A Lightweight Fault-Tolerant Mechanism for Network-on-Chip

Michihiro Koibuchi*, Hiroki Matsutani**
Hideharu Amano**. Timothy Mark Pinkston***

*National Institute of Informatics, Japan/JST, Japan
*Keio University, Japan
***University of Southern California

Background

- Improvement of the die yield
  - Circuit Level
  - Architecture Level

  e.g. Cell Brd. Eng.
    - Play Station 3：7SPE
    - HPC-Purpose: 8SPE

- Fault tolerance of the communication on multi-core systems
  - Lightweight mechanism

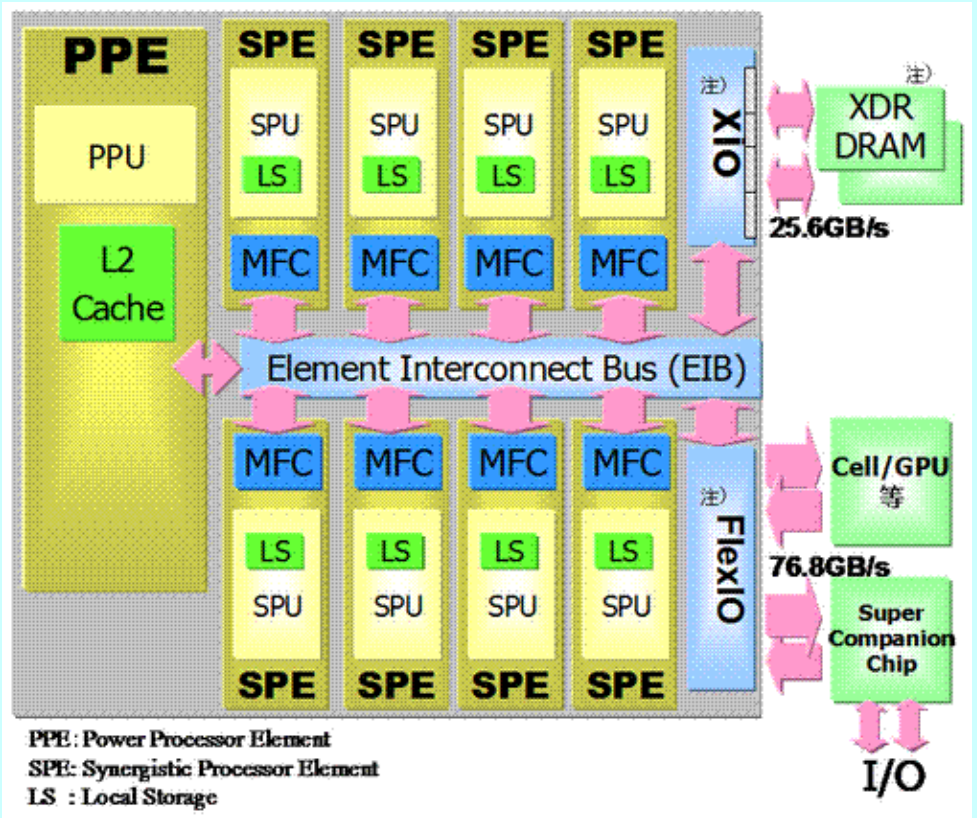

Cell Broadband Engine

Outline

- Fault patterns on Network-on-Chip (NoC)
- Default-backup path mechanism (DBP)
  - maintains the connectivity of all healthy PEs, even if the network includes hard faults

    Objective

    Provide a highly reliable network using lightweight hardware！
- Evaluation
  - Energy
  - Amount of Hardware
  - Throughput

Network-on-Chip (NoC)

- Processor Core
  - Largest component
  - Various fault-tolerant techniques
    - Resource sparing
    - Redundancy
- On-Chip Router
  - Area is not so large.
  - Infrastructure that affects on-chip communication
    - Duplication

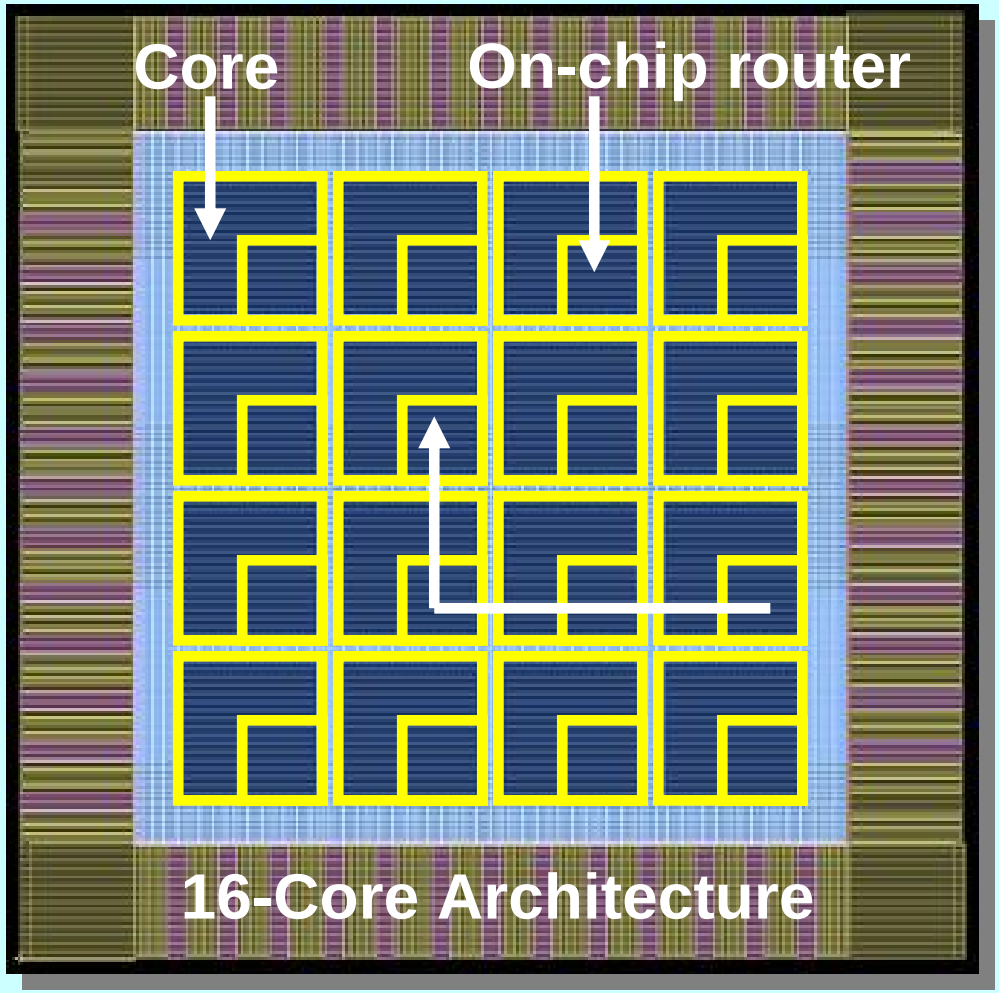

(*) Kyoto U/VDEC/ASPLA 90nm CMOS

Failures in Communication

- Transient Error (e.g. bit error)
  - Software layer is responsible, and recoverable
    - Link-to-link, and/or end-to-end [Murali,DToC05]
    - Error detection and/or error correction (e.g. CRC)
- Permanent Error (e.g. hard error)
  - System avoids using the failed modules

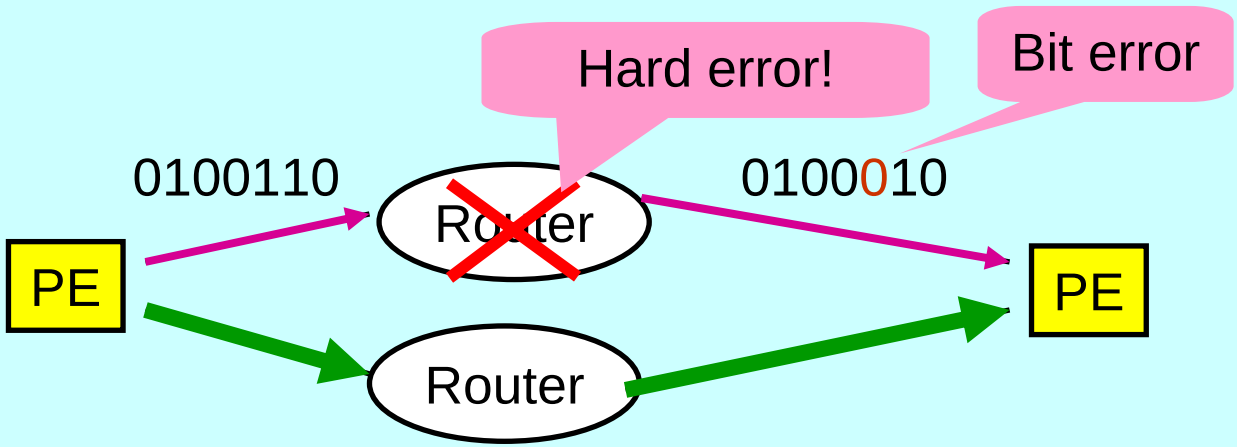

Existing NoC Fault-tolerant Techniques

Router Architecture

- Speculative Router [Kim ISCA06]
  - Providing fault-tolerance at input buffer, routing computation, and switch allocation unit.
- Dependability for misrouted packets [Thottethodi IPDPS03]
- Channel Reconfiguration[DallyText03, Soteriou ICD04]

Routing Paths

- Resource Sparing
- Dynamic Reconfiguration
- Fault-Tolerant Routing

Each Technique is resilient for portion of possible failures.

- Using them together enables high reliability! But, how about simplicity?

