TUCS

International Symposium on Networks-on-Chip

Presenter: Amir-Mohammad Rahmani

Generic Monitoring and Management Infrastructure for 3D NoC-Bus Hybrid Architectures

Amir-Mohammad Rahmani[1,2], Kameswar Rao Vaddina[1,2], Khalid Latif[1,2], Pasi Liljeberg[2], Juha Plosila[2], and Hannu Tenhunen[2]

[1]Turku Centre for Computer Science (TUCS), Turku, Finland

[2]Computer Systems Lab., Department of Information Technology, University of Turku, Finland

Outline

- Introduction
  - 3D Integration Technology
  - 3D Networks-on-Chip
- Motivation
- ARB-NET Monitoring and Management Platform
  - *AdaptiveXYZ* Routing Algorithm
  - Thermal Monitoring and Management
- Experimental Results
  - Synthetic Traffic Analysis
  - Realistic Traffic Analysis
  - Hardware Implementation Details
- Summary and Ongoing Work

Introduction

- Communication plays a crucial role in the design and performance of multi-core system[illegible]hip (SoCs).
- Recently on-chip transistor [illegible] density has been considerably increased.
  - This enables the integration of doze[illegible]f compo[illegible]nts on a single die.
  - One outcome [illegible]f grea[illegible] in[illegible]tion[illegible]at interconnection networks have[illegible]rted t[illegible]eplace[illegible]red buses.
- Network[illegible]n-ch[illegible] are [illegible]pos[illegible] to be used in complex SoCs [illegible] m[illegible]catio[illegible]tween cores, because of improve[illegible]nts[illegible]s of:
  - Scalabil[illegible]
  - Performa[illegible]
  - Power cons[illegible]ption
  - Reliability
  - …

Introduction *(cont.)*

- The advent of stacked technologies provides a new horizon for on-chip interconnect design.
- 3D integrated circuits have emerged to overcome the limitations of **interconnect scaling** by stacking active silicon layers.
- 3D ICs offer a number of advantages over 2D ICs:
  - Shorter global interconnects
  - Higher performance
  - Lower interconnect power consumption
  - Higher packing density
  - Smaller footprint
  - Support for the implementation of mixed-technology chips

3D IC

Introduction ([illegible]nt.)

NoC + 3D IC = 3D NoC

3D NoC

- The amalgamation of No[illegible]nd 3[illegible] IC a[illegible]ws [illegible] the creation of new structur[illegible] tha[illegible]nabl[illegible] significant enhancements ove[illegible]e tr[illegible]tion[illegible]ol[illegible]ns.
- With freedom in t[illegible]ension, architectures that were impossib[illegible]or p[illegible]tive due to wiring constraints in planar I[illegible] are no[illegible] possible.
- Ma[illegible]plementati[illegible] can outperform their 2D [illegible]nterpa[illegible]

Symmetric NoC Architecture

- This architecture has two major inherent drawbacks.
  - It does not exploit the beneficial attribute of a negligible inter-wafer distance.
  - A larger 7×7 crossbar is obligated as a result of two extra ports.
  - The power consumption of a 7×7 crossbar is approximately 2.25 times more than the 5×5 counterpart.

