Design Automation for Networks-on-Chip

Status and Outlook

Luca Benini

lbenini@deis.unibo.it

DEIS, Università di Bologna

Application Pull

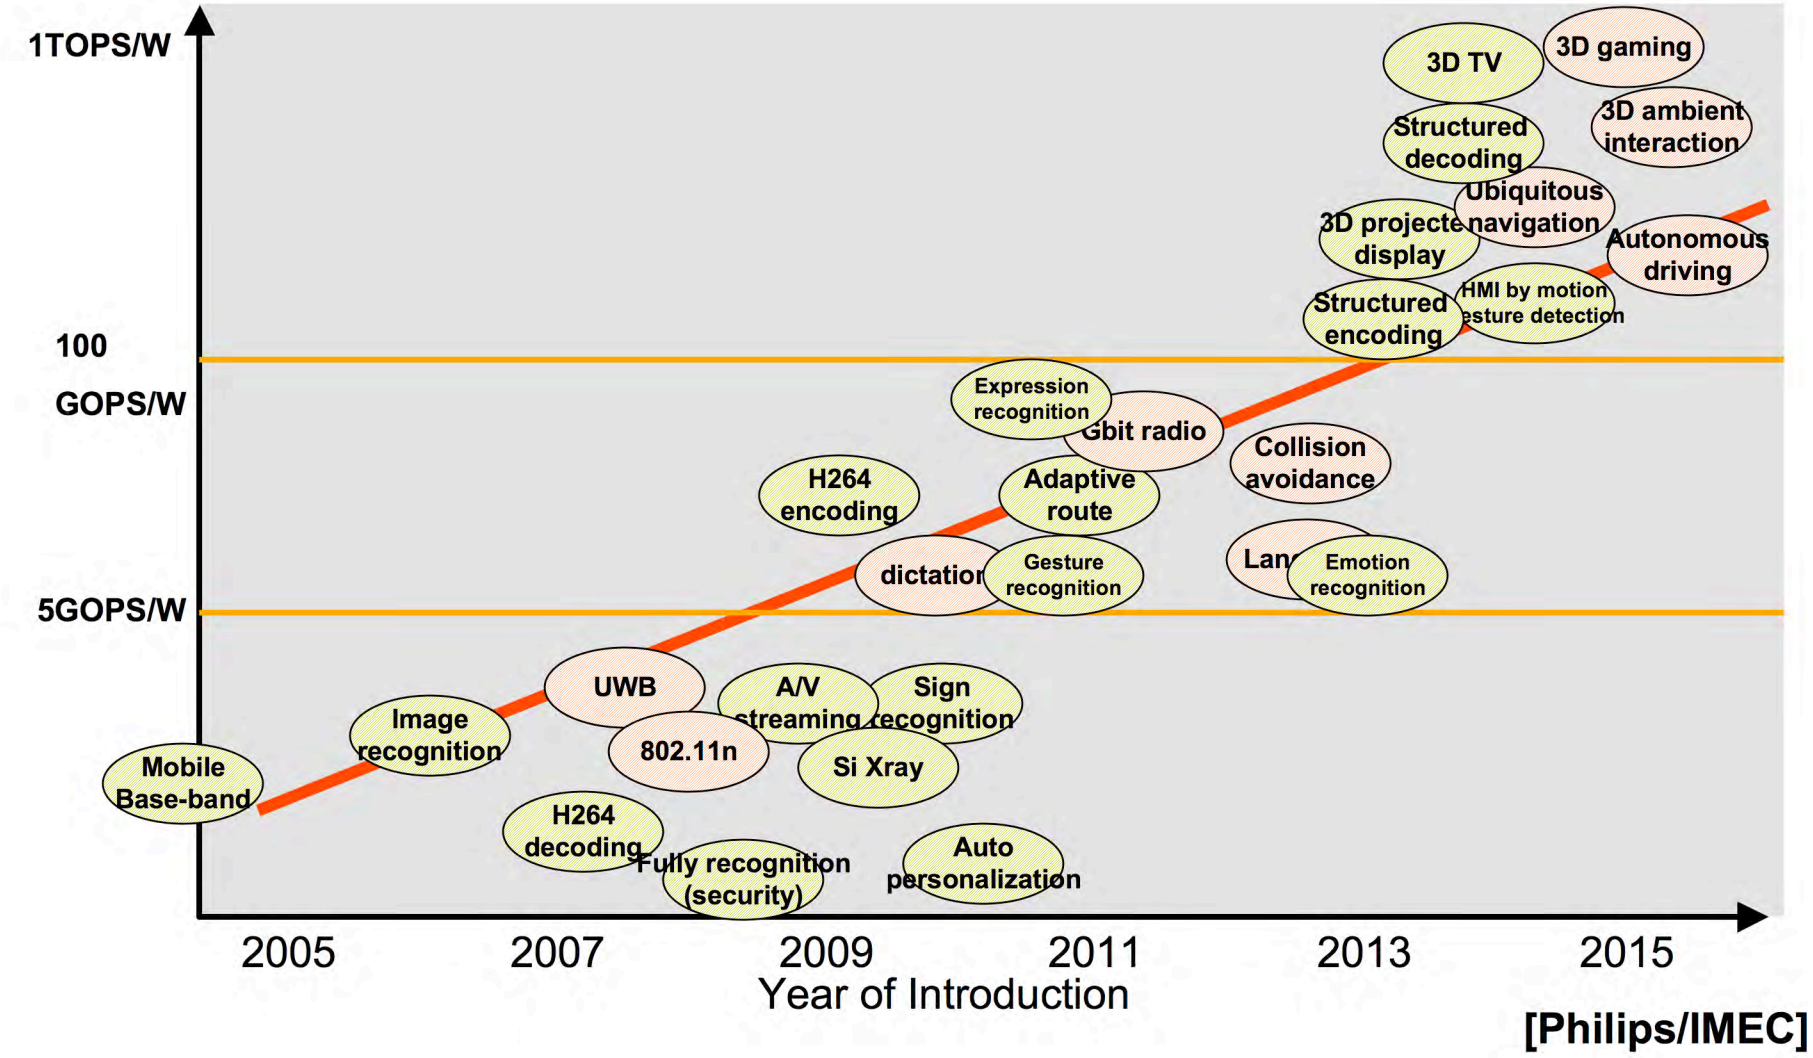

[Philips/IMEC]

Application Pull

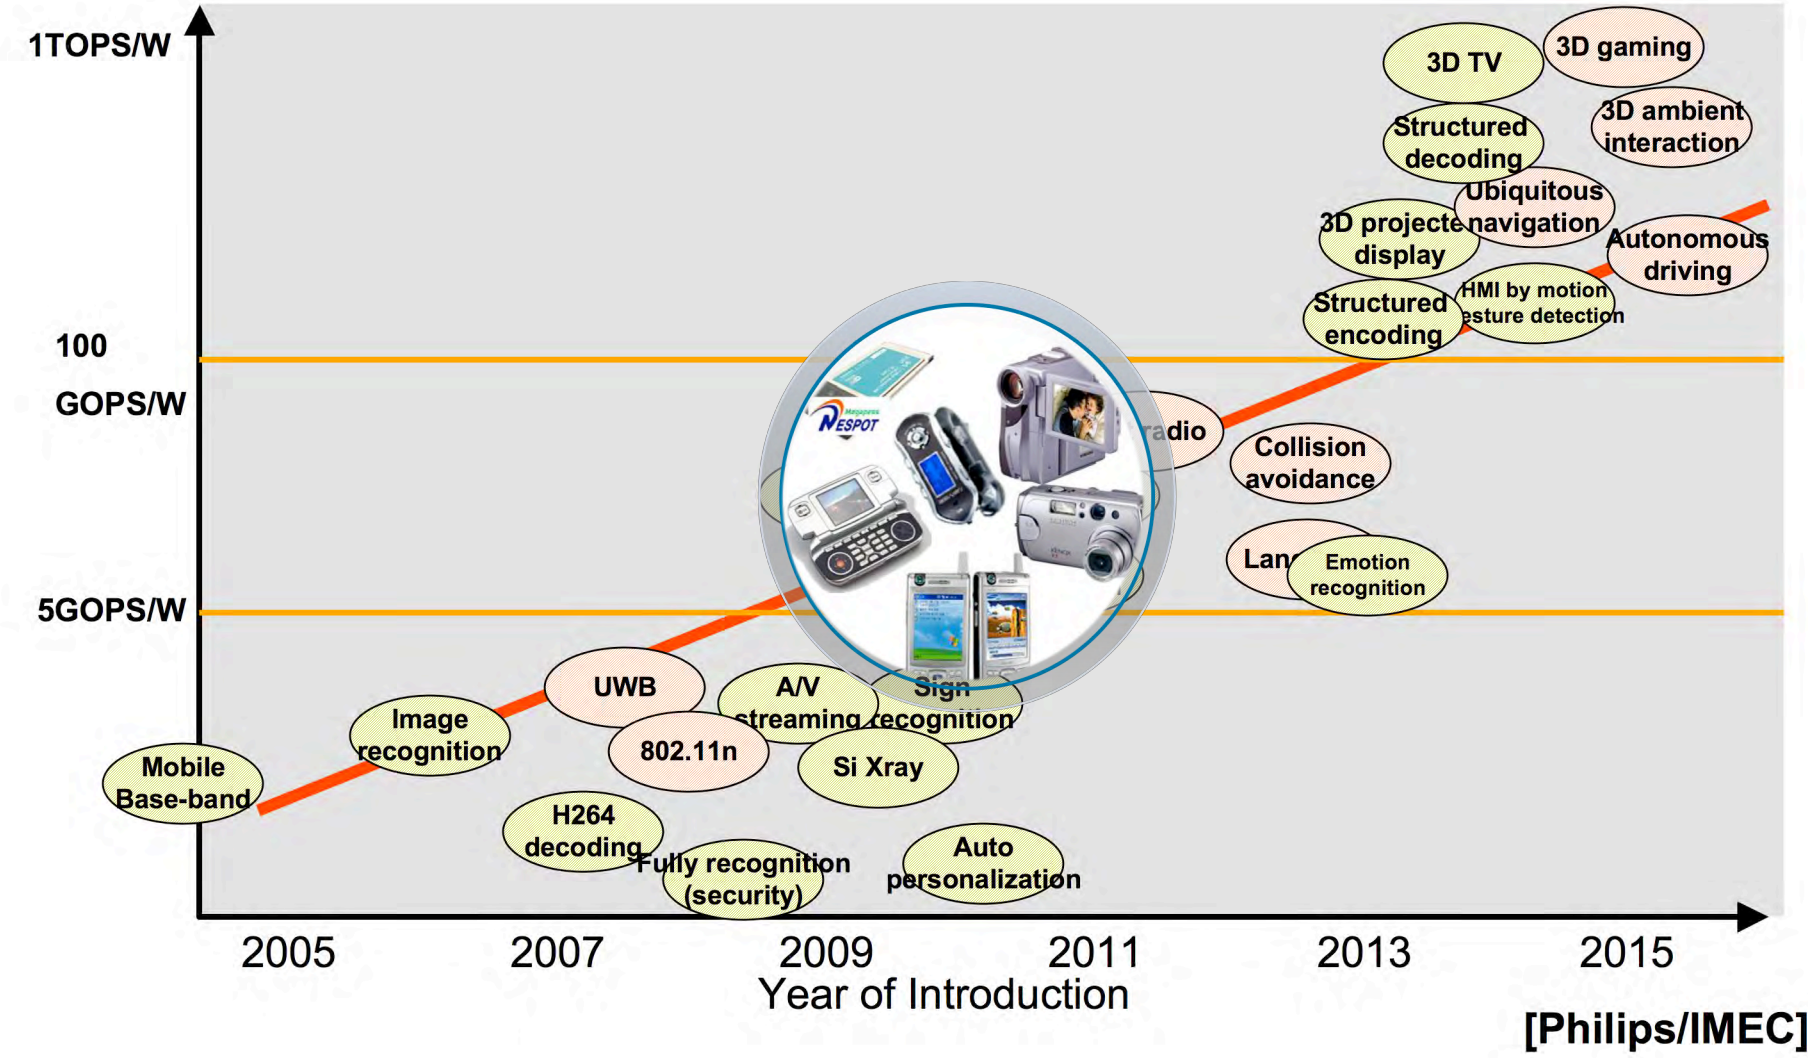

[Philips/IMEC]

