Performance Evaluation of 2D-Mesh, Ring, and Crossbar Interconnects for Chip Multi-Processors

NoCArc 09

Jesús Camacho Villanueva, José Flich, José Duato
Universidad Politécnica de Valencia

Hans Eberle, Nils Gura, Wladek Olesinski
Sun Microsystems

December 12, 2009

Index

Introduction

Network simulator

Simulation model

Performance analysis

Conclusions

Future work

Introduction

Network simulator

Simulation model

Performance analysis

Conclusions

Future work

Introduction

- Networks-on-chip (NoCs) are the critical component of a chip multiprocessor (CMP) as the number of cores increases
- CMPs with 32 cores are already on the drawing table
  - 48 cores recently announced by Intel
- Need for a full-system simulator with an accurate network simulation model
- Not considering the network component and full-system simulation may lead to **Incorrect Conclusions**

Introduction

Topology considerations for NoCs (in CMPs)

- Crossbars simplify the design, but they have a limited scalability [Micro07]
- 2D-Meshes have better scalability than crossbars and simplify the design of a tiled organization
- Rings have a simpler design than 2D-Meshes, but the average distance between nodes is higher
- The network capacity is also a critical parameter in the design of NoCs

[Micro07] Hoskote Y., Vangal S., Singh A., Borkar N., Borkar S.: 'A 5-GHz mesh interconnect for a teraflops processor', IEEE Micro Mag., 2007, 27, (5), pp. 51–61

Introduction

Goals

- To develop an accurate simulation tool for the on-chip network taking into account the target machine: coherence protocol, OS, and application

- At the network level the simulation tool needs to allow:
  - Collective communication
  - Different topologies
  - Different architectures:
    - Switch architecture
    - Switching mechanisms (WH, VCT)
    - Flit size, flow control…

- Introduction
- Network simulator
- Simulation model
- Performance analysis
- Conclusions
- Future work

Network simulator

- SIMICS + GEMS + GAPNET
- SIMICS: Full-system simulator
- GEMS: A set of modules for SIMICS that enables detailed simulation of Chip-Multiprocessors (CMPs)
  - Provides a detailed memory system simulator
  - Implements the cache coherence protocol
- GAPNET: Event-driven network simulator providing collective communication