- Hard to recover crossbar failures

Outline

- Fault patterns on Network-on-Chip (NoC)
- Default-backup path mechanism (DBP)
  - maintains the connectivity of all healthy PEs, even if the network includes hard faults

    Objective

    Provide a highly reliable network using lightweight hardware！
- Evaluation
  - Energy
  - Amount of Hardware
  - Throughput

Motivation

- NoC Component
  - Router, Link Failure
    - disabling healthy local PEs
    - Segmentation of the network
  - NI Failure
    - Disabling the healthy local PE

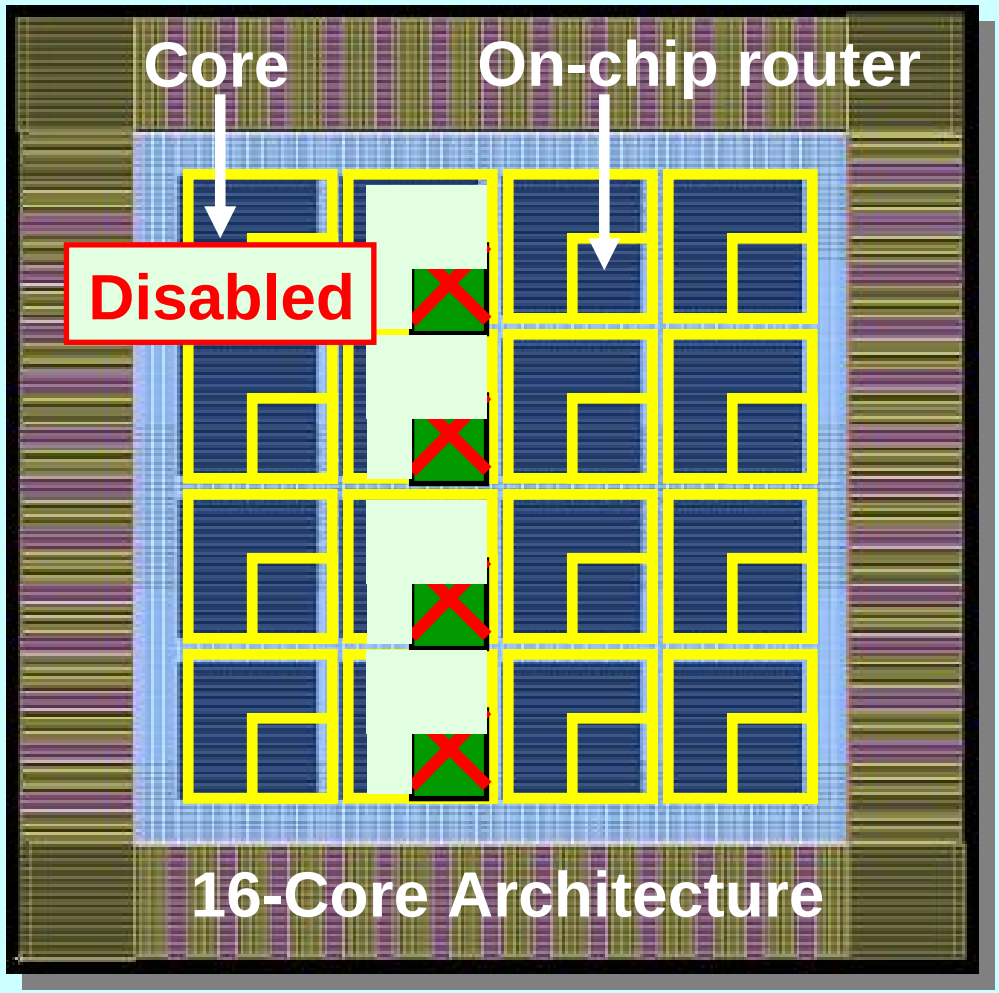

Unlike off-chip systems, a faulty module cannot be removed and replaced

The proposed lightweight fault-tolerant technique on a router maintains network connectivity of all the healthy PEs

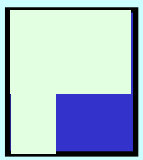

Disabled healthy PE

Conventional NoC Router (2-D mesh)

- 5-by-5 Router, channel bit-width (flit size) 64-bit

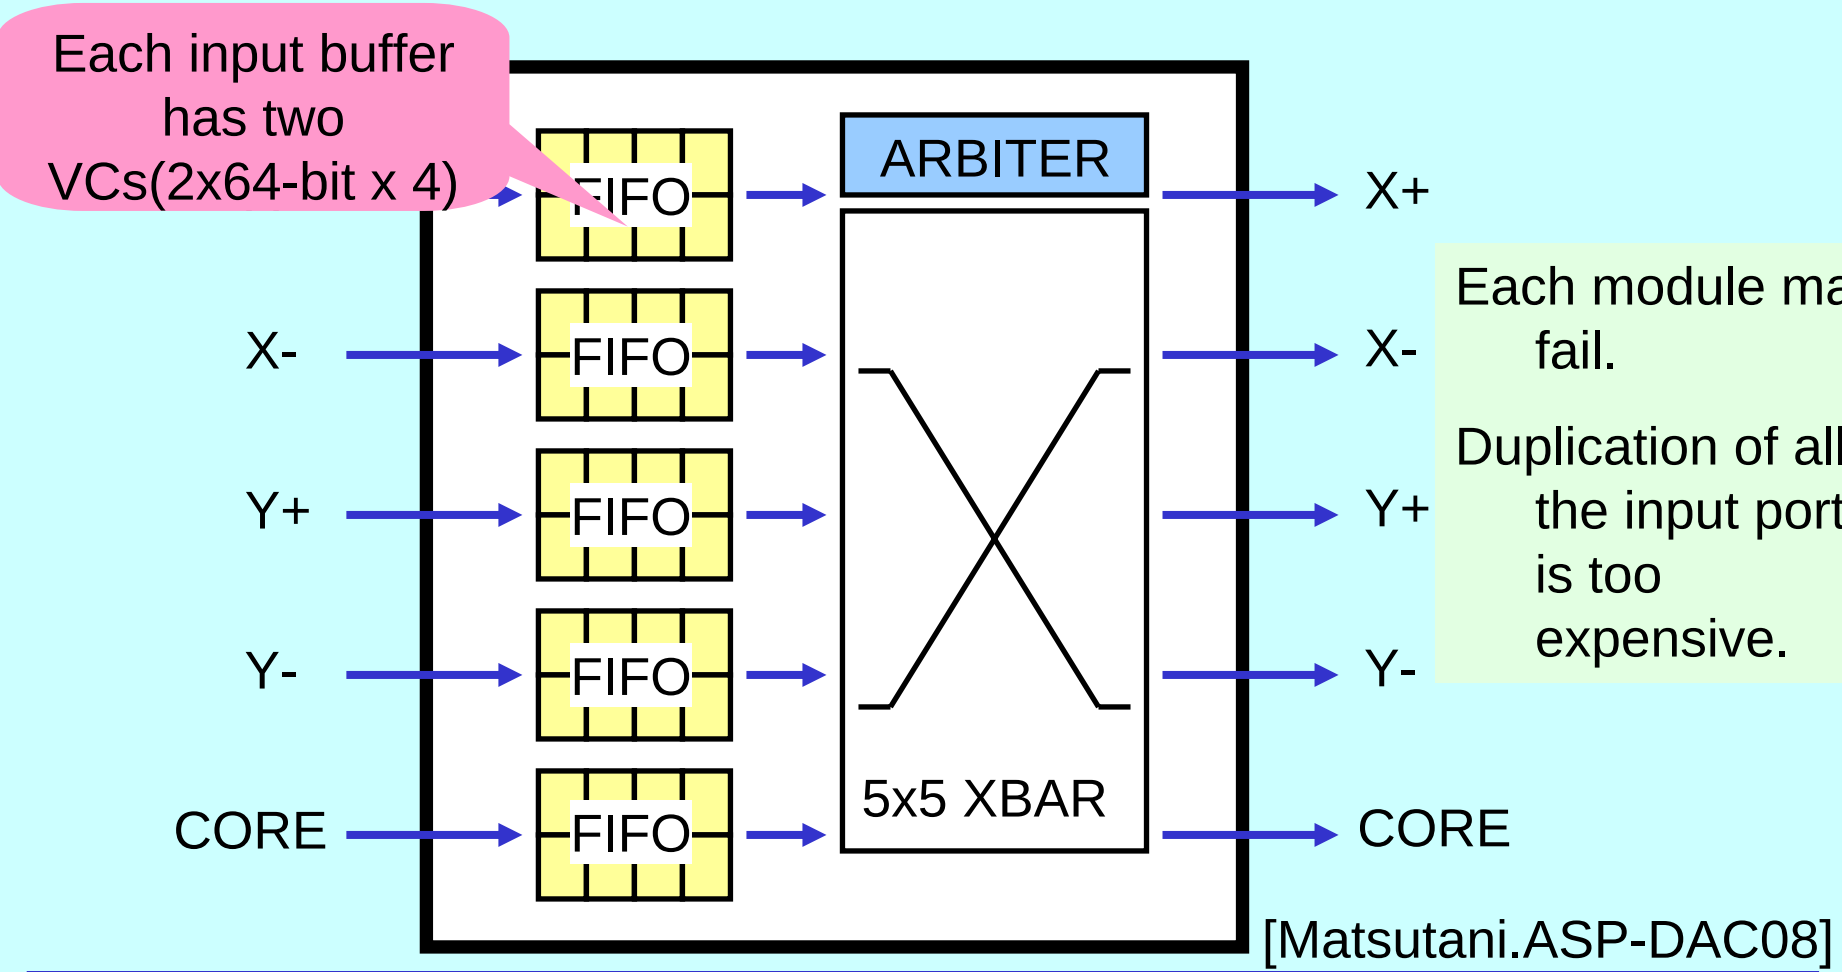

Each module may fail.

Duplication of all the input ports is too expensive.

[Matsutani.ASP-DAC08]

