Reducing the Interconnection Network Cost of Chip Multiprocessors

Pablo Abad, Valentín Puente and Jose Ángel Gregorio.

UNIVERSIDAD DE CANTABRIA

Outline

- **Motivation**
- **Reactive traffic**
  - **End-to-end deadlock**
  - **Rotary solution**
- **In order delivery**
- **Evaluation**
- **Conclusions**

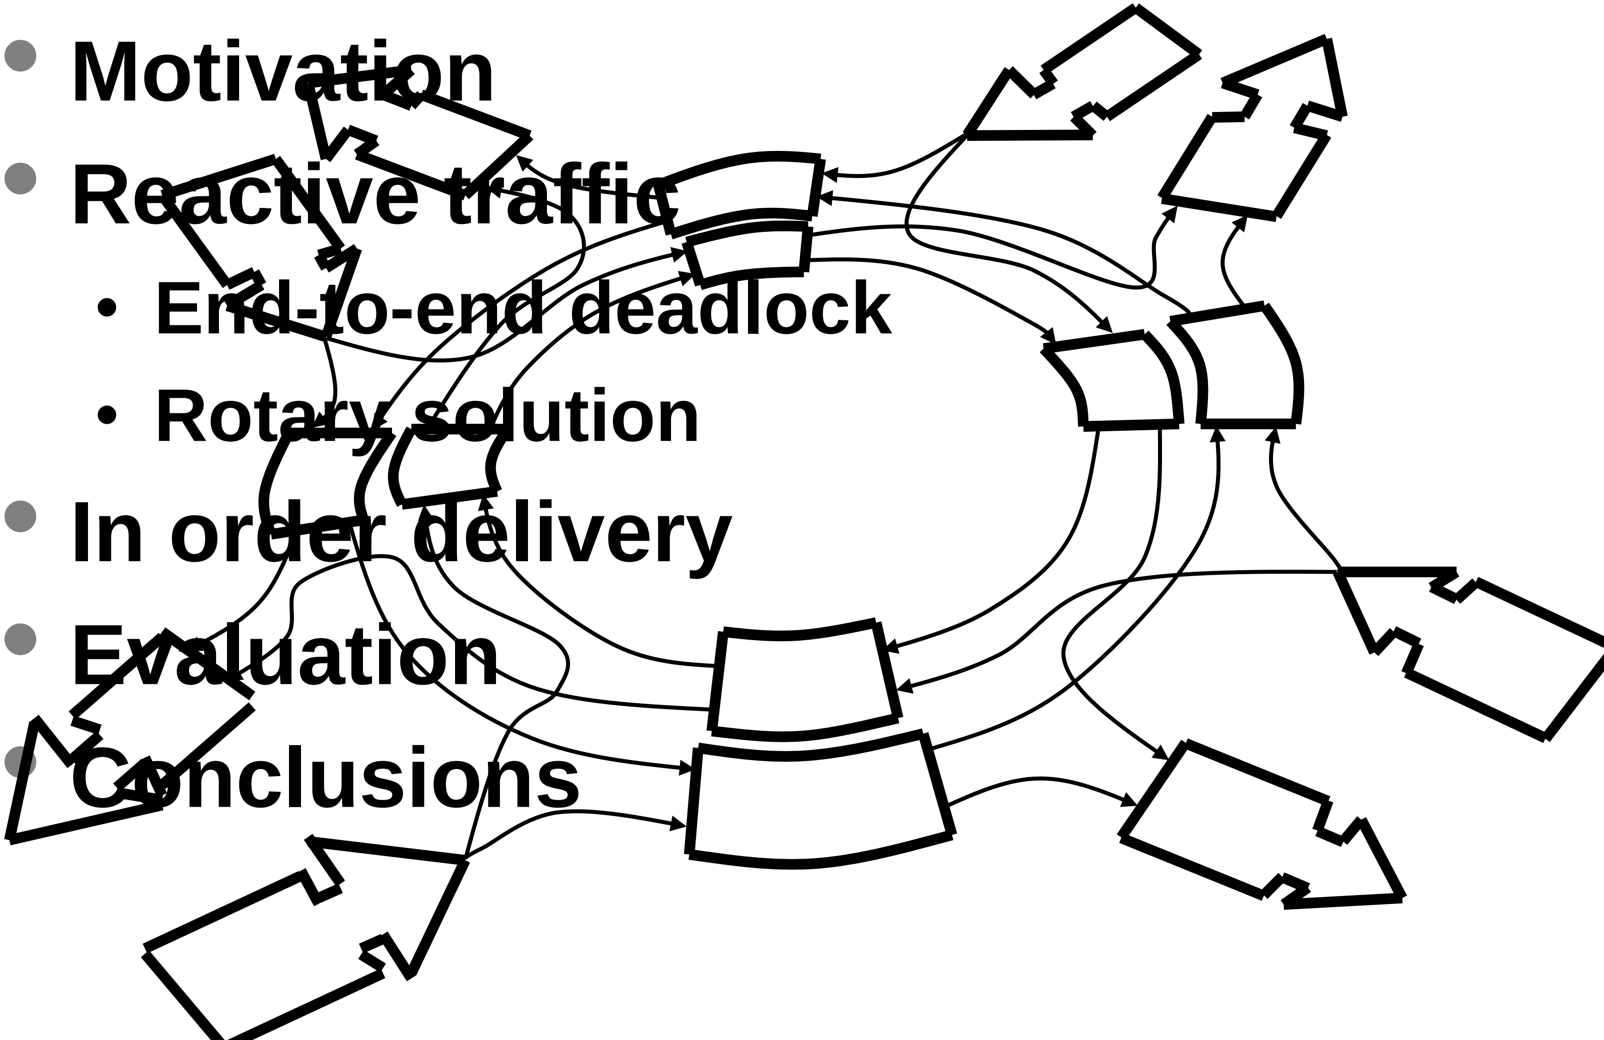

Motivation: NOCs for CMPs

- CMP systems usually assume the presence of cache coherency mechanisms.
- Cache coherence requirements for the communication subsystem:
  - Handle of reactive traffics (end-to-end deadlock).
  - In-order message delivery.
- Solutions for these requirements should have a minimal impact on NoC technological boundaries.

Outline

- Motivation
- Reactive traffic
  - End-to-end deadlock
  - Rotary solution
- In order delivery
- Evaluation
- Conclusions