Network simulator

GapNet and network interface

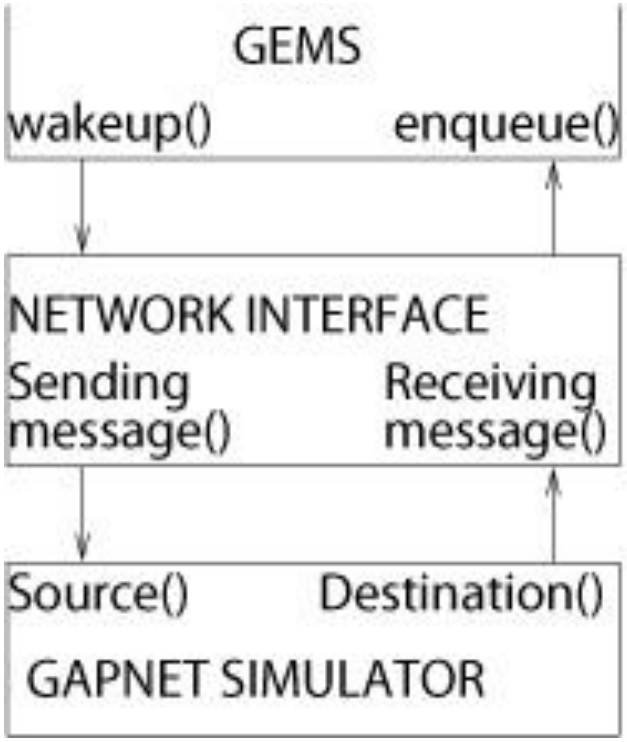

Network simulator

GapNet simulator events

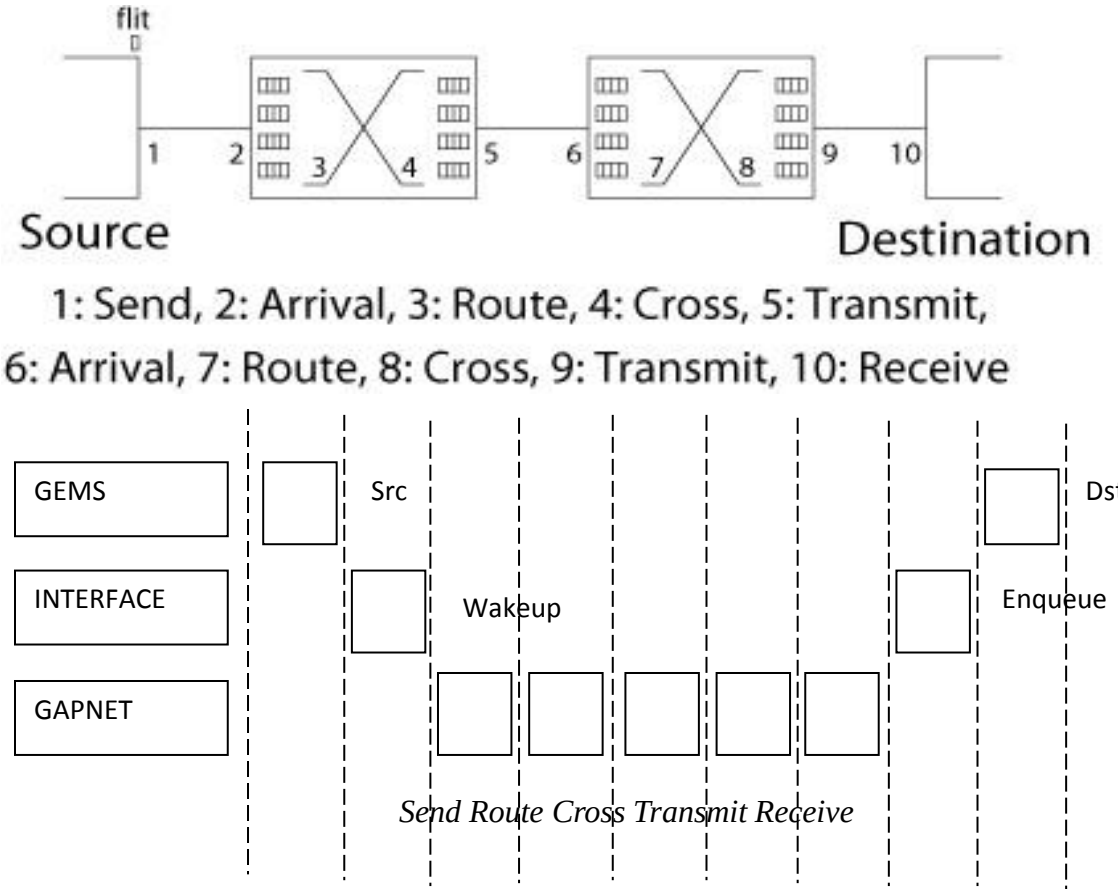

Introduction

Network simulator

Simulation model

Performance analysis

Conclusions

Future work

Simulation model

- Sarek machine (Sun Fire server) with Solaris10
- 32 cores with a SPARC CPU, private cache for the L1 and shared cache among all the processors for the L2

	L1 cache	L2 cache
Size	128 KB	8 MB
Associativity	8-way	16-way
Line Size	64 B	64 B
Hit Latency	3 cycles	6 cycles

- Cache coherency protocol is a directory protocol with non-inclusive and blocking caches

Simulation model

Interconnects

- Four interconnect types: fixed delay interconnect, crossbar, 2D-mesh and bidirectional ring
- 2D-mesh is organized as a 4x8 array and routing is based on X-Y dimension order routing. Bidirectional ring choose the shortest path

	Ideal	Crossbar	2D-Mesh	Ring
Link Latency [cycles]	-	5	1	1
Switch Delay [cycles]	1..128	2	1	1