Area (after place and route) is 40～45 [KGate]; 75% is FIFO

Minimum Requirements for Communication

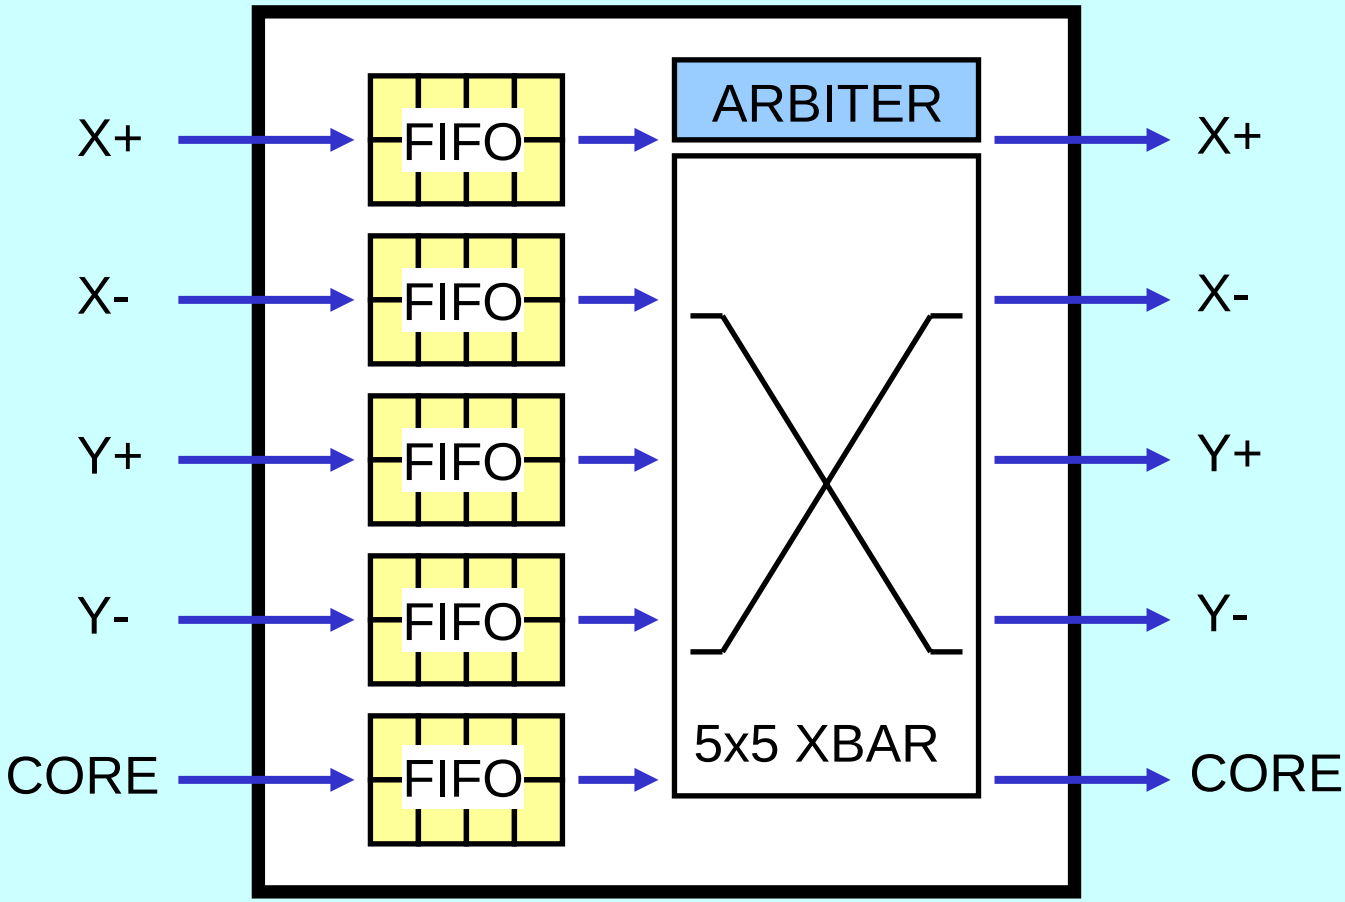

To communicate a local core with neighboring cores,

- It should send packets to at least one output port
- It should receive packets from at least one input port

Default-backup Path(DBP) Mechanism

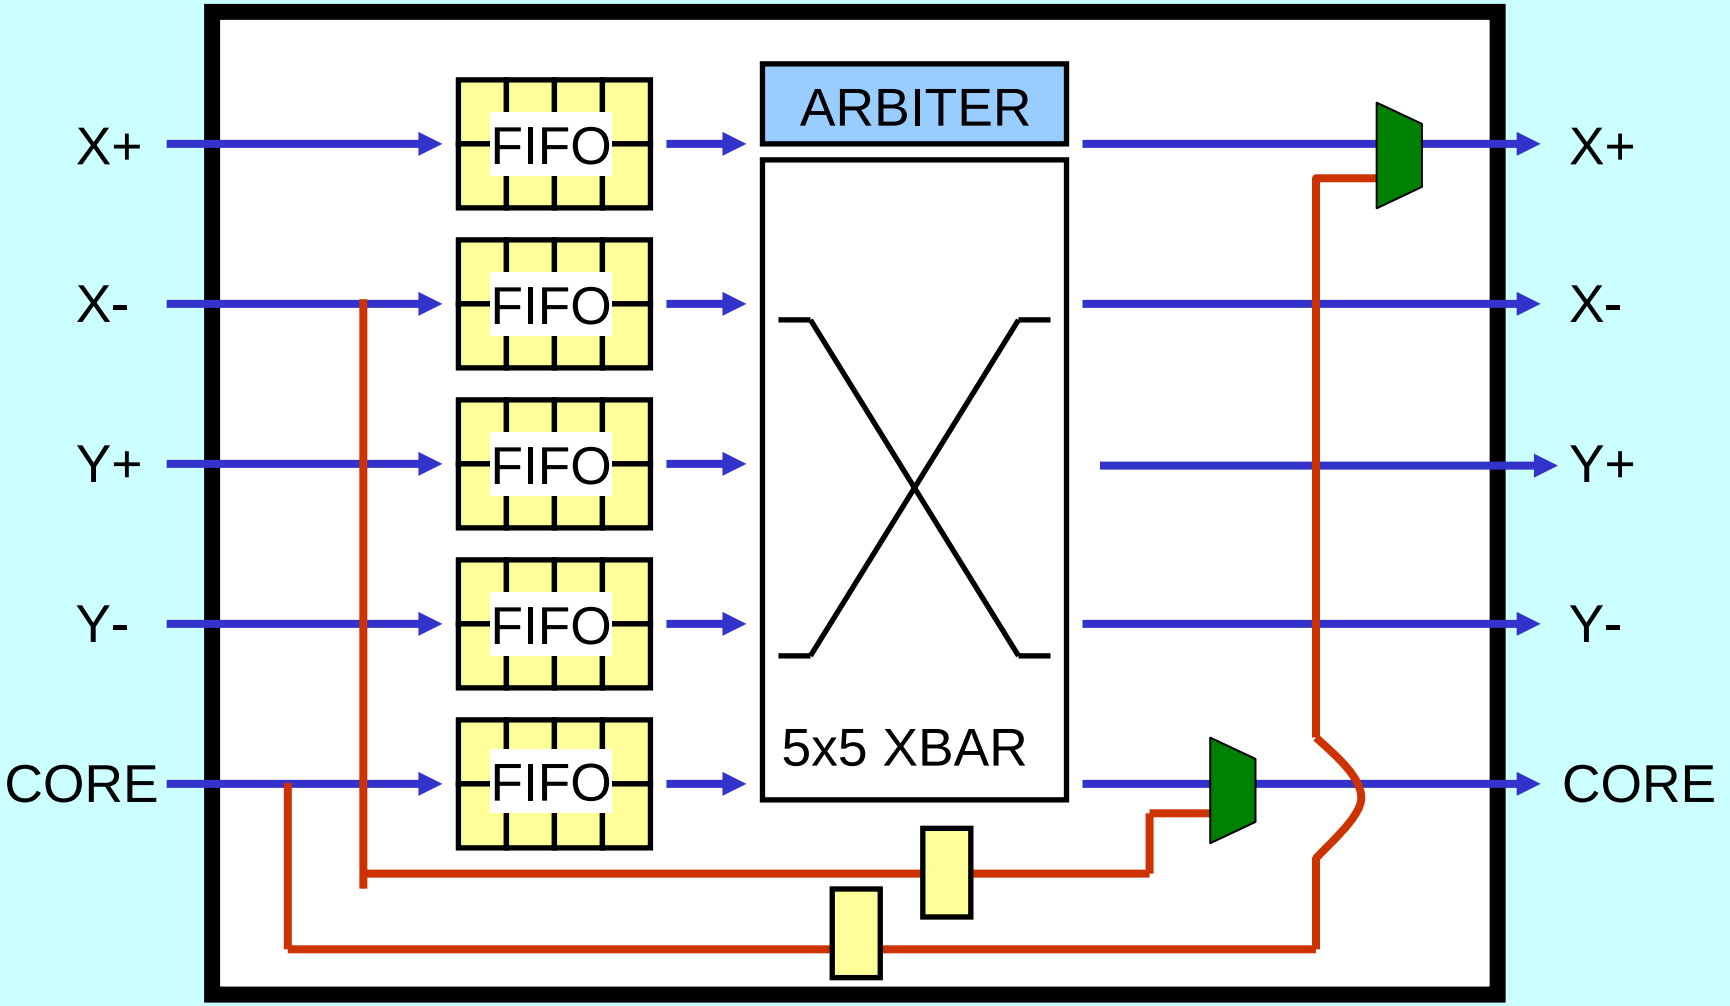

- A local core can send packets to at least one output port
- A local core can receive packets from at least one input port