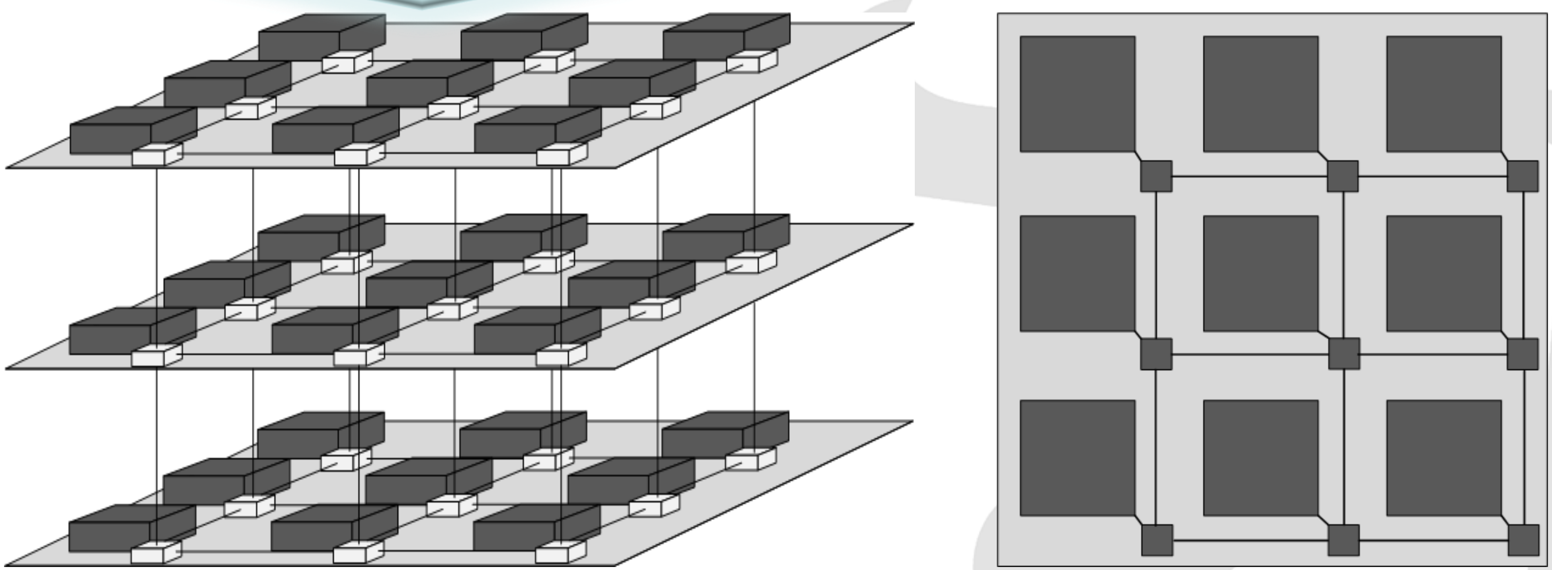

3D Mesh

2D Mesh

TUCS 3D NoC-Bus Hybrid Architecture

- It does not allow concurrent communication in the third dimension.
- In a high network load, the probability of contention and blocking critically increases.

