Exploring High-Dimensional Topologies for NoC Design Through an Integrated Analysis and Synthesis Framework

F. Gilabert⁺, S. Medardoni[‡], D. Bertozzi[‡], L. Benini⁺⁺, M.E. Gomez⁺, P. Lopez⁺ and J. Duato⁺

⁺Universidad Politécnica de Valencia.
[‡] University of Ferrara.
[†] [†] University of Bologna.







Multi-dimension topologies



2D mesh frequently used for NoC design

- perfectly matches 2D silicon surface
- high level of modularity
- controllability of electrical parameters

But its avg latency and resource consumption scale poorly with network size

Topology with more than 2 dimensions attractive:

- higher bandwidth and lower avg latency
- on-chip wiring more cost-effective than off-chip



But layout (routing) issues might impact their effectiveness and even feasibility (use of more metal layers) (links with different latencies)

Objective

Explore the effectiveness and feasibility of multi-dimensional topologies

Exploration methodology issues arise

1. Fast and accurate exploration tools required for system-level analysis and topology selection



Our approach

Abstract the behaviour of all NoC architecture-level mechanisms while retaining RTL clock cycle accuracy

(flow control, arbitration, switching, routing, buffering, injection and ejection)

Objective

Explore the effectiveness and feasibility of multi-dimensional topologies

Exploration methodology issues arise

2. Realistically capture traffic behavior

Traffic pattern usually abstracted as an average link bandwidth utilization



May lead to highly inaccurate performance predictions (traffic peaks, different kinds of messaging, synchronization mismatches)

Our approach

- Project network traffic based on latest advances in MPSoC communication middleware
- Generate traffic patterns for the NoC "shaped" by the above communication middleware (e.g., synchronization, communication semantics)

Objective

Explore the effectiveness and feasibility of multi-dimensional topologies

Exploration methodology issues arise

3. Backend synthesis flow required for assessment of layout effects



Our approach

- Silicon-aware topology exploration
- Derive physical constraints that, if met, allow to keep the better theoretical properties of multi-dimensional topologies

Topology exploration framework

- Reference NoC architecture
 - Transaction Level models
- Traffic pattern generation

Exploration of multi-dimensional topologies

- System-level performance analysis
- Implementation space exploration
- TLS driven physical synthesis

Topology exploration framework

Reference NoC architecture

- Transaction Level models
- Traffic pattern generation

Exploration of multi-dimensional topologies

- System-level performance analysis
- Implementation space exploration
- TLS driven physical synthesis

Reference NoC architecture

Xpipes-Lite switch architecture



- Input and output sampling
- Latency: 1 cycle in the switch, 1 cycle in the link
- Wormhole switching
- Round-robin arbitration on the output ports

Reference NoC architecture

The Network Interface



- Protocol conversion (from OCP to network)
 - ✓ Packetization
- Clock domain crossing
 - ✓ OCP Clock is an integer divider of NoC Clock
- Pre-computation of routing path (source routing)
- A symmetric network interface target exists

Topology exploration framework

Reference NoC architecture

Transaction level models

Traffic pattern generation

Exploration of multi-dimensional topologies

- System-level performance analysis
- Implementation space exploration
- TLS driven physical synthesis

Transaction Level Models



Our transaction level models thus achieve maximum accuracy

Each cycle, only components affected by an event require simulation time •Speed-up is dependent on the system idleness

Network Interface Master



Network Interface Master



Switch Fabric



Switch Fabric



TL Models Validation





Several OCP traffic patterns (parameters: burst length and inter-burst idle time)

TL Models Validation

- Maximum error of all the tests was: 0.03%
- Simulation speedup varied from 20x to 100x with respects to RTL simulator
 - Depends heavily on the number of idle cycles of the simulation
- 4x4 mesh test:
 - Maximum error: 0.01%
 - Speed-up: ~100x

Topology exploration framework

Reference NoC architectureTransaction level models

Traffic pattern generation

Exploration of multi-dimensional topologies

- System-level performance analysis
- Implementation space exploration
- TLS driven physical synthesis

Tile Architecture



• Processor core

Connected through a Network Interface Initiator

- Local memory core
 - Connected through a Network Interface Target
- Two network interfaces can be used in parallel

Communication protocol

- Step 4: Consumerchedds
 budific an photoconflor
 complete messages for
 the destination
 Requests proposer for
 - topdests probateer for topdests probateer for the set is not intersecting to the set is a semaphore at the



consumar tila

- Message sent only when consumer is ready to read it
- Only one outstanding message for a producer-consumer pair

Dadla Tetwerk benelwidthPutlizeuienan Efficient Communication Library for Embedded Stigentalewcyticouster Platson" the Epidlogy hop on Embedded Systems for Real-Time Multimedia, 2007.