Application Pull

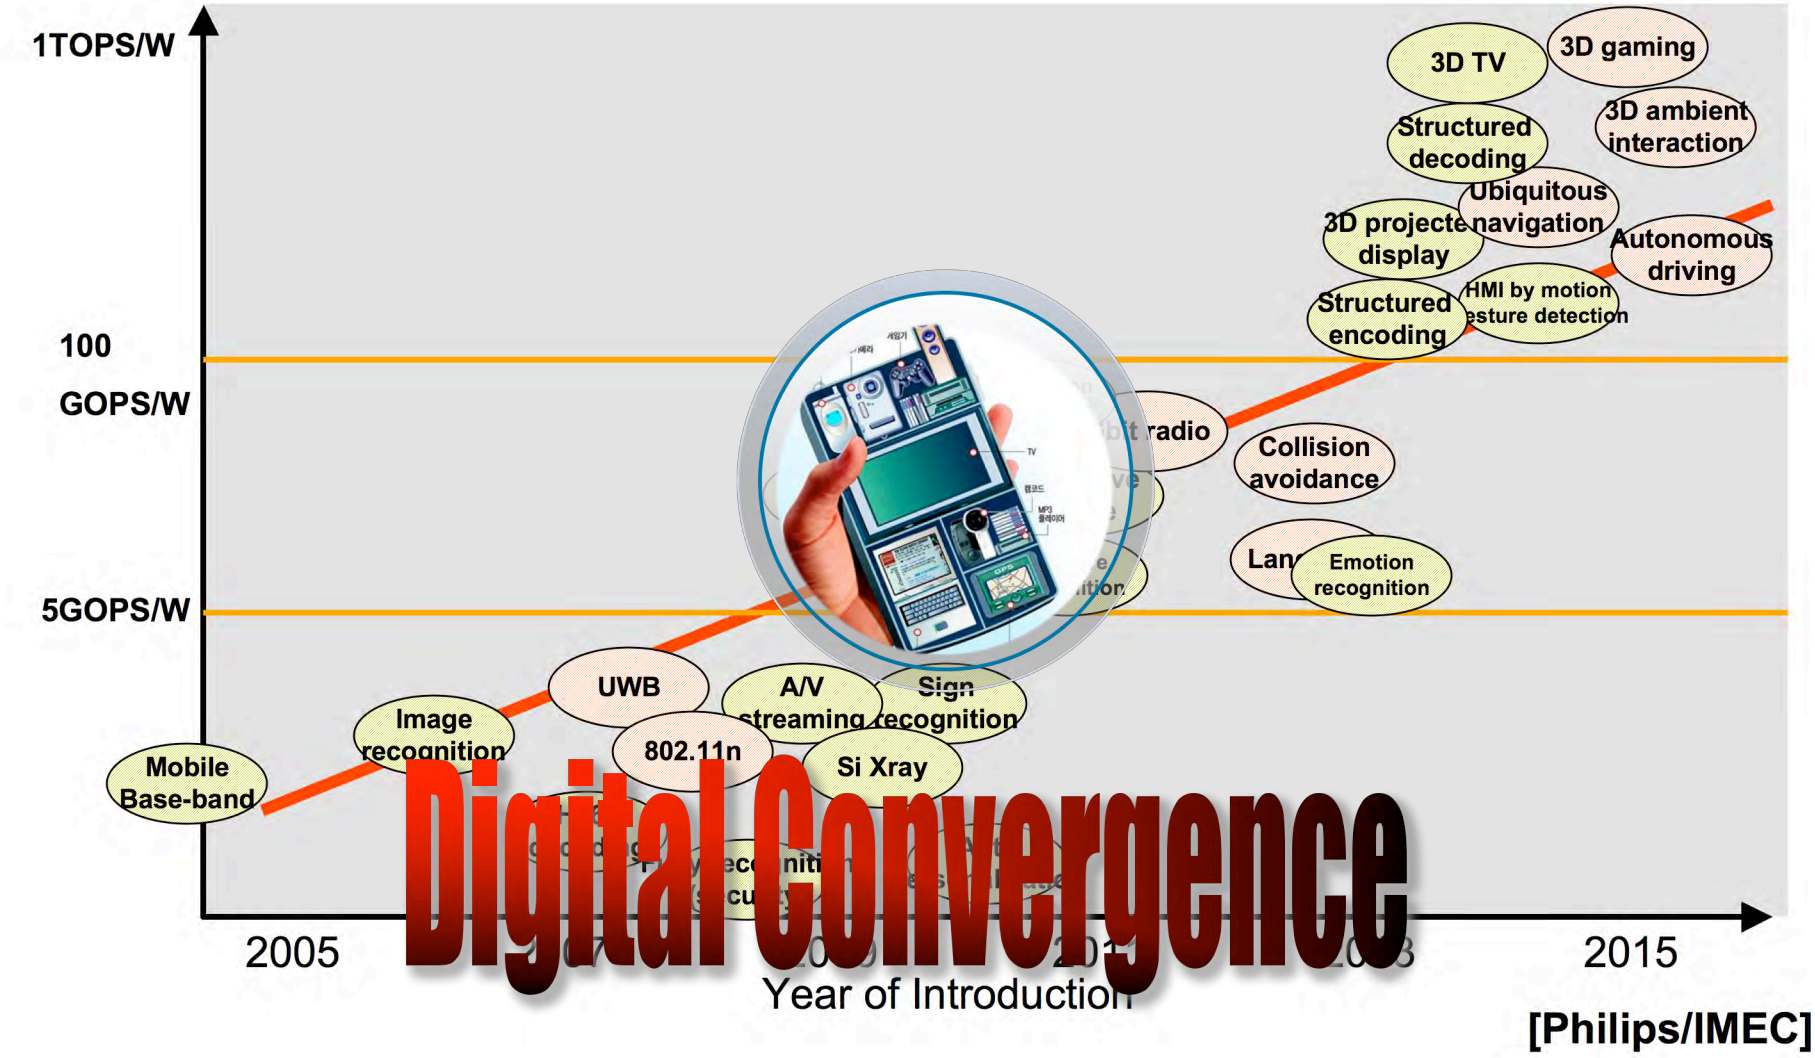

Embedded SoC Architecture Trends

Heterogeneous clusters

Embedded SoC Architecture Trends

Heterogeneous clusters

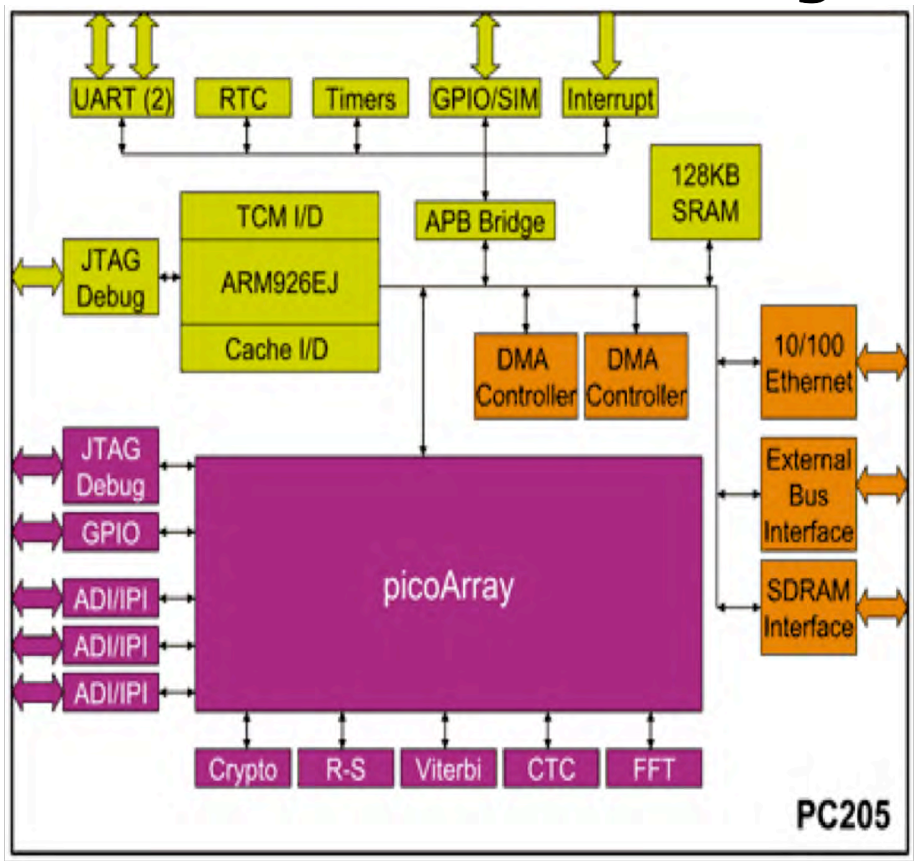

Wireless Networks Mesh nodes, Picocells

Picochip PC205 (Apr'06)

- 260MHz, 31GMAC/s, 160GIP/s
- 64KB I,D$, 128KB SRAM
- Less than 5 W, less than 1$/GMAC

Embedded SoC Architecture Trends

Heterogeneous clusters

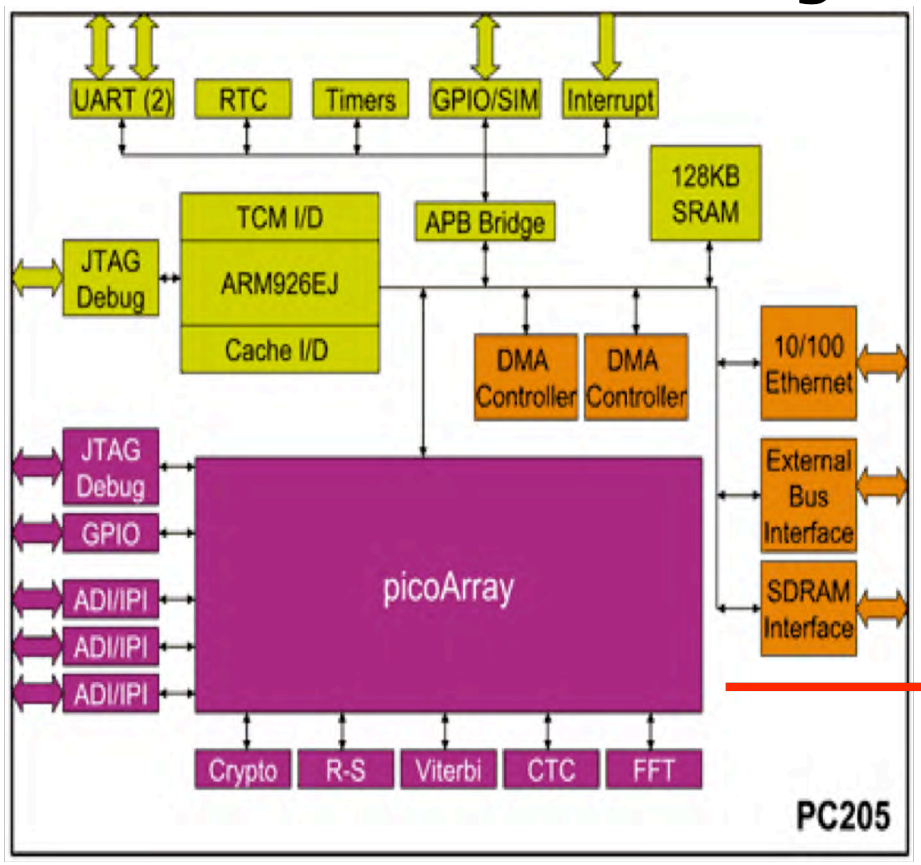

Wireless Networks Mesh nodes, Picocells

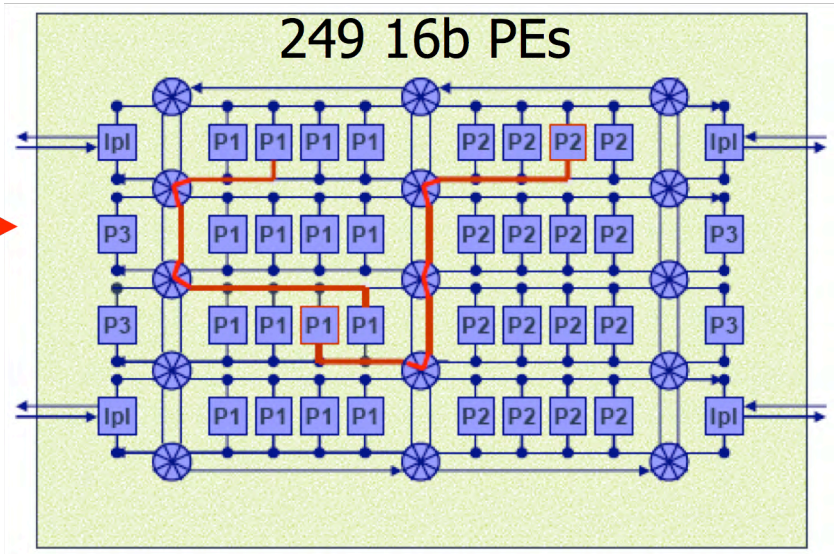

Picochip PC205 (Apr'06)

- 260MHz, 31GMAC/s, 160GIP/s
- 64KB I,D$, 128KB SRAM
- Less than 5 W, less than 1$/GMAC

Embedded SoC Architecture Trends

Heterogeneous clusters

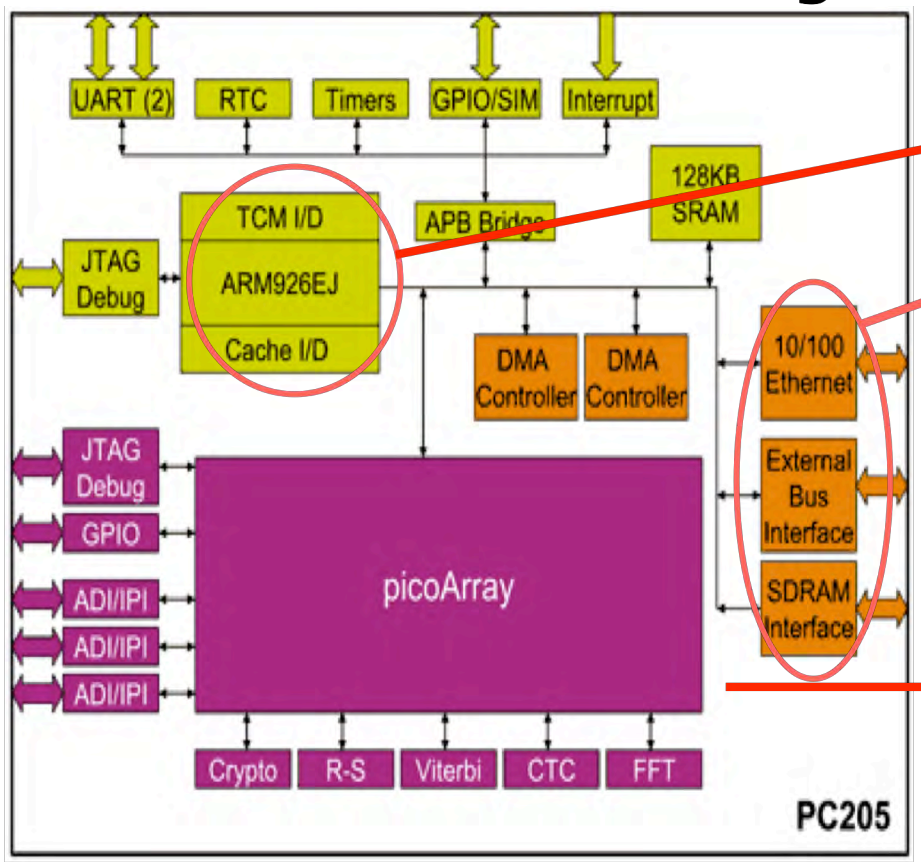

GP core

IOs

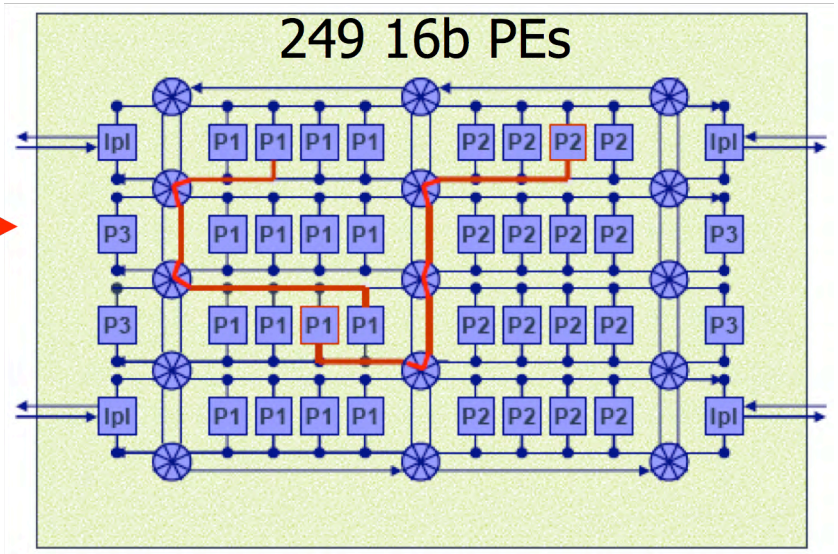

Wireless Networks Mesh nodes, Picocells

Picochip PC205 (Apr'06)