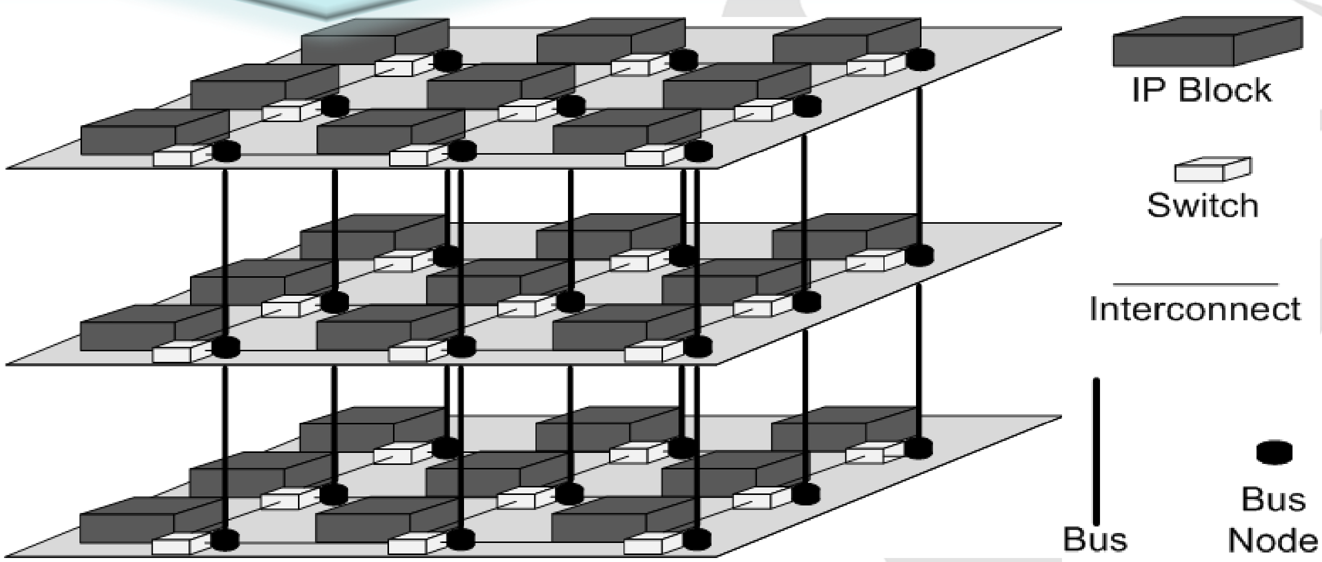

Motivation and Contribution

- The dynamic Time-Division Multiple Access (dTDMA) bus was used as a communication

- Hybridization of two different communication media necessitates new monitoring and management frameworks.
- The available frameworks cannot efficiently utilize the benefits of hybrid architectures.
- We propose a system monitoring platform called ARB-NET customized for 3D NoC-Bus Hybrid mesh architectures.

- The output buffers hinder the on-chip network from implementing adaptive routing algorithms.

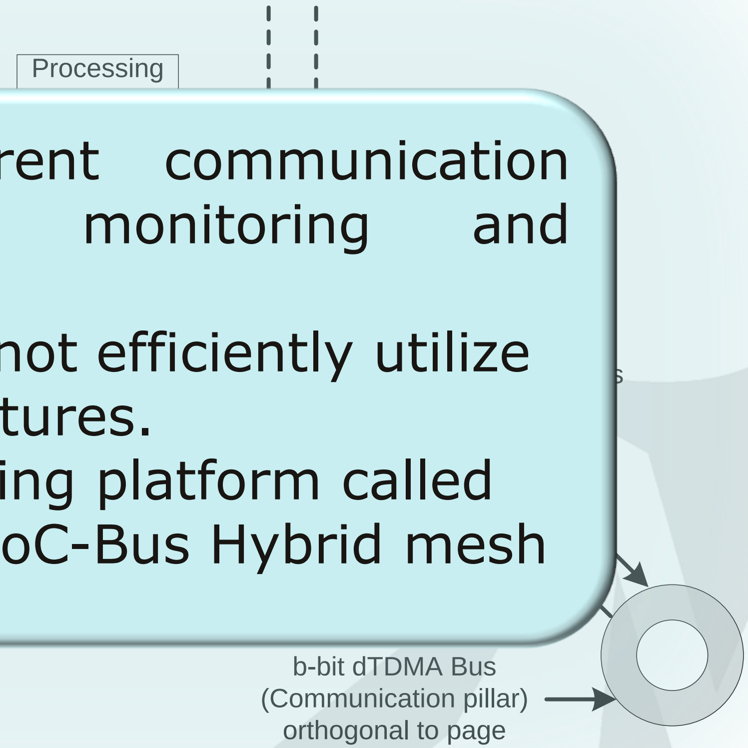

High level overview of the stacked mesh router architecture

ARB-NET Architecture

The arbiters resolve the contention between different IP blocks for bus access.

Arbiters can be prudently used by bringing them together to form a network and thereby creating an efficient monitoring and controlling mechanism.