Default-backup Path(DBP) Mechanism

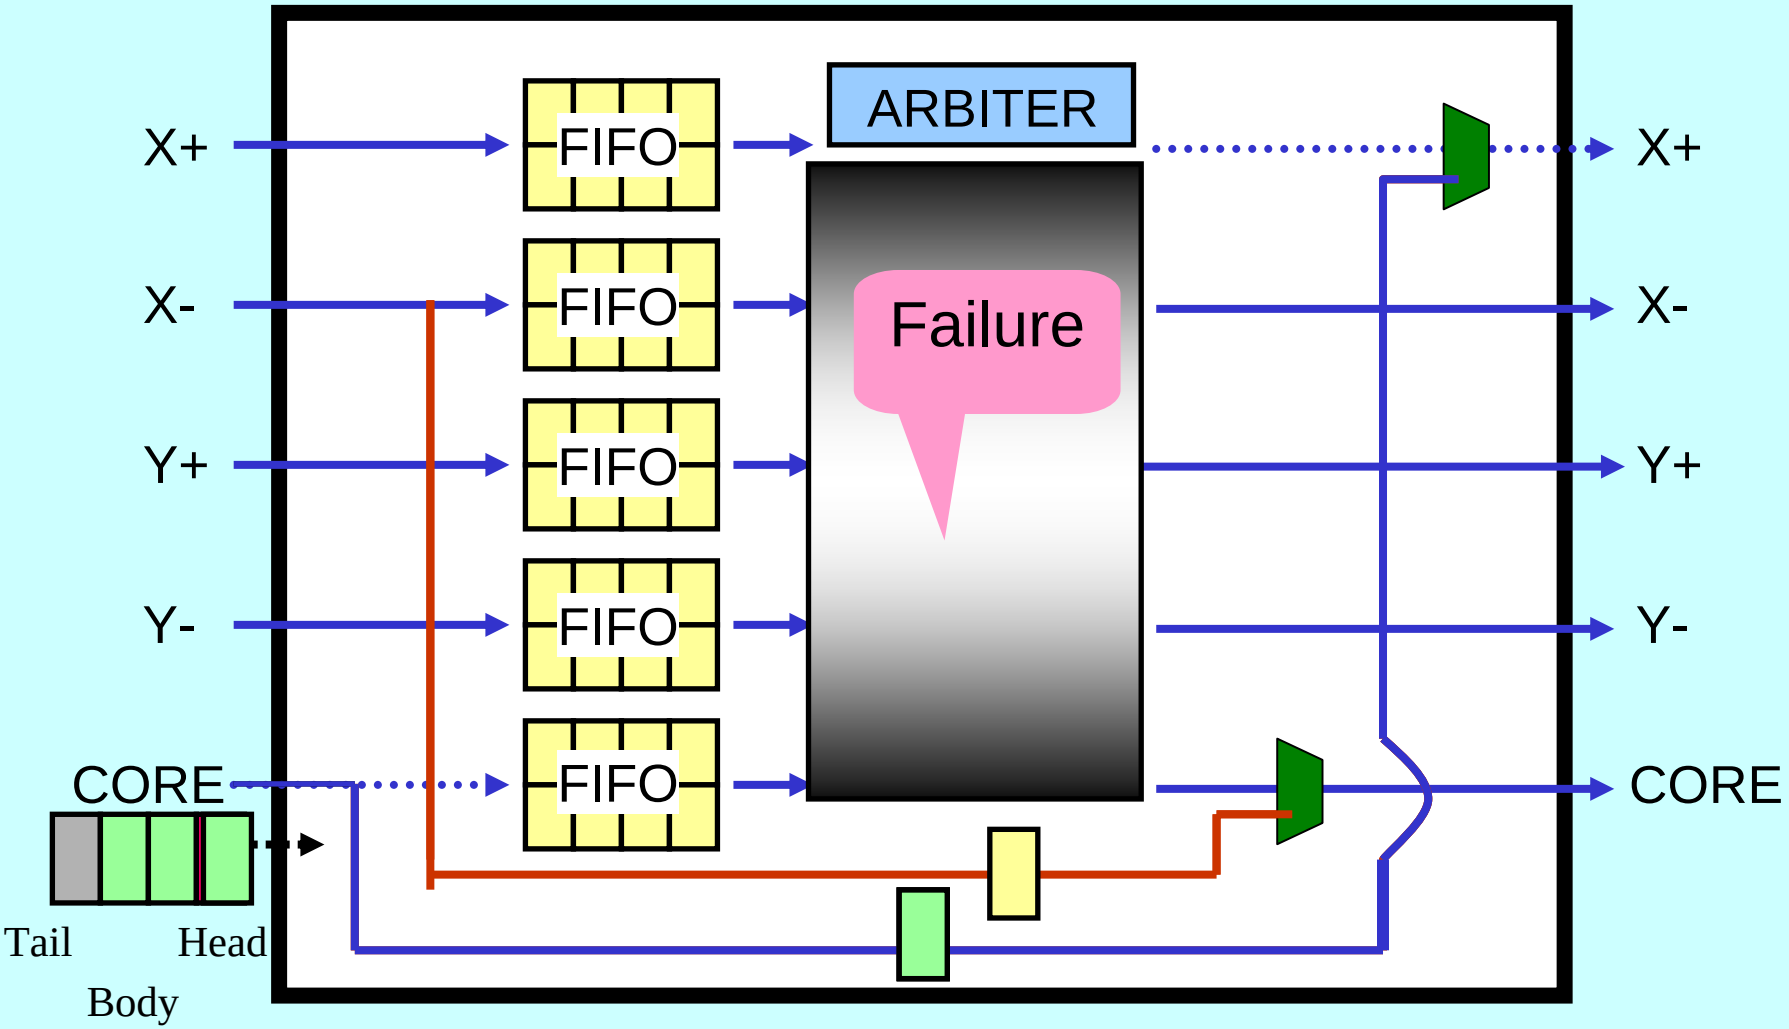

- A local core can send packets to at least one output port
- A local core can receive packets from at least one input port

Behavior of the DBP Mechanism (within a Router)

- Cores can communicate with each other, even if router modules fail
  - maintain packet transfers from X- direction, o X+ direction

