Low-Power Reconfigurable Network Architecture for On-Chip Photonic Interconnects

I. Artundo, *W. Heirman, C. Debaes, M. Loperena,*
J. Van Campenhout, H. Thienpont

New York, August 27th 2009

Iñigo Artundo, PostDoc Researcher
TONA Dept. – Vrije Universiteit Brussel (VUB)

Vrije Universiteit Brussel

Outline

1. Introduction to optical on-chip networks
2. Network reconfiguration
3. Proposed architecture
4. Full-system simulation results
5. Conclusions

Vrije Universiteit Brussel

Optical NoC

Interconnection network:

- Connects processors/cores, cache and memories
- Vital role in determining the overall performance of a multiprocessor/multicore system

Networks-on-Chip:

- Very low-level interconnection network at the on-chip size
- Latency even more important
- Critical power consumption and associated heat dissipation
- Reduced buffering

Optical NoC:

- Overcomes high frecuency electrical problems
- Very high bandwidth and interconnect density
- Now becoming a possibility via integration through silicon photonics

Optical integration (PIC)

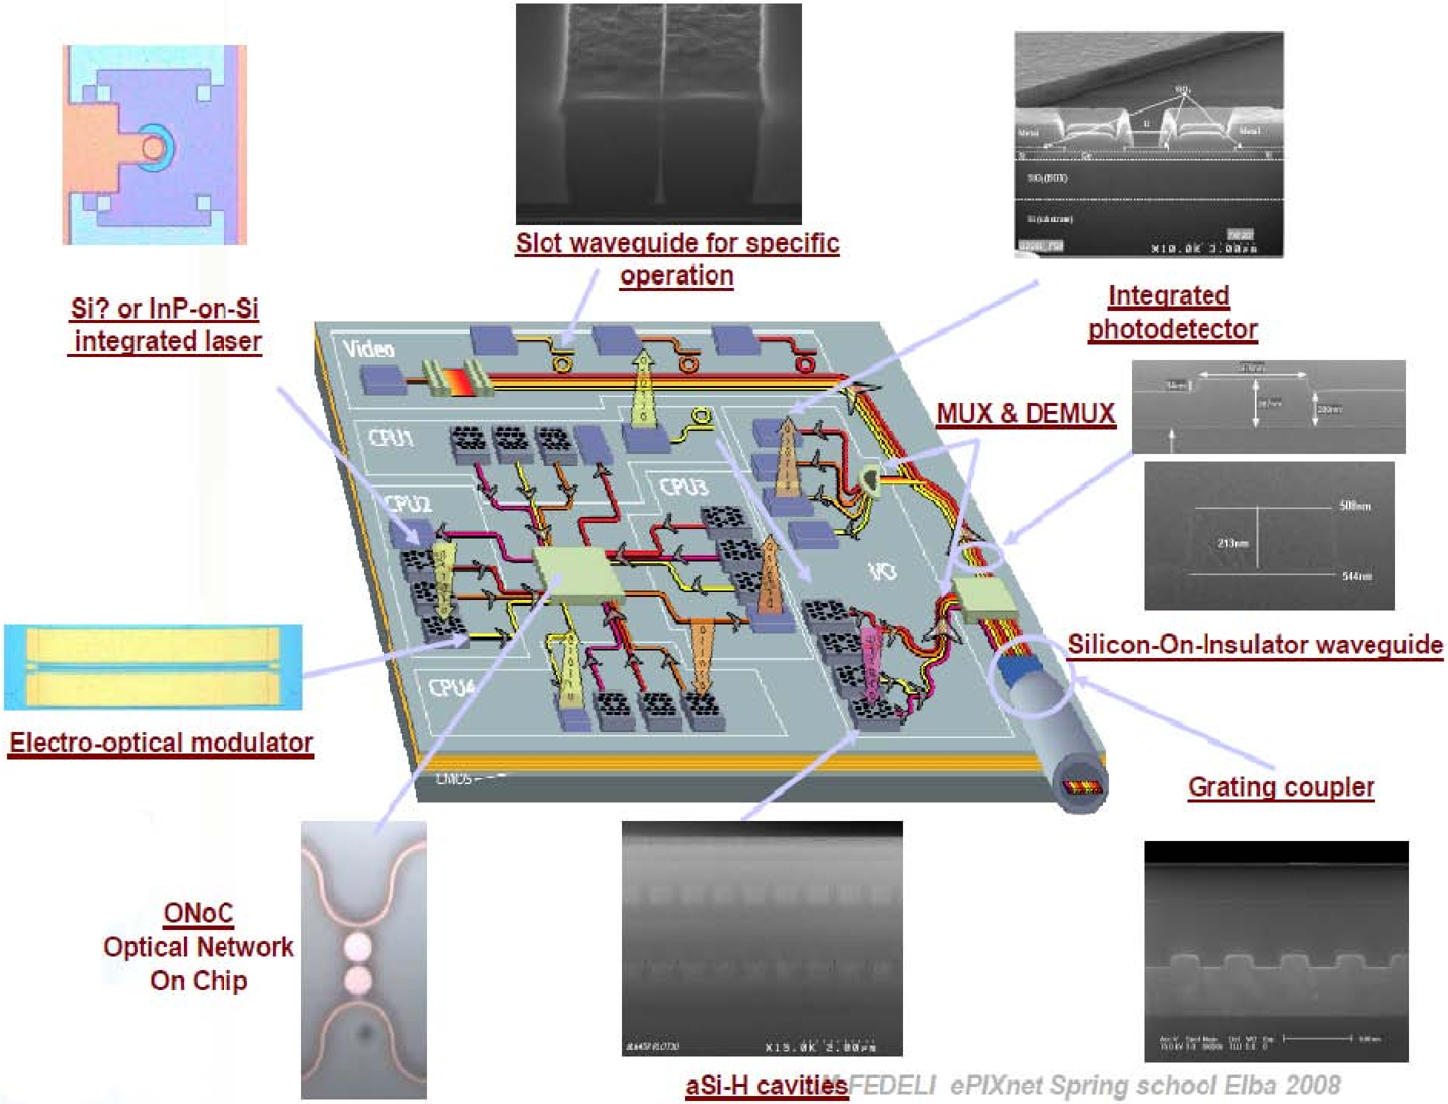

Network reconfiguration

Reconfigurability on the interconnection network can provide several benefits:

- Better adaptation of the network topology to the evolving traffic.
- Less congestion in the links, leading to faster communications.
- Redundancy on the network connections against failures of components.
- Versatility of the interconnection architecture for specific computational problems.