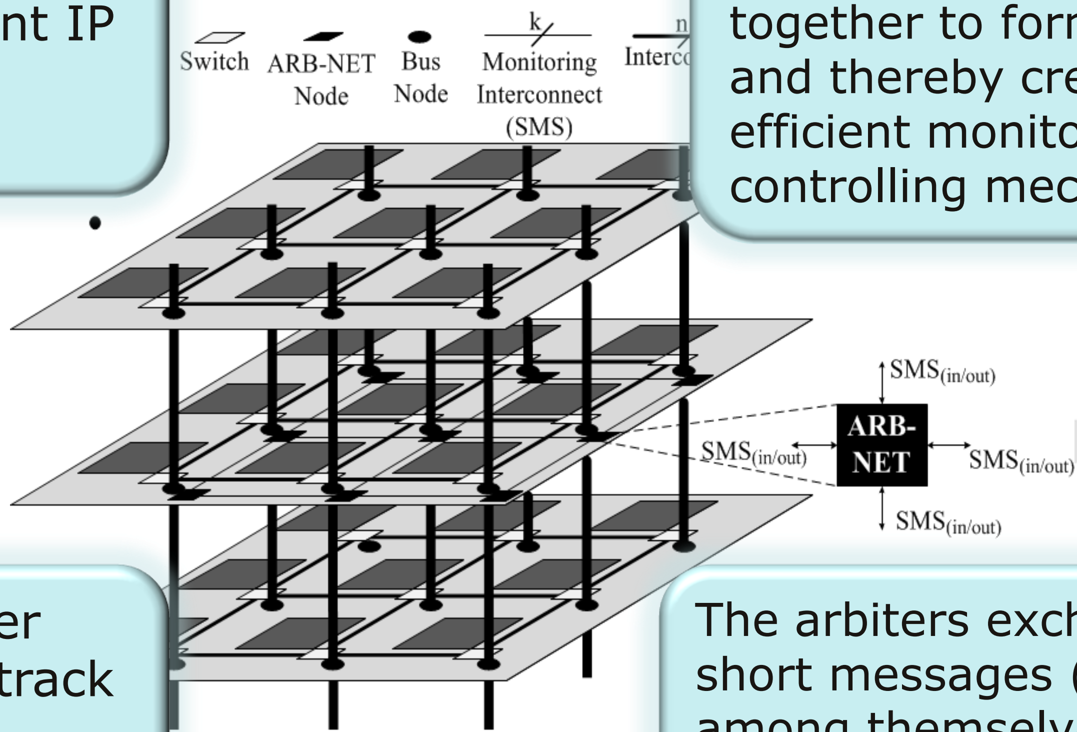

-based 3D Hybrid NoC-Bus mesh

They are a better source to keep track of monitoring information.

The arbiters exchange very short messages (SMS) among themselves regarding various monitoring services that are on offer.

ARB-NET Node Architecture

Algorithm 4 *Inter-Layer Fault Tolerant AdaptiveXYZ*

Additional Input: $Faulty\_Bus\_X$, $Faulty\_Bus\_Y$

```
if (Faulty_Bus_X = '0' and Faulty_Bus_Y = '0') then
    if (Arb_rsp_Balance = '0') then
        Stress_value_X = {Arb_rsp_X/Y, queue_length_X};
        Stress_value_Y = {Arb_rsp_X/Y, queue_length_Y};
    end if
else if (Faulty_Bus_X = '1' and Faulty_Bus_Y = '0') then
    Send the packet to Y_Axis towards the destination and exit;
else if (Faulty_Bus_X = '0' and Faulty_Bus_Y = '1') then
    Send the packet to X_Axis towards the destination and exit;
end if
```