Fixed delay interconnect means constant latency and infinite bandwidth

Simulation model

Interconnects

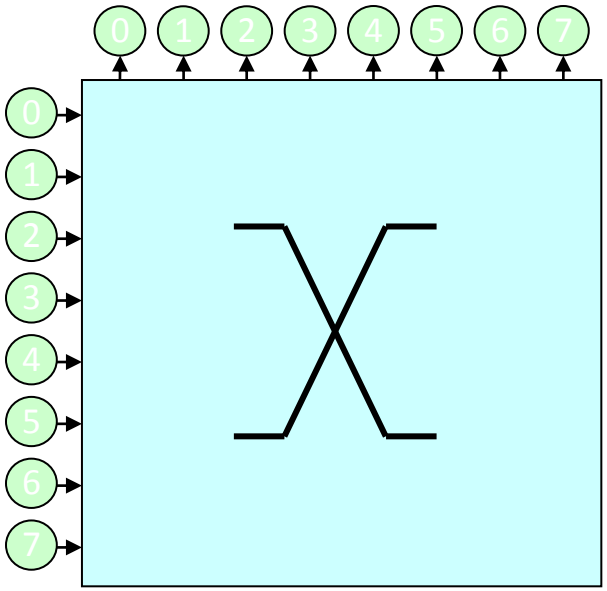

Crossbar

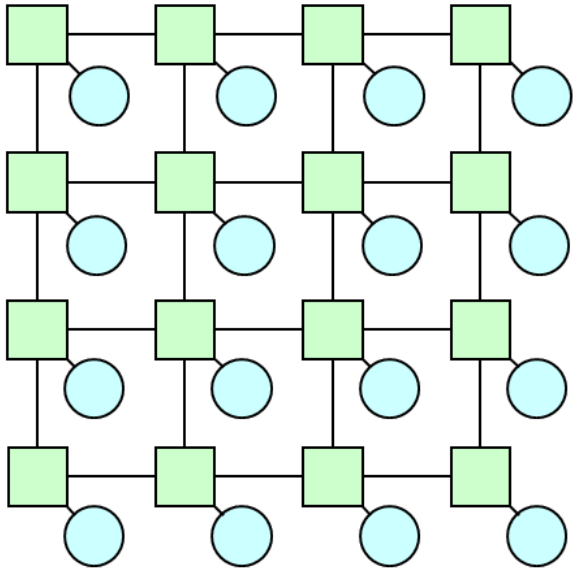

2D-mesh

Simulation model

Interconnects

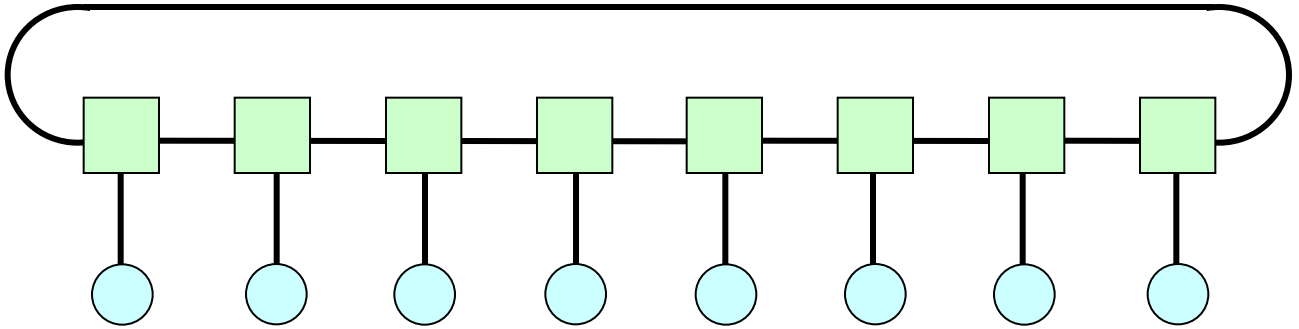

Ring

Ideal network:

- fixed delay
- free of contention
- unlimited amount of bandwidth

Simulation model

Tile based design

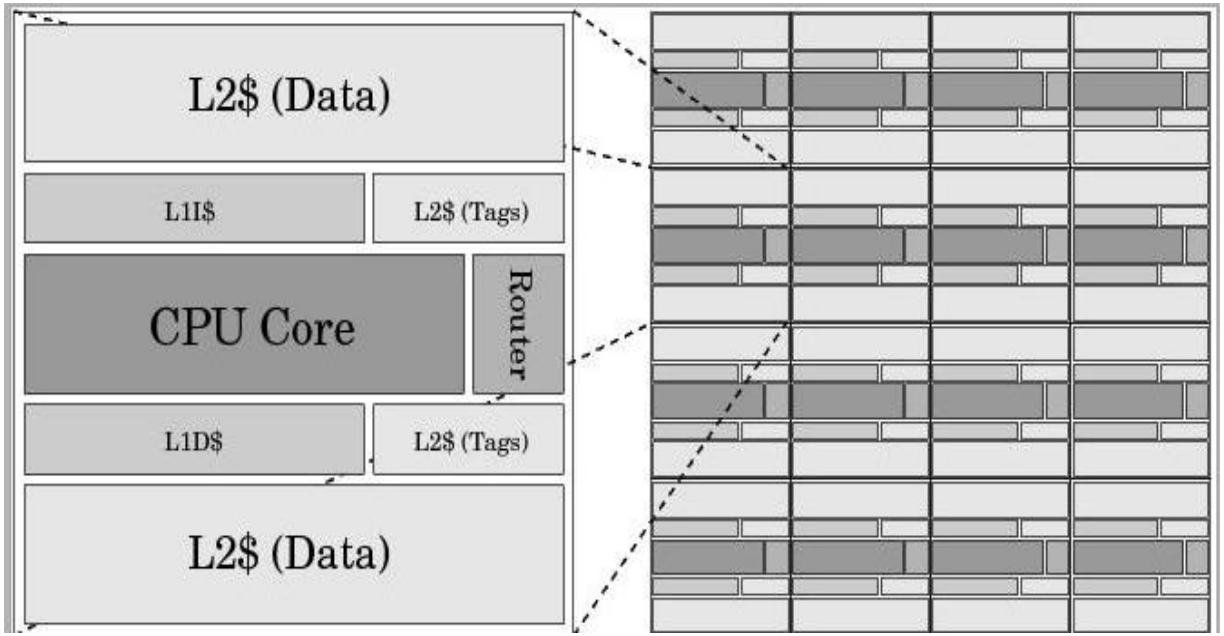

Tile based 2D-mesh 4x4

Simulation model

Network capacity

- We change the capacity of the network by modifying the flit size
- The flit is the minimum amount of data information that can be flow-controlled through a link
- The flit size is an important parameter at 2 levels:

  - Architectural level: Assuming wormhole, different flit sizes lead to different contention levels
  - Design level: Large flit size lead to more expensive router designs that consume more area and power

Introduction

Network simulator

Simulation model

Performance analysis

Conclusions

Future work

Performance analysis

- Ideal network: normalized execution time (cycles)

