# Tackling NoC Reliability and Power with Routing Reconfiguration

**Ritesh Parikh, Reetuparna Das and Valeria Bertacco**

Center for Future Architectures Research

CFAR

STARnet

**Theme 2384.005, Robust Design**

**TASK 5.3, Ultra-low cost mechanisms for fault and variability tolerant architectures**

## Tackling Faults with Routing Reconfiguration

**NoCs: significant silicon footprint and heavy activity**

—Permanent wear-out
—Device aging

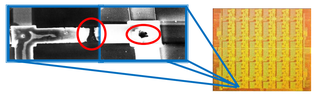

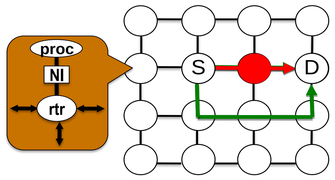


✖ cannot send on faulty path
✔ need to re-route around fault

Faulty irregular network with deadlock-free routes

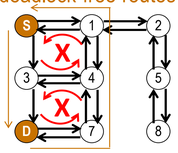


**Our reliability solution uDIREC:**
—Circumvents faults via routing-reconfiguration
—Guarantees connectivity and deadlock-freedom
—Implemented in software for low area overhead

1

## Tackling Power with Routing Reconfiguration

**NoCs: significant power footprint and heavy activity**

— Leakage is majority contributor
— Power-gating is ineffective

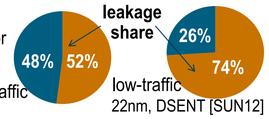


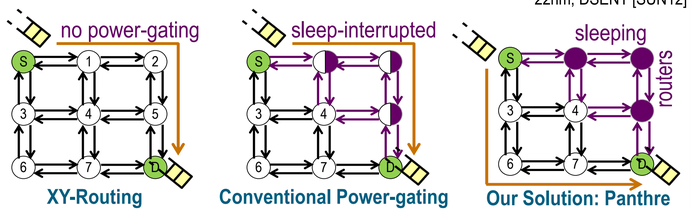


**Our power-aware solution Panthre:**
—Bypasses sleeping components via routing-reconfiguration
—Guarantees connectivity and deadlock-freedom
—Implemented in hardware for speed

2

## Contributions and Overview

1. **Panthre** (for **P**ower-**a**ware **N**oCs **th**rough **Re**configuration) incorporates:
   — **Fine-grained power-gating** of NoC components

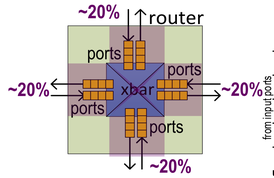


5 dpath-segments account for **99%** leakage power

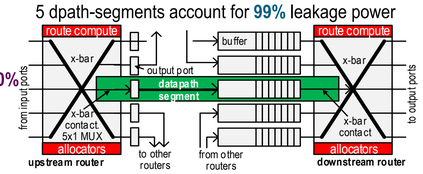


   — **Routing reconfiguration** to bypass sleeping components

2. **Adapts to application characteristics**
3. **Adds value to existing reliability hardware**
4. **Connected and deadlock-free NoC**

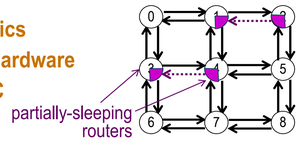


3

## Execution Flow

1. Distributed collection of component traffic statistics
2. Power-gating decisions based on traffic **A**ctivity **th**reshold ($A_{th}$)
3. Reconfiguration to update routing paths


**Adapting to Application Behavior**

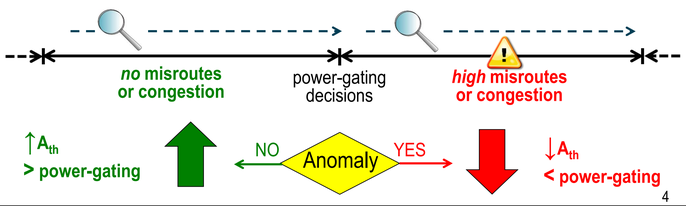


4

## Reconfiguration via Spanning Tree Construction

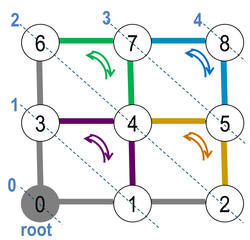


**Links added to the spanning tree at:**
**level #1: Links 0-1, 0-3**
**level #2: Links 3-6, (3-4 or 1-4), 1-2**
**level #3: Links (6-7 or 4-7), (4-5 or 2-5)**
**level #4: (7-8 or 5-8)**