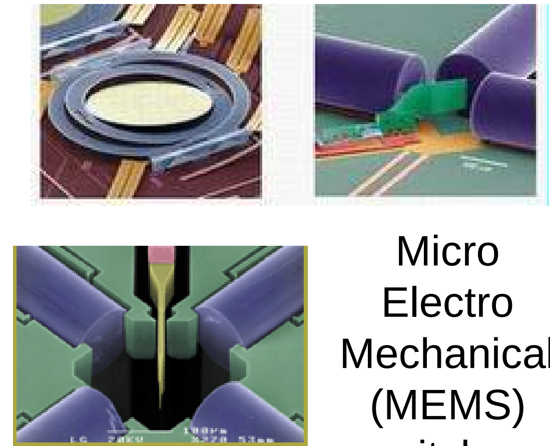

Micro Electro Mechanical (MEMS) switches

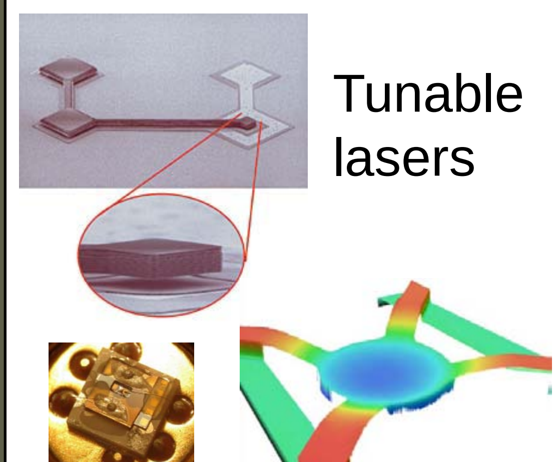

Tunable lasers

Vrije Universiteit Brussel

Network reconfiguration

Reconfigurability on the interconnection network can provide several benefits.

COST

Microring resona...

- M...actu... throu... si... oces... s
- ...ssibility ... integra... them i... very ...rge sc...e
- Ve... power (passive)
- Low coupling losses
- High switching speeds (order of ns)
- Can also be resonant to more than one wavelength

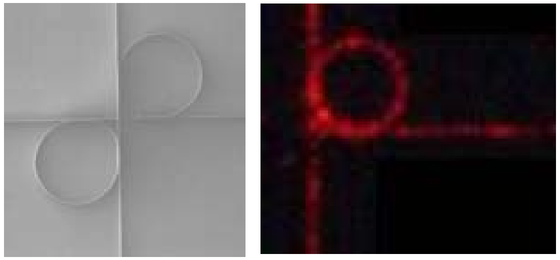

90º turn on an embedded waveguide

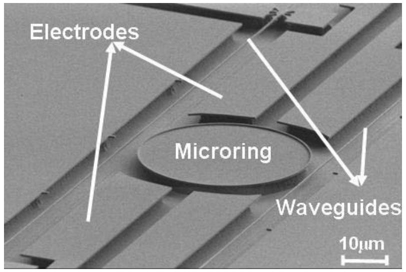

M.-C. Tien, S. Mathai, J. Yao, and M. C. Wu, "Tunable MEMS actuated microring resonators," in IEEE/LEOS International Conference on Optical MEMS and Nanophotonics, Hualien, Taiwan, 2007.

Previous optical NoC architectures proposed

Mixed interconnection network:

- Fixed electrical control network
- Optical data network
  - High-speed communication
  - Small power consumption
  - Off-chip BW = On-chip BW

Interaction between electrical and optical plane:

- Optical paths are used on a circuit-switched way in a packet basis (one packet for each circuit)
- Messages in the electrical layer establish paths
- Set-up, ACK/Path-blocked and tear-down messages

M. Petracca et Al., “Design exploration of optical interconnection networks for chip multiprocessors,” Proc. 16th IEEE Symp.on High Performance Interconnects, 2008.

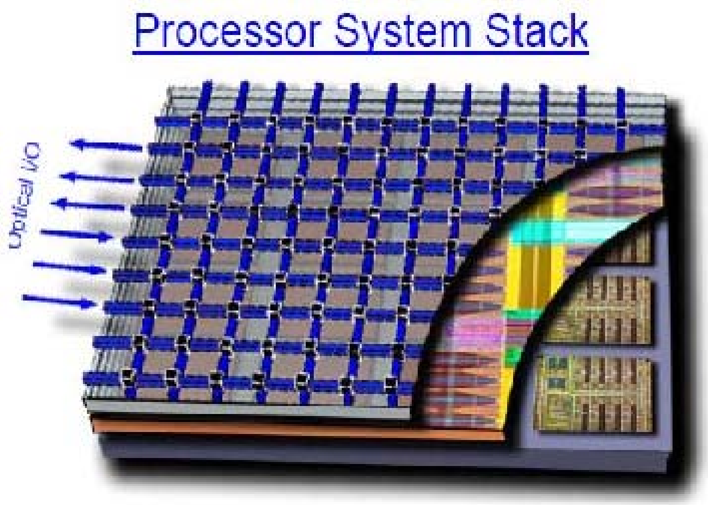

On-Chip Photonic Communications for High Performance Multi-Core Processors
Keren Bergman, Luca Carloni, Columbia University
Jeffrey Kash, Yurii Vlasov, IBM Research

Main issues:

- The set-up time is long if we use the established path only for one message
- This proposal is thus valid only for large data messages
- Switching the microrings more often means a significant increase in power

Network topology

Non-blocking torus:

- N(number of cores)=$4x^2$ (up to 36)
- $N^2/4$ switches
- N/2 warp around links
- 960 Gbps per node (using TDM-WDM)
- Routing policy:
  - Allow two nodes to inject/eject messages on each row/column through gateway
  - XY minimal routing, columns first
- Same spacial distance between every node with its neighbours

- Non-blocking router:
  - It is implemented with 8 microring non-blocking structures (0.07 mm^2)
  - 38.5 GHz bandwidth, single wavelength
  - None, two or four microrings are involved to route four lightpaths simulateously