Workload distribution

- Producer, worker and consumer tasks
- I/O devices dedicated to input OR output data
- Modeling of layout constraints (I/O devices on one side of the chip)

Topology exploration framework

- Reference NoC architecture
- Transaction level models
- Traffic pattern generation

Exploration of multi-dimensional topologies

System-level performance analysis

- Implementation space exploration
- TLS driven physical synthesis

System-level performance analysis

- Tile-based architecture
- 16 tiles system

- Up to 5 tiles used for access external I/O

- Baseline topology 4x4 mesh (4-ary 2-mesh)
 - Switch frequency: 1GHZ
 - Tile Frequency: 500 MHz
 - External I/O frequency: 500 MHz

Topologies Under Test

	4-ary 2-mesh	2-ary 4-mesh
Switches	16	16
Bis. Band.	4	8
Tiles x Switch	1	1
Switch Arity	6	6
Max. Hops	6	4



4-ary 2-mesh Baseline Topology



2-ary 4-mesh High Bandwith

Topologies Under Test

	4-ary 2-mesh	2-ary 4-mesh	2-ary 3-mesh
Switches	16	16	8
Bis. Band.	4	8	4
Tiles x Switch	1	1	2
Switch Arity	6	6	7
Max. Hops	6	4	3



4-ary 2-mesh Baseline Topology



2-ary 3-mesh Low concentration degree

Topologies Under Test

	4-ary 2-mesh	2-ary 4-mesh	2-ary 3-mesh	2-ary 2-mesh
Switches	16	16	8	4
Bis. Band.	4	8	4	2
Tiles x Switch	1	1	2	4
Switch Arity	6	6	7	10
Max. Hops	6	4	3	2



4-ary 2-mesh Baseline Topology



2-ary 2-mesh High concentration degree

Scenarios

- Performance scenarios defined by:
 - Number of producer and consumer (I/O) tiles
 - Computation time of the worker
- 4 different performance scenarios
 - Worker
 - Consumer
 - Producer
 - Balanced
- At first only execution cycles considered
- Real frequency considered in next step

Worker Scenario



Execution cycles normalized to the performance of the 2D-mesh

- Bottleneck in the workers
- Producer are fast enough to feed workers with data
- Consumer are fast enough to absorb output data from workers
- Producer and consumer experience idleness

- Waiting for workers to process data

Network is not the bottleneck, performance scenario is not very topology-sensitive
Choice among topologies should be based on physical implementation considerations

Consumer Scenario



Execution cycles normalized to the performance of the 2D-mesh

- Bottleneck in the consumers
- Producer are fast enough to feed workers with data
- Consumer are NOT fast enough to absorb output data from workers

•Workers waits almost 50% of time to send data to the consumers

•Network latency-sensitive scenario

-Concentrated topologies outperform the others (in execution cycles)

Producer Scenario



Execution cycles normalized to the performance of the 2D-mesh

- Bottleneck in the producers
- Producer are NOT fast enough to feed workers with data

Network is not the bottleneck, scenario performance insensitive to topology
Choice among topologies should be based on physical implementation considerations

Balanced Scenario



Execution cycles normalized to the performance of the 2D-mesh

- Balanced scenario.
 - Minimized idle time of all tiles in the system (producer, consumer and worker)

Highest bandwidth pressure to the network
4-hypercube provides more bandwidth
Concentrated topologies trade bandwidth for low latency
Worst performance than the 4-hypercube, but still outperform the 2D-mesh

Topology exploration framework

- Reference NoC architecture
- Transaction level models

Traffic pattern generation

Exploration of multi-dimensional topologies

System-level performance analysis

Implementation space exploration

TLS driven physical synthesis

Implementation Space Exploration

- Scenarios where topology is relevant for performance were selected
 - Exploration of possible implementation effects
 - Results will drive the synthesis process as optimization directives
- Selected scenarios
 - Consumer Scenario
 - Producer Scenario

Balanced Scenario

- Implementation space restricted to the baseline 2Dmesh compared with the best performing topology: 4-hypercube
- Switch arity is the same for both topologies
 - Same switch frequency
- Possible latency degradation at the express links





2-ary 4-mesh High Bandwith

Balanced Scenario



Breakeven point occurs with 5 cycles latency on the express links

Consumer Scenario

- Implementation space restricted to the baseline 2D-mesh compared with the best performing topology: 2-hypercube
- Possible physical degradation effects:
- Maximum achievable frequency (switch arity)
- Several layouts possible to connect more cores per switch:
 - Latency in the network links
 - Latency in the injection links

Different layouts might impact performance in a very different way

Network Link Constrained Layout

- 4 tiles placed around the switch.
 - Network links might be multi-cycle
- 2-ary 2-mesh has switches with higher radix



Under which min frequency and link latency concentrated hypercube still outperforms 2D mesh?