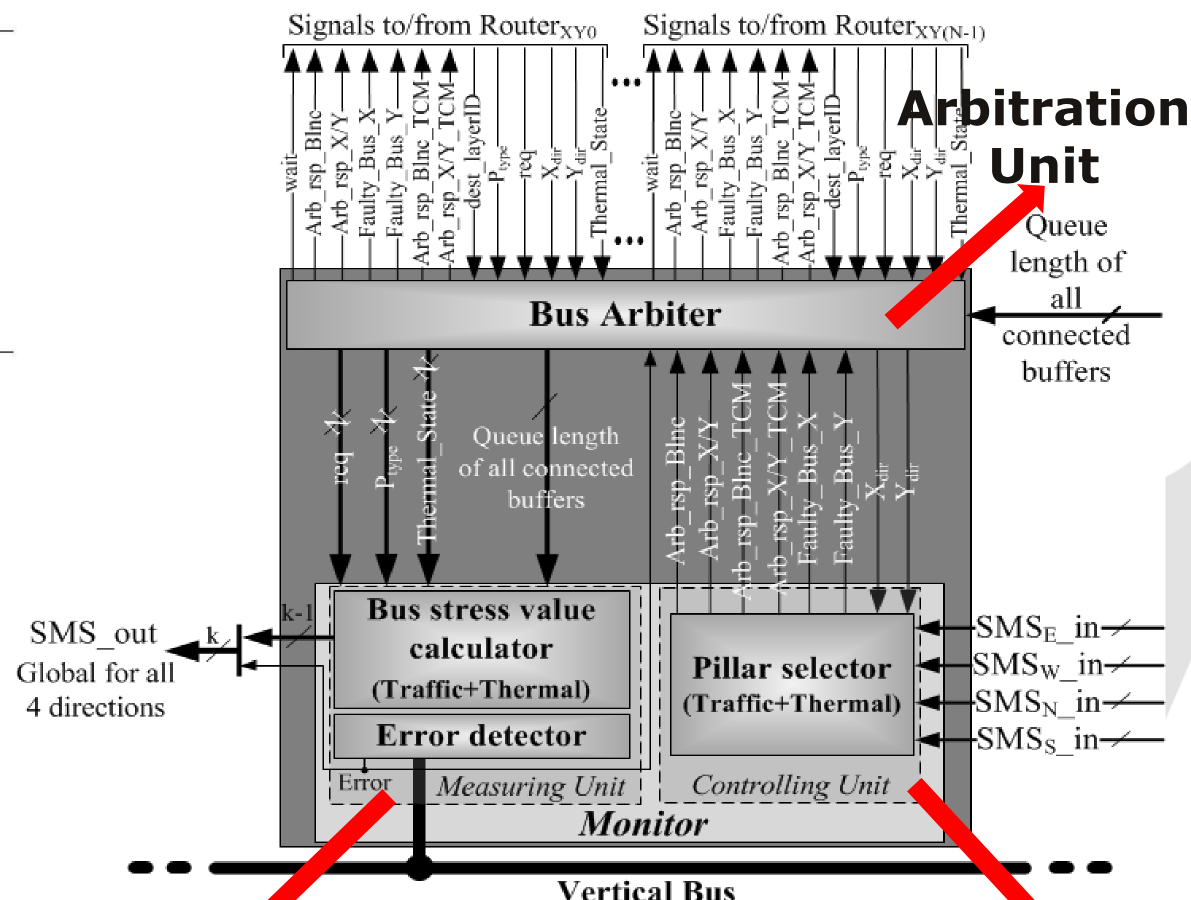

ARB-NET node architecture

Measuring Unit **Control Unit**

Thermal Monitoring and Management

- Hotspots by their very nature are localized and can lead to timing uncertainties in the system.
- There is a need to move towards run-time thermal management solutions which can effectively guarantee thermal safety.
- A thermal monitoring and management strategy on top of our ARB-NET infrastructure is proposed.
  - It responds to thermal hotspots in a 3D NoC by routing data packets which bypass the regions with greater density of hotspots.
- We assume that a distributed thermal sensor network is embedded in the 3D NoC.
  - They regularly provides thermal feedback to the routers and bus arbiter network thereby aiding and controlling temperature with our thermal control mechanism.

Thermal Monitoring and Management (*cont.*)

- We use threshold approach which when crossed, the thermal control mechanism kicks in.
- When the temperature rises above a certain threshold trip level (Thermal trip), the thermal control unit (TCU) changes the control policy to thermal control mode until the temperature drops to a certain safe zone (Thermal safe).

