



SHAPING THE NEXT GENERATION OF ELECTRONICS

**JUNE 23-27, 2024**

MOSCONE WEST CENTER  
SAN FRANCISCO, CA, USA



**JUNE 23-27, 2024**

MOSCONE WEST CENTER  
SAN FRANCISCO, CA, USA

# NoC Efficient Re Design


Jonathan Ezroni

 **mobileye**<sup>™</sup>



# What is a NoC

- a packet switched platform
- single chip system
- scales well to an arbitrary number of processor like resources



*S. Kumar et al., "A network on chip architecture and design methodology," Proceedings IEEE Computer Society Annual Symposium on VLSI. New Paradigms for VLSI Systems Design. ISVLSI 2002, Pittsburgh, PA, USA, 2002, pp. 117-124*




# Design Process


- (Allegedly) Two phases of design:
  1. Concrete Architecture of network
  2. map application to concrete architecture

*S. Kumar et al., "A network on chip architecture and design methodology," Proceedings IEEE Computer Society Annual Symposium on VLSI. New Paradigms for VLSI Systems Design. ISVLSI 2002, Pittsburgh, PA, USA, 2002, pp. 117-124*

# 1D Subsystem Example - design phase 1

- Simple Mesh
- 4 processing cores (P)
- 1 L2 logical memory






# Map to physical domain


## Routing Hotspot

- 2 Buses for EW
- 2 Buses for NS



# Map to physical domain


- [Jira-0001] transform mesh to hierarchical topology
- 1 Bus for EW
- 1 Bus for NS






# Map to physical domain

- [JIRA-0002] additional buffering required for throughput bottlenecks
- you could increase BW but don't want to risk routing congestion again



# Sensitivity to processor scaling


- As we scale number of processors
- [Jira-0003] add pipelining to close timing\*



(\*) Partial solutions from IP vendors. Will eventually be turn-key solution.

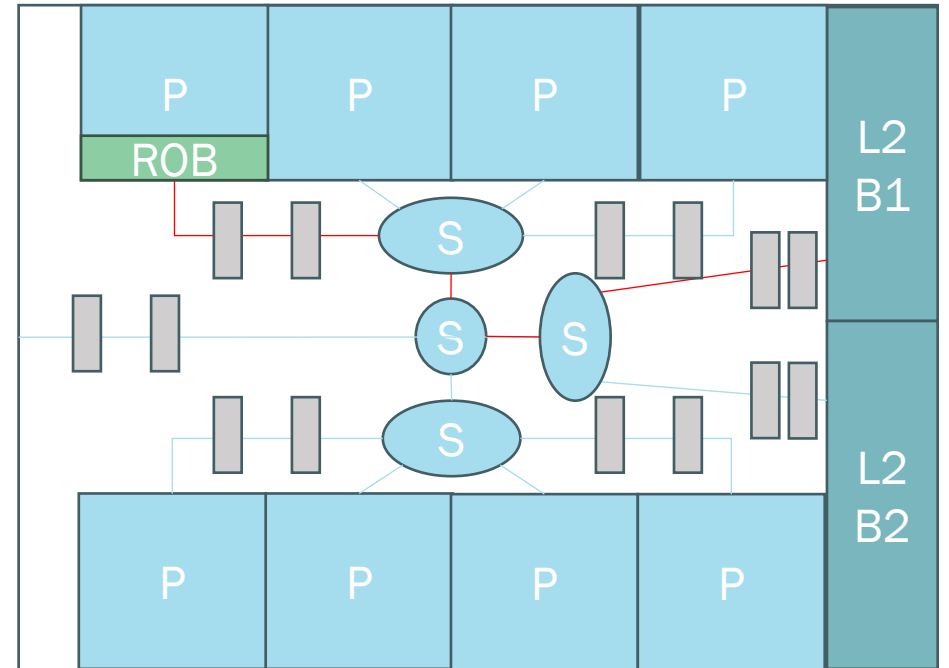
# Map to physical domain

- Not all EDA tools support reorder buffers outside of NoC master abstraction



# Generality has its costs

- [JIRA-004] Add buffers for worst case latency to accommodate for pending transactions
- Not all EDA tools take workload into account
- All Processors pay same penalty
- internal node buffer size is based on worst case analysis (think GBA vs. PBA)



# Design Process


- (Allegedly) Two phases of design:
  1. Concrete Architecture of network
  2. map application to concrete architecture