Design Exploration of Optical Interconnection Networks for Chip Multiprocessors
Michele Petracca, Benjamin G. Lee, Keren Bergman, Luca P. Carloni

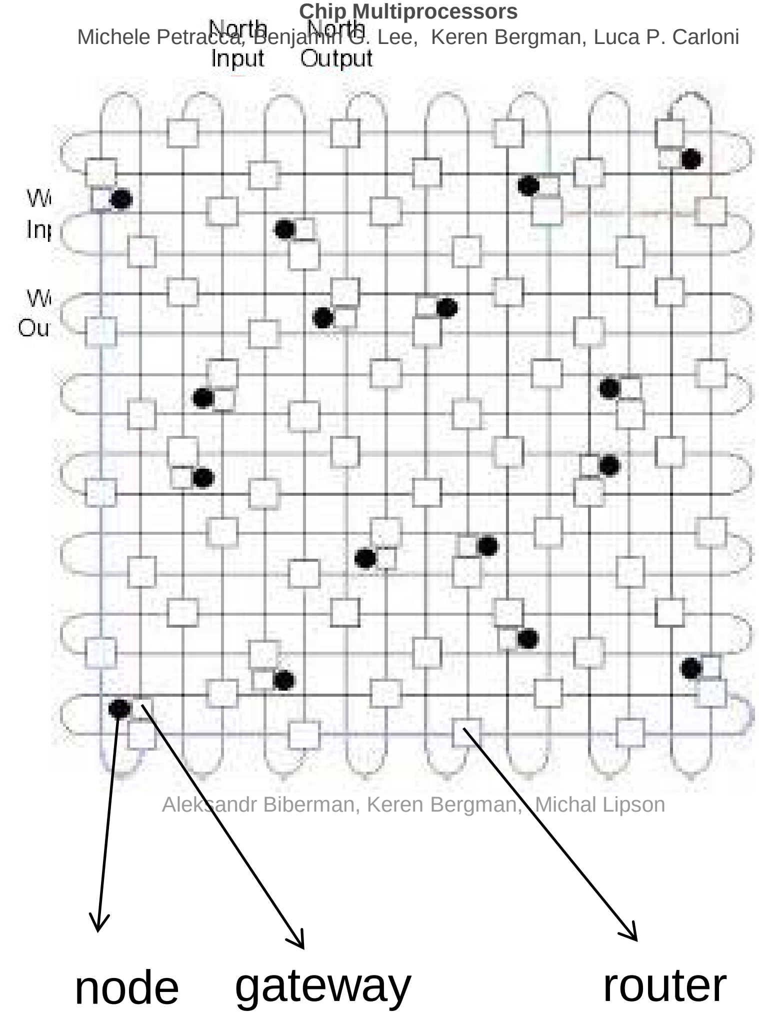

Vrije Universiteit Brussel

Proposed reconfiguration

• Based on same microring router approach, but with a "slow"-switching optical layer that allows for reconfiguration of longstanding traffic patterns (mem locality).

• Photonic circuits (*elinks*) are used as short cuts between the most communicative and distant node pairs.

• Topology changes follow the dynamics of the traffic:

  • Measurement of the busiest node pairs in the reconfiguration interval (1 μs)

  • Assignment of the elinks:

    t_{se}: selection time of possible circuits (elinks) to be established

    t_{sw}: switching time to create the optical elinks

  • Transmission of data through the elink and new measurements for the next interval.

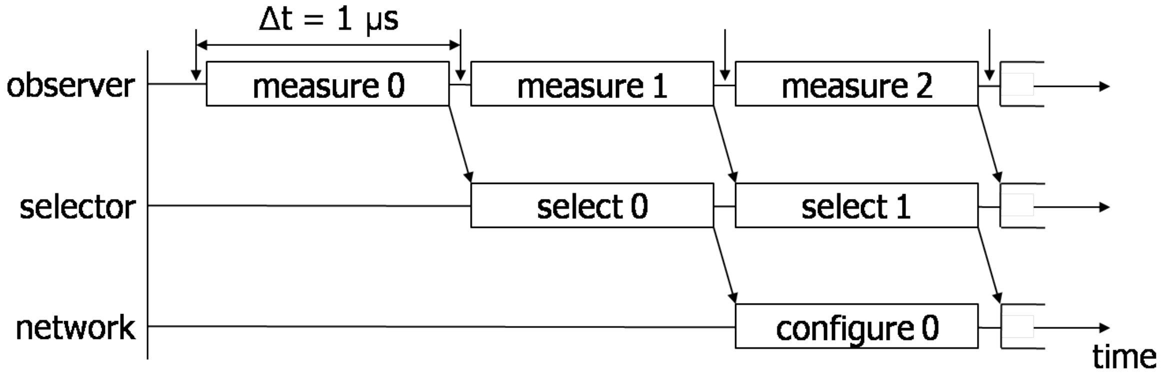

Power consumption

- Reconfigurable layer:
  - Based on slow-switching devices (nanoseconds instead of 30 ps)
  - Only active microrings on the router will be powered