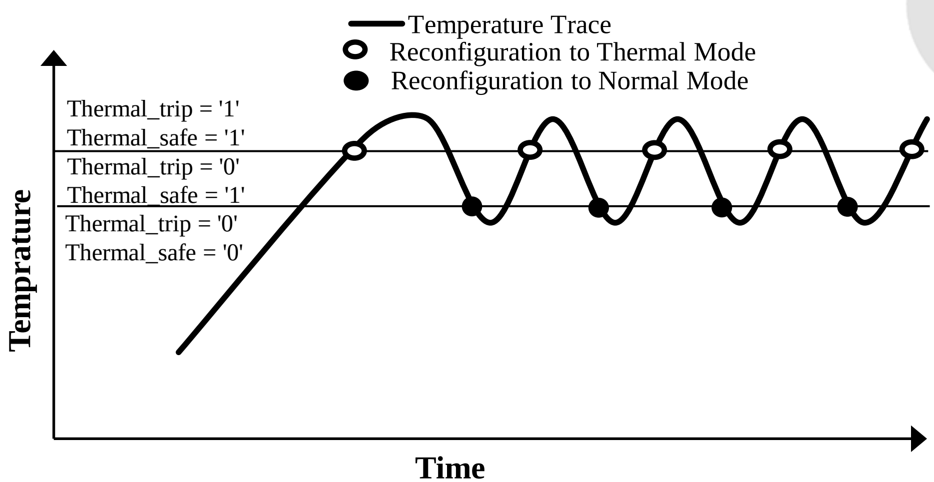

Temperature profile using run-time thermal management

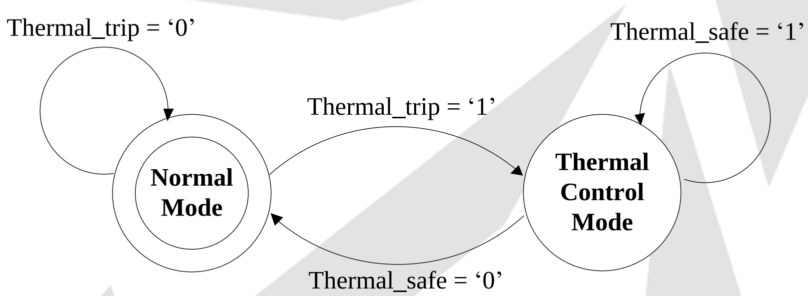

State diagram of the proposed thermal control unit

Thermal Monitoring and Management (*cont.*)

- If the tile's temperature increases beyond the predefined thermal trip state then a signal called *Thermal State* is set which will be sent to the respective bus arbiter for further processing.
- We measure the thermal state of the bus in terms of its thermal stress value.
  - The total thermal stress value of the bus is the sum of the Thermal State values of the respective routers connected to the bus.
- It takes into account the total thermal stress, traffic and fault stress values of the neighboring buses.

Algorithm 7 *Thermal-Aware and IL Fault Tolerant AdaptiveXYZ*

Additional Input: Arb_rsp_X/Y_TCM, $Arb\_rsp\_Balance\_TCM$

```
if (Faulty_Bus_X = '0' and Faulty_Bus_Y = '0') then
    if (Arb_rsp_Balance_TCM = '0') then
        if (Arb_rsp_X/Y_TCM = '0') then
            Send the packet to X_Axis towards the destination and exit;
        else
            Send the packet to Y_Axis towards the destination and exit;
        end if
    else
        if (Arb_rsp_Balance = '0') then
            Stress_value_X = {Arb_rsp_X/Y, queue_length_X};
            Stress_value_Y = {Arb_rsp_X/Y, queue_length_Y};
        end if
    end if
else if (Faulty_Bus_X = '1' and Faulty_Bus_Y = '0') then
    Send the packet to Y_Axis towards the destination and exit;
else if (Faulty_Bus_X = '0' and Faulty_Bus_Y = '1') then
    Send the packet to X_Axis towards the destination and exit;
end if
```