- 260MHz, 31GMAC/s, 160GIP/s
- 64KB I,D$, 128KB SRAM
- Less than 5 W, less than 1$/GMAC

Embedded SoC Architecture Trends

Heterogeneous clusters

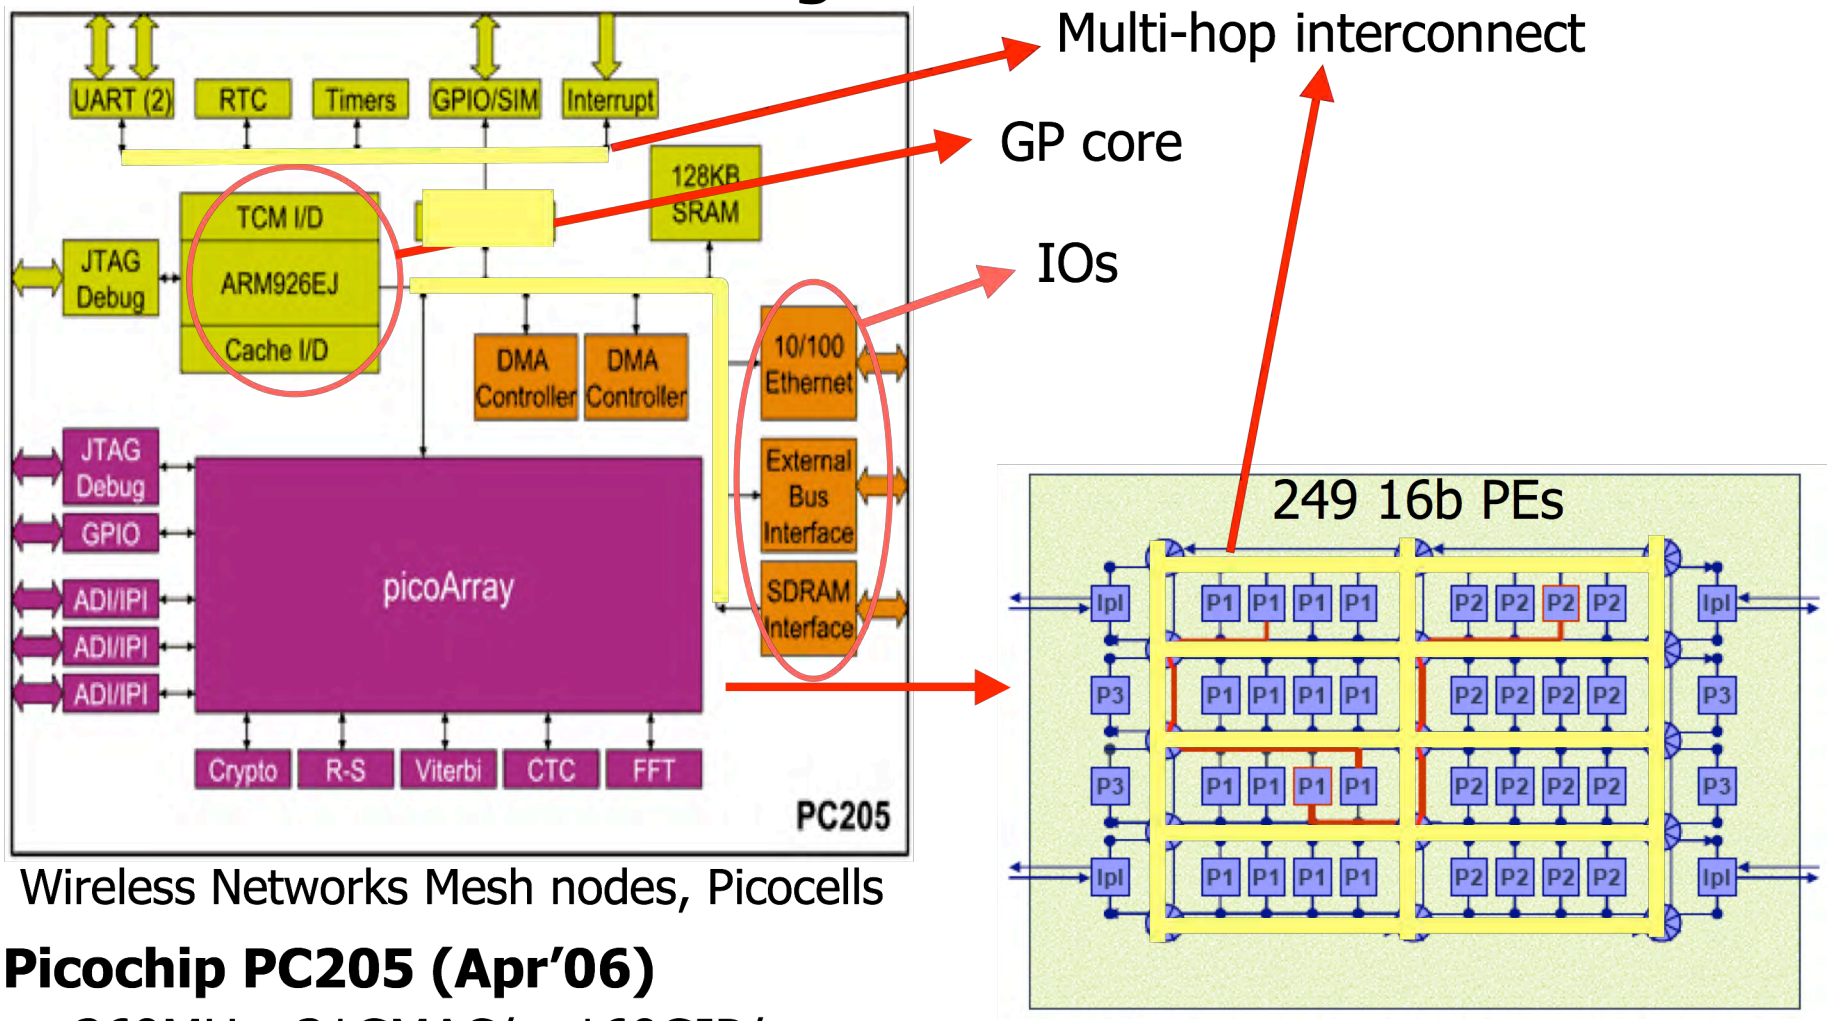

Wireless Networks Mesh nodes, Picocells

Picochip PC205 (Apr'06)

- 260MHz, 31GMAC/s, 160GIP/s
- 64KB I,D$, 128KB SRAM
- Less than 5 W, less than 1$/GMAC

Architecture Evolution

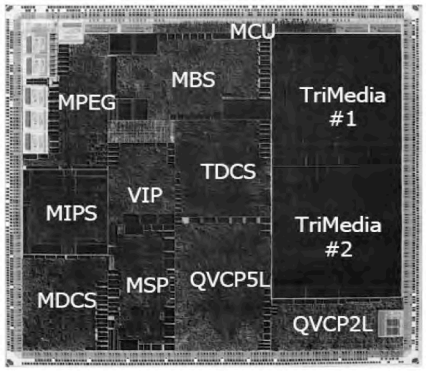

- Roadmap continues: 90→65→45 nm
- "Traditional" Bus-based SoCs fit in one tile !!

Architecture Evolution

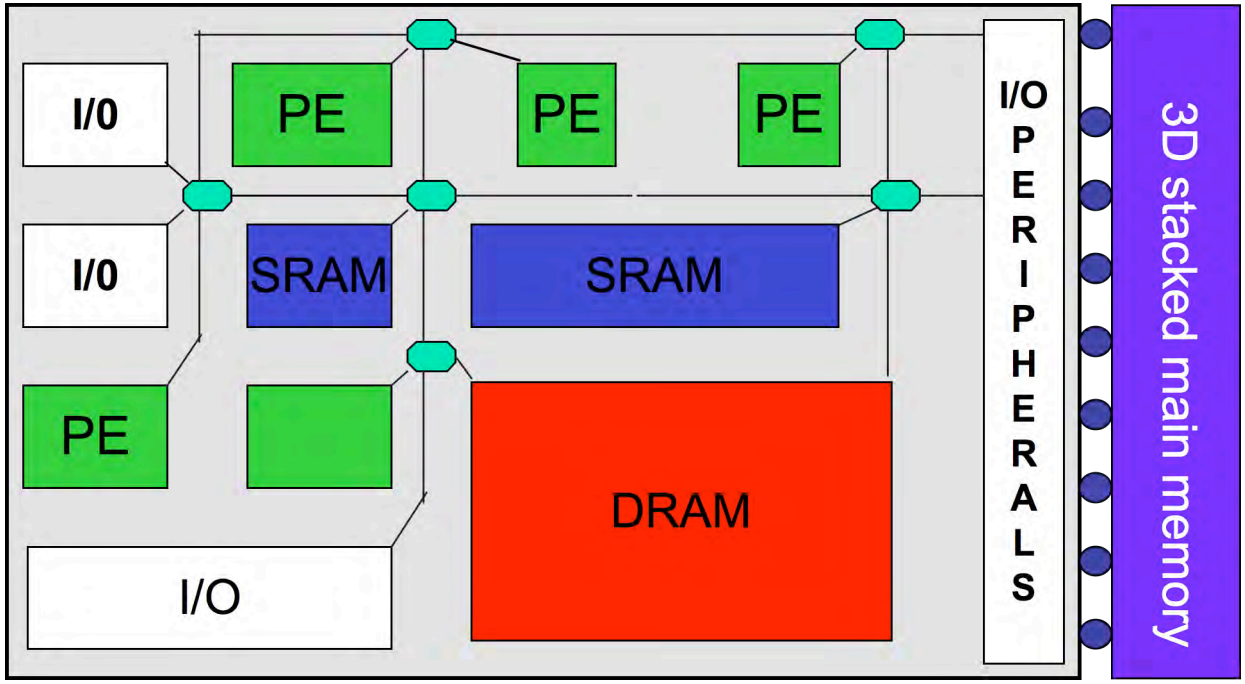

- Roadmap continues: 90→65→45 nm
- "Traditional" Bus-based SoCs fit in one tile !!

Architecture Evolution

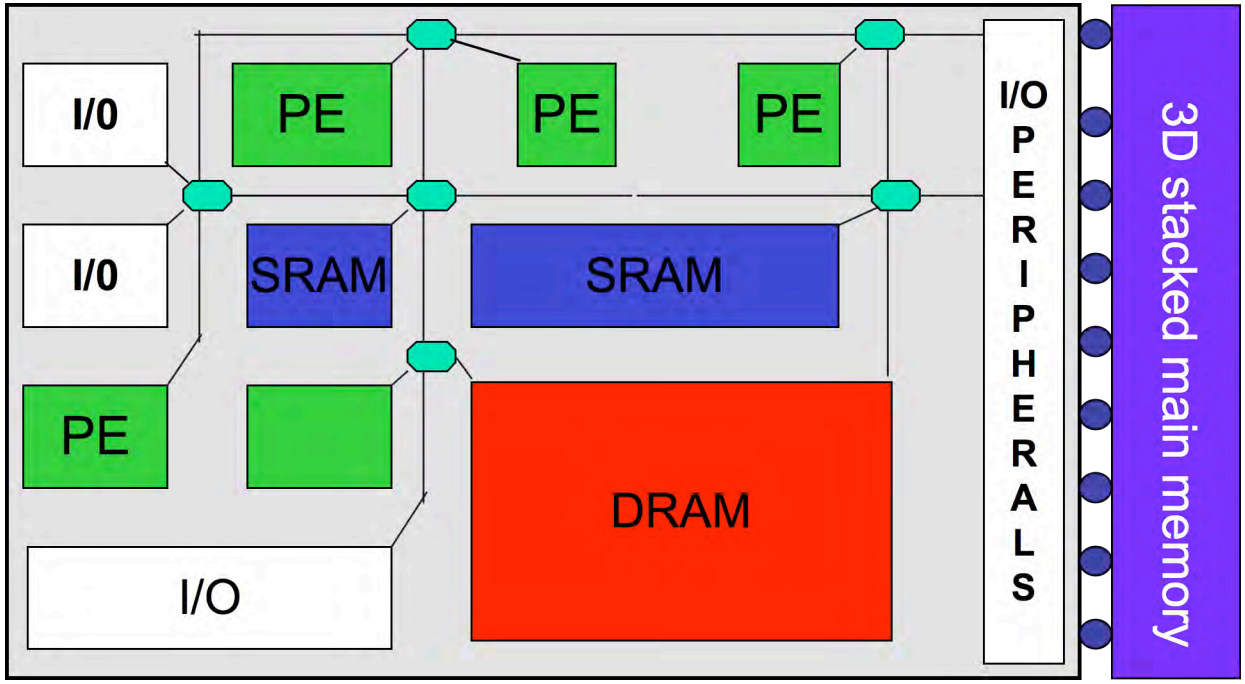

- Roadmap continues: 90→65→45 nm
- "Traditional" Bus-based SoCs fit in one tile !!
- Communication demand is staggering, but unevenly distributed, because of architectural heterogeneity

Interconnect Bottleneck

Power consumption
Unidirectional link
(38 bits+flow control)