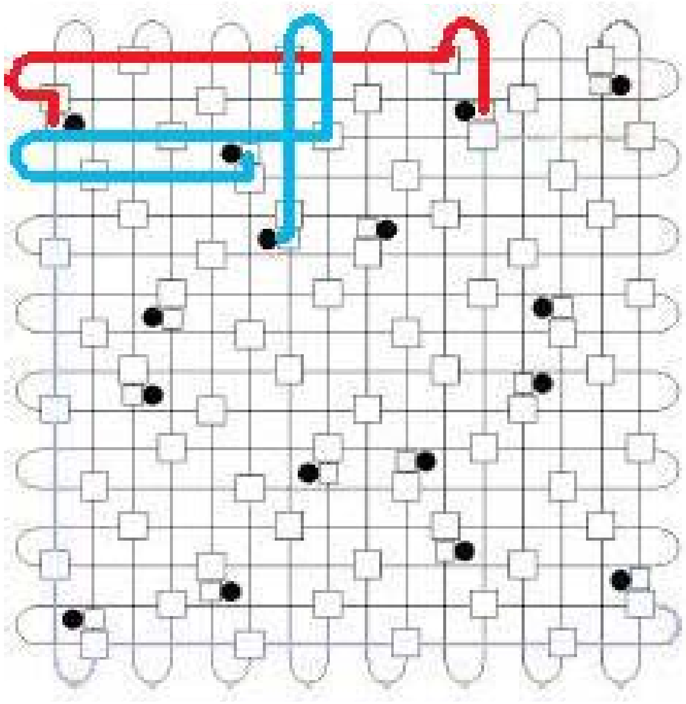

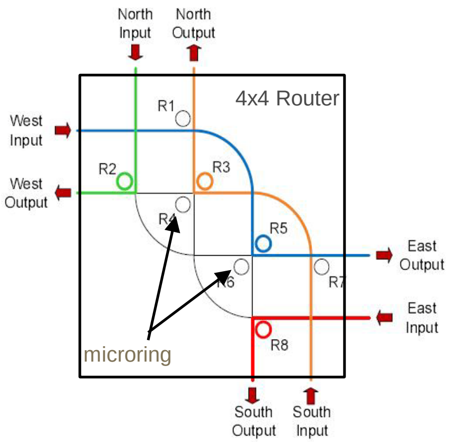

Power consumption (II)

Electric part:

Num elinks * (Dynamic + Static)

Num elinks:

16
64

Dynamic:

Base link activity
Tile length
(1.67 mm)
Electrical link power
(0.34 pJ/bit/mm)
Electrical buffering
(0.12 pJ/bit)
Electrical cross of switches
(0.36 pJ/bit)

Static:

Base link activity
Electrical routing
(0.35 pJ/bit)
Electrical link static power
@10-40 Gbps
(0.25-2 mW)

Vrije Universiteit Brussel

Power consumption (III)

Photonic part:

Num elinks * (Dynamic + Static)

Num elinks:
16
64

Dynamic:
Base link activity
Optical link power
(0.5 pJ/bit)

+ Microrings power
(400 μW)

Static:
Optical link static power
(500 μW)

Simulation environment

Vrije Universiteit Brussel

virtutech **SIMICS**

UNIVERSITEIT GENT

PARIS

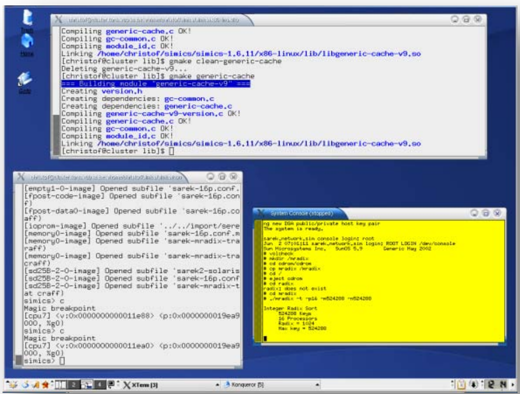

Slow down factor of ~10000 times

⇩

Limited number of simulated nodes

Sunfire servers

Sun microsystems

	UltraSPARC III	UltraSPARC T2
T_reconf	1 us	1 us
T_sel	= 1 us	= 1 us
T_switch	≈ 0 us	≈ 0 us
Cache Size (L2)	512 KB	32 KB
Cache Size (L1i/d)	32 KB	64 B
Processor speed	2.5 GHz	2.5 GHz
Routing Latency	= 1ps	= 1 ps
Bandwidth	10 /2.5 Gbps	10 /2.5 Gbps
Elink Bandwidth	10 Gbps	10 Gbps

Benchmark applications

SPLASH-2
Scientific parallel algorithms

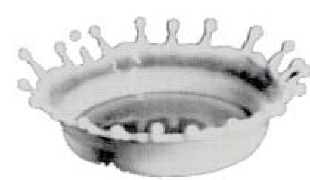

Vrije Universiteit Brussel

Simulation environment (II)

- Architectures simulated:
  - 16 USIII cores, single threaded
  - 16 USIII cores with 4 threads running per core (64 threads total)
  - 64 UST2 cores, single threaded

- Electrical links unning @ 10 Gbps and optical @ 40 Gbps

- Networks compared: “slow” and “fast”, reconfigurable and not.

- We measured the average hop distance between node pairs in the topology, the average memory access latency and the total power consumption of the interconnect.