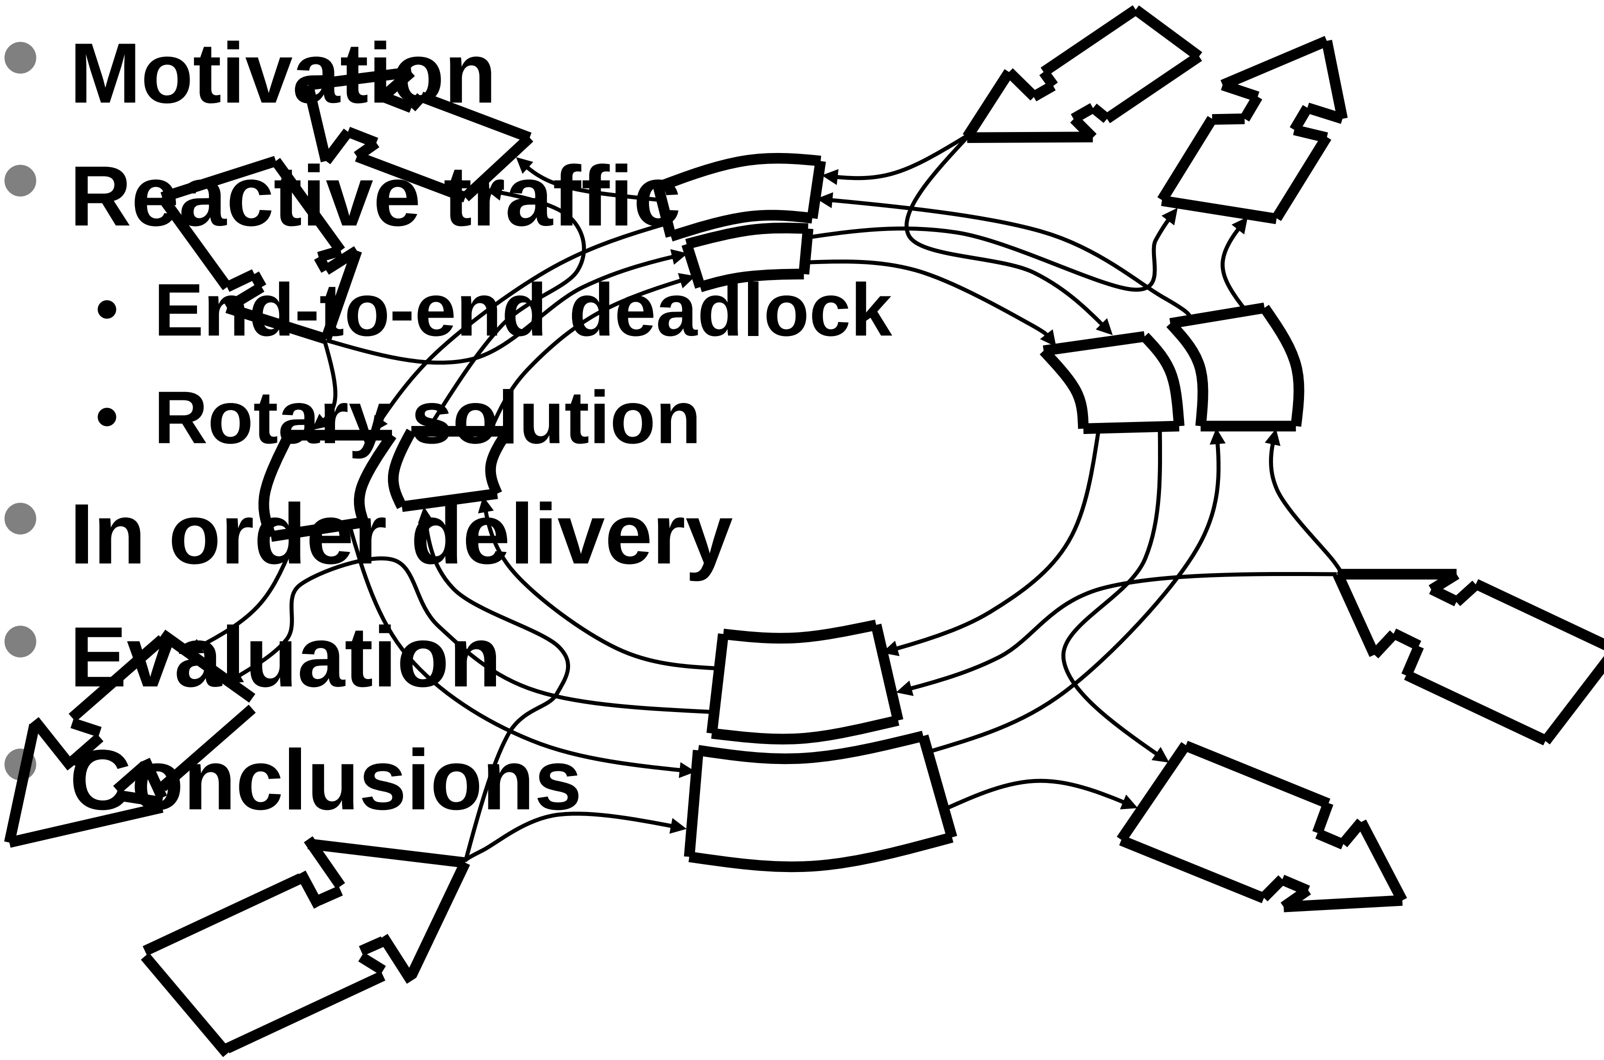

Reactive Traffic

Messages involved in a memory transaction depend one upon the other

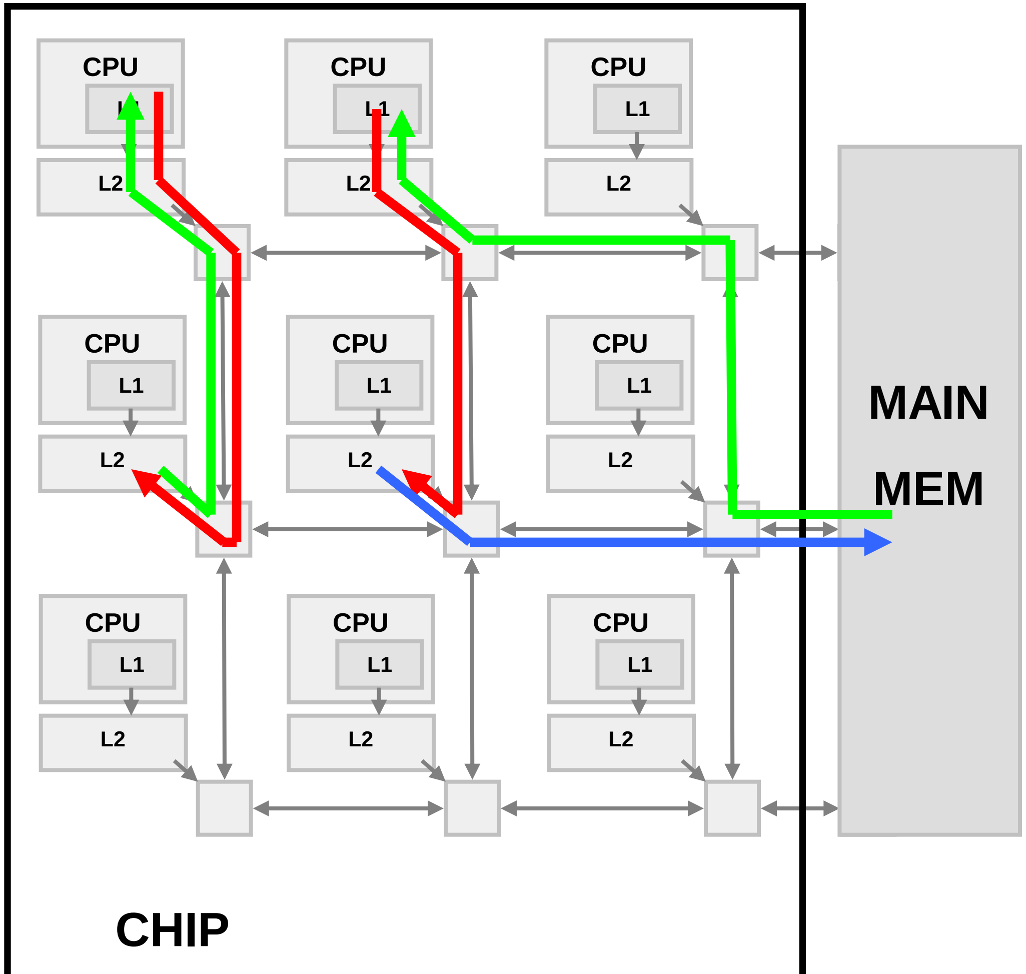

- Minimal 2 messages:
  - CPU-A requests a cache line.
  - CPU-B L2 provides the block.
- Longer Dependencies:
  - CPU-A requests a cache line.
  - The line is not in CPU-B L2, to memory.
  - Memory provides the block.

Reactive Traffic

This kind of communication can cause message-dependent deadlocks.

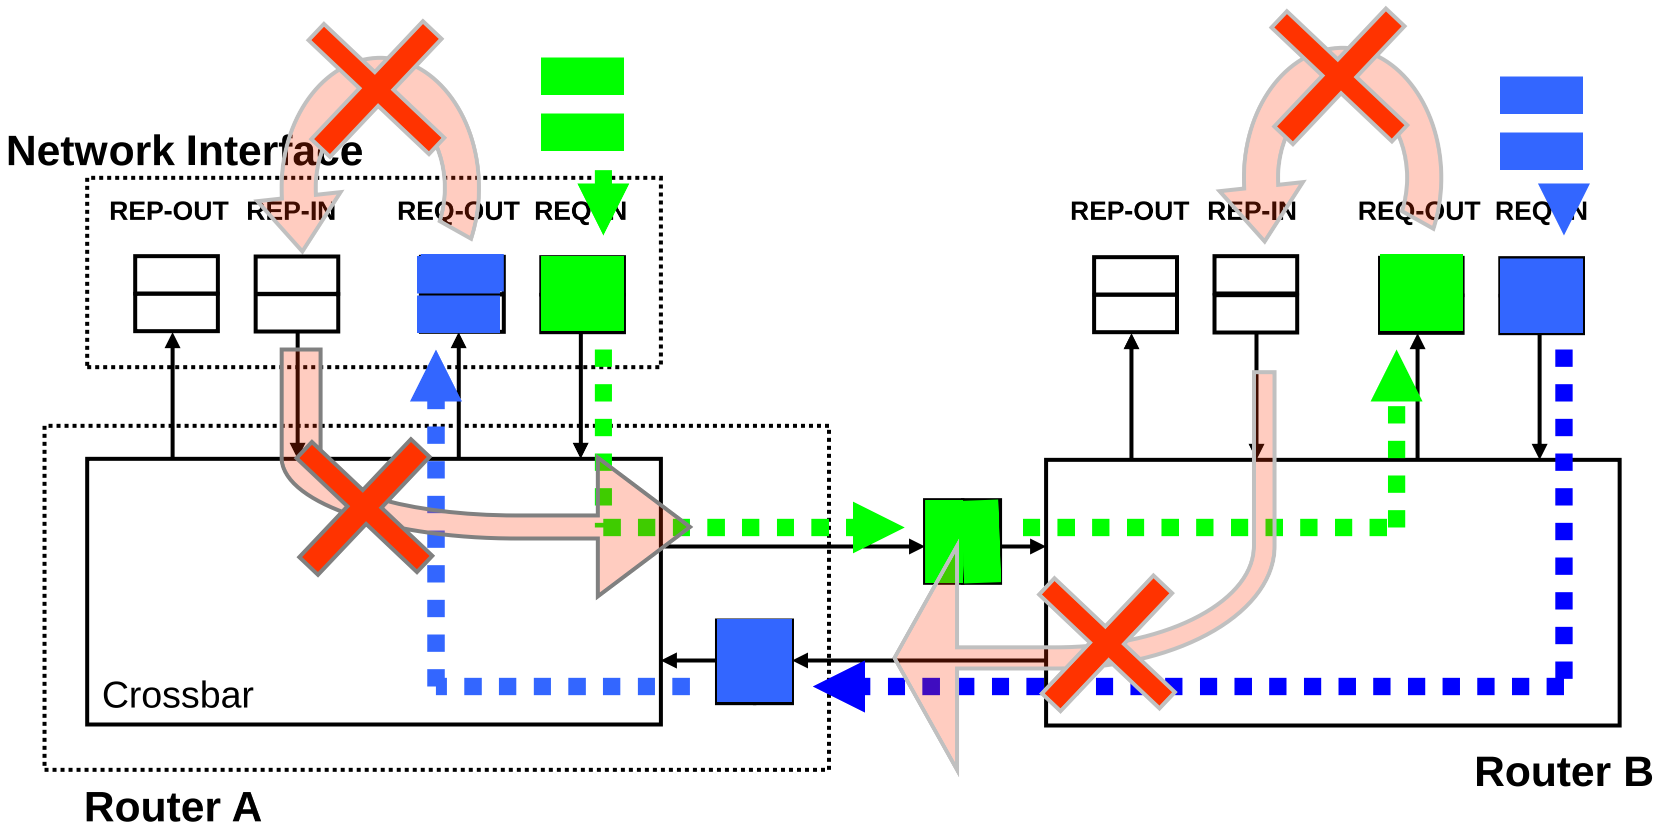

1 Router A and Router B flood the network with REQUEST messages
2 REQUEST messages are only attended if a REPLY can be generated

The hole leaved by an attended REQUEST is occupied by another REQUEST.
3 DEADLOCK: No more REQUESTS can be attended and REPLIES cannot reach destination.

Reactive Traffic

A widely utilized solution to avoid this problem is buffer replication. REQ and REP travel through different buffering resources (virtual networks).