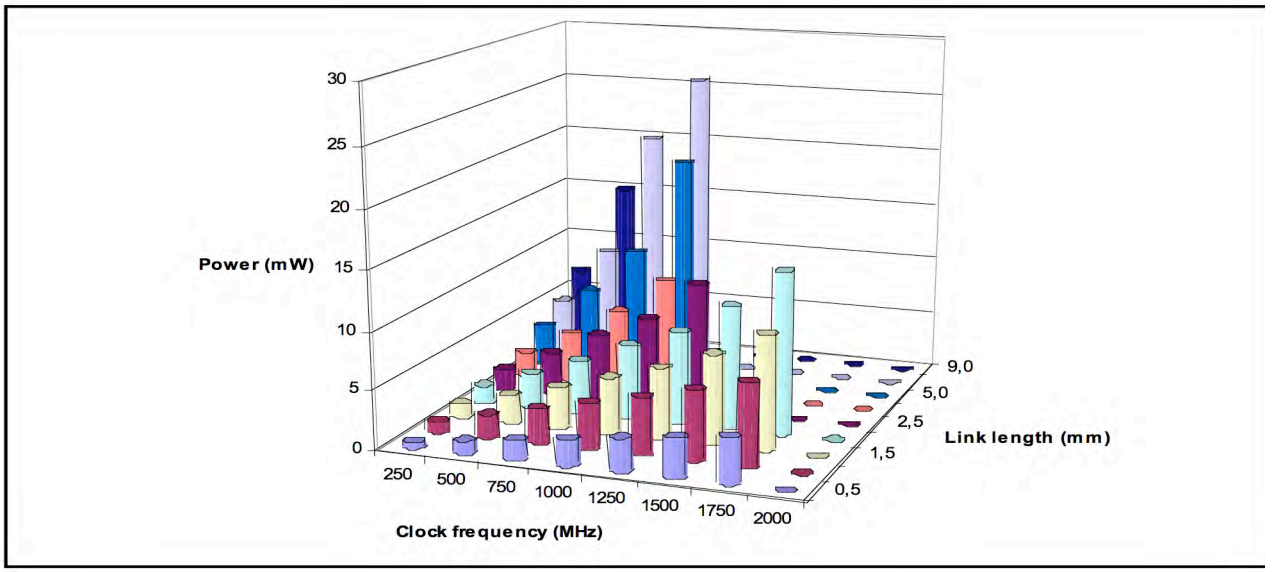

- 65 nm low-power library
- low V_t library, high V_{DD} – power/perf tradeoff
- very high frequencies or very long links unfeasible
- but even some feasible links burn up to 30 mW!!
  - heavy buffer insertion

Interconnect Bottleneck

Power consumption
Unidirectional link
(38 bits+flow control)

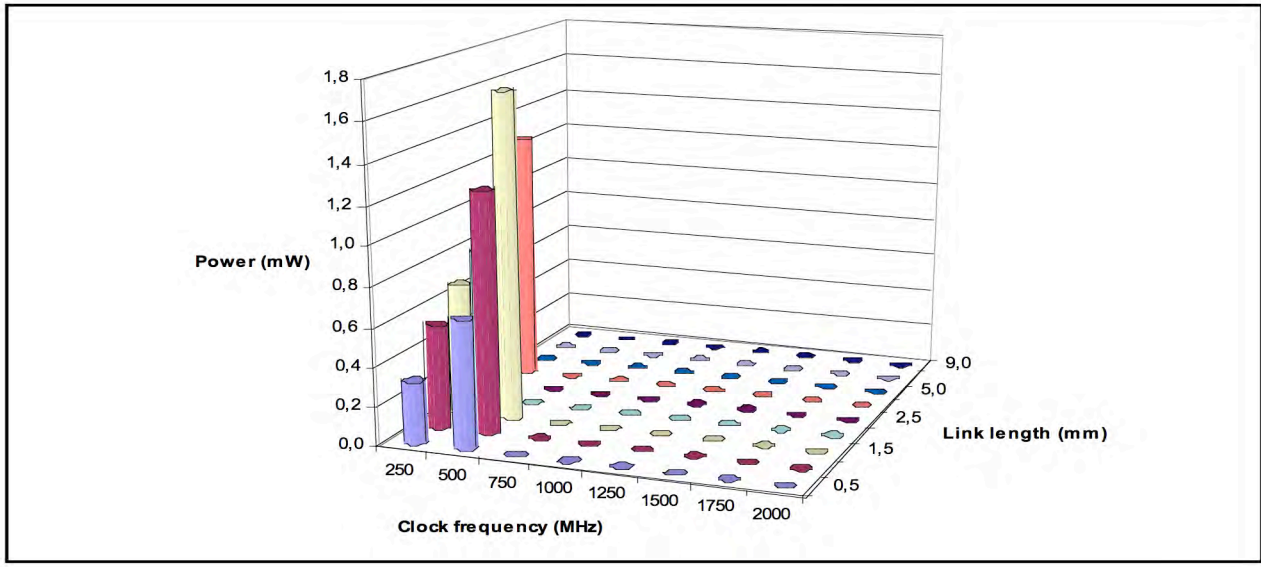

- 65 nm low-power library
- high V_t library, low V_{DD} – absolute min power
- even at 250 MHz, > 2 mm link length unfeasible

Addressing Interconnect Issues

- High-end industrial solutions:
  - Evolutionary path from shared busses

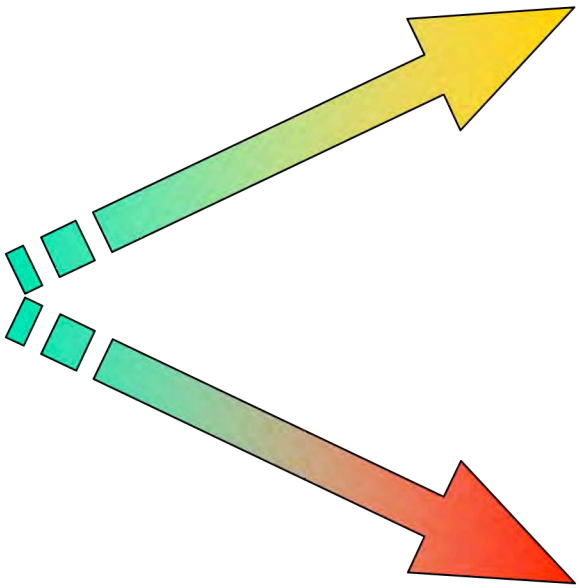

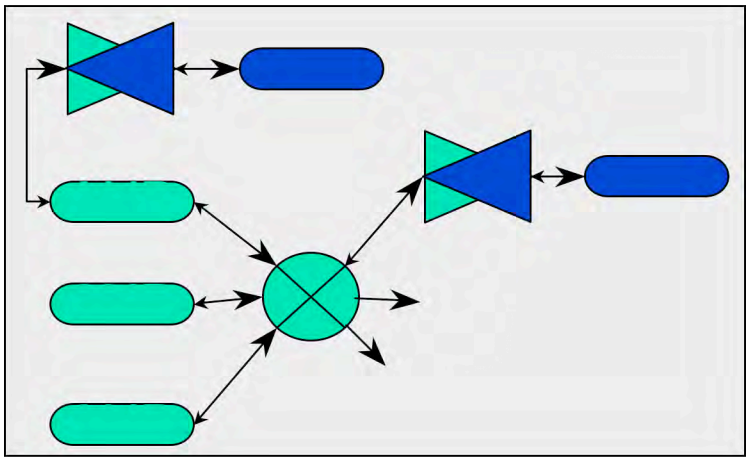

Addressing Interconnect Issues

- High-end industrial solutions:
  - Evolutionary path from shared busses

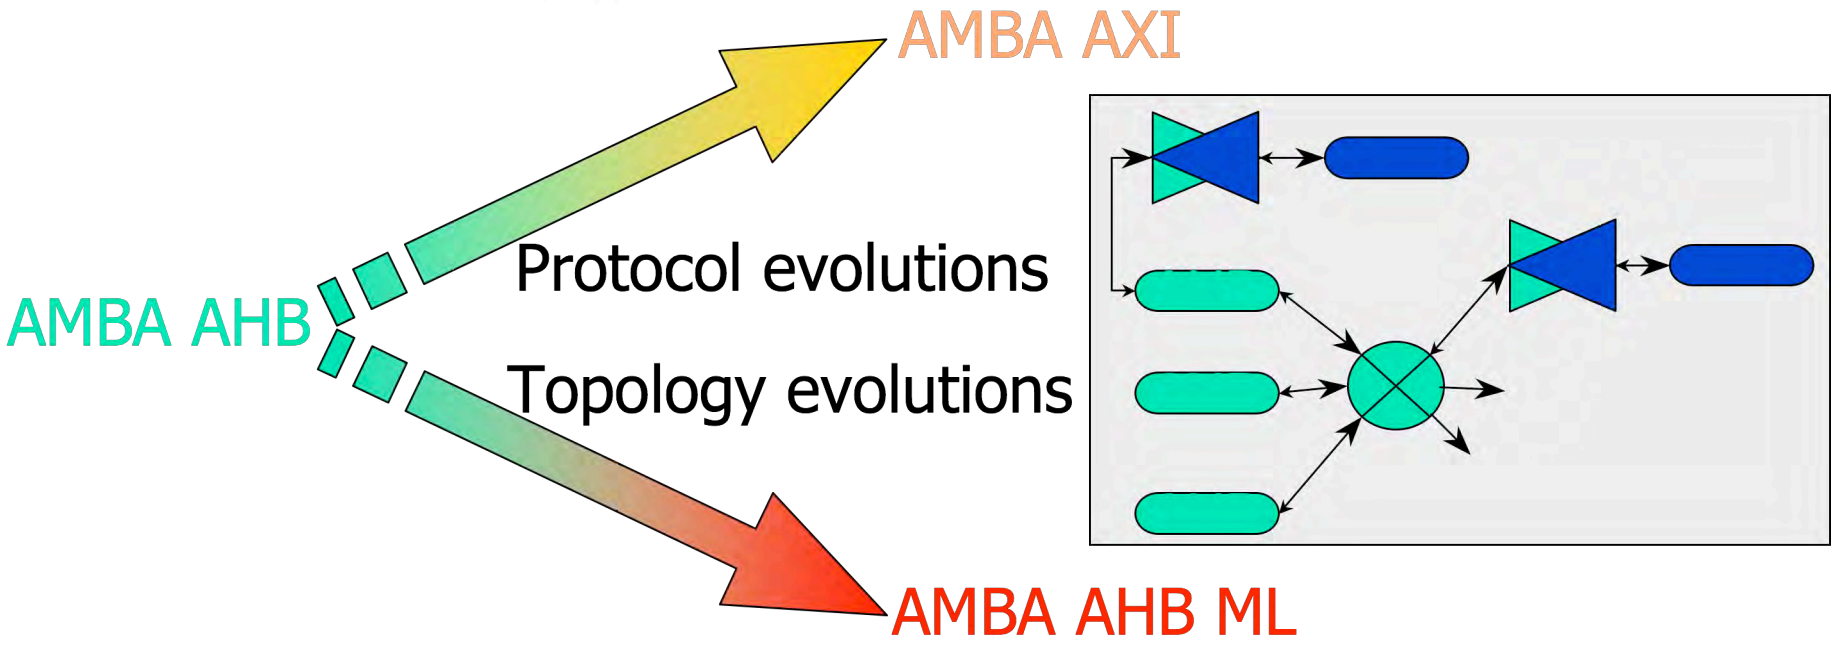

Addressing Interconnect Issues

- High-end industrial solutions:
  - Evolutionary path from shared busses

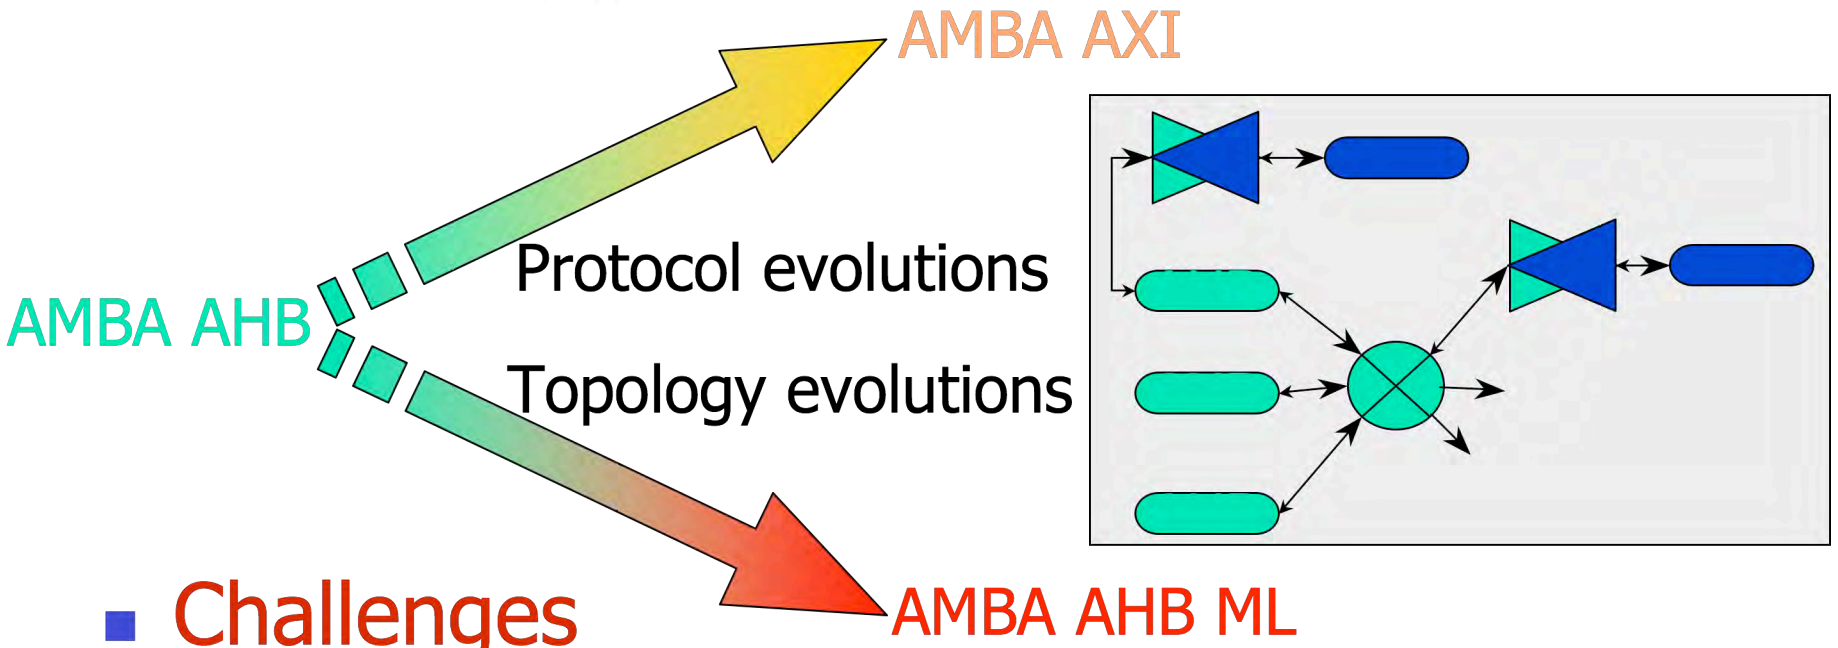

- Challenges
  - Complexity (e.g. 4-SHB + 2XBar, 75 actors): how to analyze and verify "spaghetti interconnects"?
  - Scalability: bus is bandwidth-limited, Xbar is size-limited
  - Predictability: how to tie interconnects with floorplanning