*S. Kumar et al., "A network on chip architecture and design methodology," Proceedings IEEE Computer Society Annual Symposium on VLSI. New Paradigms for VLSI Systems Design. ISVLSI 2002, Pittsburgh, PA, USA, 2002, pp. 117-124*



# Map Application

- System says:
  - All core parallel access to L2  $\geq X$  [Mbps]
  - 2 cores parallel access to same resource  $\geq Y$  [Mbps]
  - SW requires Memory Size X2 increase



# Reality NoC-ing on our door

- Verification says:
- [Jira-005] Test Fail 1 - All core parallel access to L2  $\geq$  X[Mbps]
- [Jira-006] Test Fail 2 - cores parallel access to same resource  $\geq$  Y[Mbps]
- [Jira-007] SW requires Memory Size X2 increase
  - Hey Timmy, Go tell Backend to go back to floorplanning




# Water Falls, Engineers Redesign

- Present Status of Design Phase 1
  - Explain that the current BE model isn't worth the db size its written on.
- split blame in equal portions and assign to each EDA vendor.
  - sprinkle something extra on timmy, BE hates him anyway.
  - don't forget to mention NoC accounts for 1/3 of stdcell area.
- Tell your boss to update program committee to increase original schedule by a constant factor of  $e$ .
  - Your Boss says that is unacceptable, come up with another solution.
    - Work double shifts to fix JIRA-00[1 through 7]
    - apologize to wife
    - Tell said boss to increase your RSU allocation by a factor of  $e$ .



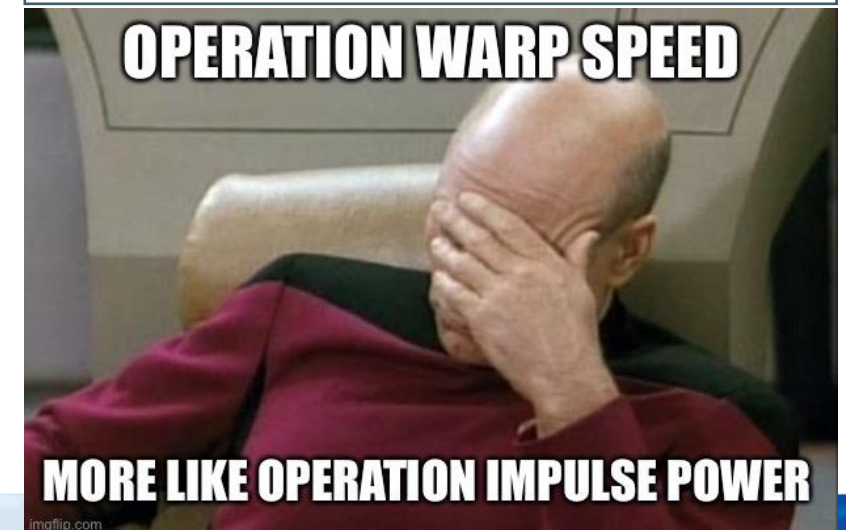

# Design Phase 2 - physical domain

- BE report +10% utilization increase due to performance bug fixes
  - [JIRA-008] replace large Reorder Buffers with area efficient memories
  - Forget to add ECC/Parity on them as a gift to Jimothy – the functional Safety lead – he has a mild heart condition.
- Timmy, tell BE we fixed Utilization and they need to add new memories and increase the size of the existing L2.
- Add +2 to layout complexity, you are no longer the simplest block in the chip.
- Wait for new floorplan
  - add pipe stages.
  - rinse and repeat.
- System update they forgot to tell you outgoing flows are also important
- reconsider the meaning of life



# 1D SoC Example

- clone yourself – you also design SoC level NoC
- roundtrip to DDR and back is 64 clock cycles
  - Derrick - the SW engineer, not jimothy, just suffers from a heart attack.
- Listen patiently as system team change Subsystem architecture of “Non-critical” control bus → KPI of a gazillion TFPS is science fiction.
- BE lead says you need to reduce I/F width by half → go back to Subsystem Design Phase 3





# End of Current Chapter

- after gathering multiple failures you have successfully established yourself as a NoC Ninja.
  - You have become indispensable -> your boss realizes your compensation increase rate is converging to  $e$
  - you will have to be in the same role until are 99 years old.

# Next Chapters...

- Unique addressing Part I, II, III
  - symmetric designs
  - systolic store and forward networks
  - Relative parametric target IDs in IPXACT
- NoC Physical Design Tools: Adding Routing/Placement blockages
- Automation: Change I/F width magic wand tool.
- Performance: Top Level performance Estimation
- FuSa: Implementing End-to-End data and control safety.