Network Link Constrained Layout

2-ary 2-mesh —— 4-ary 2-mesh ----+---



- Switch frequency of 900 MHz allows no additional delay
- Switch frequency of 1 GHz allows up to 4 cycles of additional delay

Injection Link Constrained Layout

- Central network around which all cores are placed
 - Injection and ejection links might be multi-cycle
- Test performed by scaling clock frequency together with latency in injection and ejection links



2-ary 2-mesh

Injection Link Constrained Layout

4-arv 2-mesh

Concentrated hypercube Time (ns) 1.4e+07 1.3e+07 1.2e+07 1.1e+07 1e+07 2D mesh 9e+06 8e+06 7e+06 6e+06 0.5 Frequency (Ghz) Latency

- Switch frequency of 900 MHz allows no additional delay
- Switch frequency of 1 GHz allows up to 2 cycles of additional delay
- Injection links latency has higher penalization over performance

Topology exploration framework

- Reference NoC architecture
- Transaction level models
- Traffic pattern generation



- System-level performance analysis
- Implementation space exploration

TLS driven physical synthesis



4-hypercube

2-D mesh

4-hypercube should require twice the wiring resources of the 2D-mesh

Asymmetric tile size makes the traditional assumptions on mesh and hypercube wiring questionable

Placement-aware logic synthesis and place-and-route on a STMicroelectronics 65nm Low Power library

Computation tiles are rendered as *hard obstructions Fences* are defined to limit the area where the cells of each module of the interconnect can be placed



worst case: we prevented over-the-cell routing for the hard obstructions, thus posing more constraints on network link routing



NI critical path: **0.65ns** Switch maximum delay: **0.84ns**

2D-mesh :

Network maximum link delay: **0.45ns** Network can sustain the target **1GHz**.

4-hypercube:

Maximum network link delay : **0.91ns**. Network can sustain the target **1GHz**.

Both topologies can address the target frequency (1GHz) without pipelining in the express links of the hypercube



Wire length report



Total wire length increases in the 4-hypercube only by 25%

- Limited size of the system
- Asymmetric physical size of computation tiles
- High percentage of the wiring in the component (Switch, NI, etc..)

Conclusions

- Development of an analysis framework for NoC topologies
 - Fast and accurate TL simulation
 - Exploration of the implementation space
 - Drive the synthesis process
 - Conservative BW utilization of the communication middleware taken into account
- Methodology is applied to a 16-tile system
 - There is already performance differentiation despite the limited system scale
- Automatic routing of these topologies might provide surprising results
 - Physical degradation of hypercubes can be relieved by
 - Smart switch placement
 - Asymmetric tile size
 - Performance-wise hypercubes can be an efficient and feasible solution

Future Work

- Extend analysis framework to highly integrated MPSoC

 Hundreds of tiles
- Set up a power characterization flow
- Analysis of layout feasibility of concentrated topologies
 - Possible trade-off between power/area and performance

Backup Slides

Reference NoC architecture

Stall/go flow control



Simplified realization of an ON/OFF flow control protocol

- A repeater is a simple two-stage FIFO
- Two control wires:
 - ✓ forward : flagging data availability
 - ✓ backward : signaling either a condition of buffers filled
 (STALL) or of buffers free (GO)

Sender needs two buffers to cope with stalls in the very first link repeater

Transaction Level Models



Not fast enough for topology exploration

Network Interface Slave



Flow Control



TL Models Validation



Network frequency: 1GHz OCP Core frequency: 500 MHz Newort/OCP Clock ratio: 2

Tests designed to check intra-switch communication mechanisms

TL Models Validation



Tests designed to check clock domain crossing mechanism



Tests designed to check inter-switch communication mechanisms

TL Models Validation

- Every test evaluated with different OCP traffic configurations
- Maximum error of all the tests was: 0.03%
- Simulation speedup varied from 20x to 100x with respects to RTL simulator
 - Depends heavily on the number of idle cycles of the simulation

Validation of TL simulator





Maximum error: **0.01%** Speedup of about 100x



NI critical path: **0.65ns** Switch maximum delay: **0.84ns**

2D-mesh :

Network maximum link delay: **0.45ns** Network can sustain the target **1GHz**.

4-hypercube:

Maximum network link delay : **0.91ns**. Network can sustain the target **1GHz**.

Link delay overhead is due to logic gates along the link (buffers, flow control cells)

Both topologies can address the target frequency (1GHz) without pipelining in the express links of the hypercube