The Network-on-Chip Paradigm

The "power of NoCs":

- Clean separation at session layer
  - Cores issue end-to-end transactions
  - Network deals with transport, network, link, physical
- Modularity at HW level: only 2 building blocks
  - Network interface
  - Switch (router)
- Physical design aware (floorplan global routing)

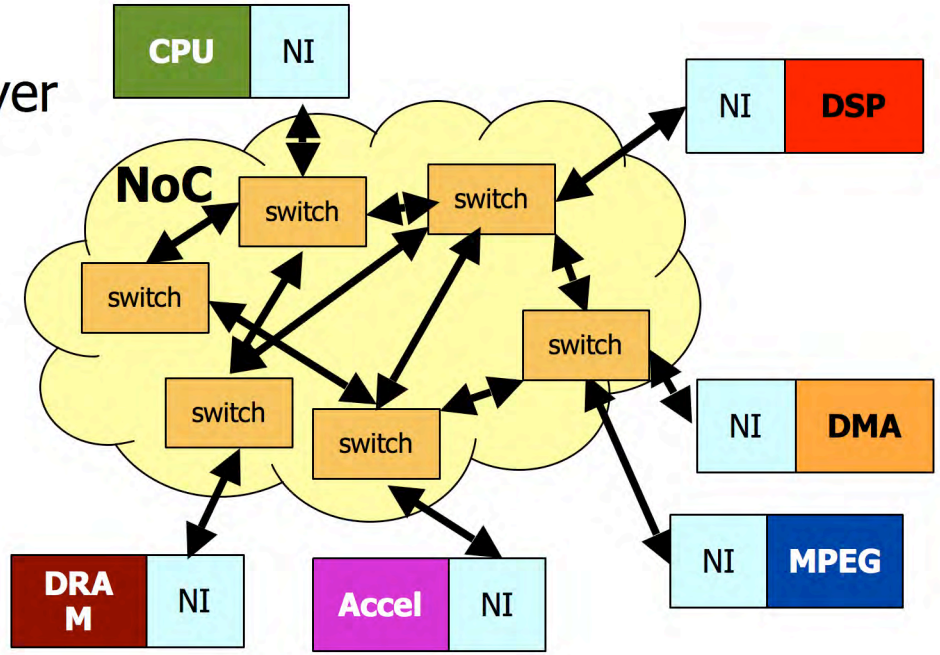

Scalability is supported from the ground up!

NoC Synthesis Project

- Started in 2002
- UNIBO, UNICA, Stanford, EPFL
- Objective: develop a complete EDA flow for NoC synthesis from application to P&R

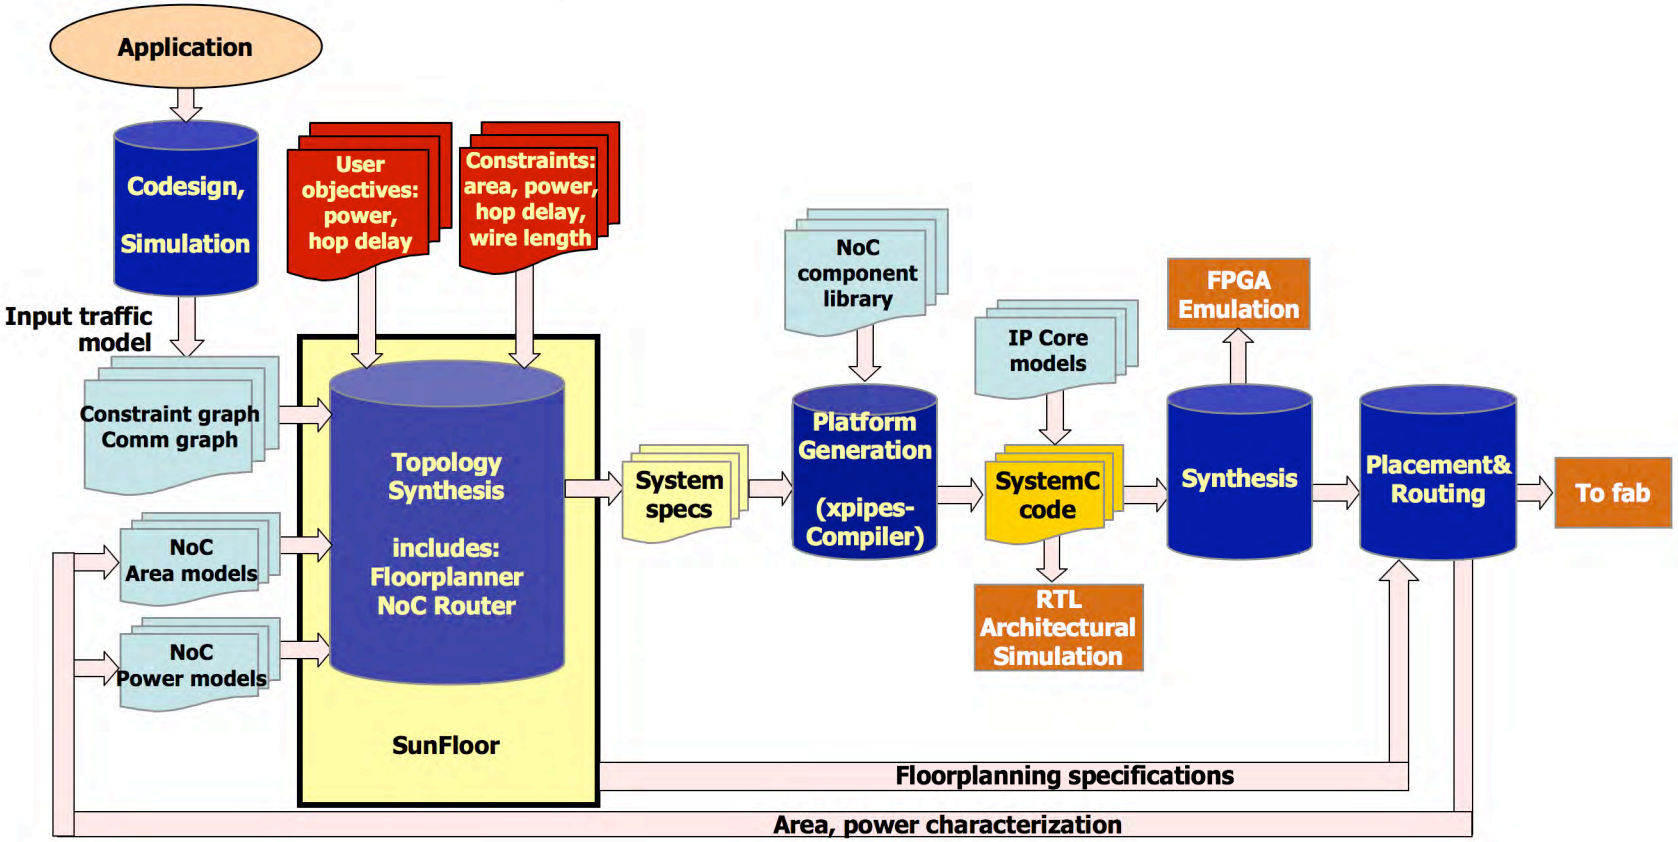

NoC Synthesis Project

- Started in 2002
- UNIBO, UNICA, Stanford, EPFL
- Objective: develop a complete EDA flow for NoC synthesis from application to P&R

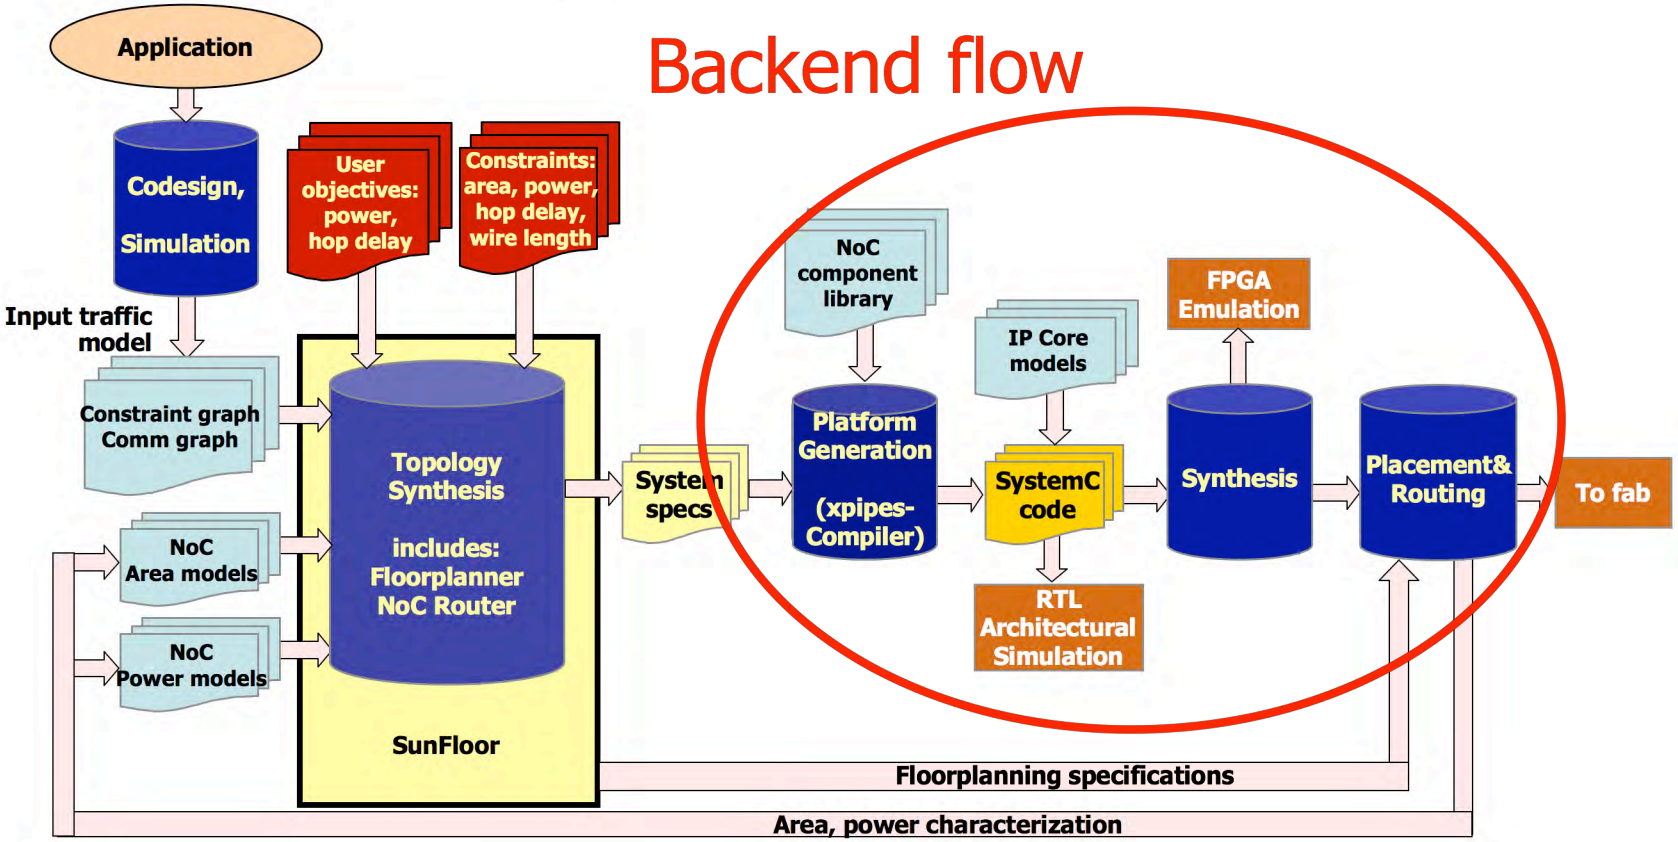

The xpipes NoC

- A soft macro library:

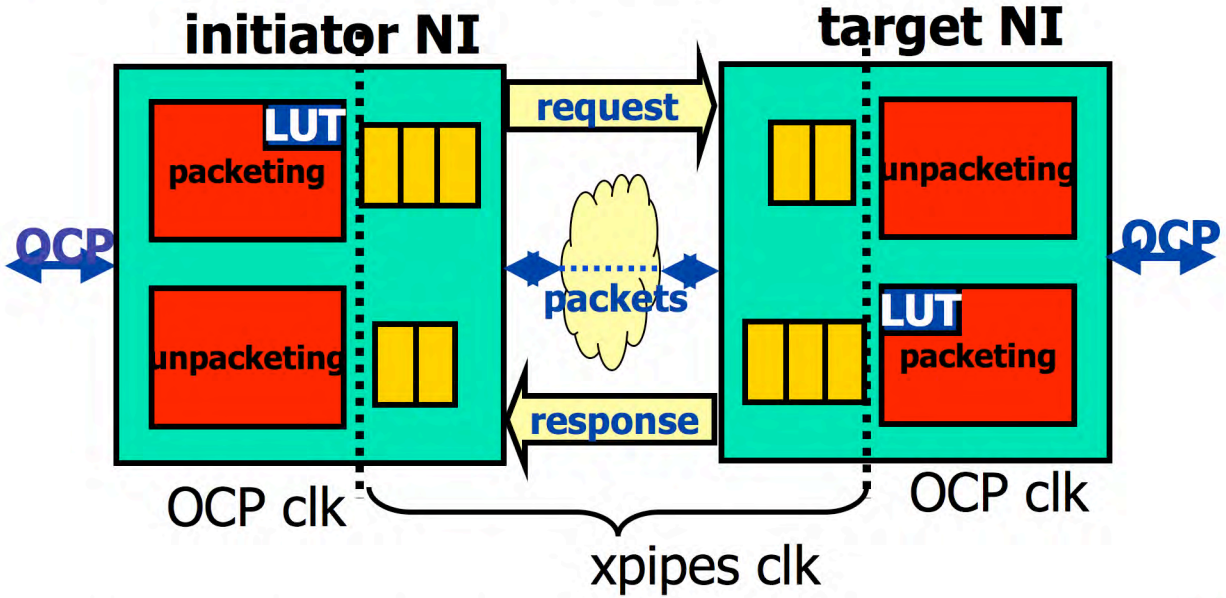

- Packeting/unpacketing
- OCP 2.0 protocol to connect to IP cores
- Source routing
- Dual Clock
- 2 Stage Pipeline

