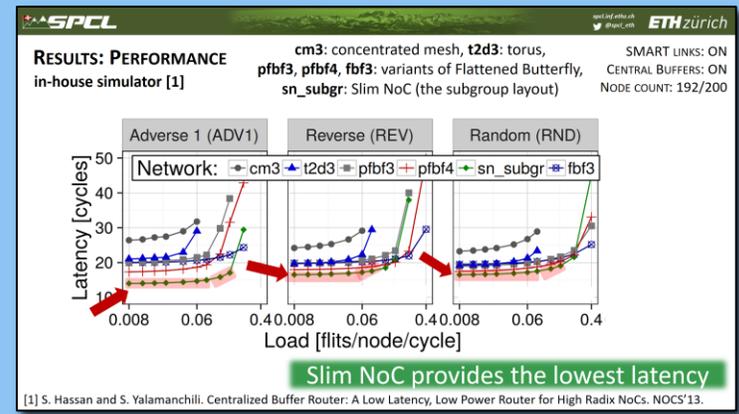
## A POWER-PERFORMANCE-SWEETSPOT TOPOLOGY



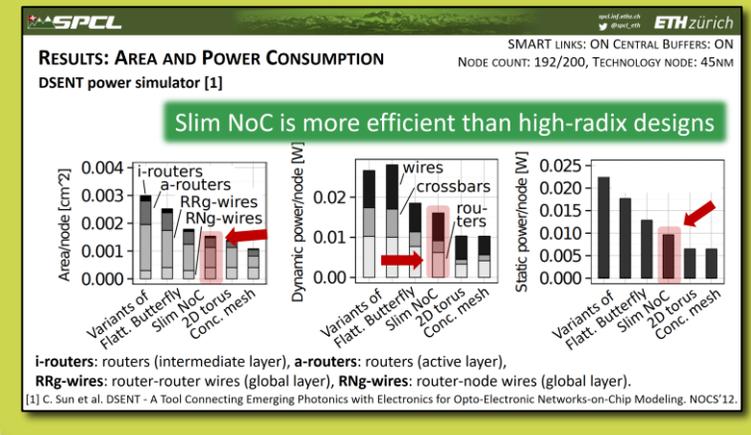
	bar.	can.	cho.	de.	fer.	fft	fl.	oc.	radio.	radi.	str.	vip.	vol.	wat.
fbf3	7.7	8.1	6.6	7.3	7.3	8.5	7.3	8.5	9.1	8.2	7.2	8	7.4	6.9
pfbf3	9.2	8.7	6.8	7.5	7.6	9.2	7.3	7.8	9.8	8.6	7.4	8.3	7.6	7.3
cm3	0	0	0	0	0	0	0	0	0	0	0	0	0	0
sn	13.2	11.7	9.8	10.5	10.6	12.6	10.4	10.9	13.6	11.8	10.6	11.6	10.7	10.3