Average hop distance

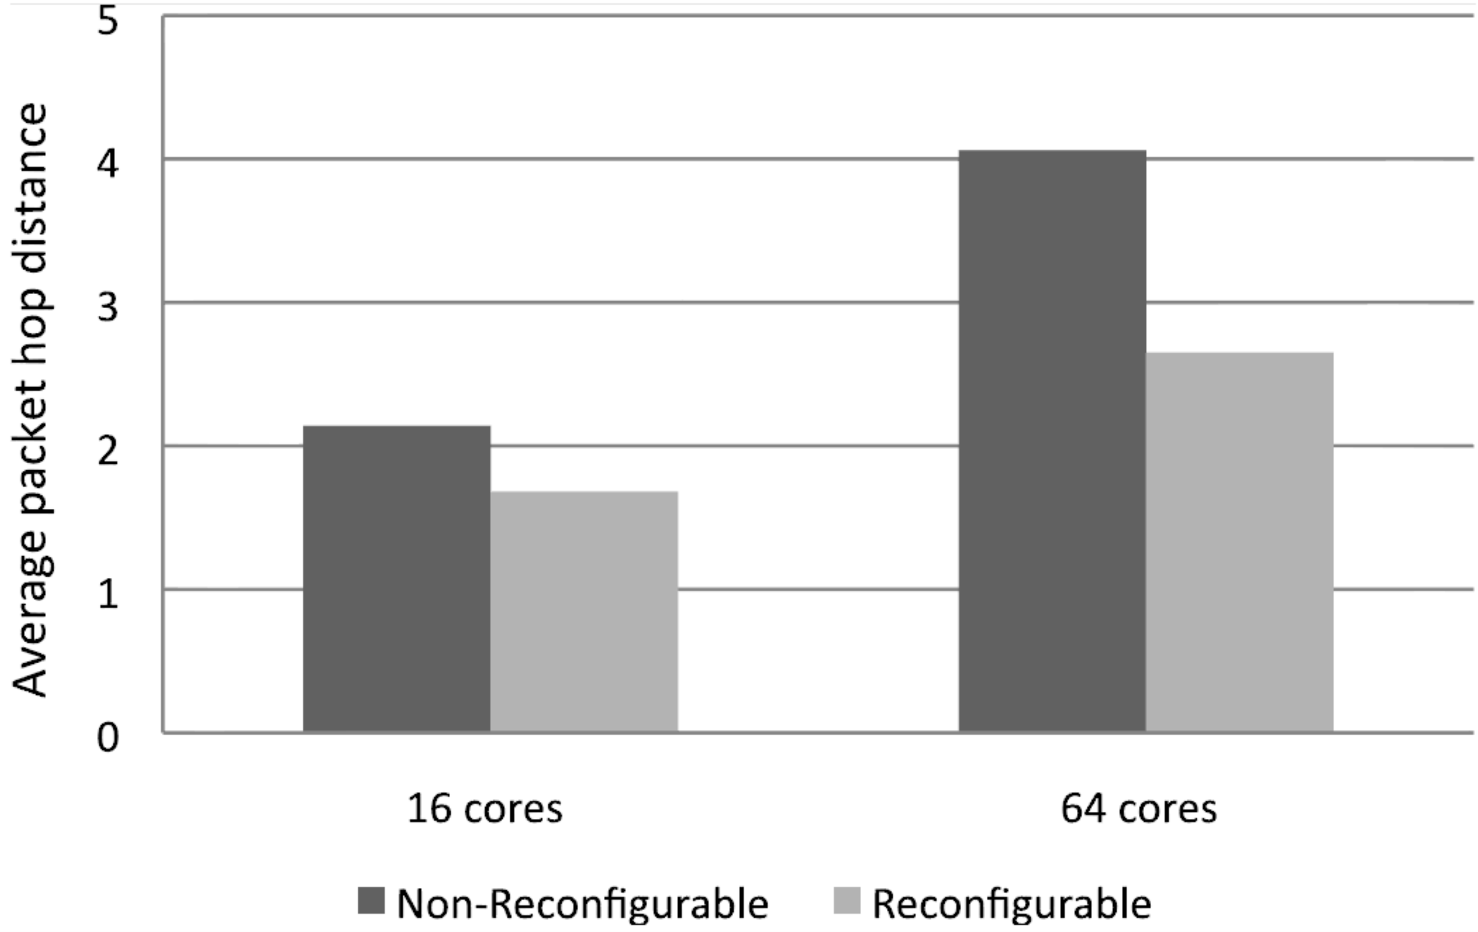

Memory access latency

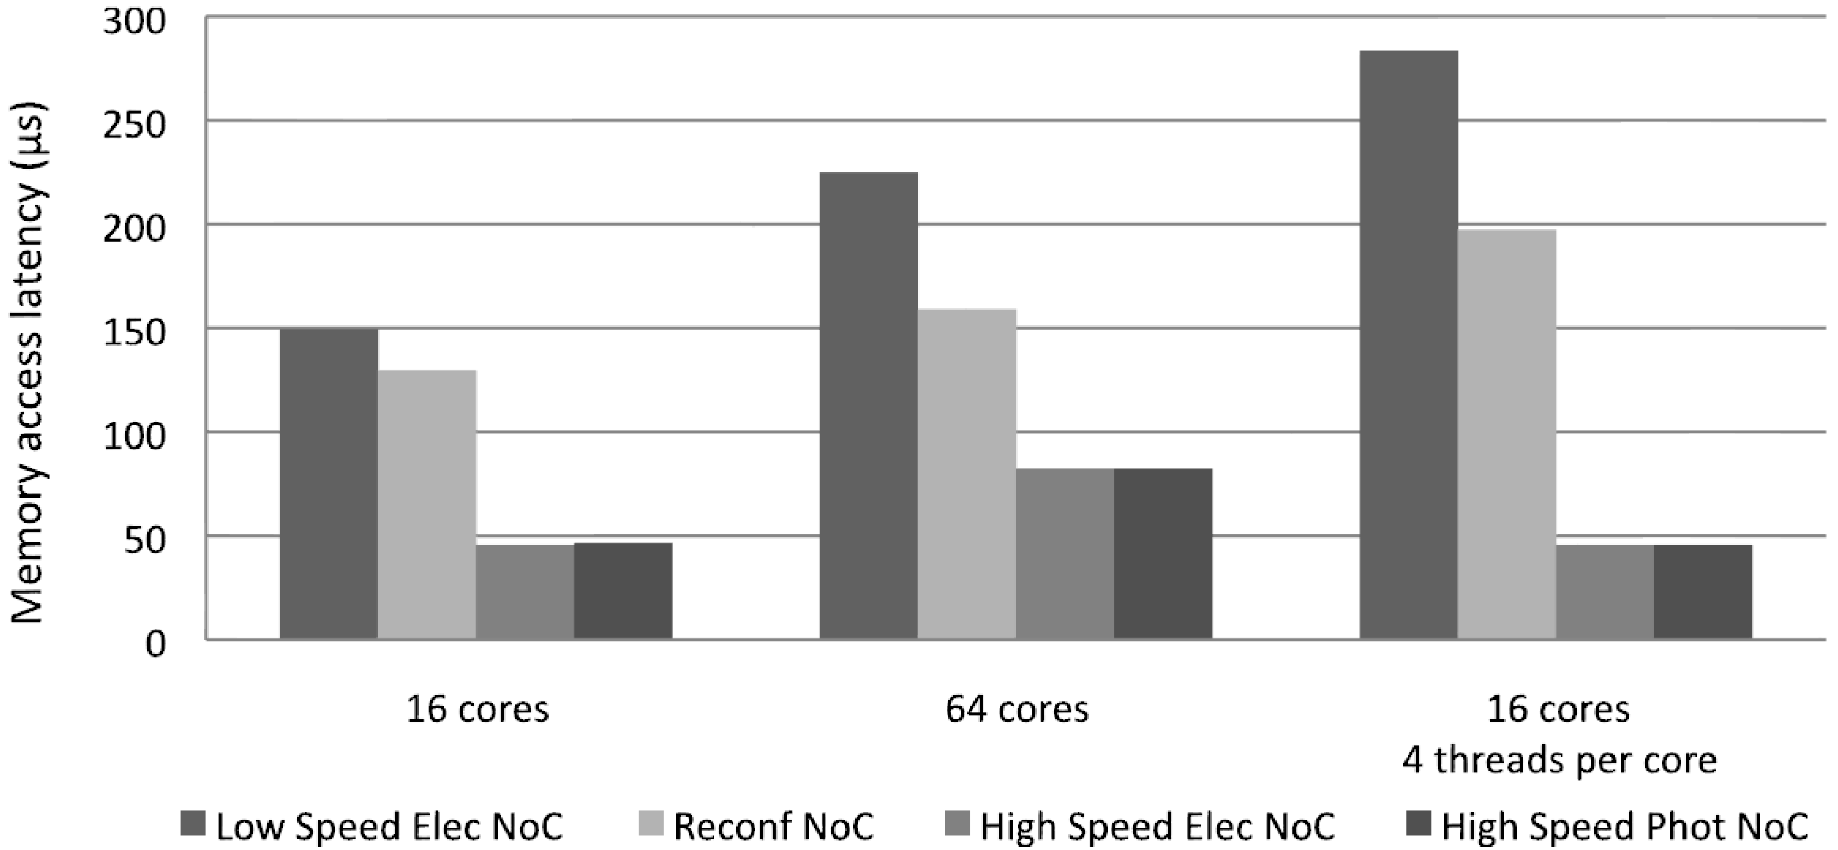