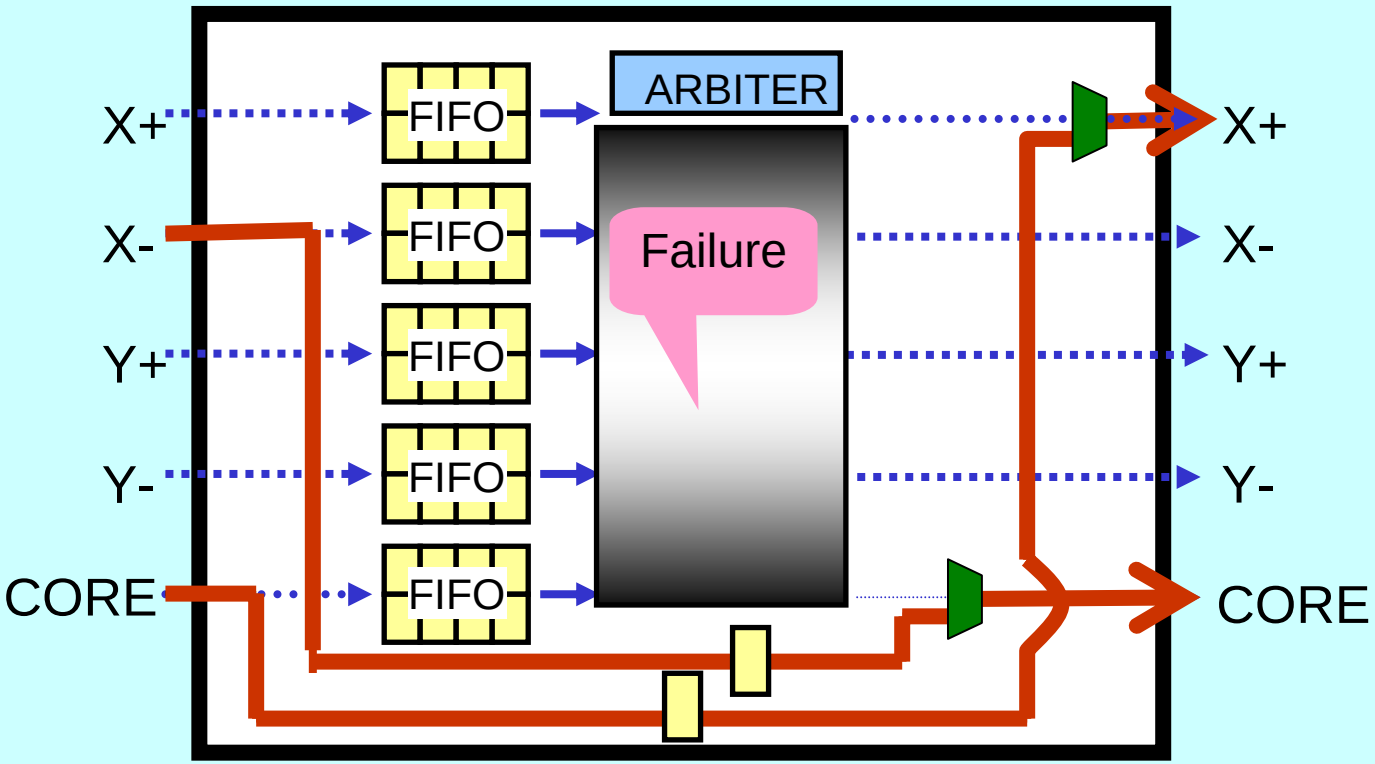

Behavior of the DBP Mechanism

（bypassing Xbar and NI faults）

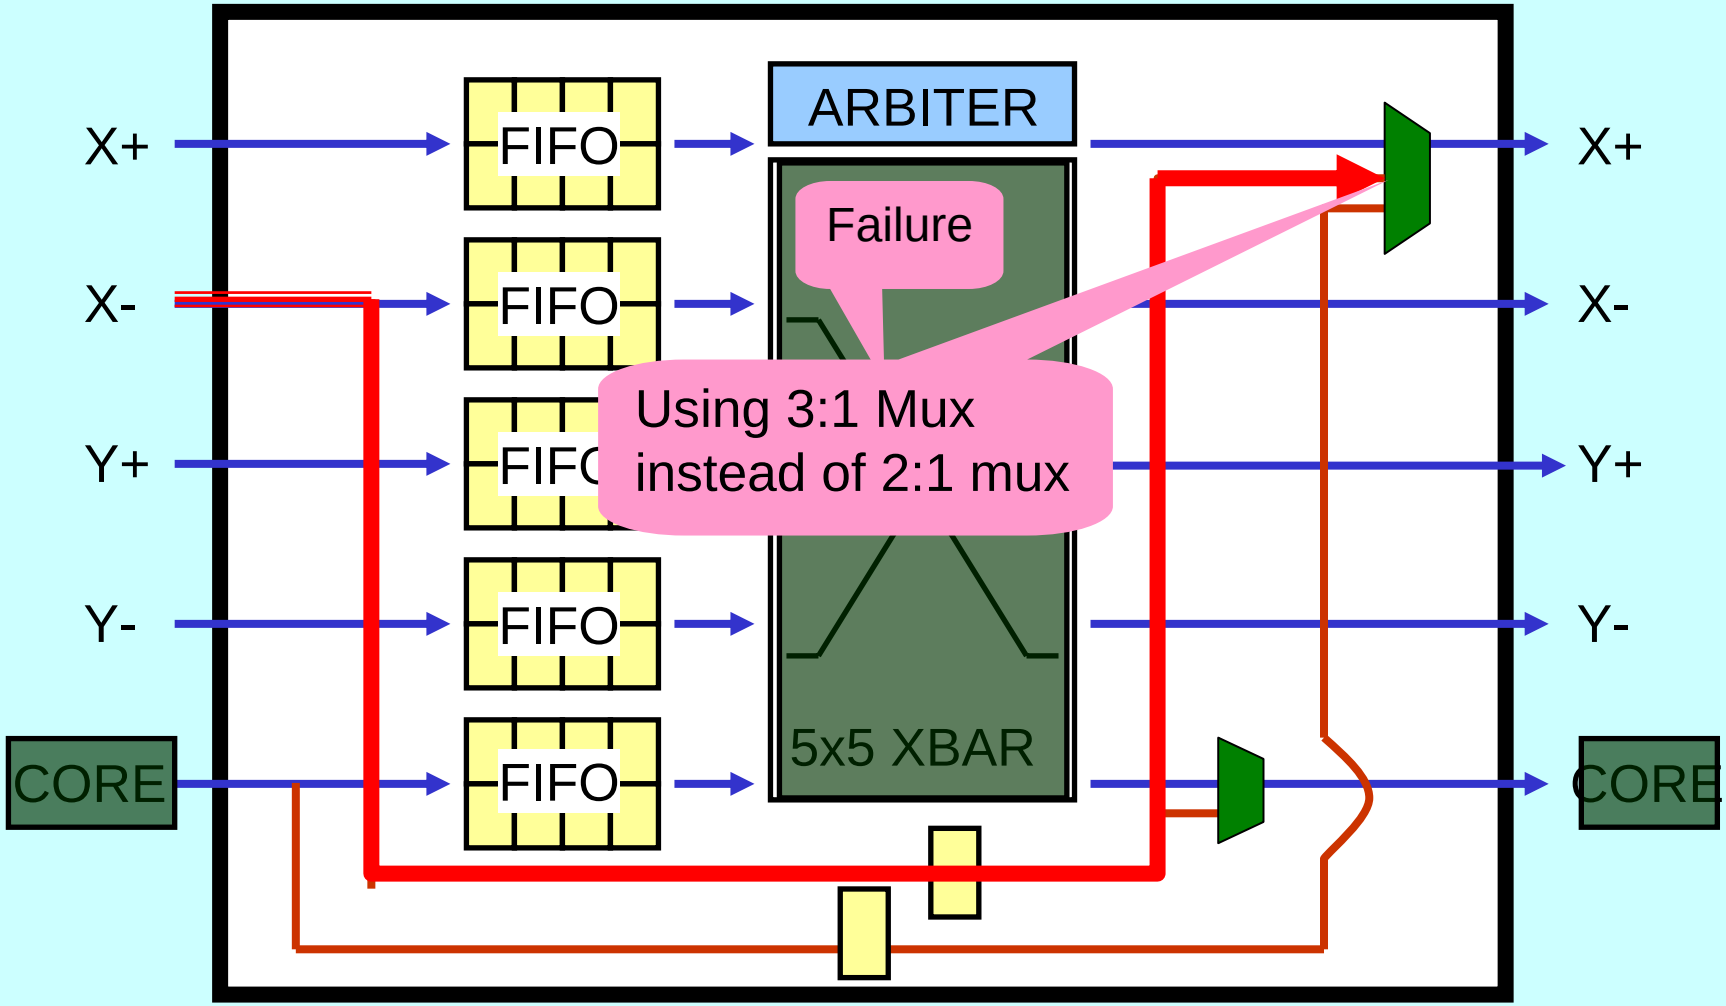

Another Issue: Network Connectivity

- Router, link failure
  - Disabling healthy local PEs
  - Segmentation of the networks
    - may disable all the PEs

The DBP mechanism provides reliability not only on intra-router datapath but also on routing paths

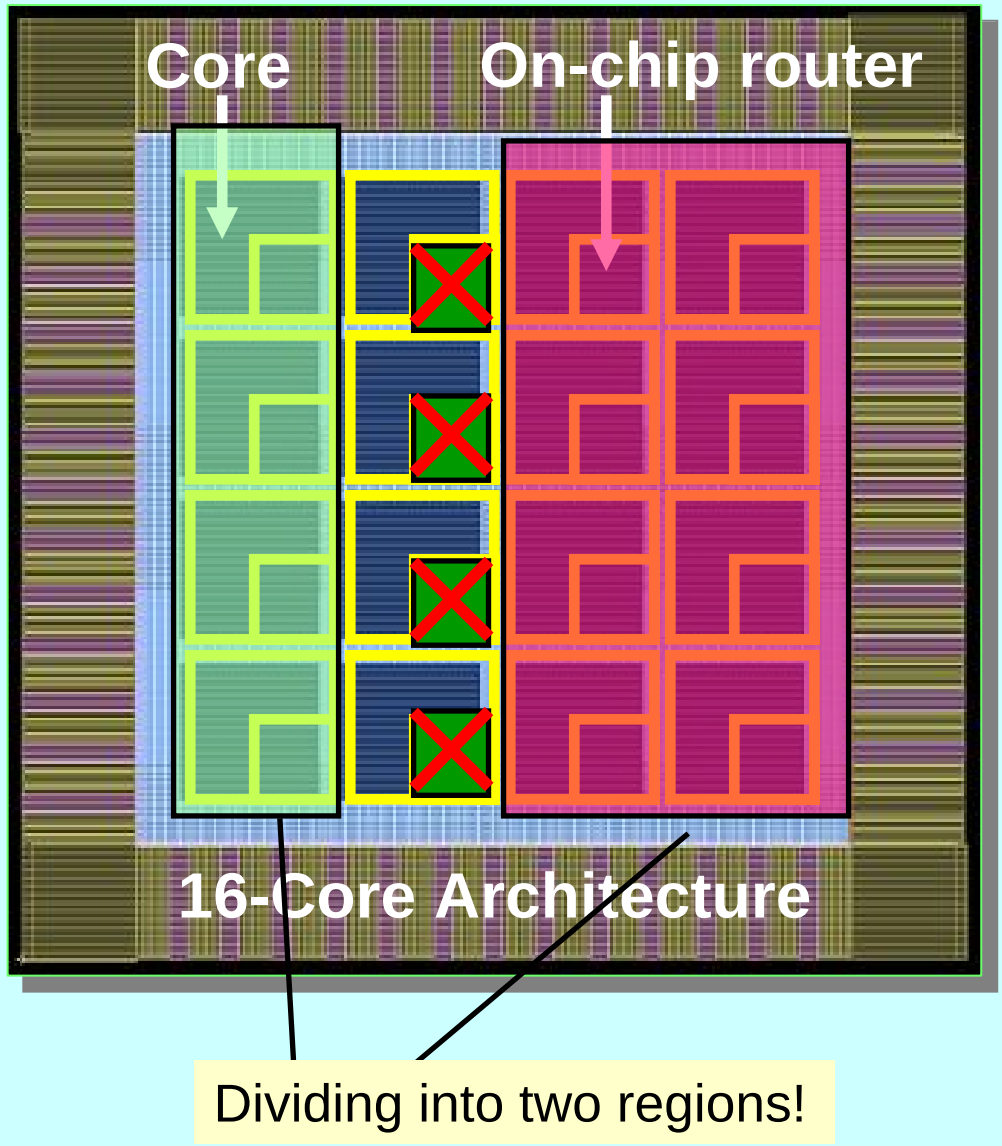

DBP Mechanism (inter-router behavior)

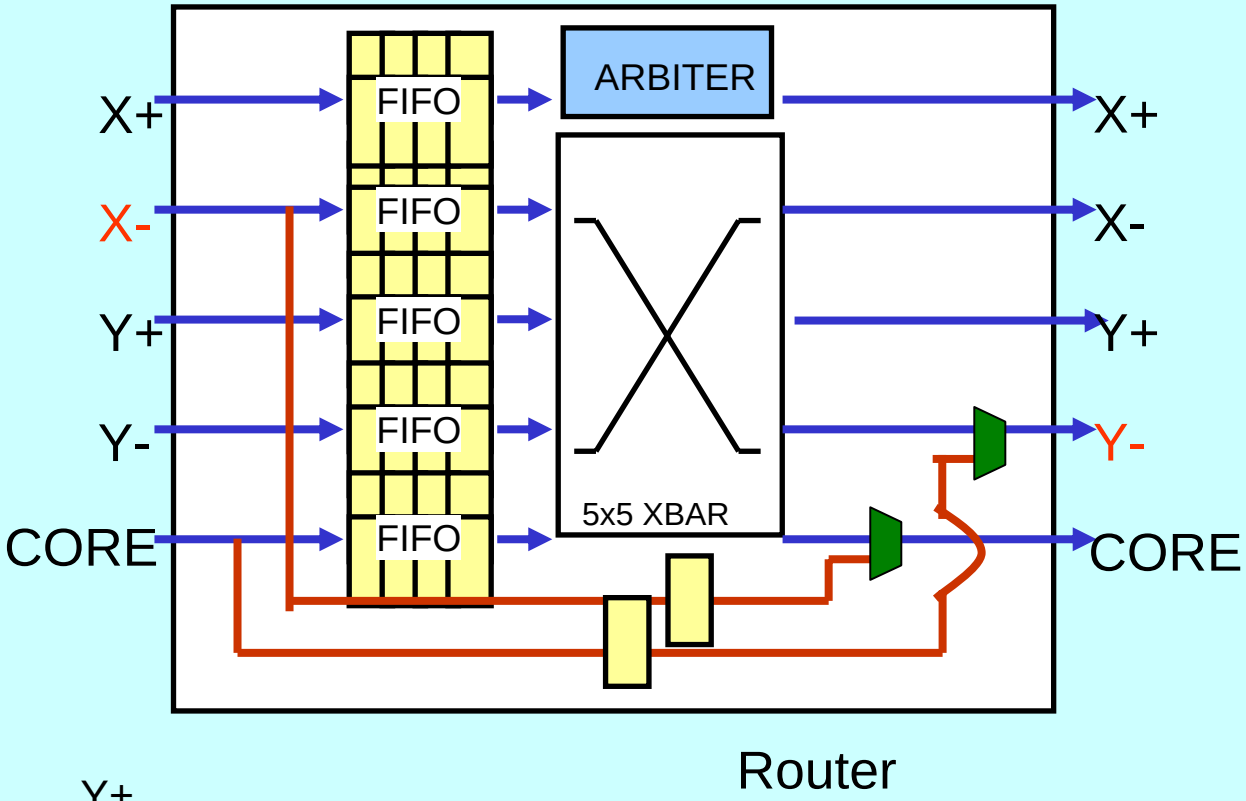

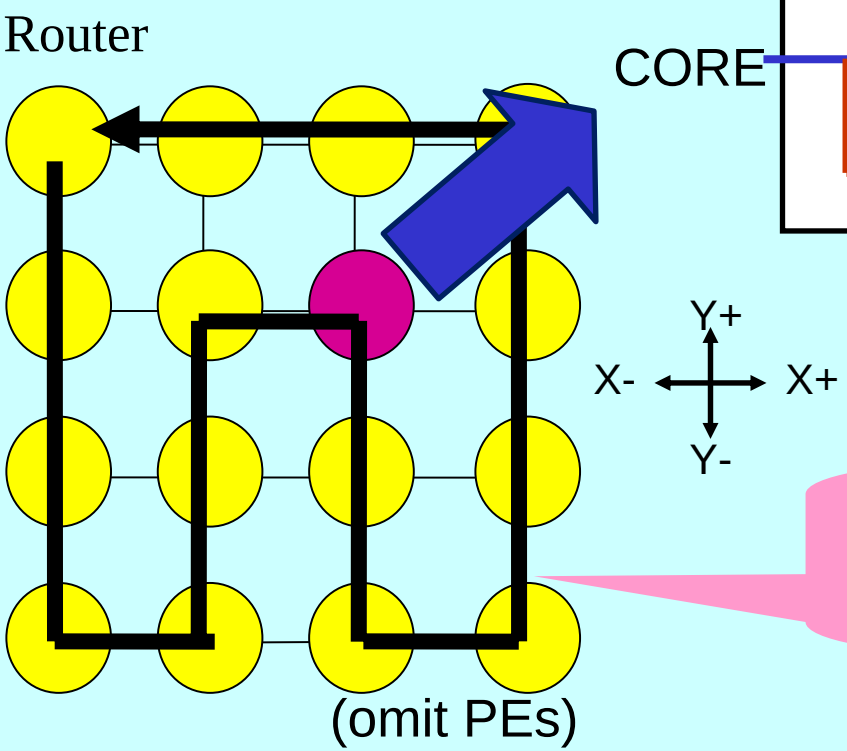

Set the DBP ports along a unidirectional embedded ring topology

Routing Bypasses Faults (e.g., failed crossbar)