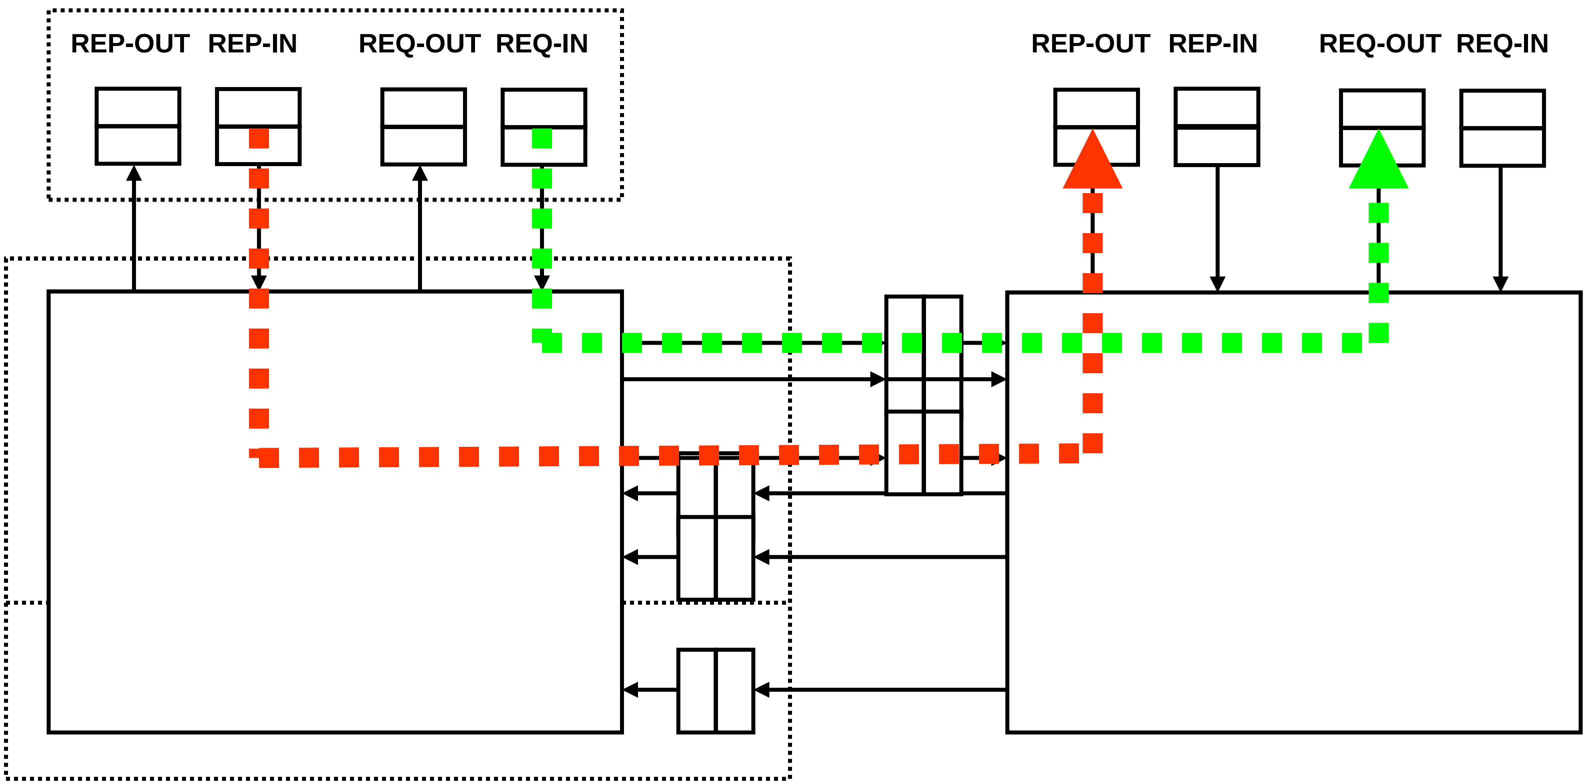

Reactive Traffic

Path replication solves end-to-end deadlock problem, but can seriously affect other relevant design aspects, such as area, complexity, power.

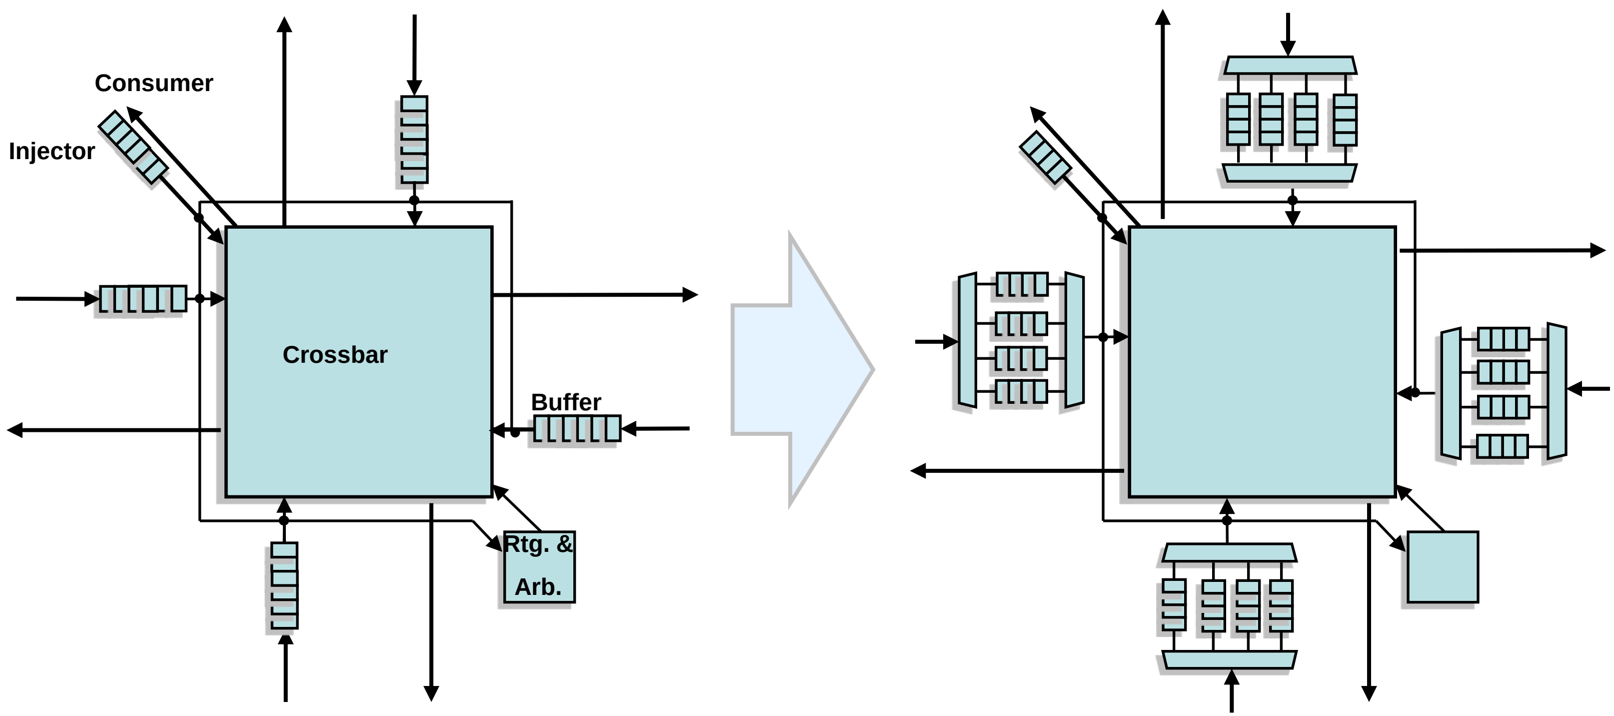

1 Message Type **4 Message Types**

Alpha 21364 router: 7 message types

Previous work: The Rotary Router

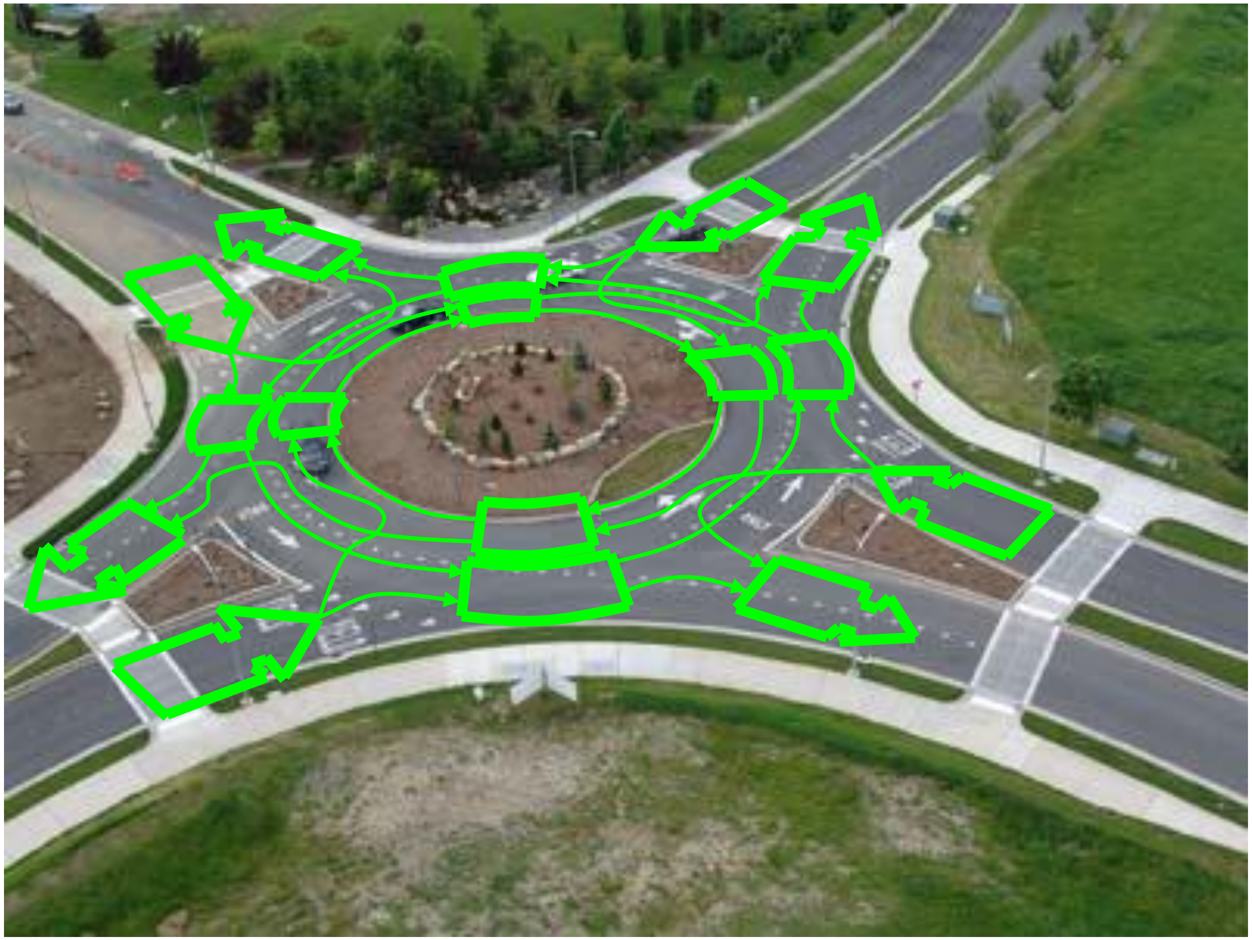

[REF] P. Abad, V. Puente, P. Prieto, J.A. Gregorio, "Rotary Router: An Efficient Architecture for CMP Interconnection Networks", International Symposium on Computer Architecture (ISCA), 2007.