10 Gbps   10+40 Gbps   40 Gbps   40 Gbps

Power consumption (16c)

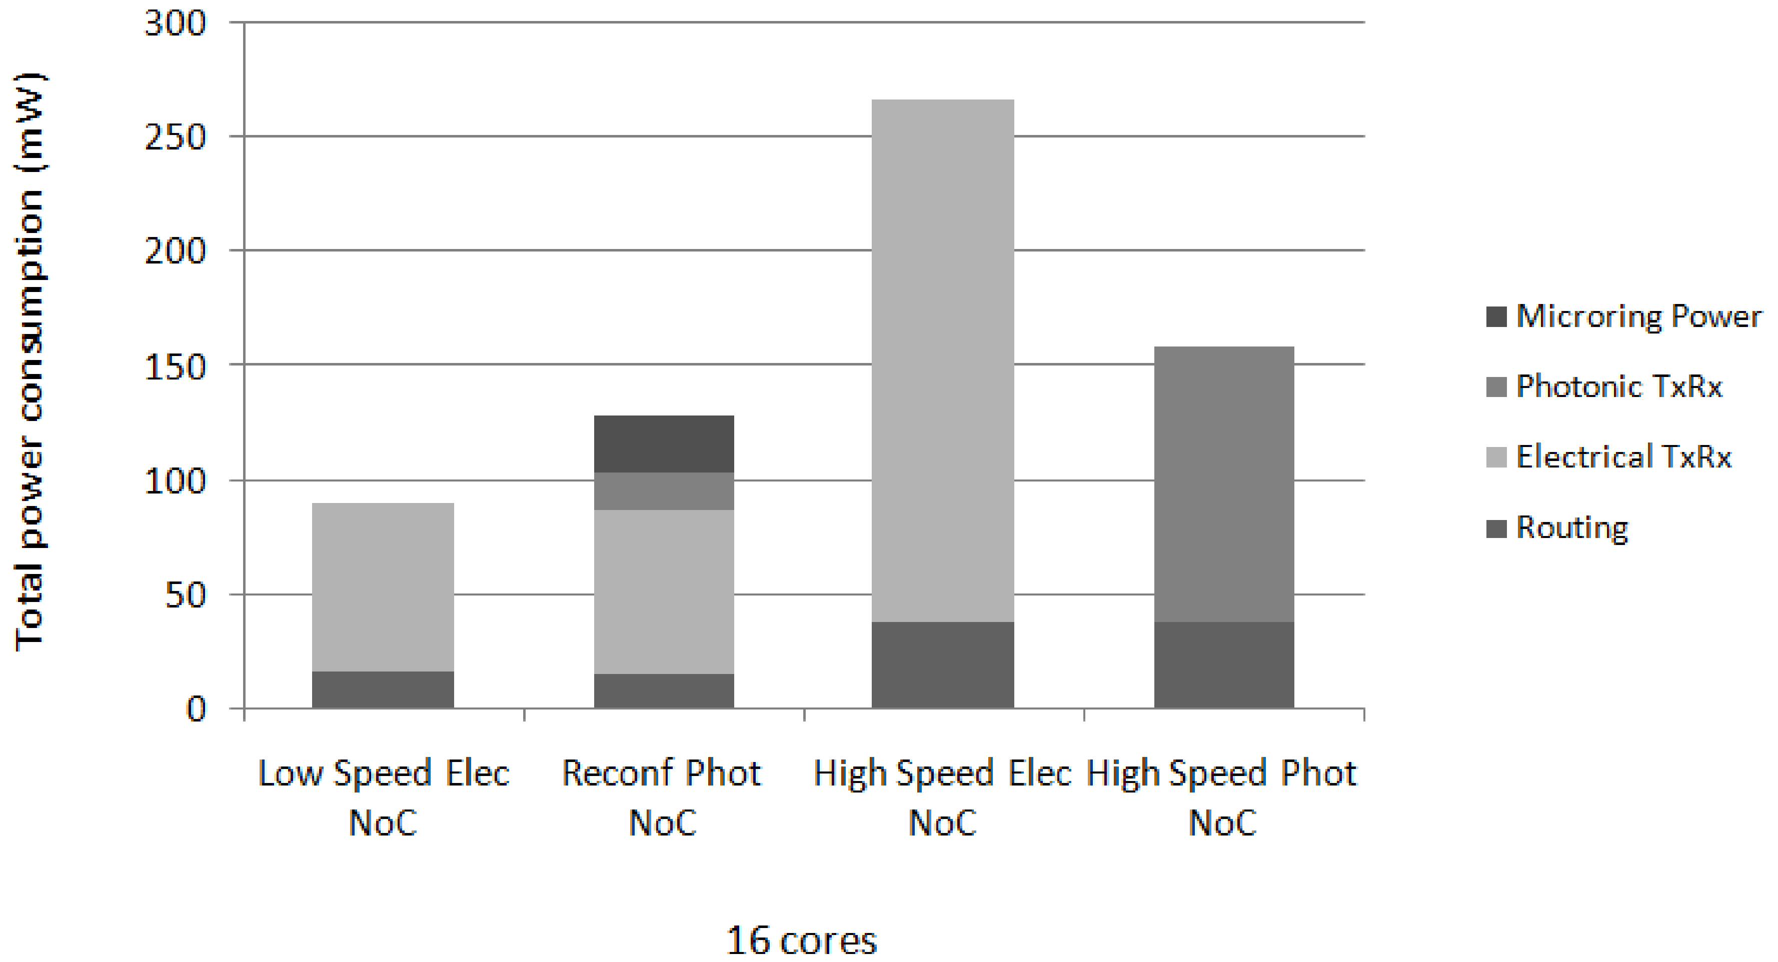

Vrije Universiteit Brussel

Power consumption (16c/4thr)