**disable turns: 1-4-3, 4-7-6, 2-5-4, 5-8-7**

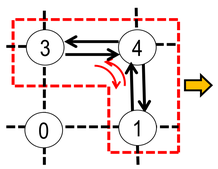

+ Switch off two datapath-segments per router
+ Potential 45% savings for 8x8 mesh
+ All configurations follow same turn restrictions

5

## Complete Router Shutdown Optimization

Conditions:
— Local-node is not producing or accepting traffic
— Router is at the leaf of the spanning tree

**Equivalent to power-gating 10 datapath-segments at once**

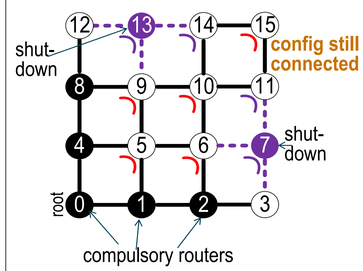


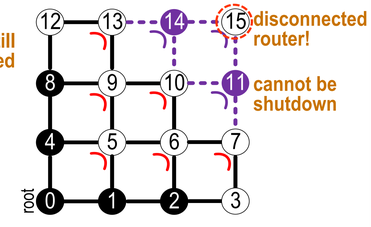


6

## Continuous Operation via Background Updates

**Continuous operation** even during the route update
— Always deadlock-free and connected!
— Backup routing using all links
— Leverage reliability-specific reconfiguration hardware

**Reconfiguration for fault-tolerance [Aisopos11]**

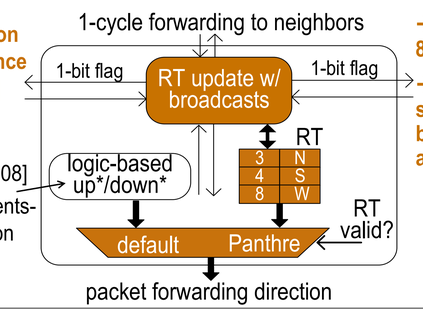


→ LBDR [Flich08]
→ All-components-ON configuration

→4K cycles for 8x8 network
→ Involves synchronized broadcasts by all nodes

7

## SPEC CPU 2006 Multi-Programmed Workloads

| 8x8 mesh NoC, 3-stage 4-VC routers | | 64 2-wide cores with shared L2 | |
|---|---|---|---|
| baseline | No power-gating | PG_conv | Conventional router-level Power-Gating |
| Panthre | Panthre w/ fine-grained power-gating | PG_fg | Conventional fine-grained Power-Gating |

**Panthre's** IPC slow-down is within **2%**

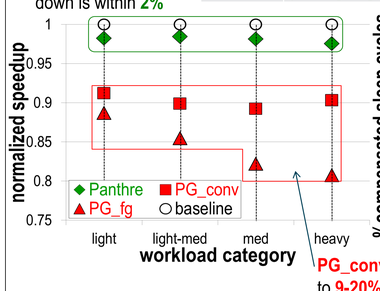


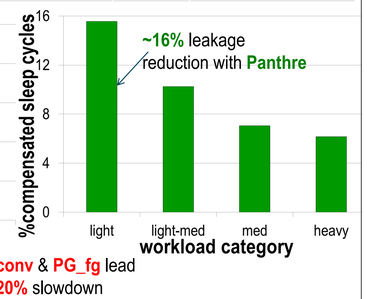


**PG_conv** & **PG_fg** lead to **9-20%** slowdown

**Panthre provides considerable power savings with minimal slowdown**

8

## Results with 1-16 Idle Cores – Synthetic Traffic

**Panthre_RS:** Panthre w/ complete router shutdown optimization

Latency increase is **less than 30%**

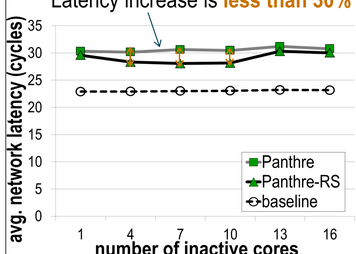


**37%** leakage power can be saved for **10-16 idle nodes**

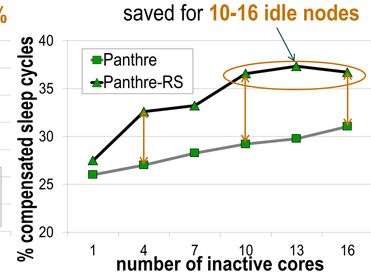


**Panthre_RS exhibits lower latency and greater leakage power savings**

9