Rotary Router Sketch

Input Stage

- ✓Packet Pre-Routing.
- ✓Ring Selection.

Buffering Segment Stage

- ✓Packet movement.
- ✓Output arbitration.

Output Stage

- ✓Flow control.
- ✓Packets storage.

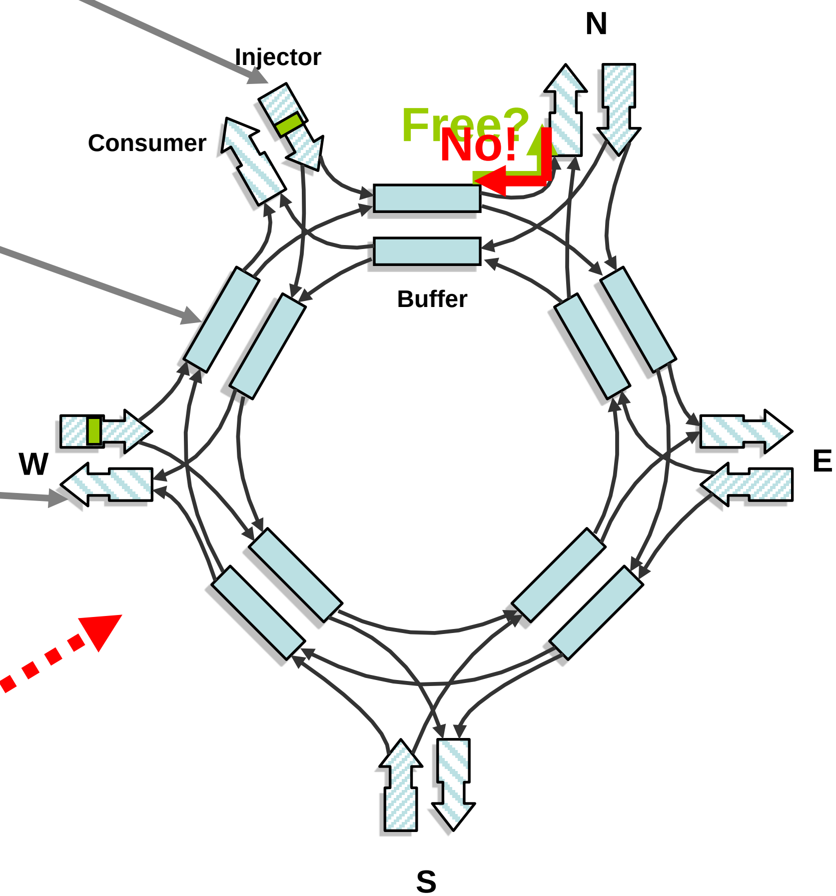

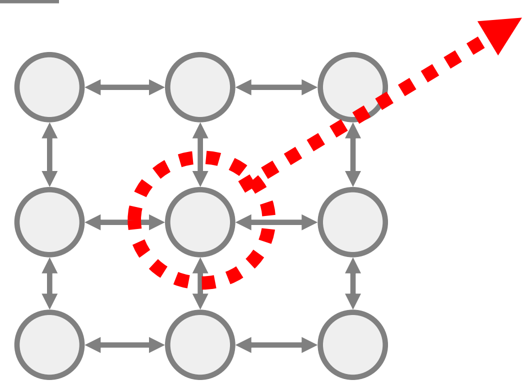

Rotary Router Advantages

- **Head of Line Blocking Avoidance.**
- **Improved Buffering utilization.**
- **Adaptive routing without virtual channels.**
- **Centralized structures avoidance (Xbar, Arbiter).**
- **Topology agnostic Deadlock avoidance Mechanism.**

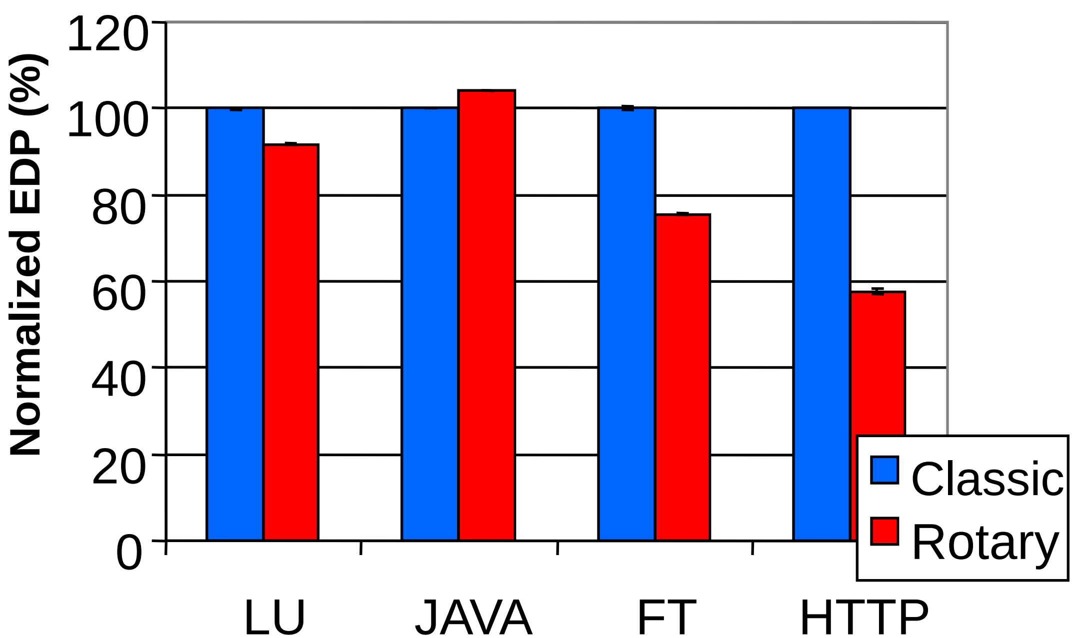

Reactive Traffic

Continuous movement of packets inside the router rings allows the Rotary Router to implement a solution to end-to-end Deadlock without requiring path replication.

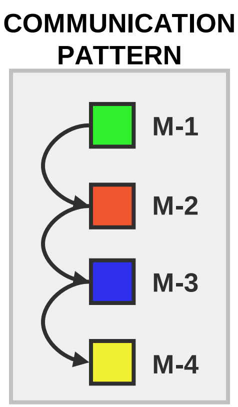

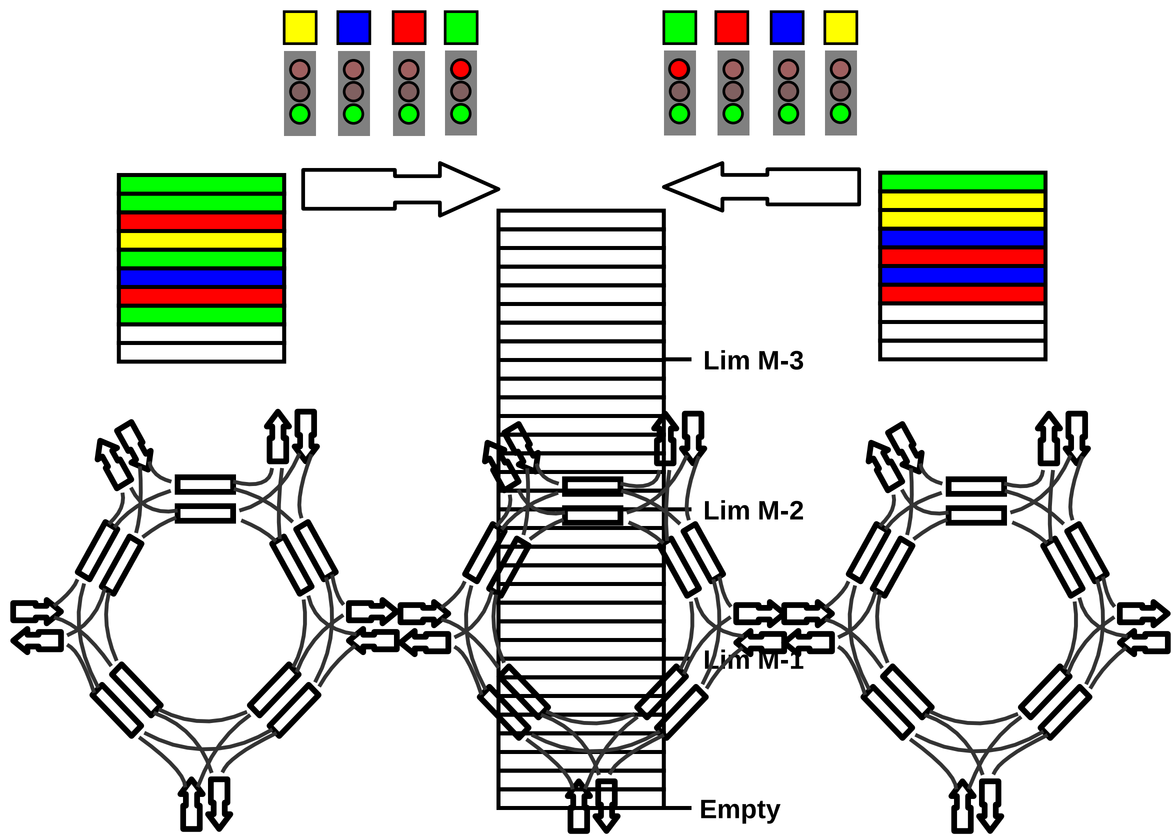

Reactive Traffic

Continuous movement of packets inside the router rings allows the Rotary Router to implement a solution to end-to-end Deadlock without requiring path replication.