- Wormhole switching
- Round-robin & fixed priority allocator
- Supports ACK/NACK & STALL/GO flow control
  - ACK/NACK: Output buffered, 2 stages
  - STALL/GO: Input buffered, 1 stage
- Link pipelining fully supported

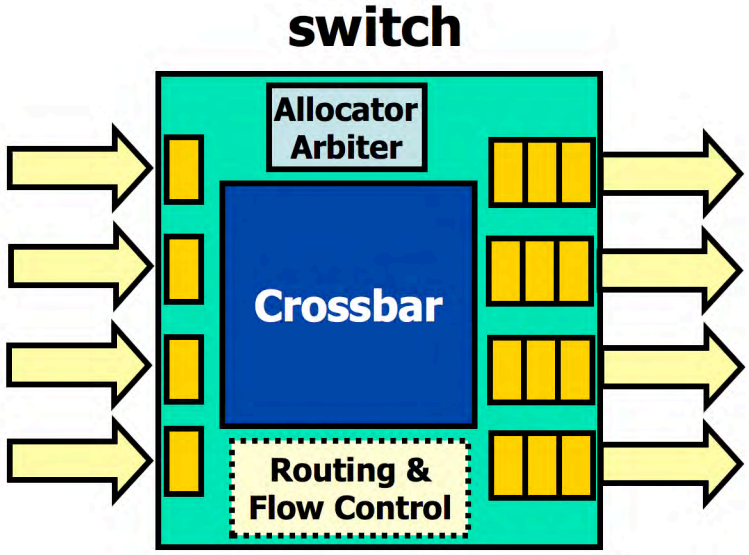

The xpipes NoC

- A soft macro library:

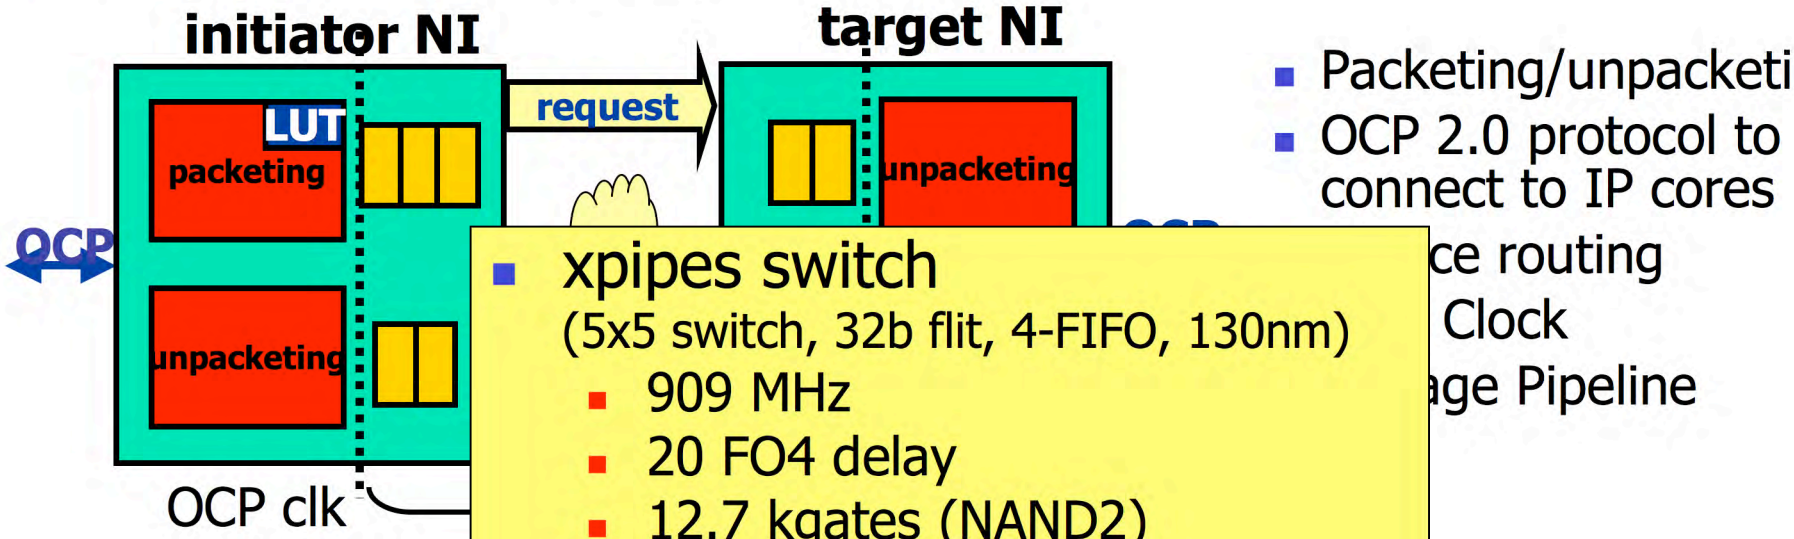

- Packeting/unpacketing
- OCP 2.0 protocol to connect to IP cores
- ce routing
- Clock
- age Pipeline

- xpipes switch
  (5x5 switch, 32b flit, 4-FIFO, 130nm)
  - 909 MHz
  - 20 FO4 delay
  - 12.7 kgates (NAND2)
  - 0.087 mm^2
  - 35.2 uW/MHz (32mW@909MHz)
- xpipes NI
  - 0.8 switch area

- Wormhole switch
- Round-robin & fix
- Supports ACK/NA control
  - ACK/NACK: Output buffered, 2 stages
  - STALL/GO: Input buffered, 1 stage
- Link pipelining fully supported