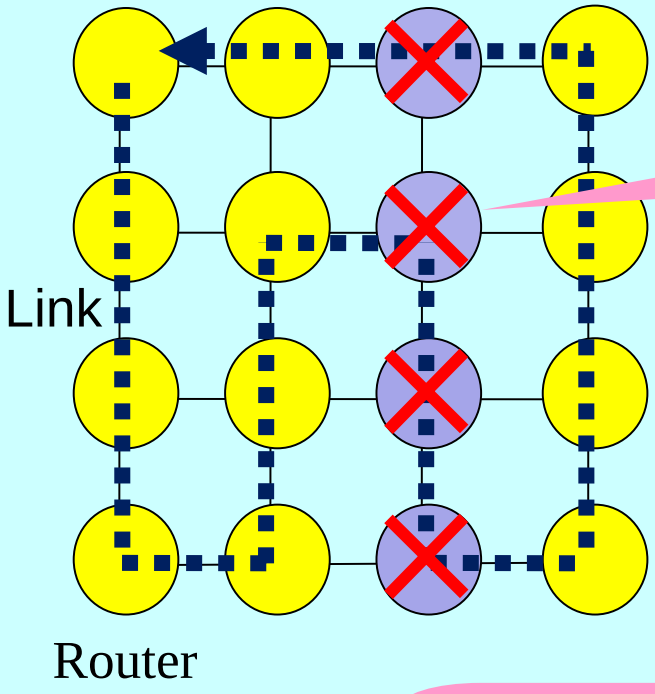

Default-backup path is used only at the faulty port

The corresponding network graph

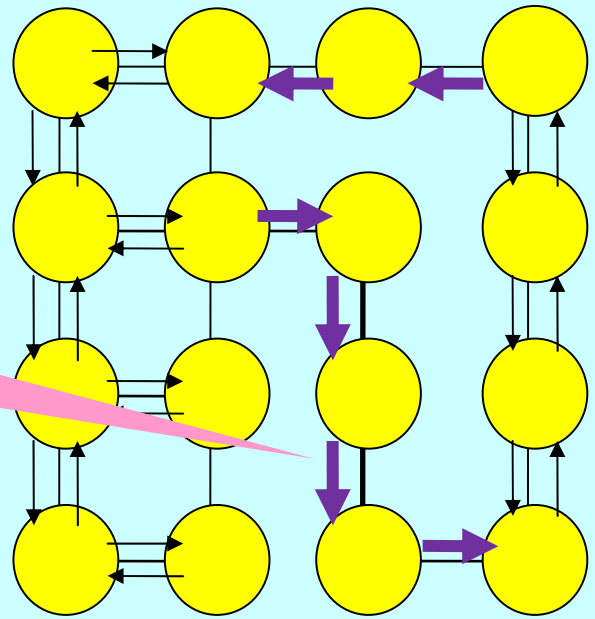

A unidirectional channel on a link

DBP Applied to Up*/Down* Routing

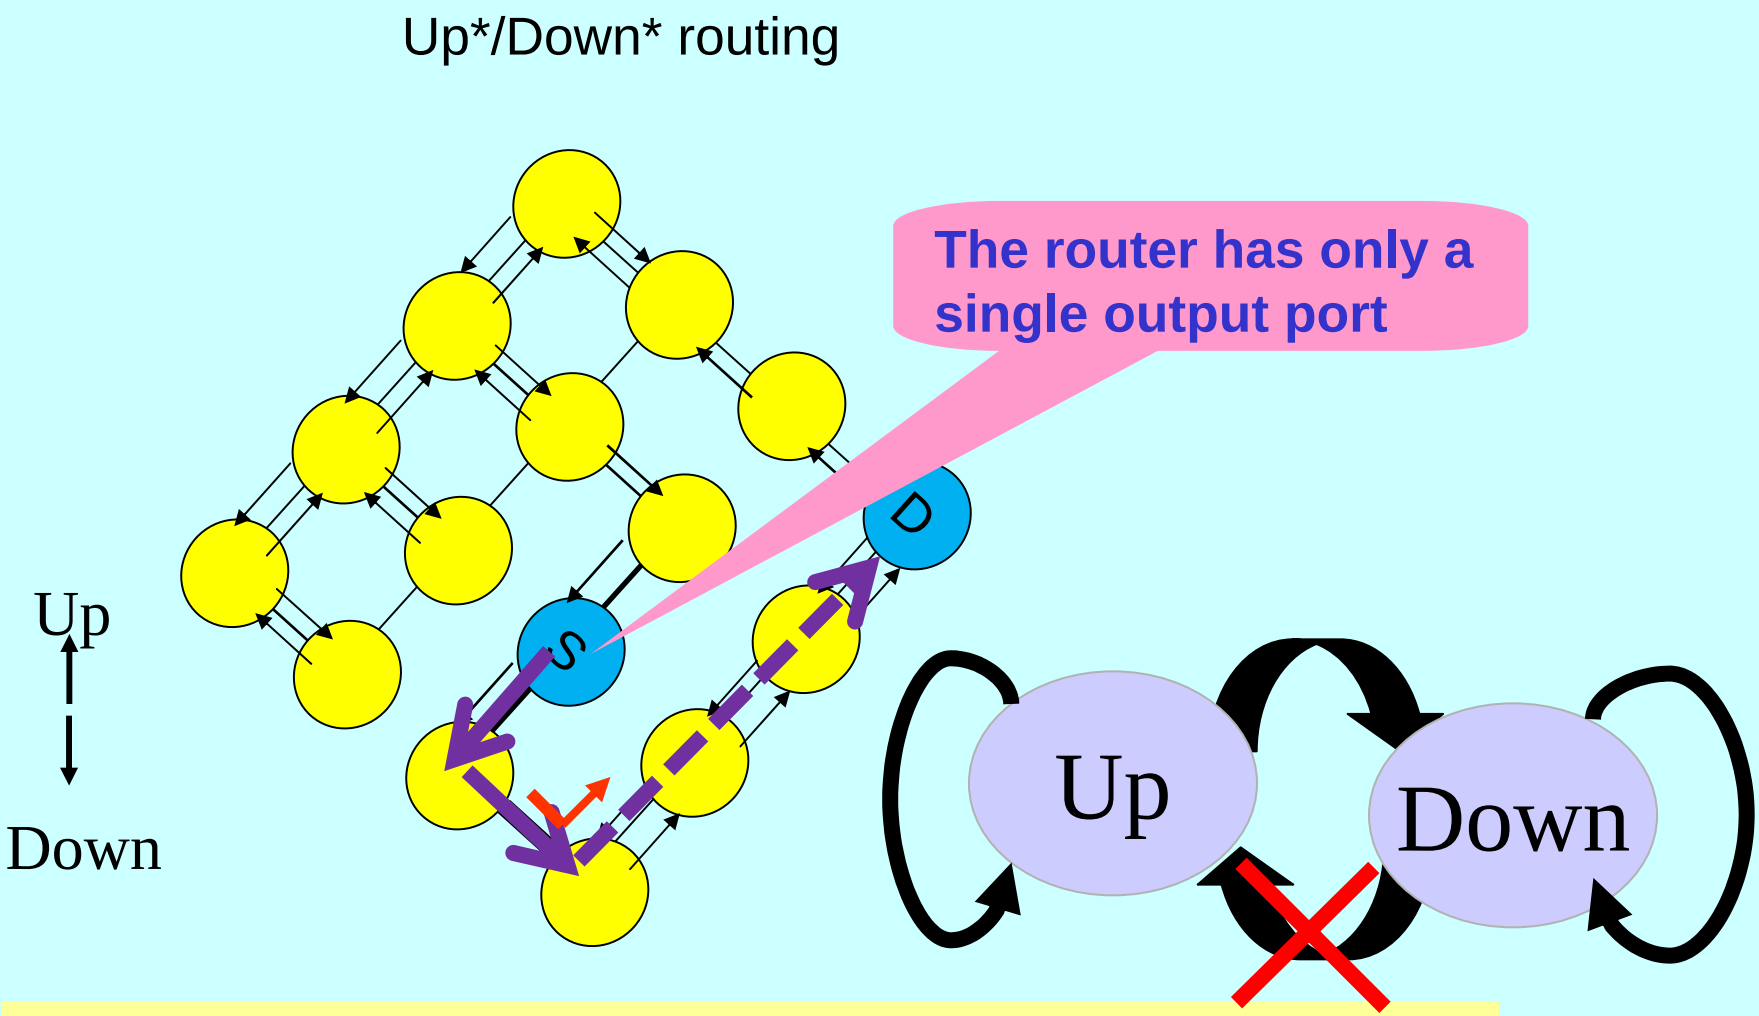

Existing deadlock-free routing cannot provide the network connectivity, due to the directional routing restrictions

DBP Routing Mechanism

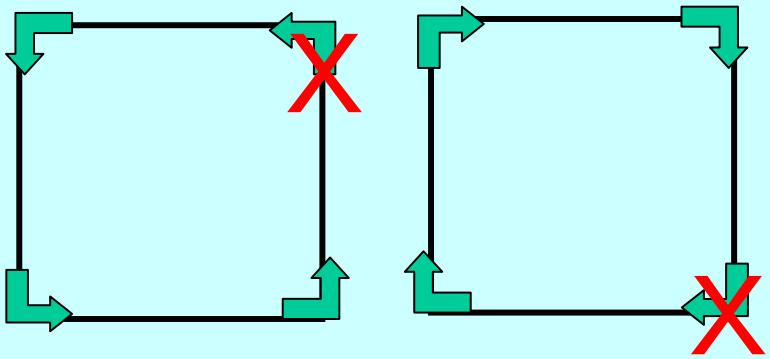

Turn Model[Glass,1992]

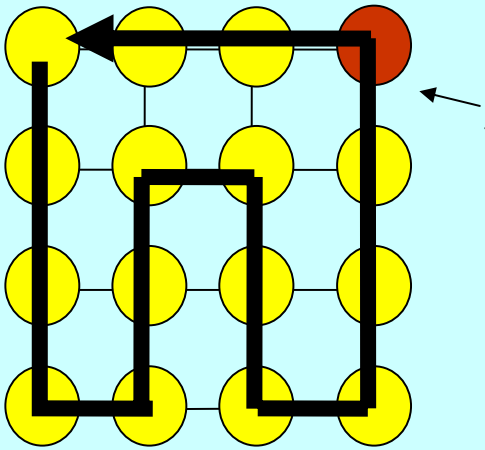

Virtual channel (VC) transition

- Guaranteeing deadlock-freedom and connectivity by imposing routing restrictions
  - Allows packet transfer along the DBP ring
  - Allows VC transitions in increasing order
  - Uses existing deadlock-free routing within every virtual-channel network

The Idea is similar to the SAN routing [koibuchi,ICPP03]

We propose a new routing strategy for NoCs with directional routing restrictions!