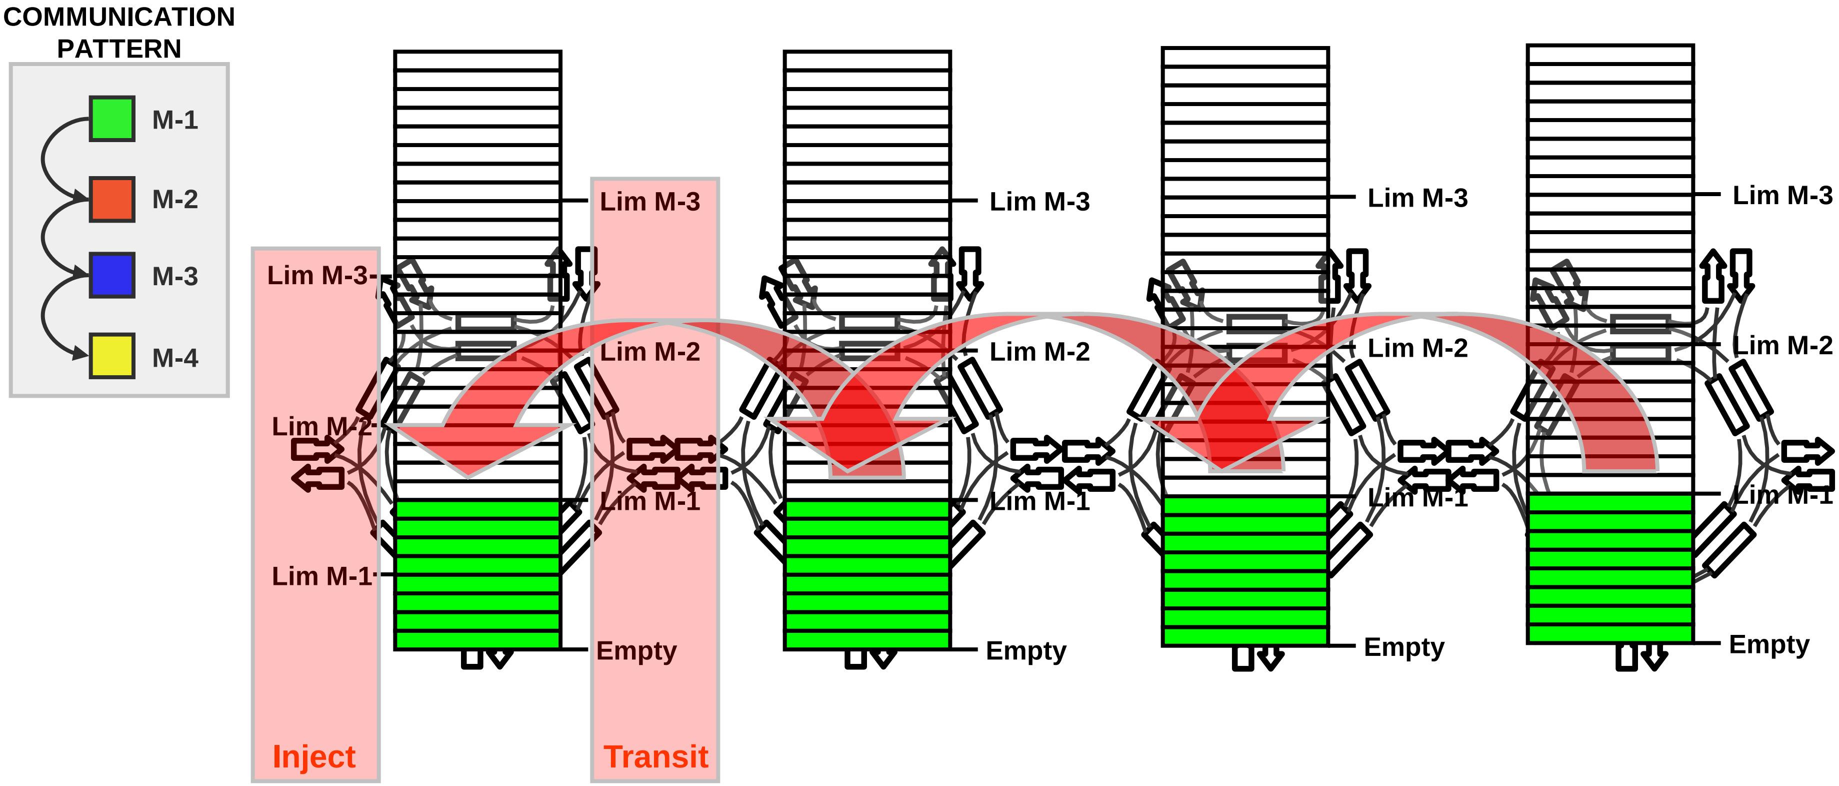

Outline

- Motivation
- Reactive traffic
  - End-to-end deadlock
  - Rotary solution
- In order delivery
- Evaluation
- Conclusions

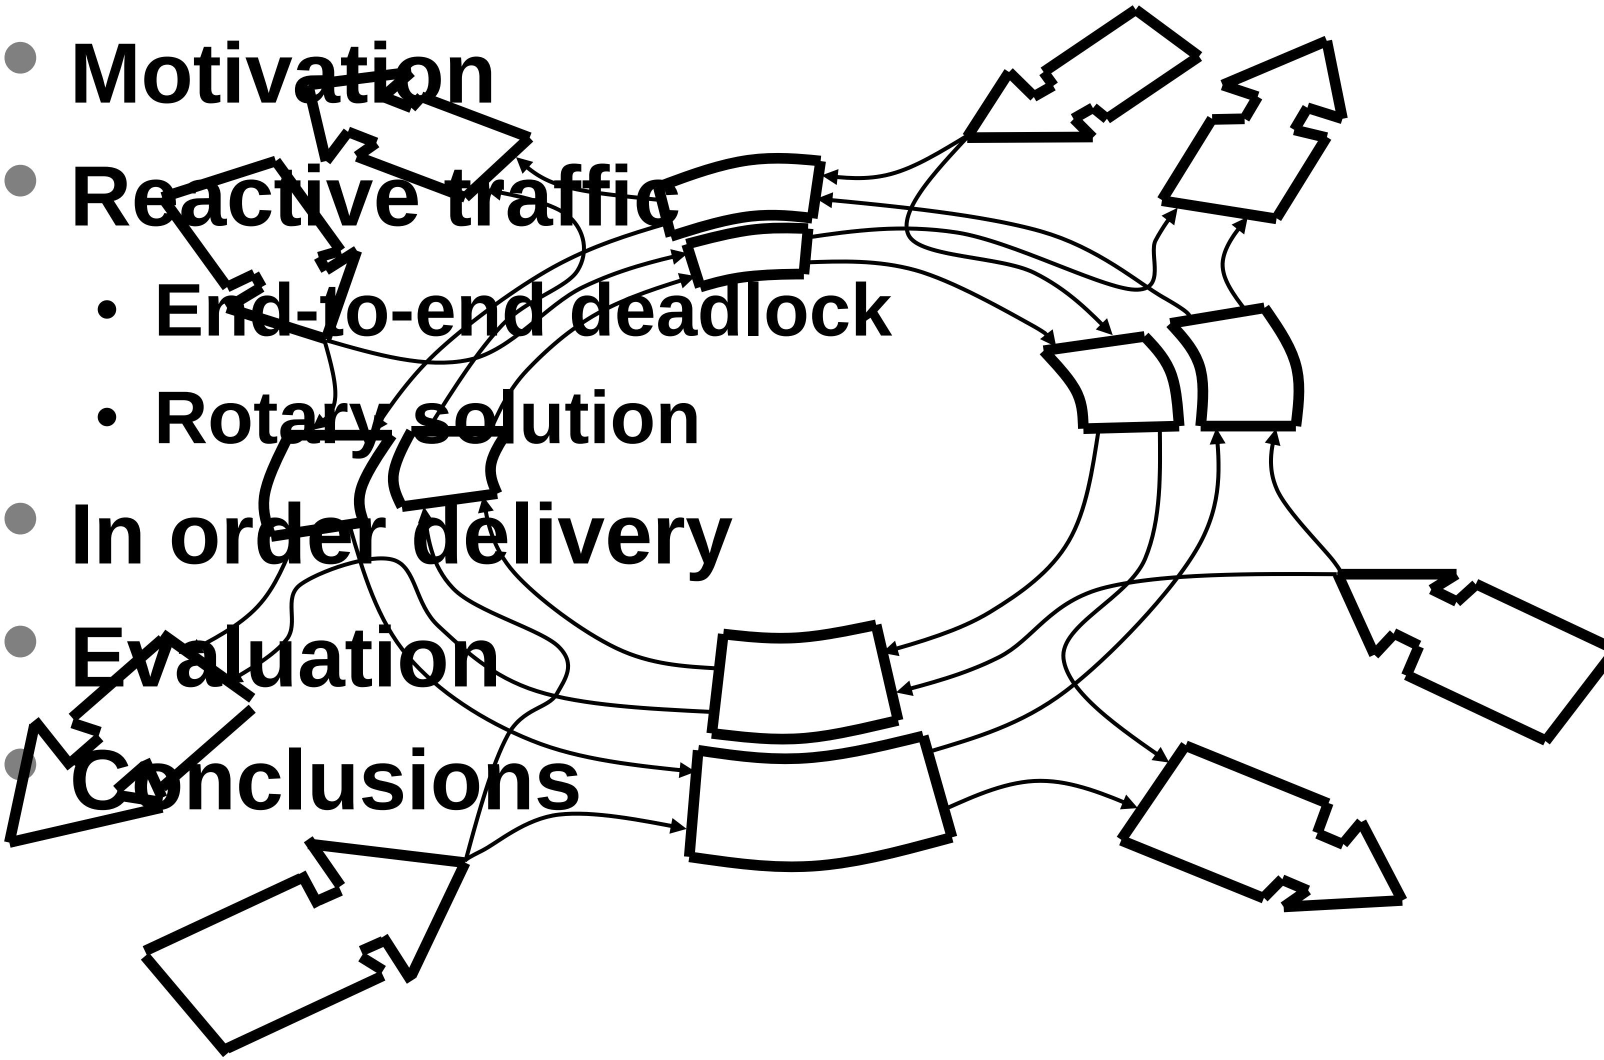

In-Order Delivery

- **This requirement is imposed by some memory coherence protocols (v.gr. Token coherence protocol) or maintenance tasks.**
- **In these cases, only specific transactions need to be ordered (v.gr. Persistent request deactivation)**
- **Ordered messages represent only a small portion of total network traffic (~5% of total traffic).**

In-Order Delivery

Fulfilling this requirement is extremely simple for input buffered routers.
It becomes a challenge for the Rotary Router:

-Adaptive routing allows inter-router packet overtaking.
-Internal router rings allow intra-router overtaking.