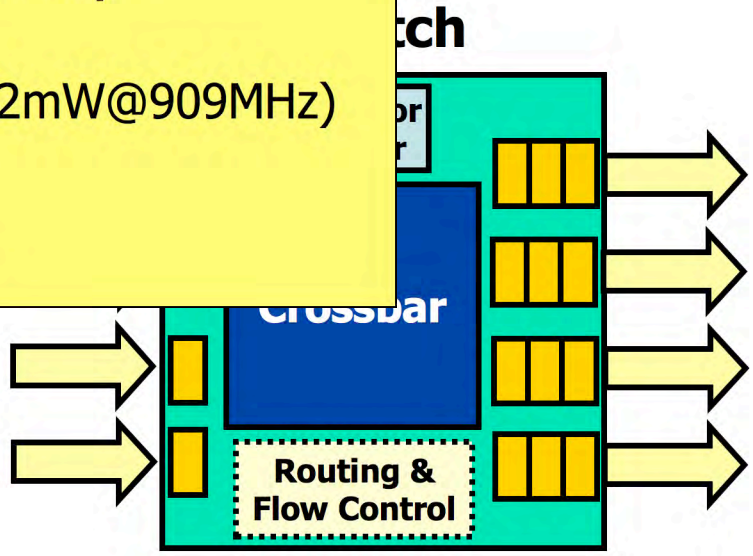

Backend Design Flow for Xpipes

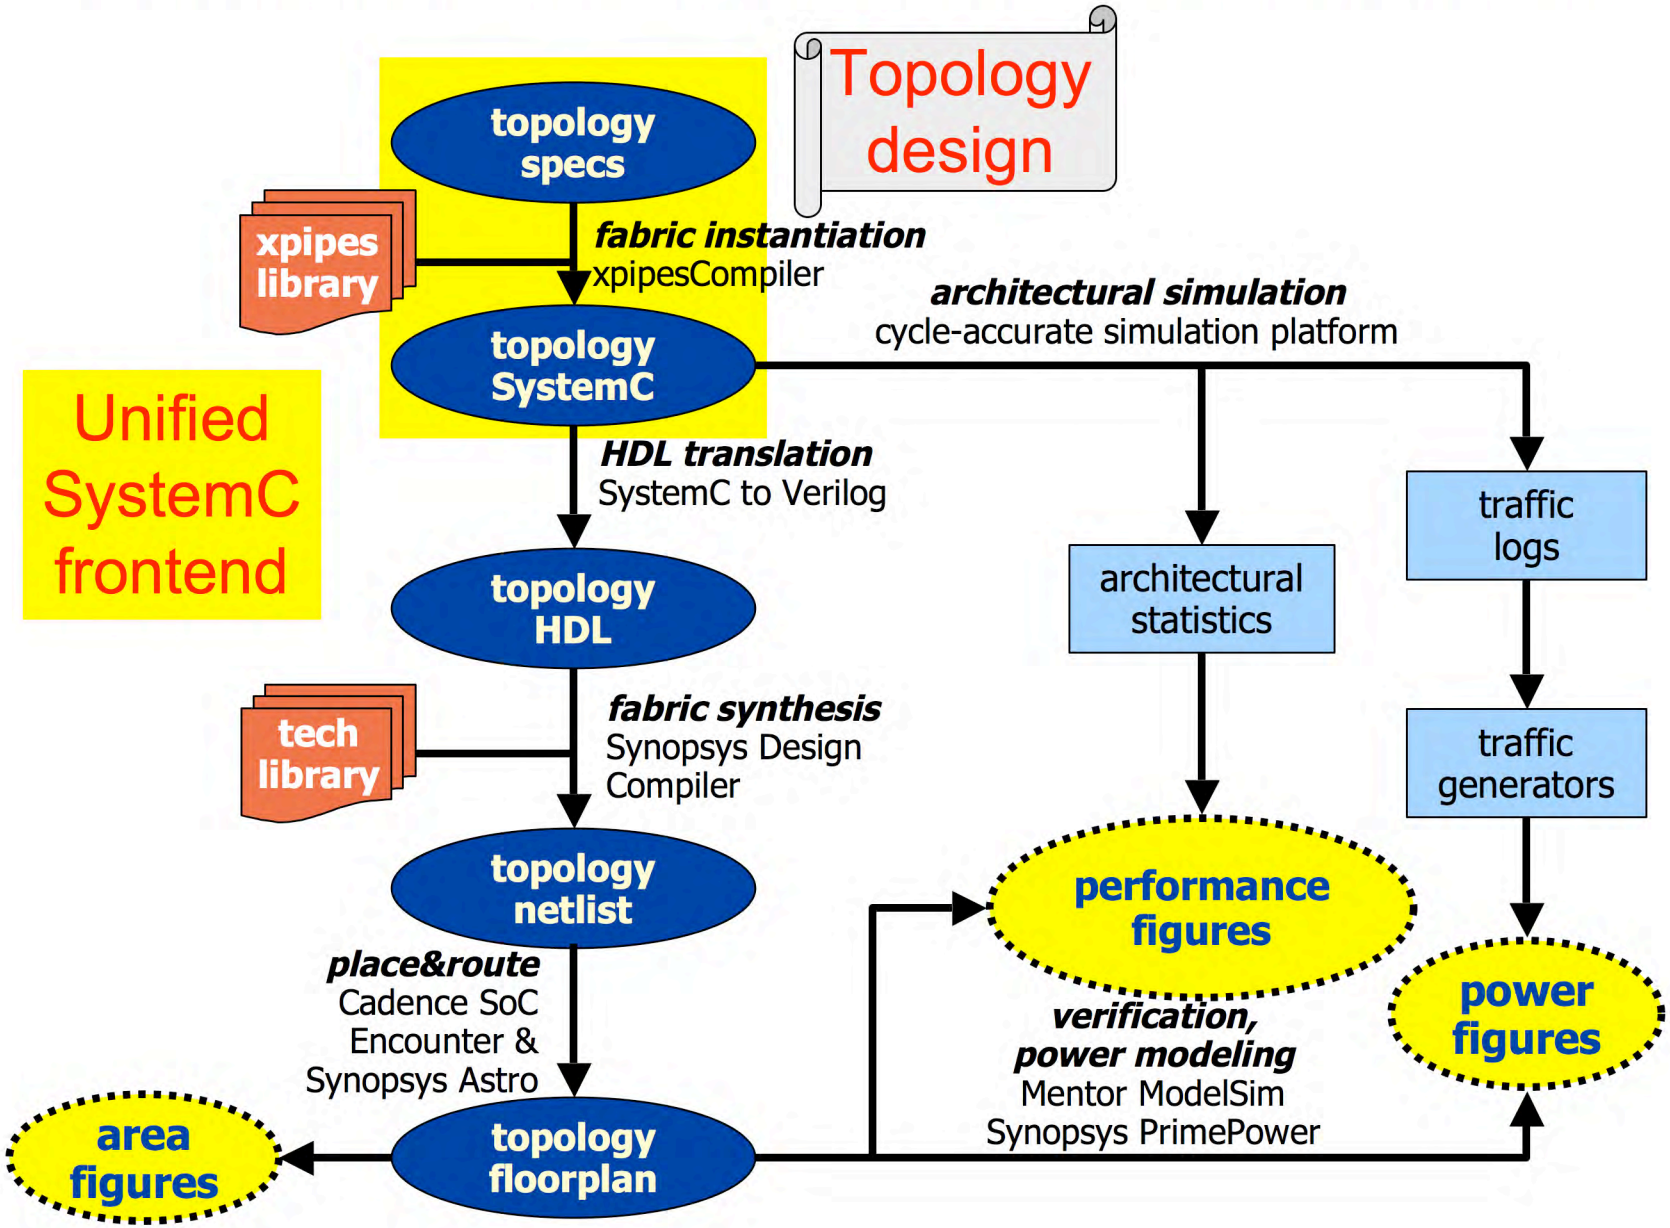

Case study in 130nm CMOS AMBA AHB Multilayer

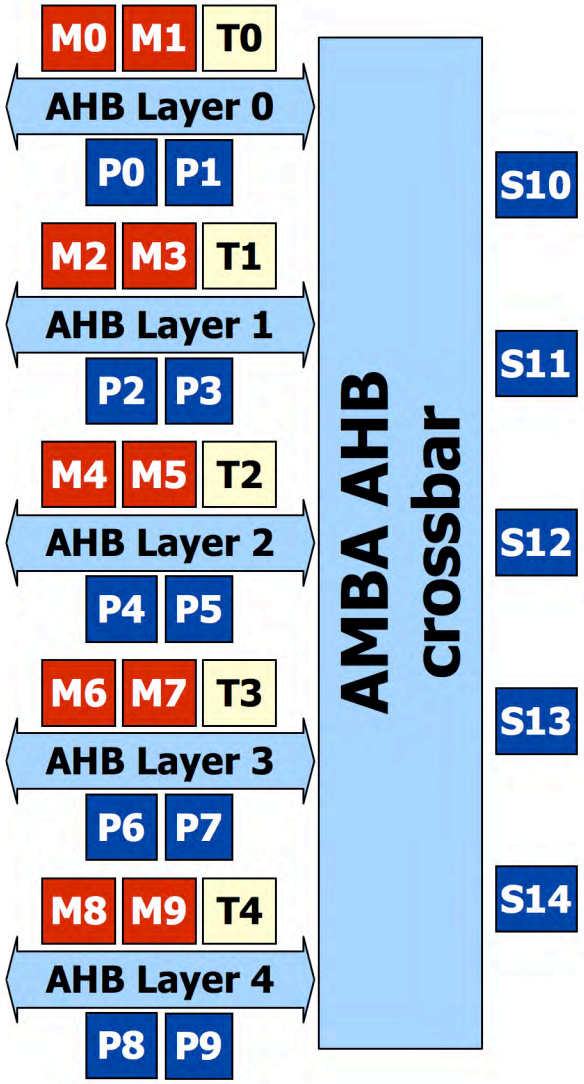

- Realistic 130nm SoC (10 cores, 5 accelerators 15 memory slaves)
- Shared bus fully saturated!
- Multilayer solution required for performance:
  - Intra-cluster traffic to *private slaves* (P0-P9) is bound within each layer, reducing congestion
  - *Shared slaves* (S10-S14) can be accessed in parallel
- Representative 5x5 Multilayer configuration (up to 8x8)
- Synthesized for max operating frequency

Xpipes Quasi-Mesh

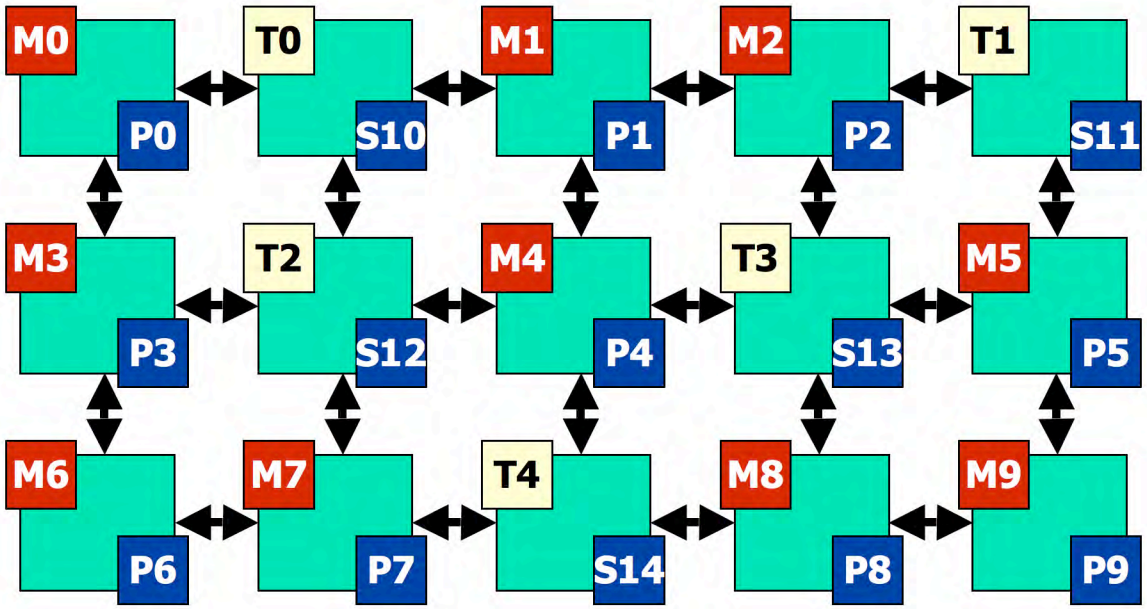

- Balanced architecture, no bottlenecks
- Regular topology: easy to floorplan
- Optimally placed private slaves, scattered shared slaves

Comparison

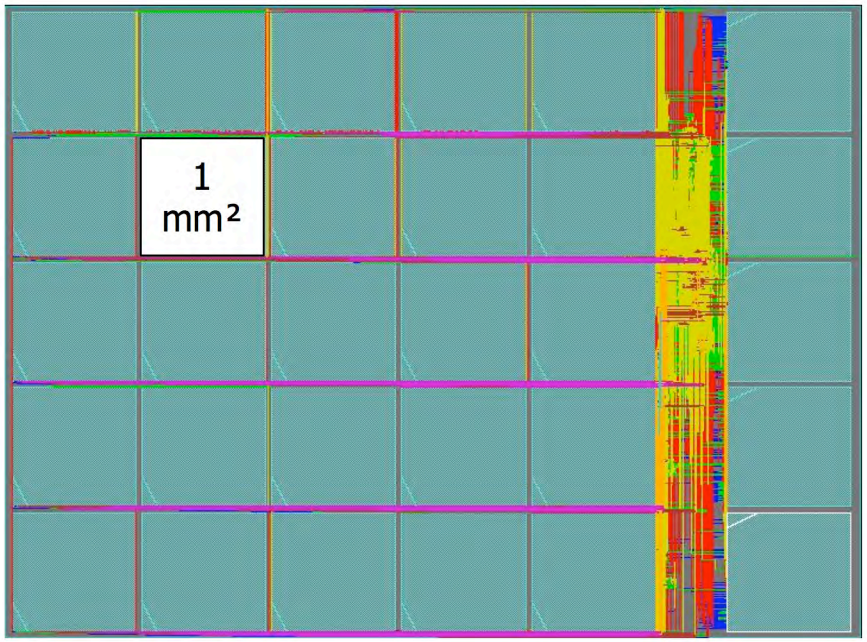

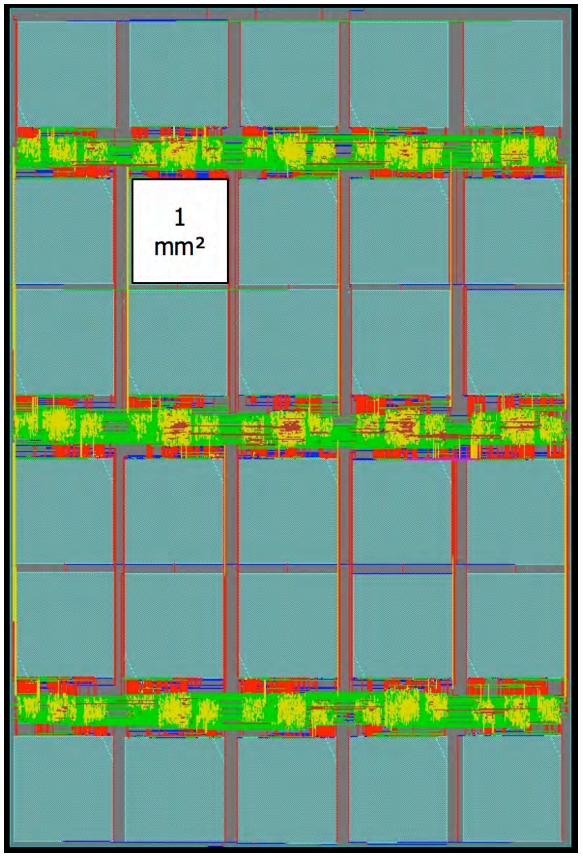

- 130nm UMC library
- IPs: 1mm² obstructions (ARM cores, 32kB SRAM)
- Wire routing over the cores was forbidden
- Minimum delay synthesis

Comparison

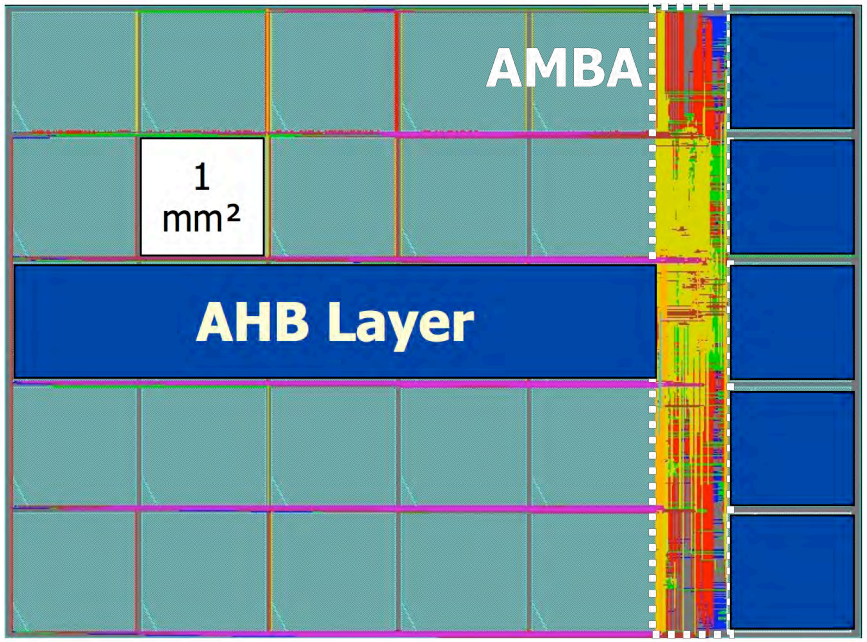

Shared Slaves

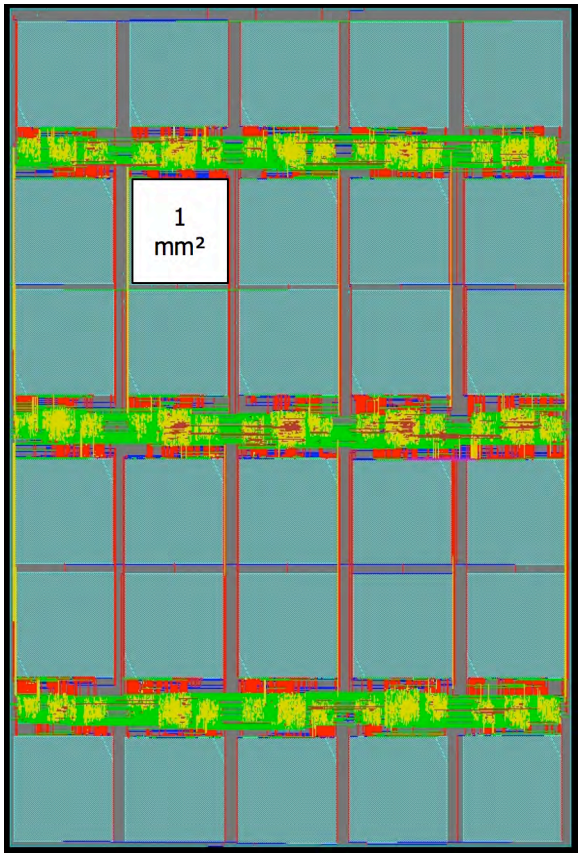

- 130nm UMC library
- IPs: 1mm² obstructions (ARM cores, 32kB SRAM)
- Wire routing over the cores was forbidden
- Minimum delay synthesis

Comparison

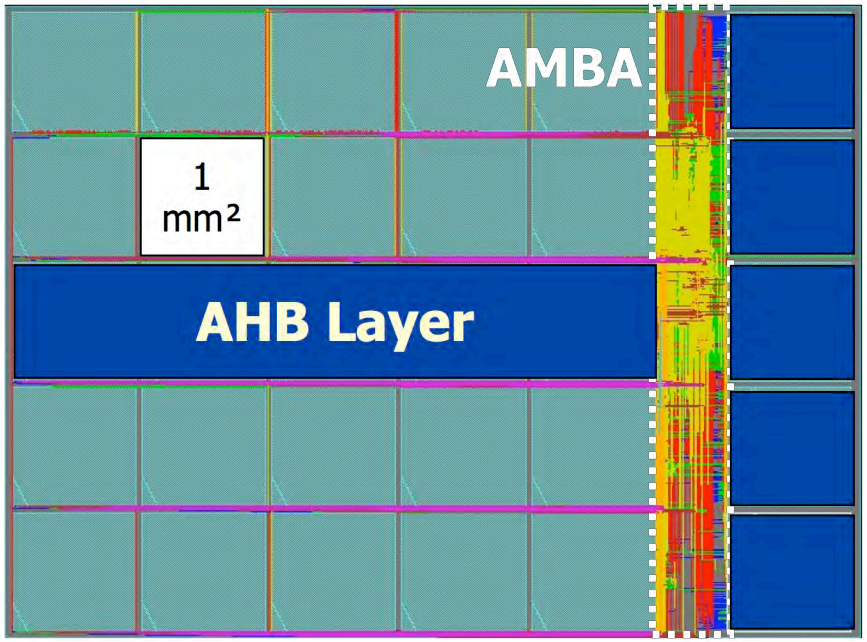

Shared Slaves

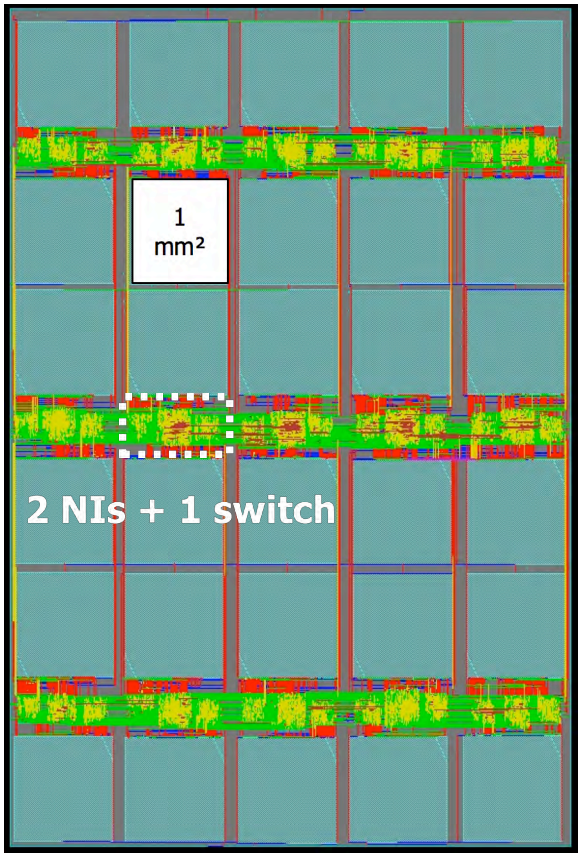

- 130nm UMC library
- IPs: 1mm² obstructions (ARM cores, 32kB SRAM)
- Wire routing over the cores was forbidden
- Minimum delay synthesis