Experimental Results

- To demonstrate the efficiency of the proposed monitoring platform in network average packet latency and power, a cycle-accurate NoC simulation environment was implemented in HDL.
- The **proposed architecture**, **Symmetric 3D-mesh NoC** and ***AdaptiveZ*-based 3D NoC-Bus Hybrid mesh** and the **proposed architecture** were analyzed for synthetic and realistic traffic patterns.

Synthetic Traffic Analysis

- The 3D NoC of the simulation environment consists of 3×3×4 nodes.
  - The performance of the network was evaluated using latency curves as a function of the packet injection rate.
  - There were two packet types (1-flit and 5-flit packets).
  - The buffer size was four flits.
  - The data width was set to 128 bits.
- To perform the simulations, we used following traffic patterns:
  - Uniform
  - Hotspot 10%
  - Negative Exponential Distribution (NED)
- The packet latencies were averaged over 50,000 packets.

Synthetic Traffic Analysis

- For all the traffic patterns, the network with proposed architecture saturates at higher injection rates.
- The reason is that the *AdaptiveXYZ* routing algorithm for inter-layer communication increases the bus utilization and makes the load balanced.

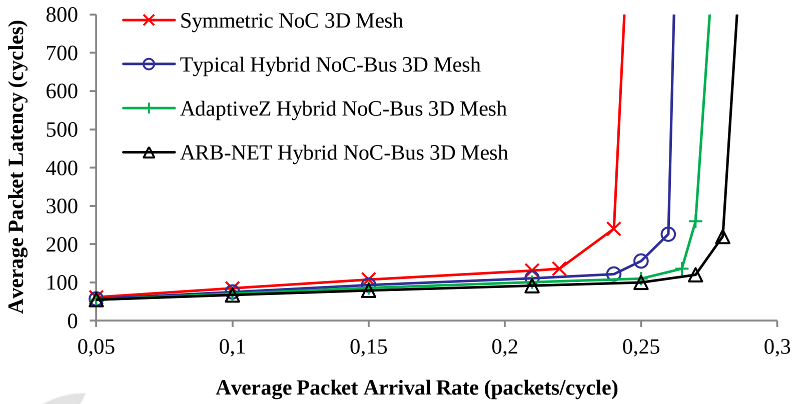

Latency versus average packet arrival rate results under NED traffic

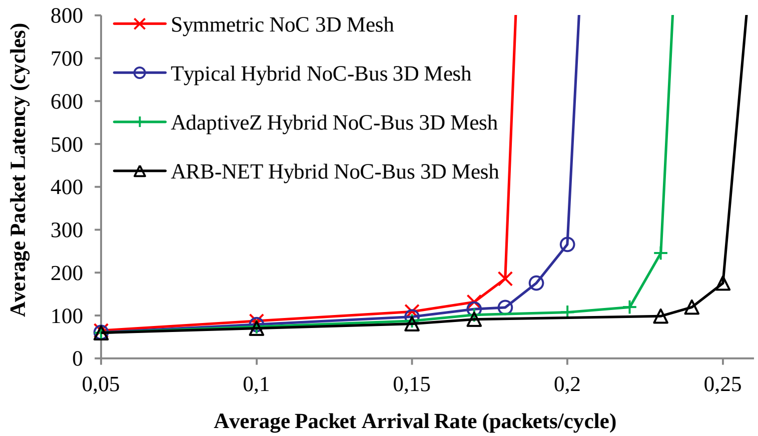

Latency versus average packet arrival rate results under uniform traffic

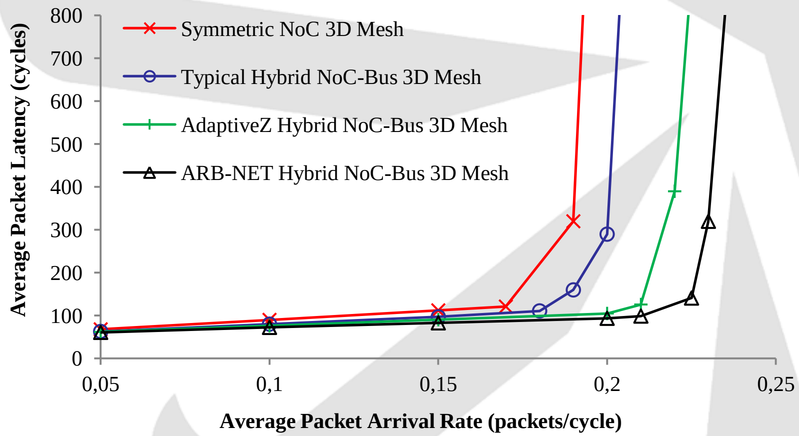

Latency versus average packet arrival rate results under hotspot 10% traffic

Realistic Traffic Analysis

- For realistic traffic analysis, the encoding part of video conference application with sub-applications of H.264 encoder, MP3 encoder and OFDM transmitter were used [*Rahmani et al. NOCS'11*].
- To demonstrate the efficiency of the ARB-NET monitoring platform for system reliability, the network running the video conference application with one faulty vertical bus was simulated.

3D NoC Architecture	**Avg. Power Cons. (W)**	**APL (cycles)**
Symmetric NoC 3D Mesh	3.195	117
Hybrid NoC-Bus 3D Mesh	2.832	100