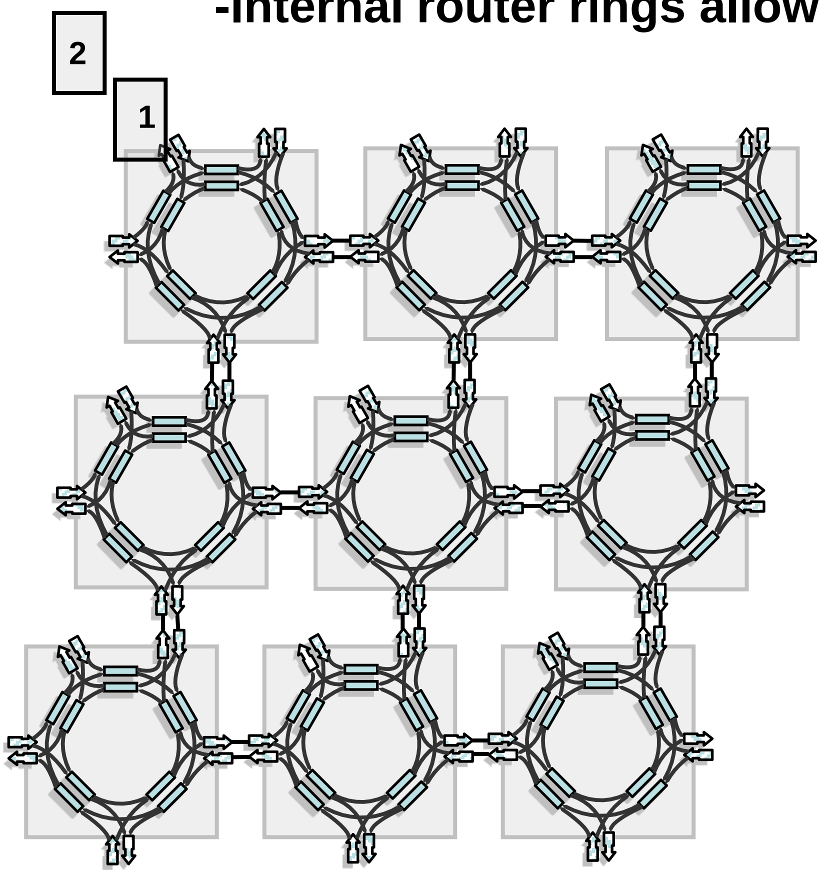

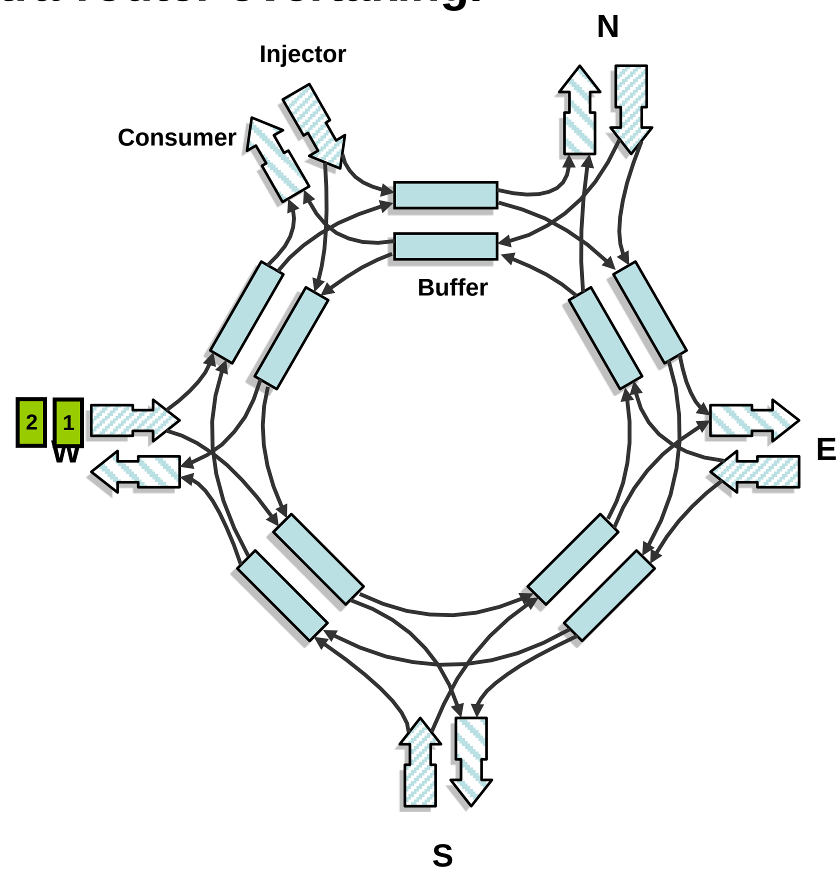

In-Order Delivery

Inter-router overtaking is avoided through specific Routing decisions for in-order messages:

-wraparound links will be avoided (Mesh)

-Adaptive routing will not be allowed (DOR).

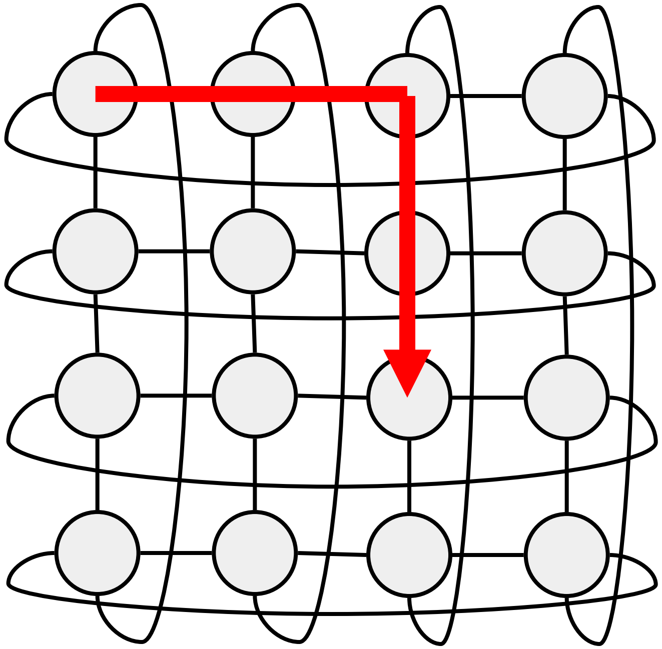

In-Order Delivery

Intra-router overtaking needs a special mechanisms to be avoided.

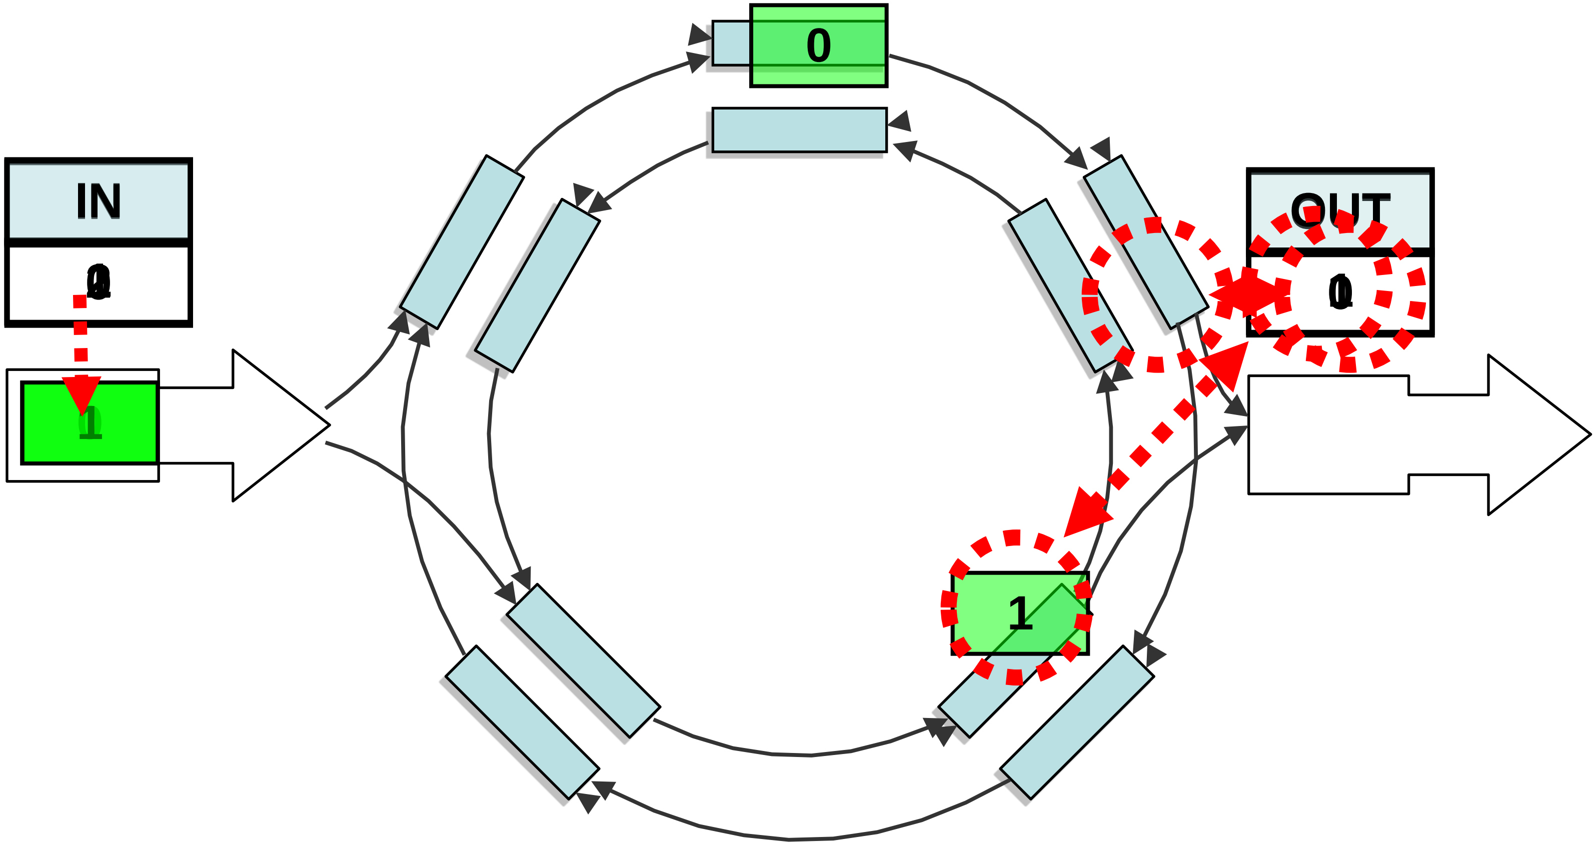

Outline

- Motivation
- Reactive traffic
  - End-to-end deadlock
  - Rotary solution
- In order delivery
- Evaluation
- Conclusions