Outline

- Fault patterns on Network-on-Chip (NoC)
- Default-backup path mechanism (DBP)
  - maintains the connectivity of all healthy PEs, even if the network includes hard faults

    Objective

    Provide a highly reliable network using lightweight hardware！
- Evaluation
  - Energy
  - Amount of Hardware
  - Throughput

Energy: NoC Energy Model

- Ave. flit energy: E_{flit}
  - Send 1-flit to destination
  - How much energy[J] ?
- Simulation parameters
  - 6/12mm square chip (16/64 cores)
  - 90nm CMOS

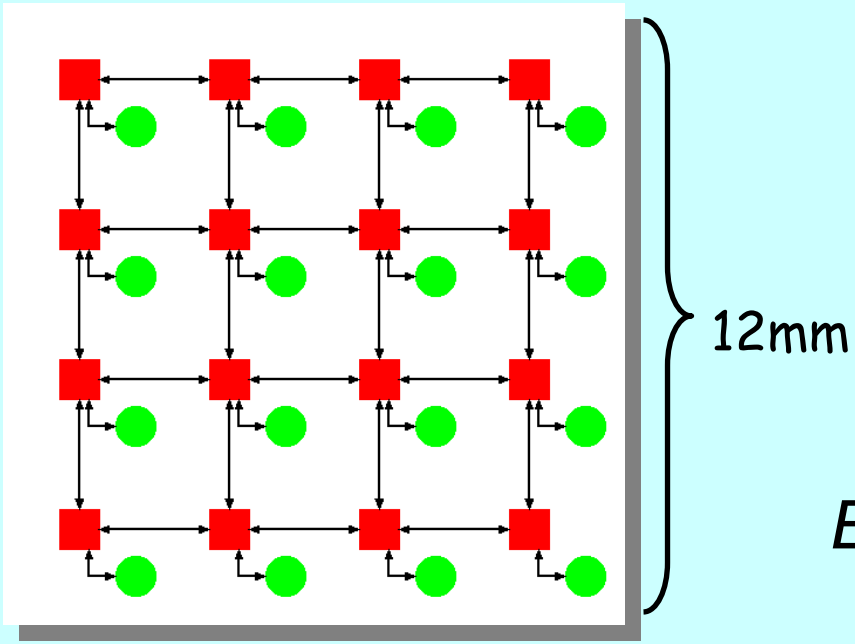

$$E_{flit} = w \cdot H_{ave}(E_{sw} + E_{link})$$

[Wang, DATE'05]

Energy Consumption

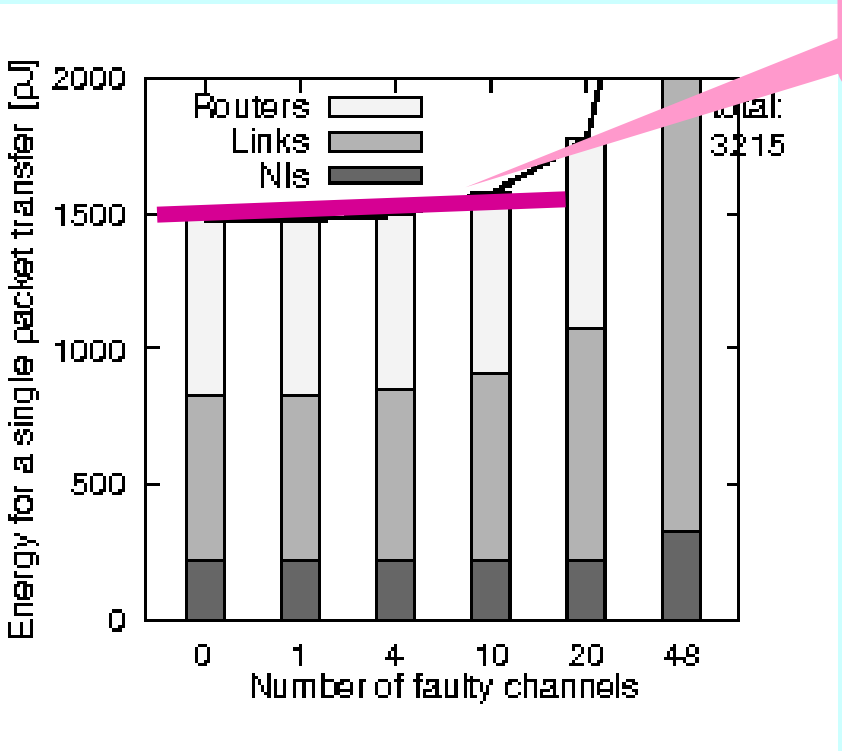

16 cores

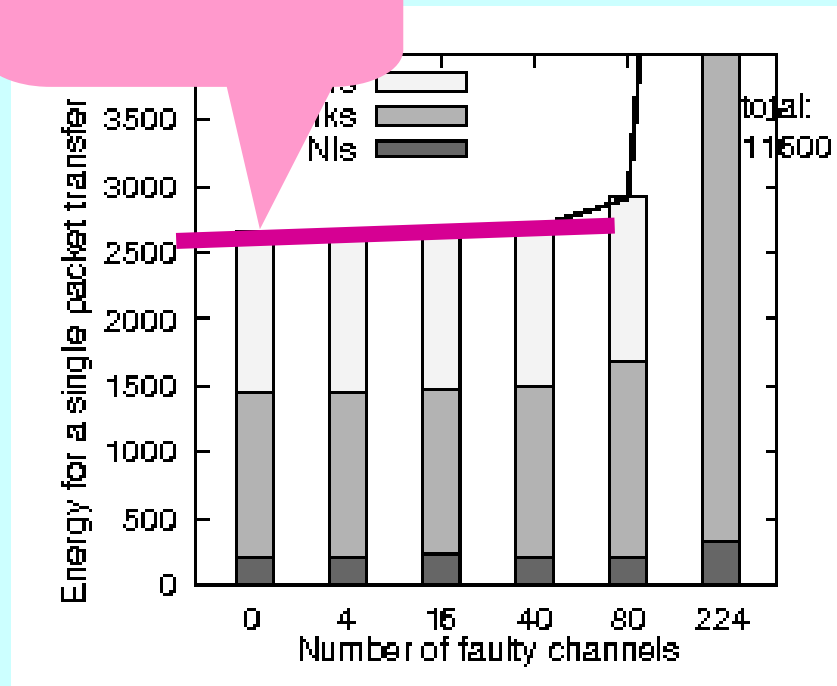

64 cores

almost constant!

As the number of faulty links increases, DBP gracefully increases the energy, due to the increased hop counts

Amount of Hardware

The ratio of additional HW is decreased, as # of ports increases.

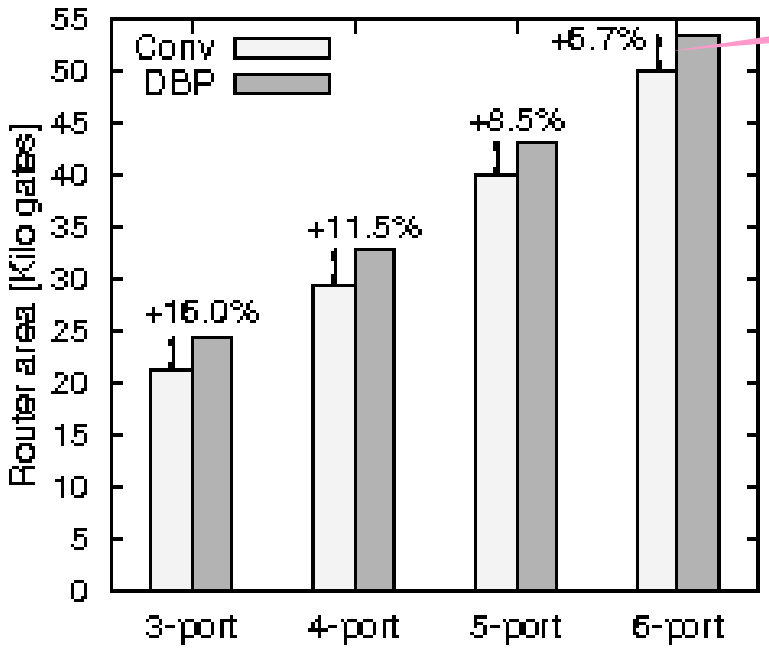

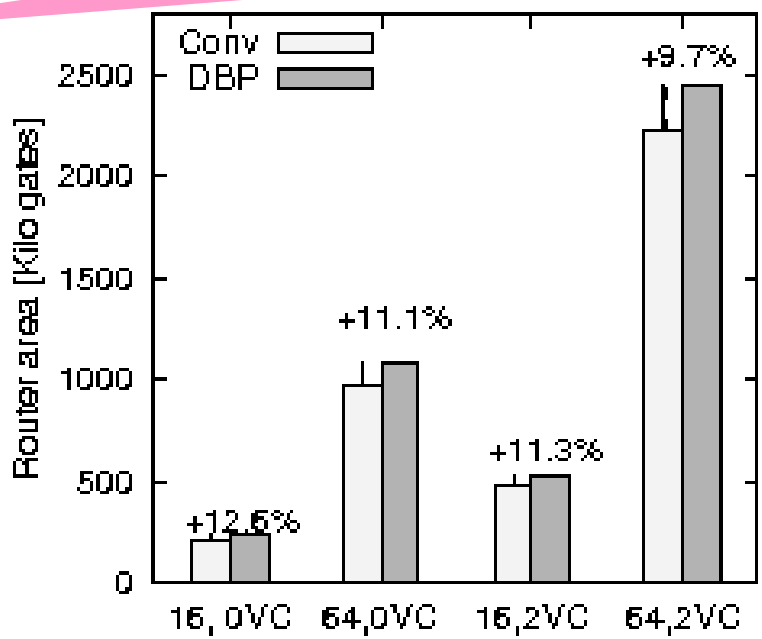

Router area with various # of ports.