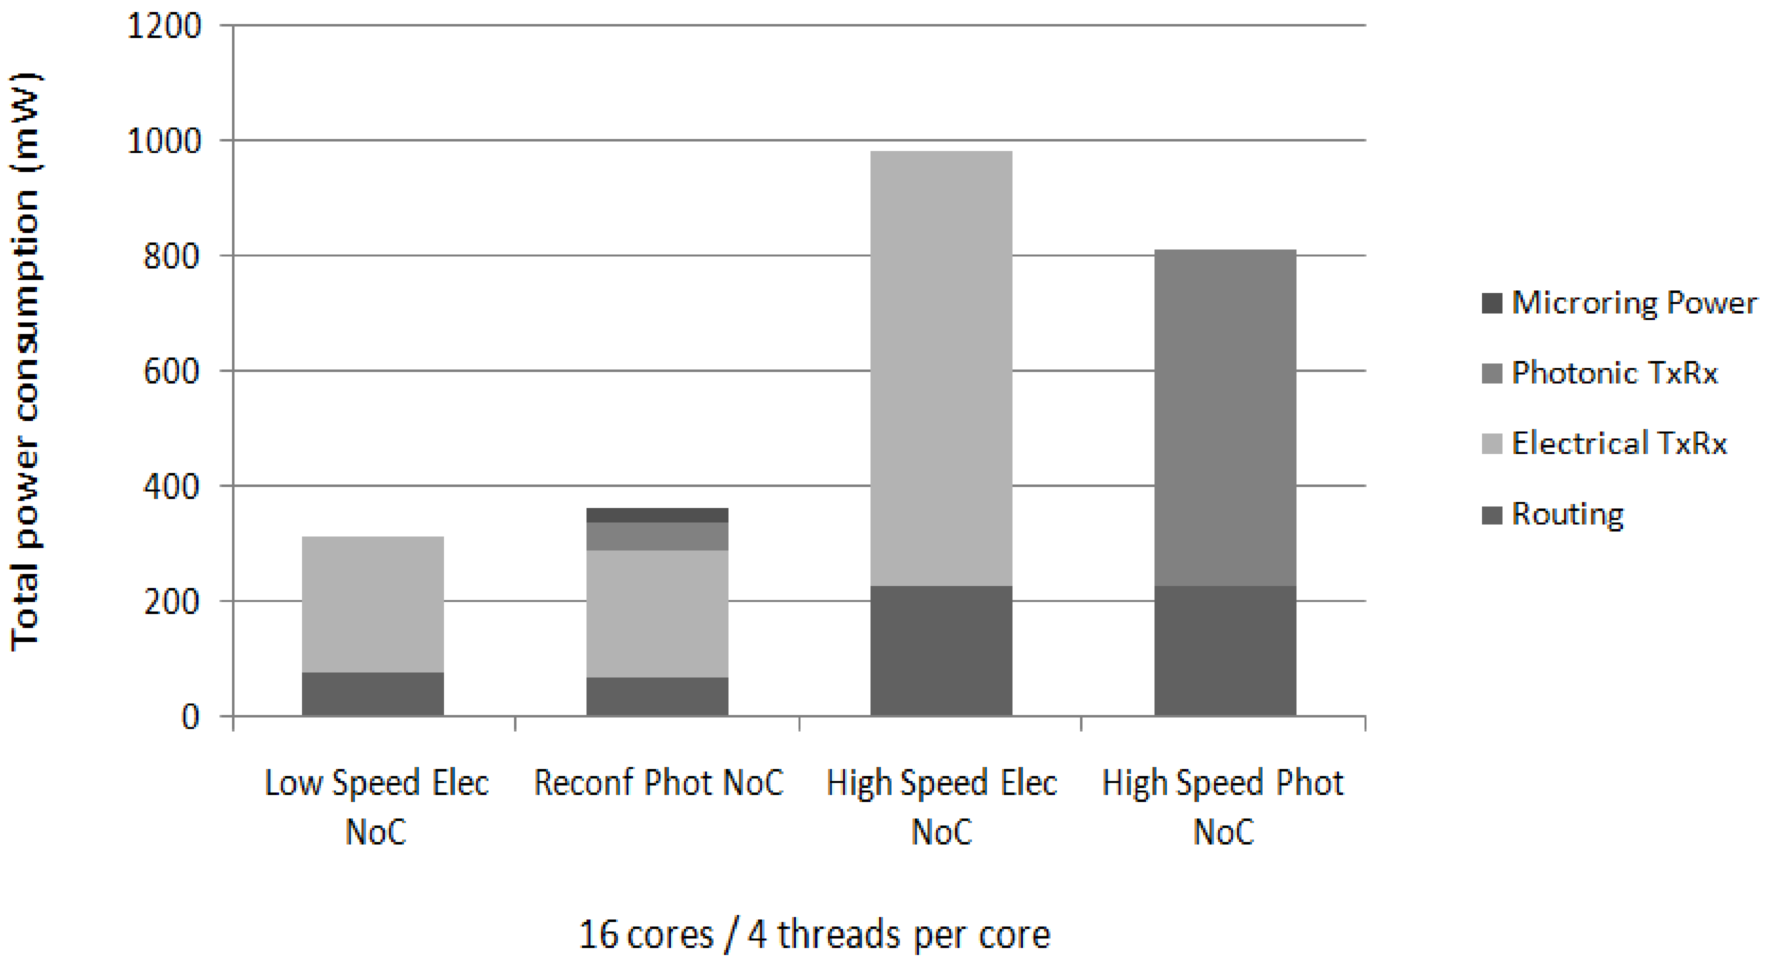

Vrije Universiteit Brussel

Power consumption (64c)