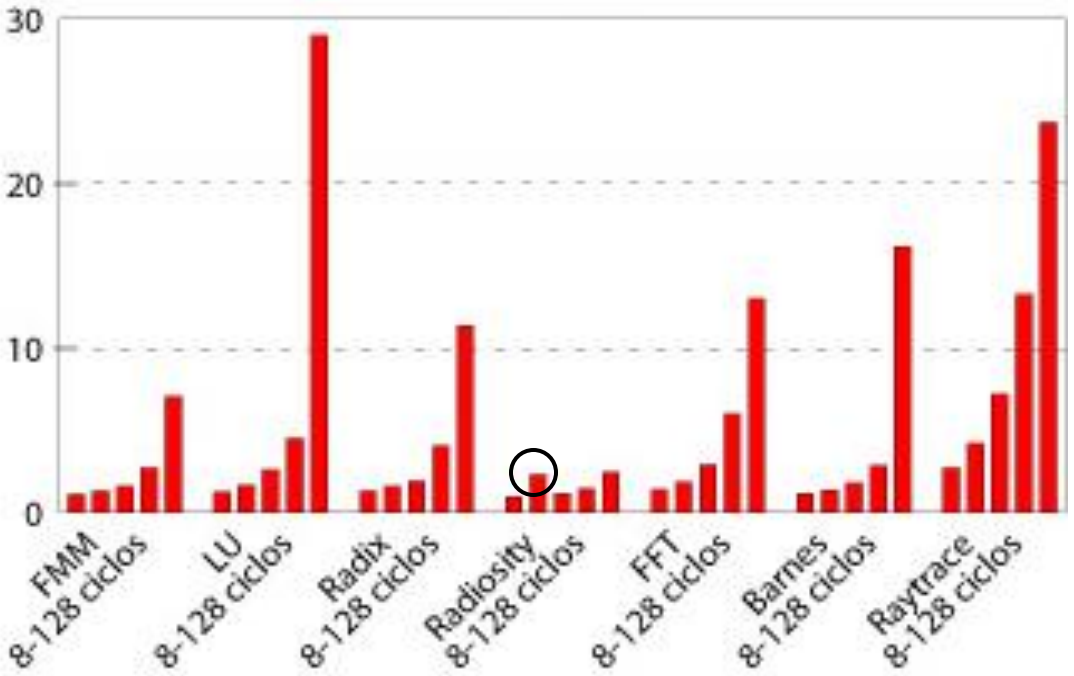

delay spectrum for each benchmark

The system (for most applications) is very sensitive to network latency. E.g. 41% increase for FFT, 171% for Raytrace, 32% for Radix (8c vs 1c delay)