Total router area of 2-D mesh

Area is increased by at most only 11.1% (the 2-VC case)

Performance Evaluation

- Network simulation
  - Throughput and latency
  - 16 cores and 64 cores
- Topology
  - 2-D mesh
- Traffic pattern
  - Random (as a baseline)

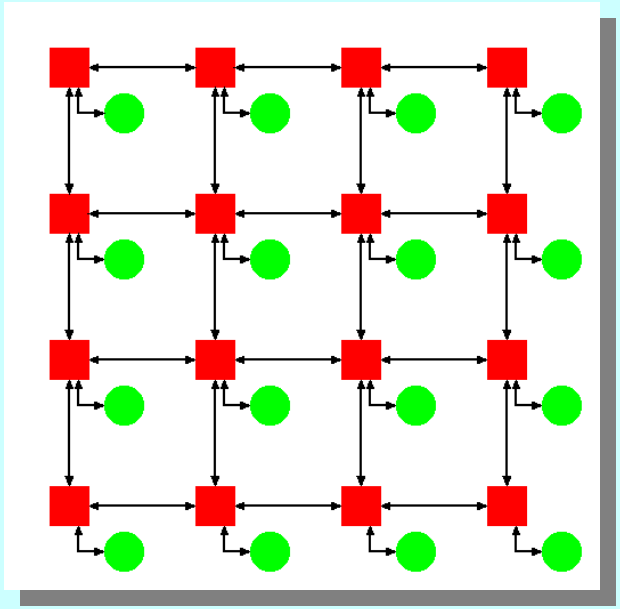

Packet size	16-flit (1-flit header)
Buffer size	1-flit per channel
Switching	Wormhole switching
# of VCs	2
Min latency	3-cycle per router

Throughput and Latency

Throughput is decreased by the increased path hops.

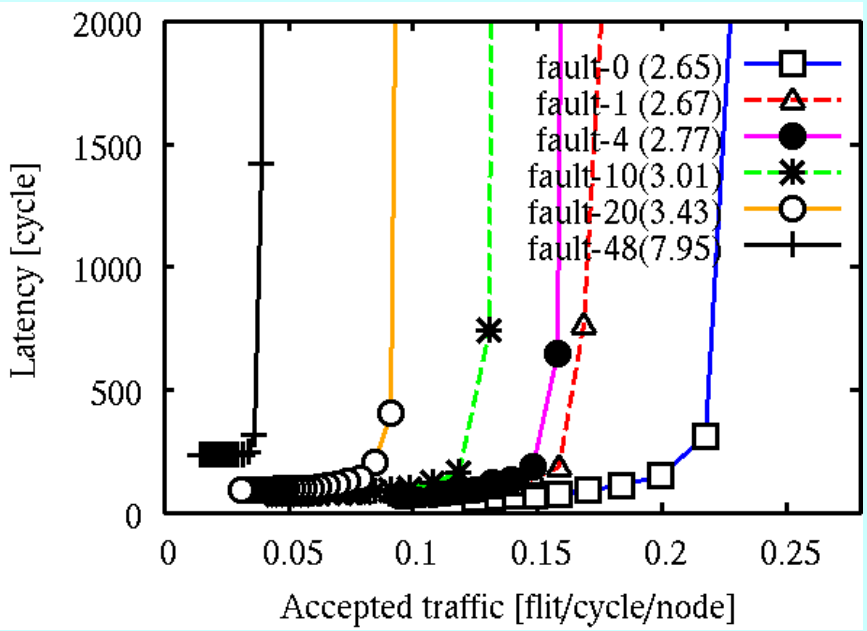

16 cores

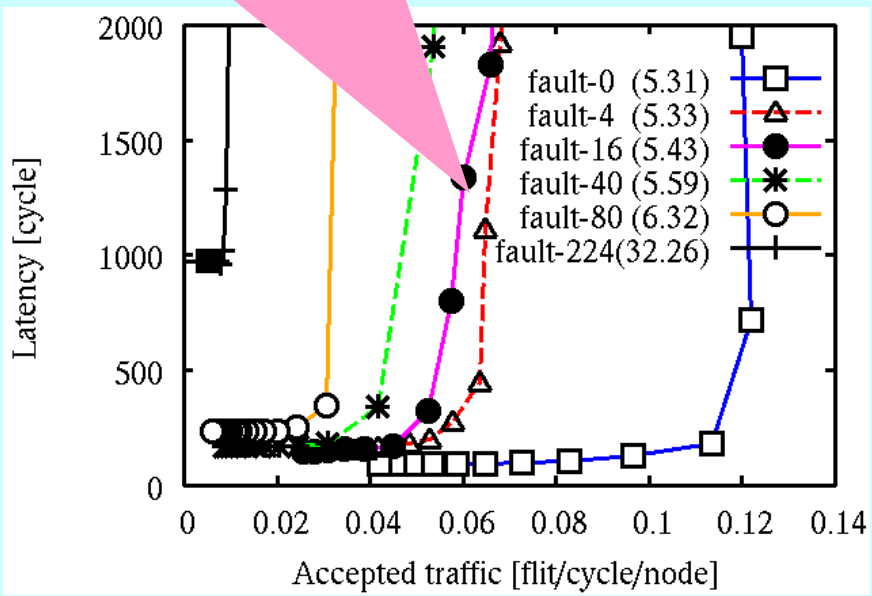

64 cores

Topology is changed from 2-D mesh (no faults) to ring at 48/224 faults on 16/64 cores

Extensions of DBP Mechanism

- Faults within the DBP itself and various ports
  - Partially duplication
  - Multiple embedded DBP rings

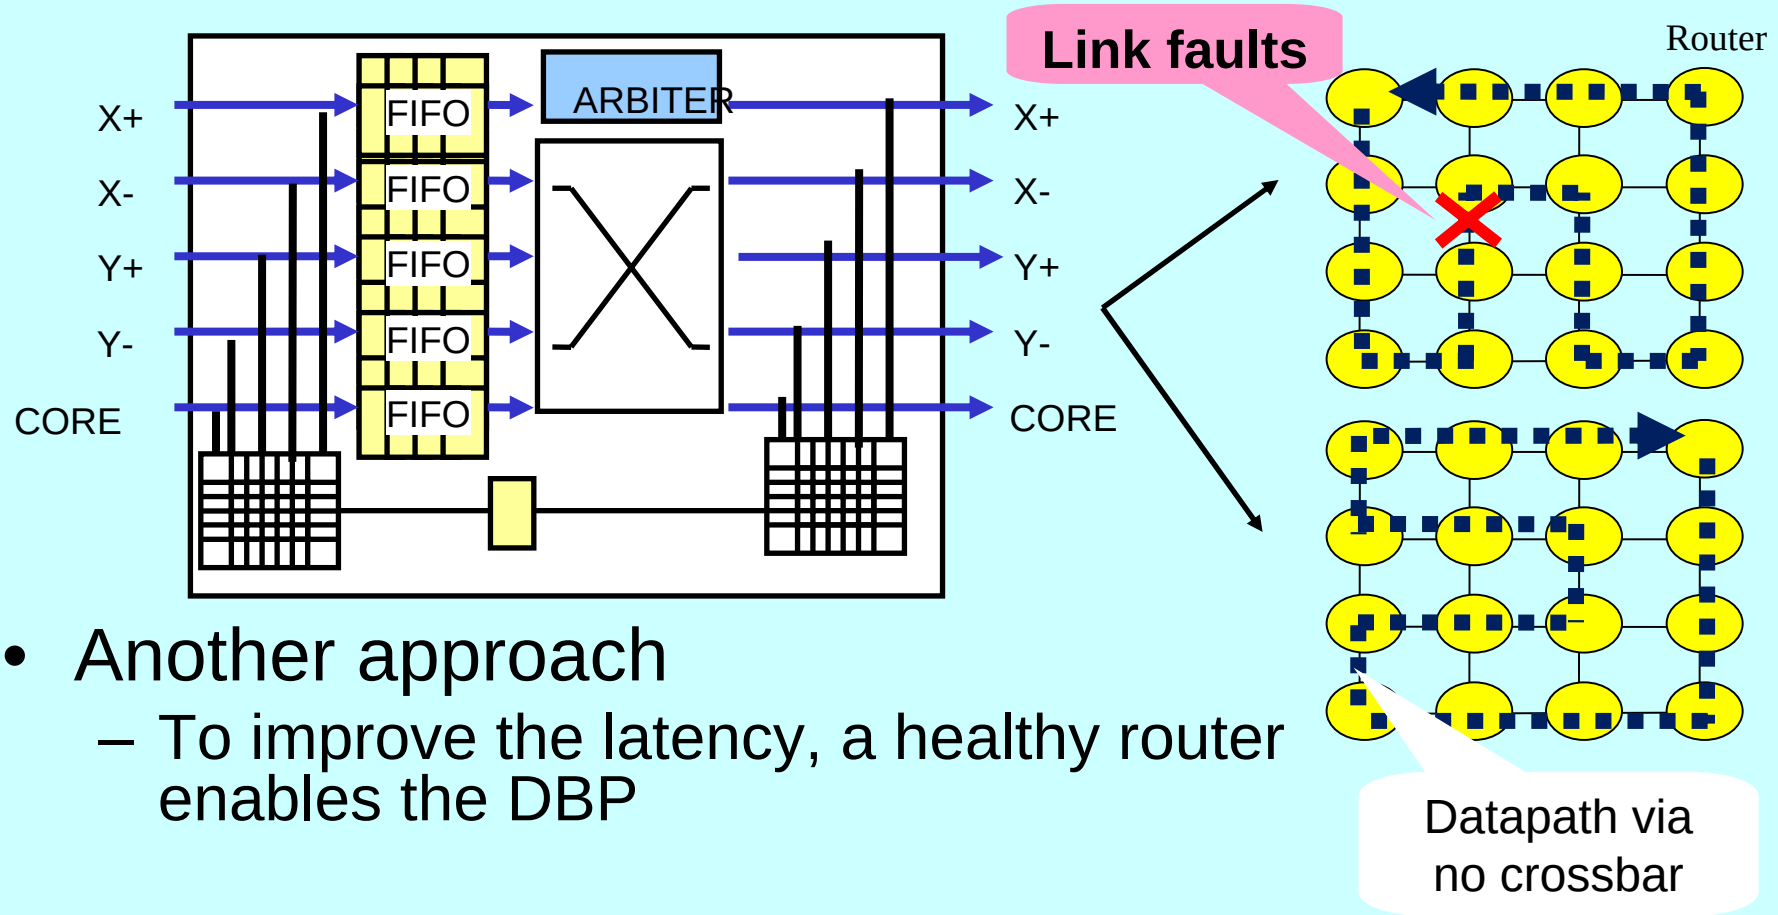

- Another approach
  - To improve the latency, a healthy router enables the DBP

Conclusions

- We proposed a lightweight fault-tolerant mechanism, DBP, for NoCs (architecture level)
  - Resilient for hardware faults of both intra-router modules and routing paths
  - A new routing strategy was developed
  - The idea is applicable to various NoC architectures
    - As well as regular topologies
- Evaluation
  - Energy consumption
    - almost constant by up to 40 faults (64 cores)
  - Amount of Hardware
    - increasing by at most only 11.1%
  - Throughput performance
    - decreasing by the increased path hops
- The DBP serves the role of “lifeline” to increase the lifetime of NoCs