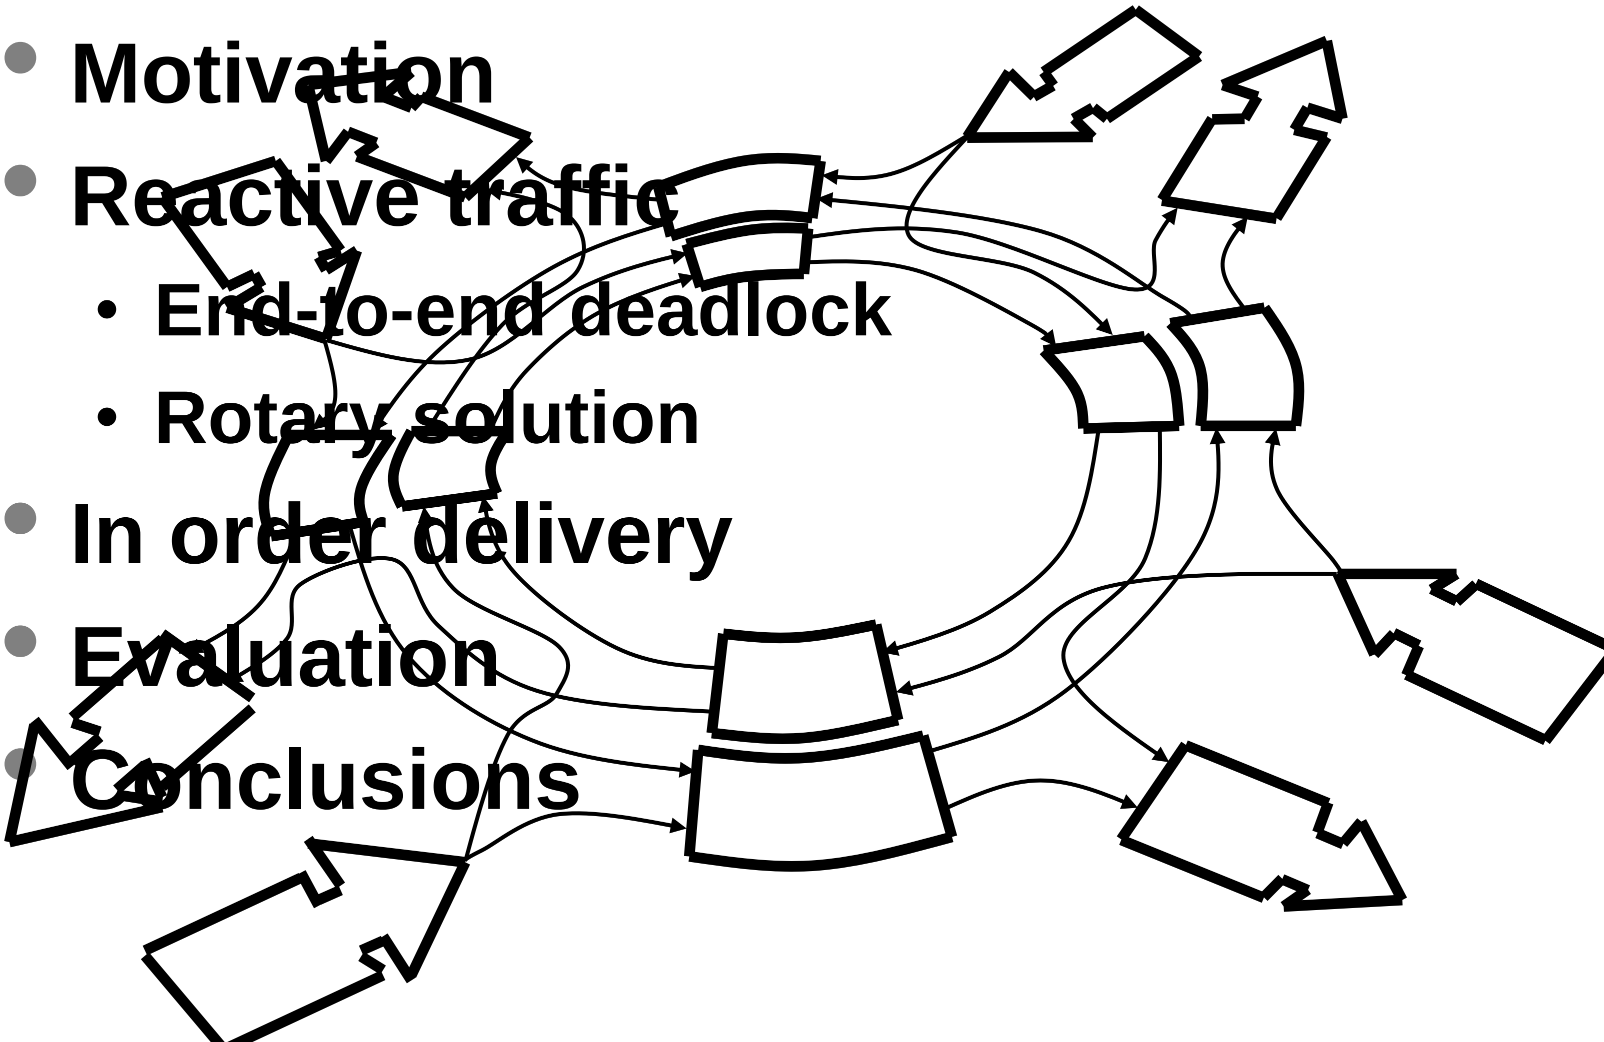

Performance Evaluation

- Compared to three different routers

Adaptive Bubble Router **Deterministic Router** **Low Latency Router**

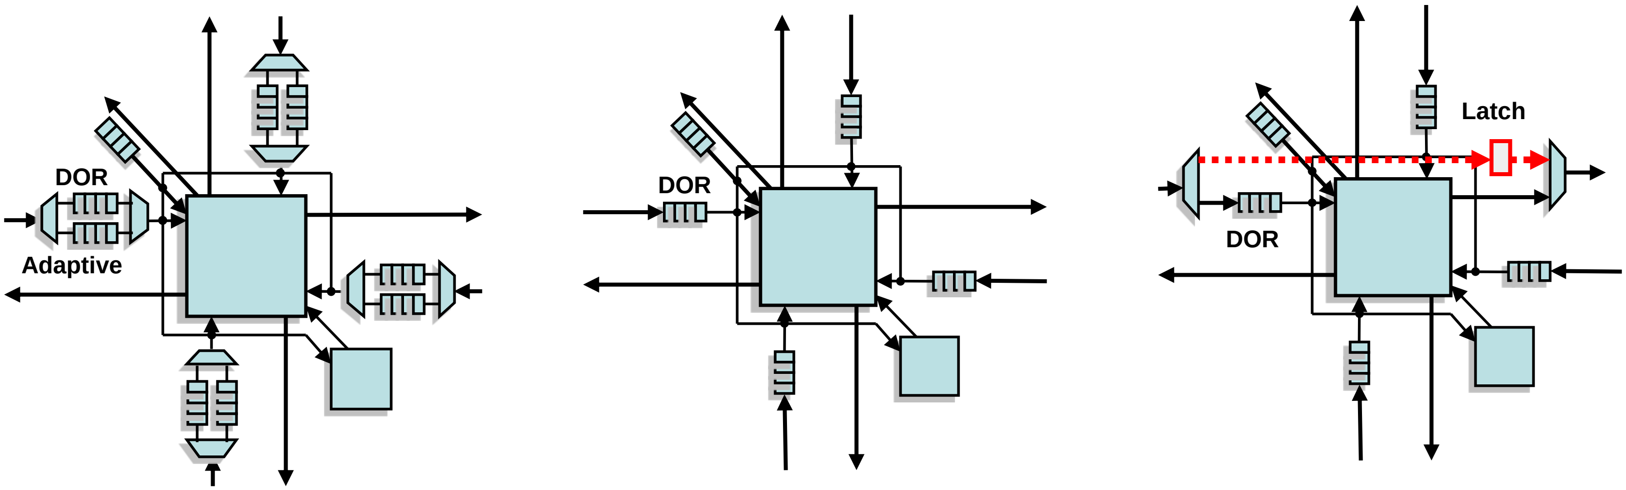

Performance Evaluation

- Synthetic Traffic Patterns
- Real Workloads
  - GEMS + SICOSYS.

Number of cores	16	**Main Memory**	4GB, 260 cycles, 320 GB/s
L1 I/D cache	Private, 32KB, 2-way, 64Bytes block, 1-cycle	**Command size**	16 bytes
L2 cache	SNUCA, 16x16 banks, 4 per router	**Network Topology**	8x8 Torus
L2 cache bank	128KB, 16-way, 3-cycles, Pseudo LRU, 64 Bytes block	**Network Link**	128 bits / 1 cycle latency

Performance Evaluation

- Synthetic Traffic Patterns
  - 5 message types.
  - 32.000 messages of each type delivered.
  - Low-lat topology: Mesh.