[*Rahmani et al. NOCS'11*] Hybrid NoC-Bus 3D Mesh	2.671	92
ARB-NET Hybrid NoC-Bus 3D Mesh using *AdaptiveXYZ* routing	2.603	86
[*Rahmani et al. NOCS'11*] Hybrid NoC-Bus 3D Mesh (1 faulty bus)	2.847	103
ARB-NET Hybrid NoC-Bus 3D Mesh using IL Fault Tolerant *AdaptiveXYZ* routing (1 faulty bus)	2.663	89

Hardware Implementation Details

- The area of the different routers was computed once synthesized on CMOS 65nm LPLVT STMicroelectronics standard cells using Synopsys Design Compiler.
- The figures given in the table reveal that the area overheads of the proposed routing unit and the ARB-NET node are negligible.

HARDWARE IMPLEMENTATION DETAILS

Component	**Area** (μm^2)
Typical 6-Port Router with 2 VCs (*Z-DyXY*)	92194
Rahmani et al. [9] 6-Port Router with 2 VCs (*AdaptiveZ-DyXY*)	93591
Proposed 6-Port Router with 2 VCs (*AdaptiveXYZ*)	93914
Typical Bus Arbiter for a 3-layer NoC	267
Rahmani et al. [9] Bus Arbiter for a 3-layer NoC	694
ARB-NET Bus Node for a 3-layer NoC (Arbiter + Monitoring)	1534

Summary and Ongoing Work

- An low-cost monitoring platform called ARBNET for 3D stacked mesh architectures was proposed which can be efficiently used for various system management purposes.
- A fully congestion-aware adaptive routing algorithm named *AdaptiveXYZ* is presented taking advantage from viable information generated within bus arbiters.
- A thermal monitoring and management strategy on top of our ARB-NET infrastructure was presented.
- Our extensive simulations reveal that our architecture using the *AdaptiveXYZ* routing can help achieving significant power and performance improvements compared to recently proposed stacked mesh 3D NoCs.
- We have performed some preliminary implementations of our thermal monitoring and management strategy which guide us to believe about the reductions in on-chip peak temperatures.
- The future work would include supplementing and verifying our preliminary work on thermal monitoring and management strategy by simulating our network using a set of realistic workloads.

TUCS

Thank you for your attention

Any Questions?