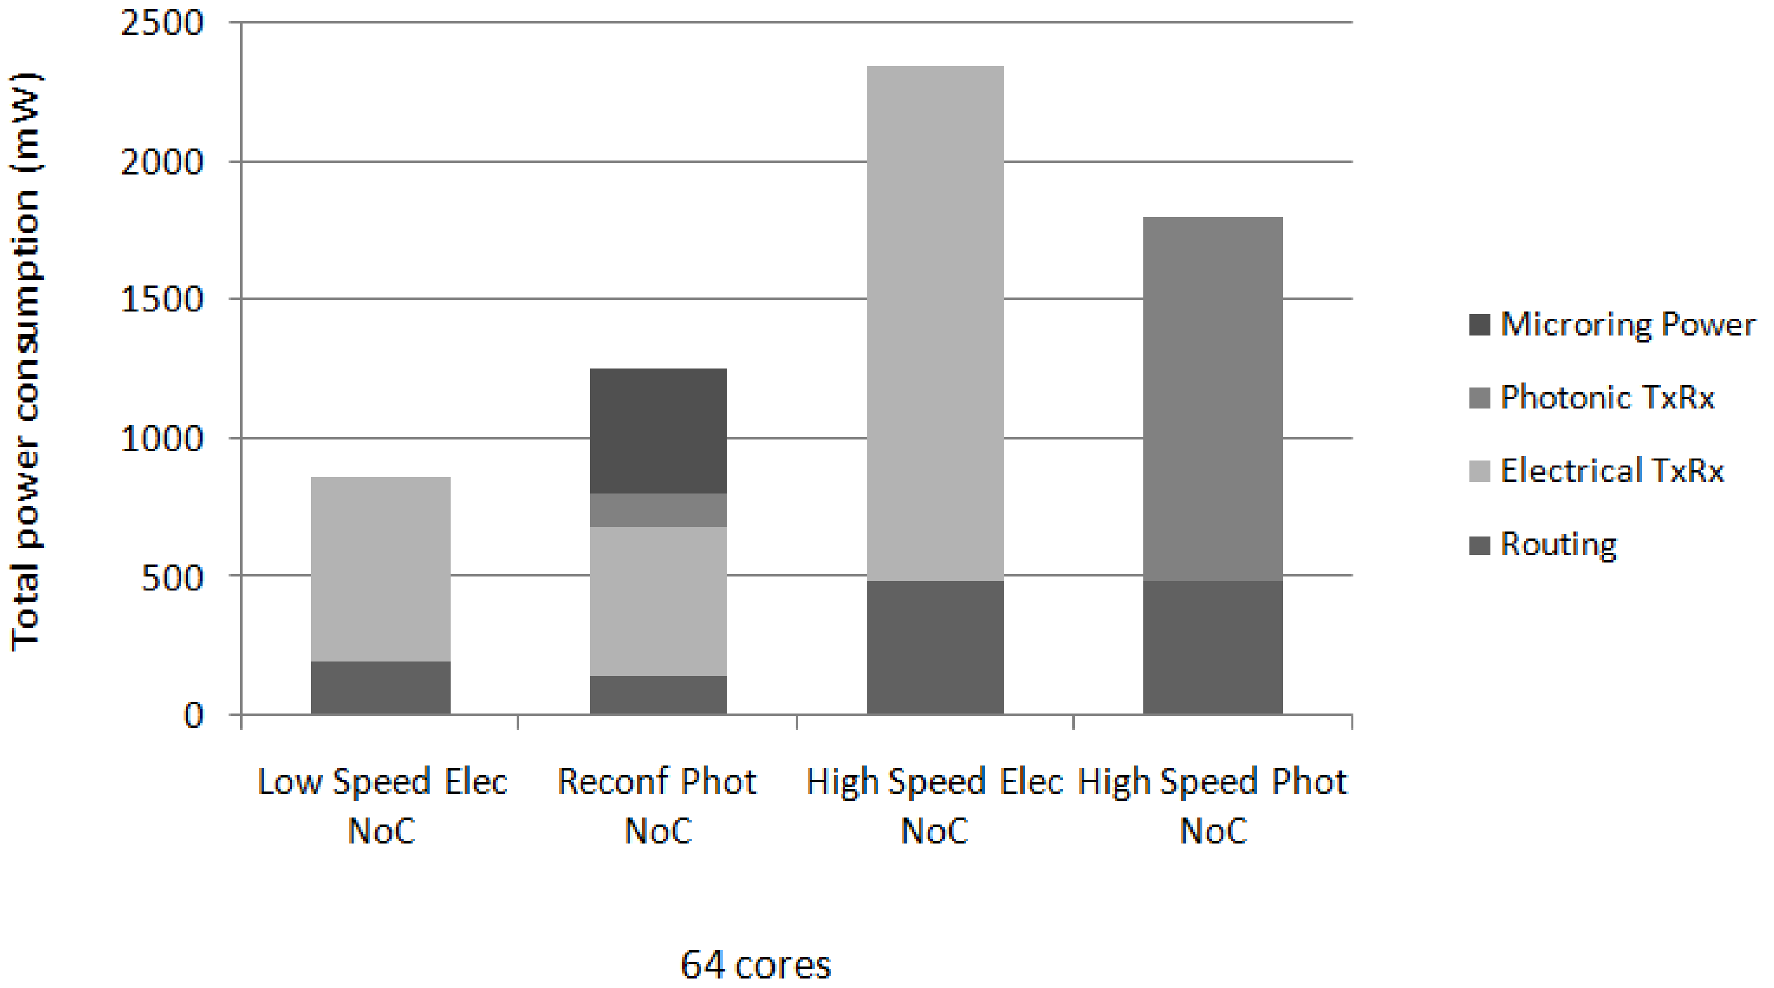

Vrije Universiteit Brussel

Summary

Introduced a reconfigurable optical interconnect for a NoC multicore system.

It makes use of low-power photonic switches, based on microring resonators, to route messages over a reconfigurable optical layer.

Electrical packet switching is used on the underlying control network, and optical circuit switching is used on the data network for all packet sizes.

Significant reduction in power usage and aditional performance boost.

Acknowledgements

Vrije Universiteit Brussel

UNIVERSITEIT GENT

PARIS

NEMO
Network of Excellence on Micro-Optics

Thank you for your attention

iartundo @ tona.vub.ac.be

More info at **http://tona.vub.ac.be**