# RESULTS

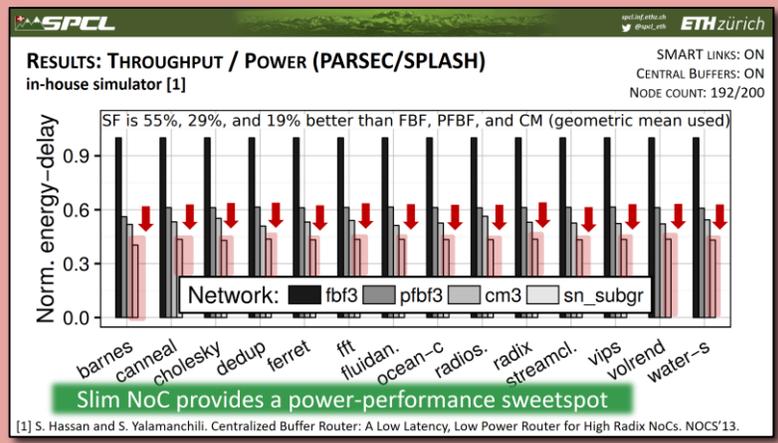
## A LOW-LATENCY TOPOLOGY



## AN AREA- AND ENERGY-EFFICIENT TOPOLOGY

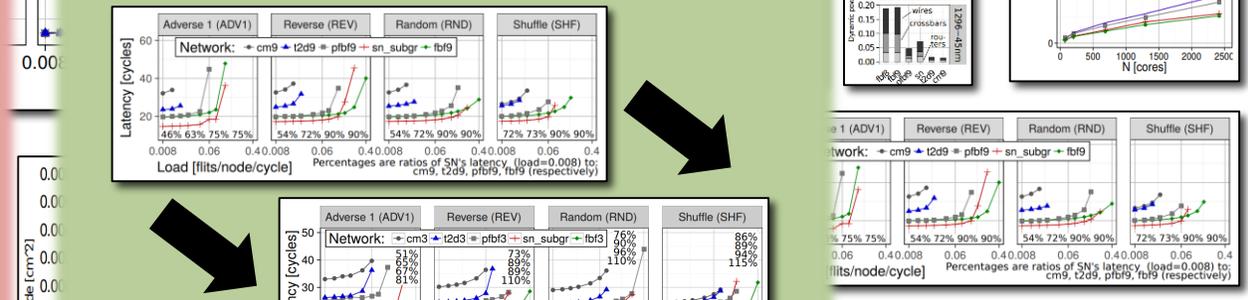


## A POWER-PERFORMANCE-SWEETSPOT TOPOLOGY



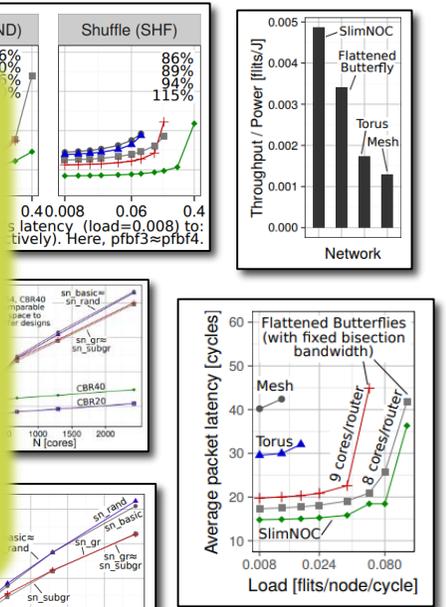
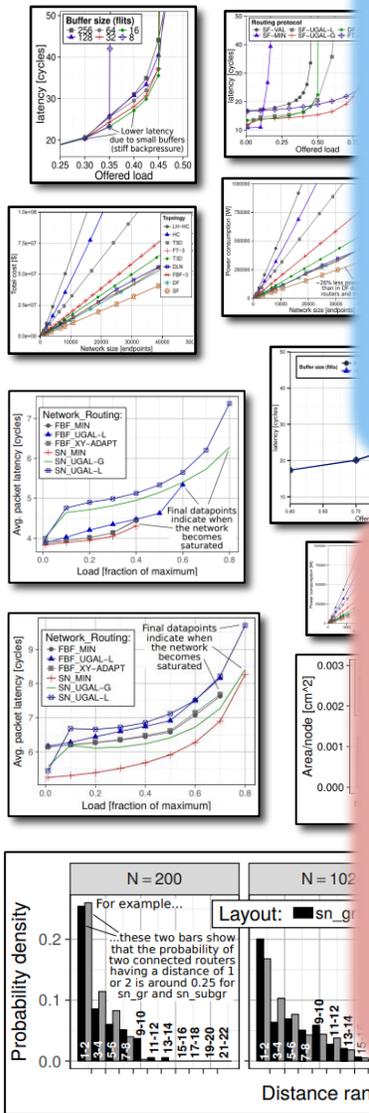
## A HIGHLY-SCALABLE TOPOLOGY

### 200 nodes



### 1296 nodes

Network	sn	cm3	t2d3	pfbf3	pfbf4	fbf3	oc.	radio.	radi.	str.	vip.	vol.	wat.
Area/node [cm <sup>2</sup> ]	13.2	11.7	9.8	10.5	10.6	12.6	0	0	0	0	0	0	0
Dynamic power [W]	8.5	9.1	8.2	7.2	8	7.4	0	0	0	0	0	0	0
Static power [W]	7.3	7.8	9.8	8.6	7.4	8.3	7.6	7.3	0	0	0	0	0
Latency [cycles]	10.5	10.6	10.6	10.5	10.6	10.6	11.6	10.7	10.3	10.3	10.3	10.3	10.3



# Slim NoC: A Low-Diameter On-Chip Network Topology for High Energy Efficiency and Scalability

MACIEJ BESTA, SYED MINHAJ HASSAN, SUDHAKAR YALAMANCHILI,  
RACHATA AUSAVARUNGNIRUN, ONUR MUTLU, TORSTEN HOEFLER

The word "SAFARI" is written in a bold, orange, sans-serif font on a white rectangular background.The Georgia Tech logo, consisting of the words "Georgia Tech" in a yellow, sans-serif font on a dark blue background, with a small yellow tower icon to the right.The SPCL logo, featuring the letters "SPCL" in a large, white, stylized font with a green outline, set against a green background with white mountain peaks.