Performance analysis

- Ideal network: normalized number of L1 misses

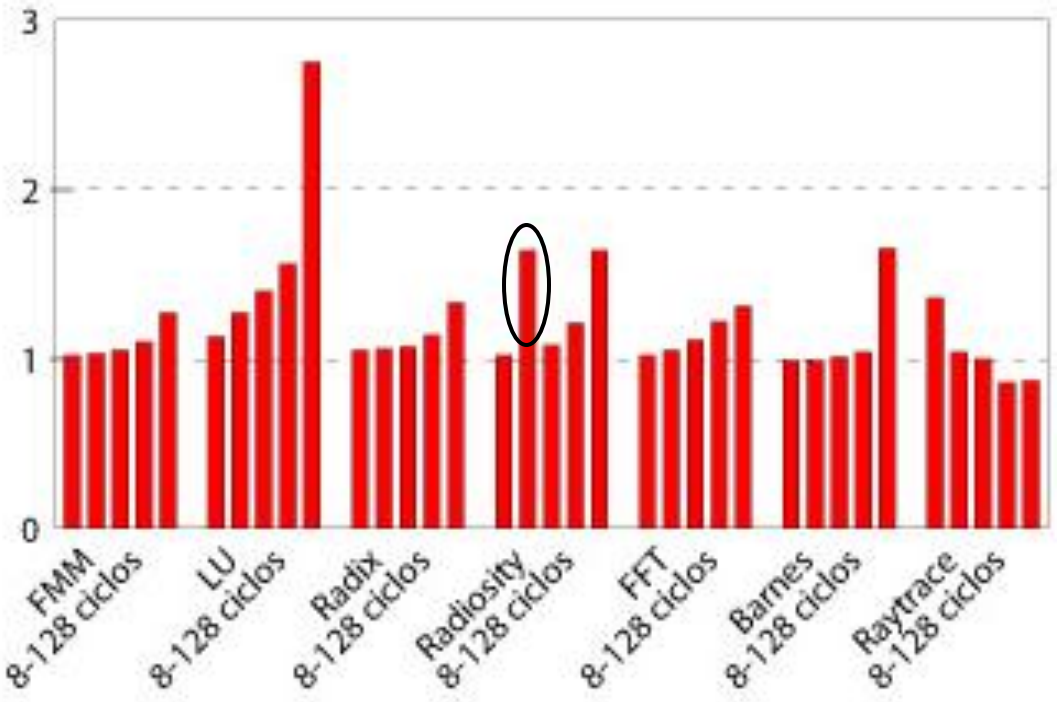

delay spectrum for each benchmark

Performance analysis

- Ideal network: normalized number of messages

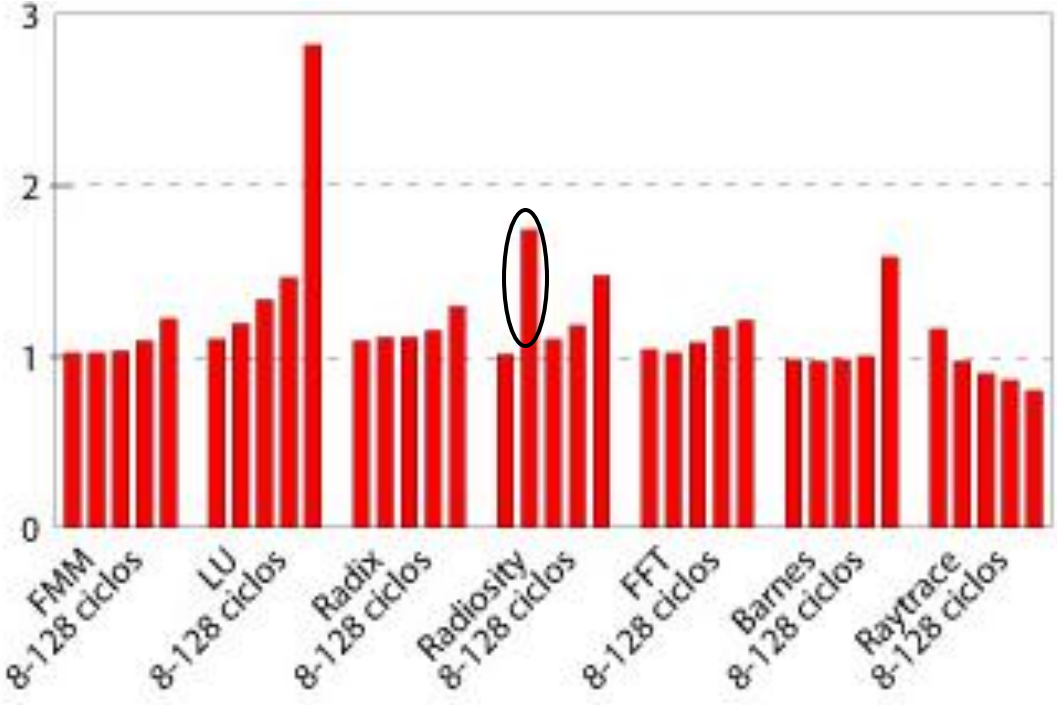

delay spectrum for each benchmark

Performance analysis

- 2D-Mesh achieves the best performance. The average savings for narrow flits:
  - 19% when compare with ring
  - 26% when compare with crossbar
- Crossbar for wide flits perform better than ring in FMM, LU, FFT and Barnes and similar than the others.
  - As we shrink the flit size, the behavior change and the crossbar becomes worse.
- Ring with wide flits achieve similar performance than 2D-Mesh with narrow flits.
- Narrow flits tend to delay execution time, regardless of the topology, however 2D-Mesh is less affected.
- A good trade-off would be a 2D-Mesh with moderate flit sizes (for example 8B), for this CMP configuration.