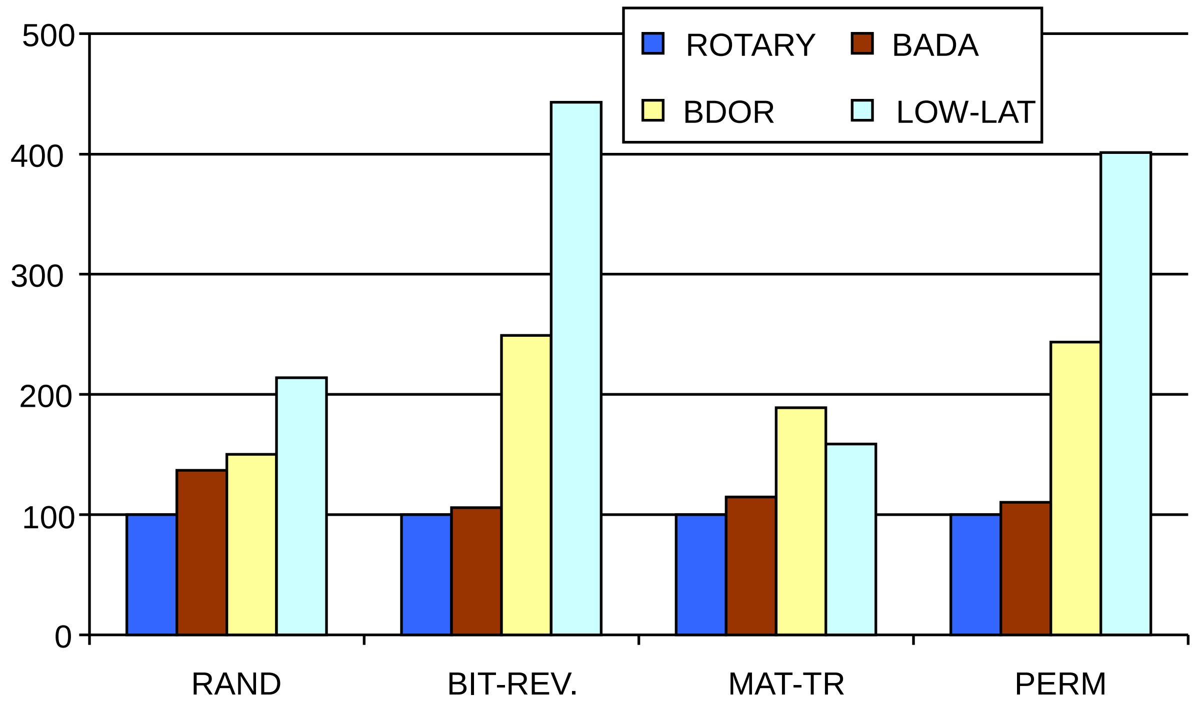

Performance Evaluation

- Real Workloads
  - Transactional & Scientific applications.

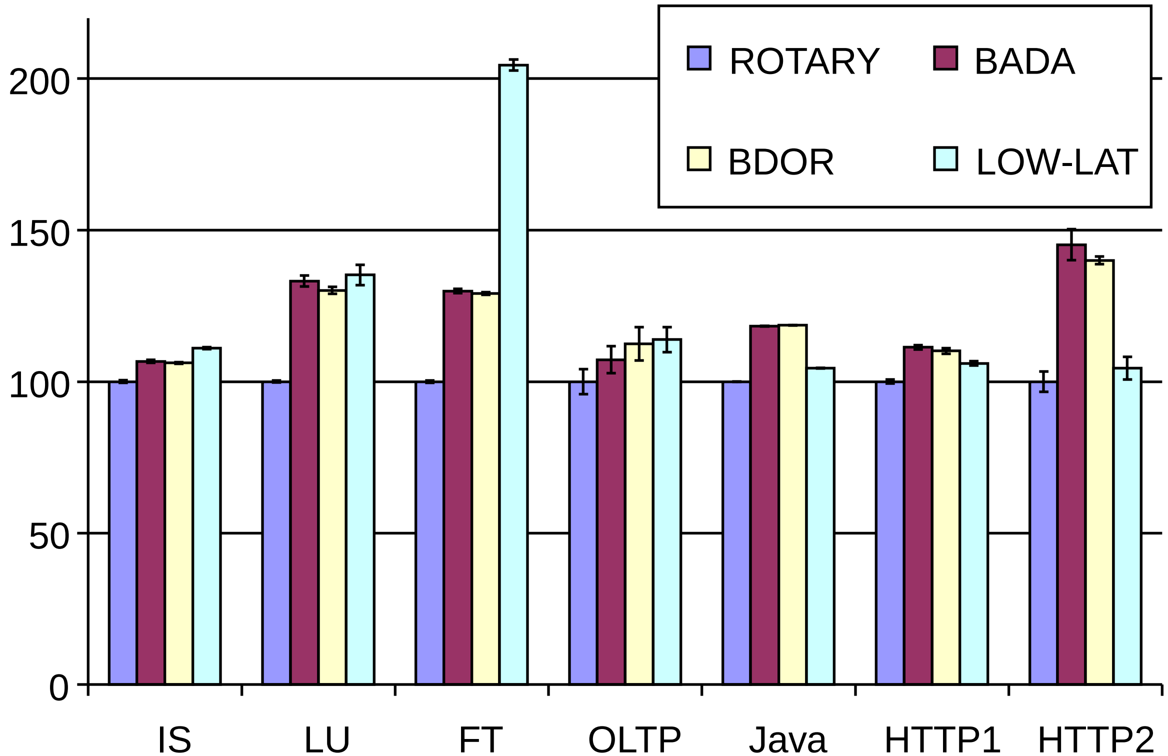

Outline

- **Motivation**
- **Reactive traffic**
  - **End-to-end deadlock**
  - **Rotary solution**
- **In order delivery**
- **Evaluation**
- **Conclusions**

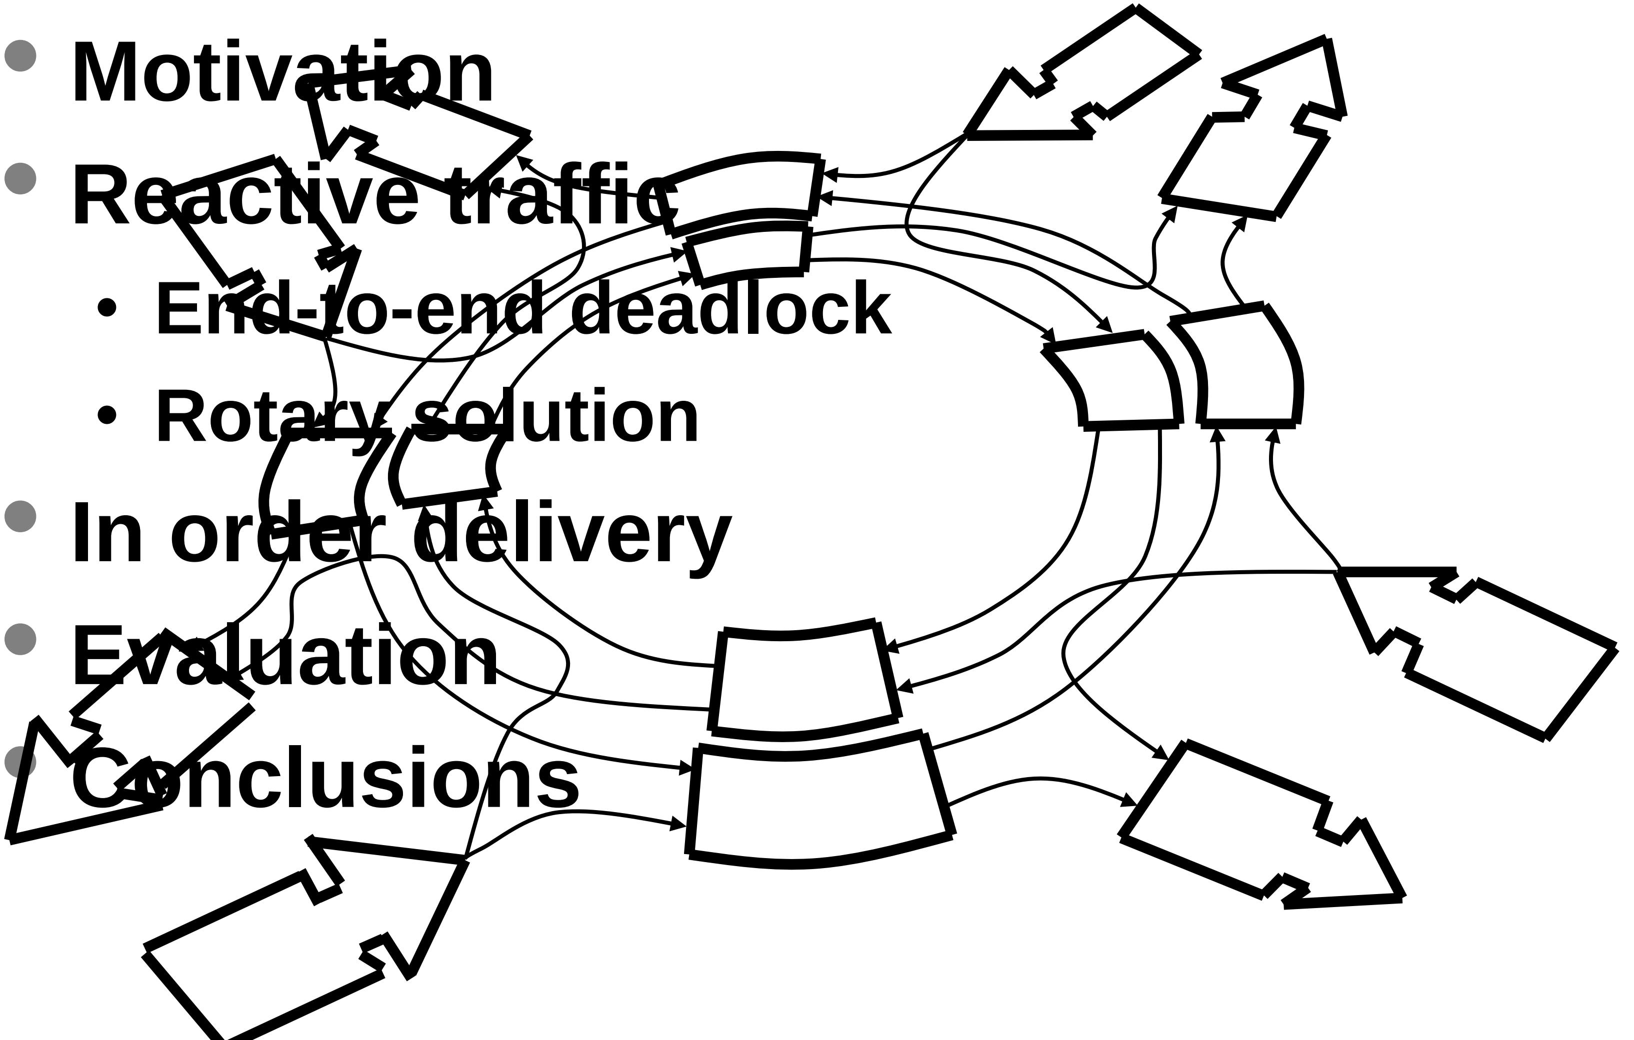

Conclusions

- The Rotary Router has been the base to implement a mechanism able to deal with end-to-end deadlocks.
- This mechanism does not require path replication.
- We solve in-order delivery with a simple method which requires few extra hardware.
- Flexible buffer utilization allows our router to obtain better performance results.

Questions?