Performance analysis

Comparison between 2D-mesh, ring and crossbar

FMM

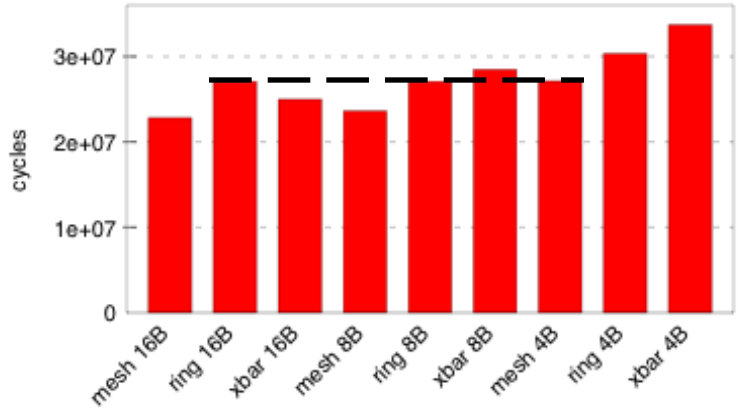

LU

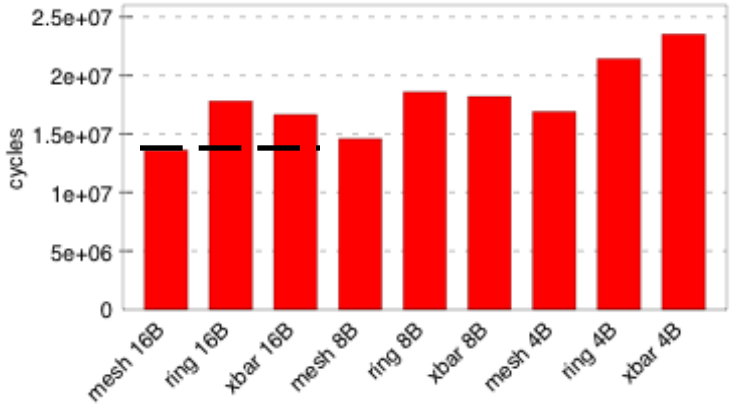

Radix

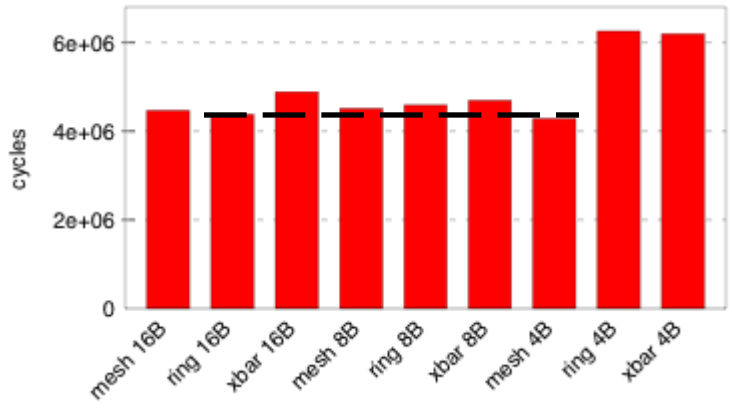

Raytrace

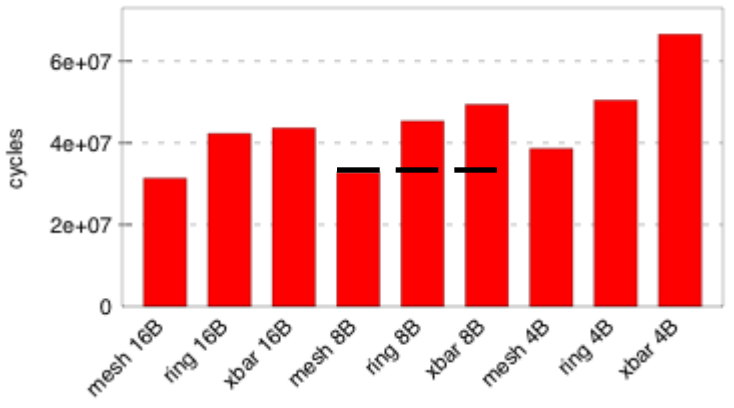

Performance analysis

Comparison between 2D-mesh, ring and crossbar

FTT

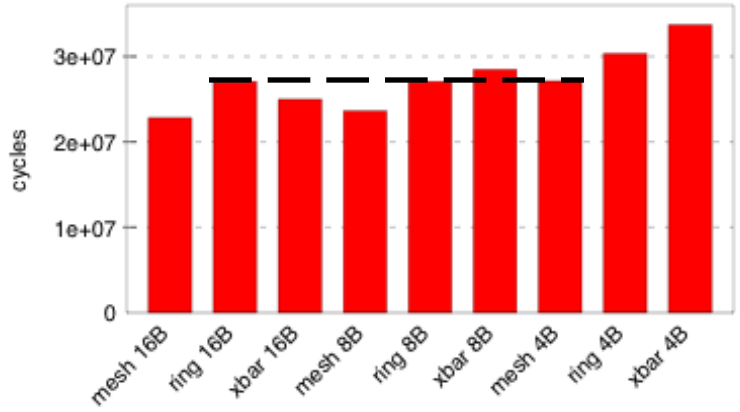

Barnes

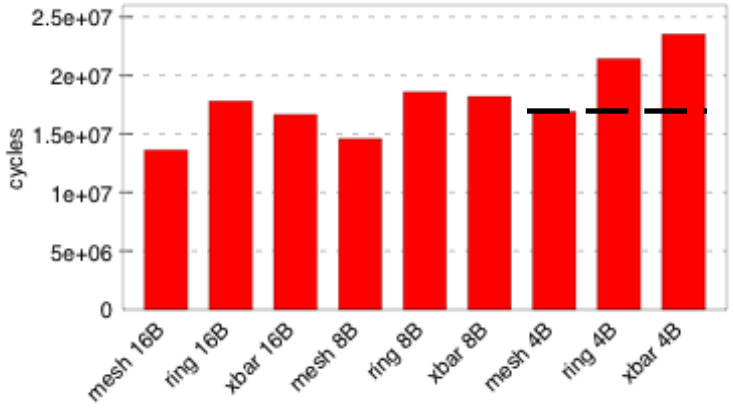

Radiosity

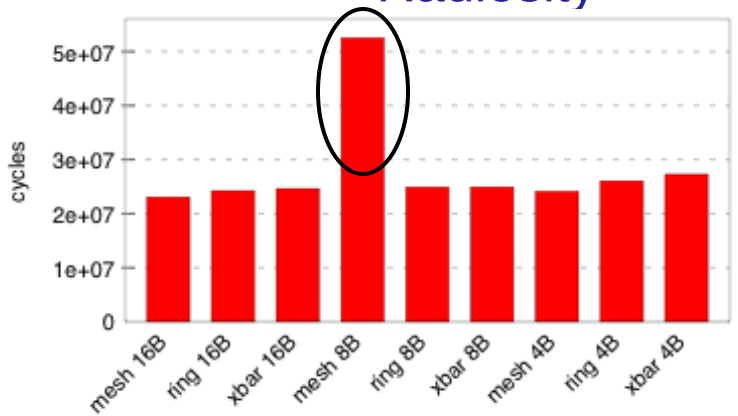

L1 Miss Types in Radiosity

	16	8	4
User	537,136	541,517	539,187
Supervisor	198,764	480,737	201,876
Total	735,901	1,022,255	741,063

Performance analysis

L1 miss rates (%): low network load

		mesh	ring	xbar
Radix	16B	0.35	0.33	0.33
	8B	0.35	0.37	0.36
	4B	0.38	0.30	0.29
Radiosity	16B	0.09	0.09	0.09
	8B	0.07	0.09	0.09
	4B	0.09	0.09	0.09
FFT	16B	0.36	0.29	0.32
	8B	0.36	0.28	0.29
	4B	0.31	0.26	0.22
Barnes	16B	0.15	0.13	0.14
	8B	0.15	0.13	0.13
	4B	0.14	0.13	0.13
Raytrace	16B	0.84	0.52	0.38
	8B	0.82	0.53	0.32
	4B	0.68	0.38	0.20

Congestion is not an issue (in this CMP configuration)

Introduction

Network simulator

Simulation model

Performance analysis

Conclusions

Future work

Conclusions

- Developed and interfaced a detailed on-chip network simulator to GEMS/SIMICS
- Analyzed the impact of topology and flit sizes on real application's execution time
- Results:
  - - Applications are very sensitive to network latency
  - - Application + system behavior may change because of the network (unpredicted behavior captured by our simulation tool)
  - - 2D-Meshes always outperforms rings and crossbars
- For this CMP configuration, 2D-Mesh with moderate flit sizes is the best option

Introduction

Network simulator

Simulation model

Performance analysis

Conclusions

Future work

Future work

- The tool will enable us to:

  - Evaluation of other cache coherence protocols (token and hammer) with strong requirements for collective communication

  - Impact of multicast traffic on application's execution time

  - Impact of memory controllers on application's execution time

  - Evaluation of commercial workloads

Thank you!

Jesús Camacho Villanueva
e-mail: jecavil@gap.upv.es

December 12, 2009