Comparison

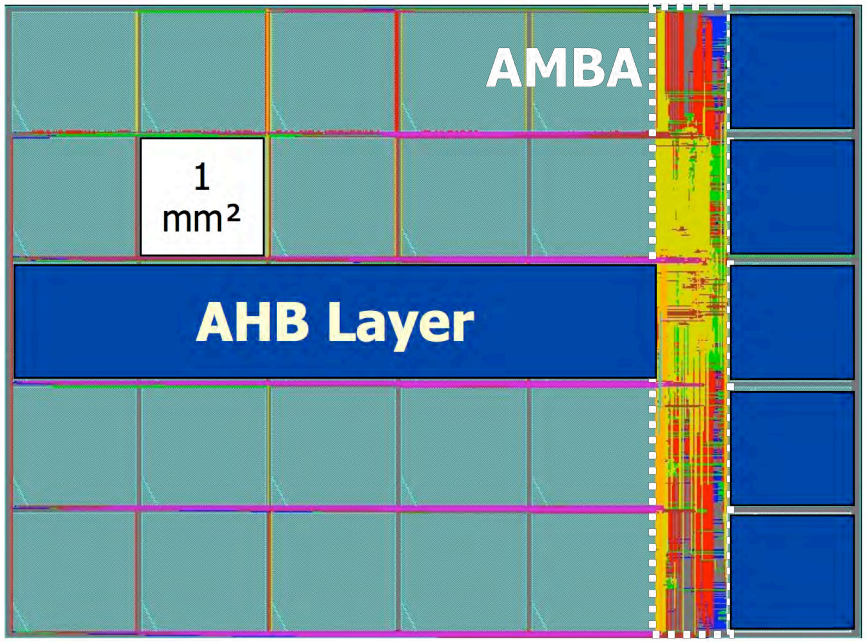

Shared Slaves

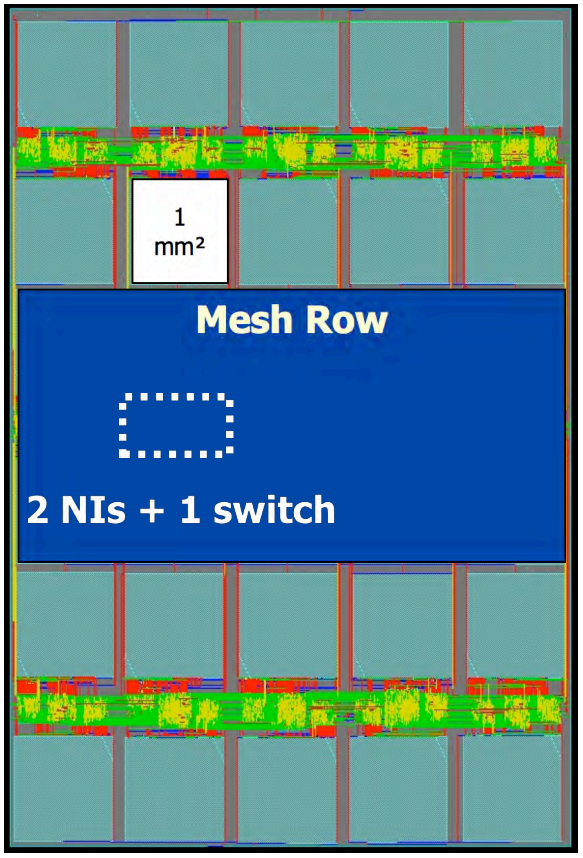

- 130nm UMC library
- IPs: 1mm² obstructions (ARM cores, 32kB SRAM)
- Wire routing over the cores was forbidden
- Minimum delay synthesis

Comparison

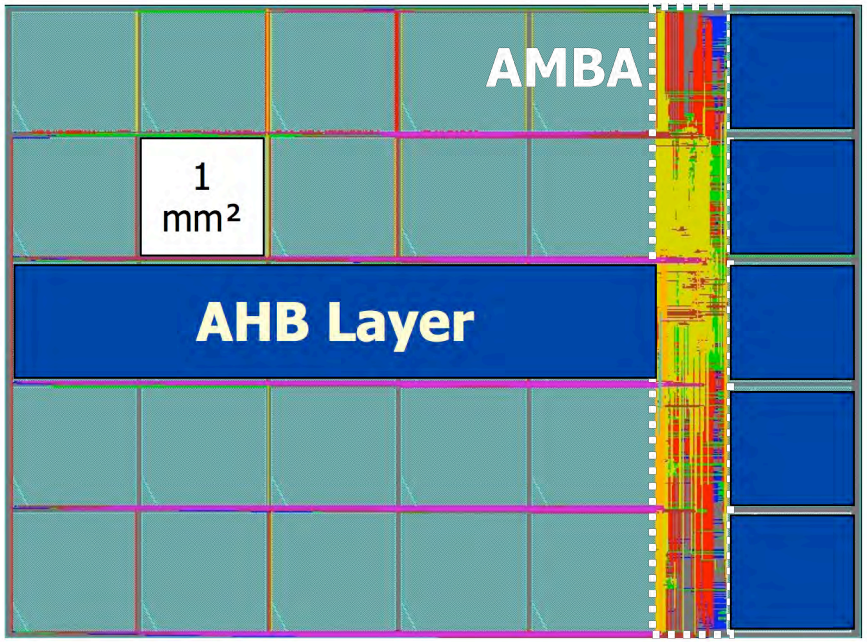

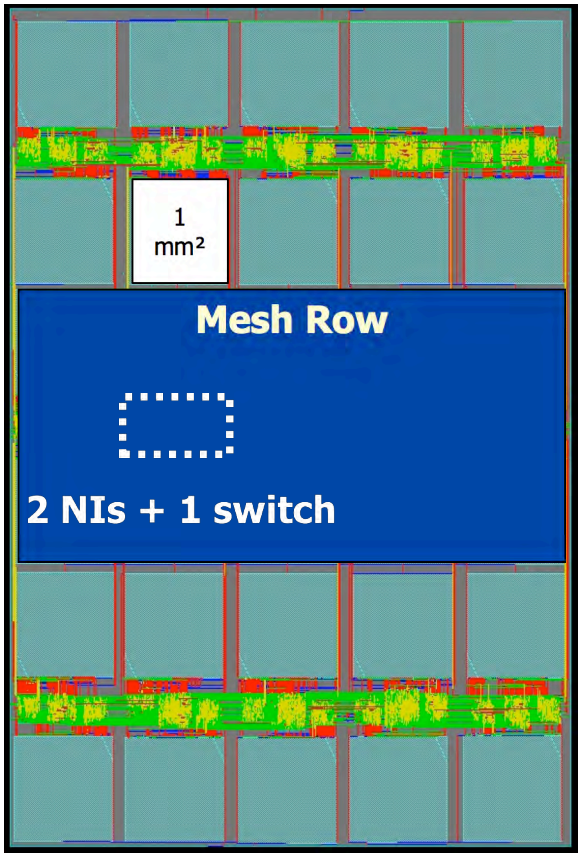

- 130nm UMC library
- IPs: 1mm² obstructions (ARM cores, 32kB SRAM)
- Wire routing over the cores was forbidden
- Minimum delay synthesis
- Summary of results:
  - 6.4% vs. 22.9% post P&R timing degradation → Much improved predictability
  - Clock frequency 793 vs. 370 MHz post P&R → Much faster
  - 15% avg application speedup (longer latency, but more effective bandwidth)
  - Higher power & area (higher frequency, flip-flops in buffers + ck tree)

Comparison

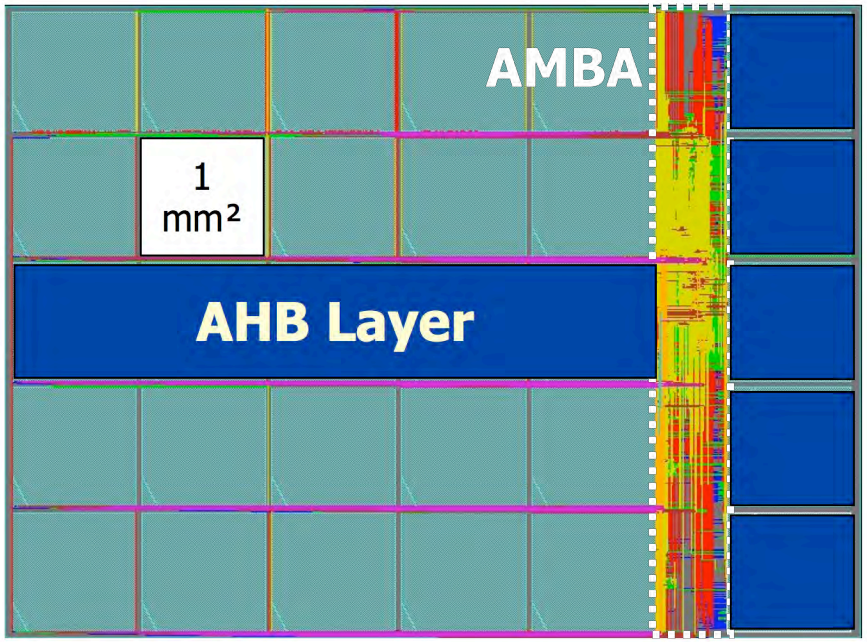

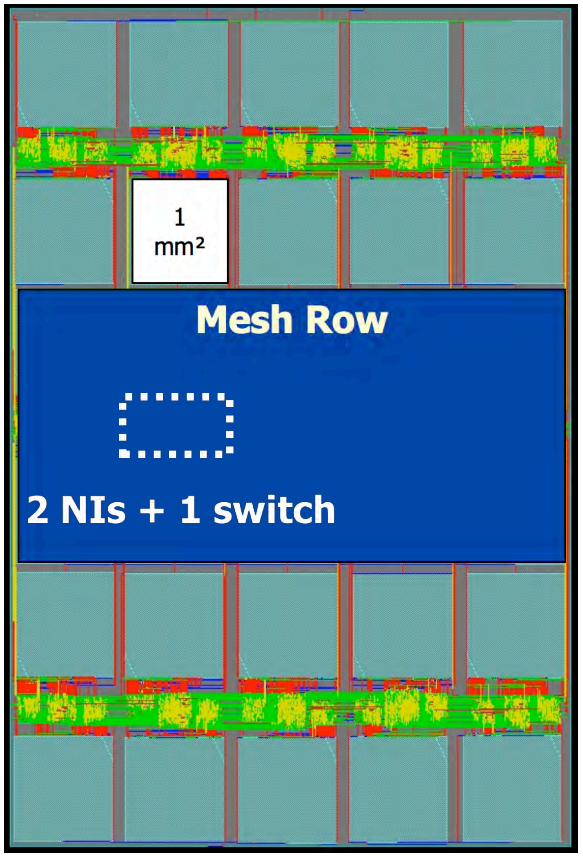

- 130nm UMC library
- IPs: 1mm² obstructions (ARM cores, 32kB SRAM)
- Wire routing over the cores was forbidden
- Minimum delay synthesis

Application specific NOC is competitive with state-of-the-art interconnect even in 130nm technology

Technology Scaling on a Topology

- Three designs for max frequency:

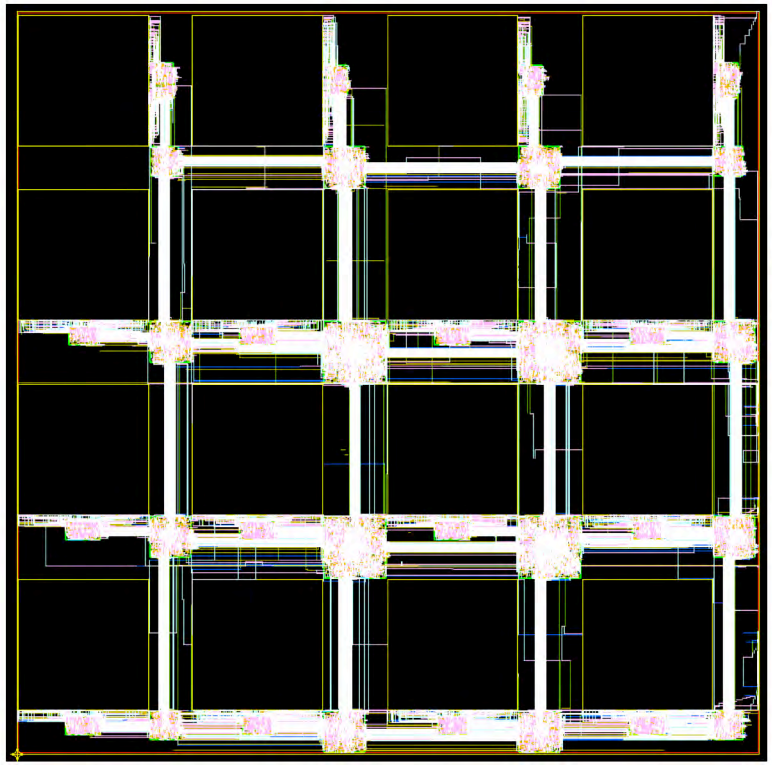

90 nm, 1 mm^2 cores

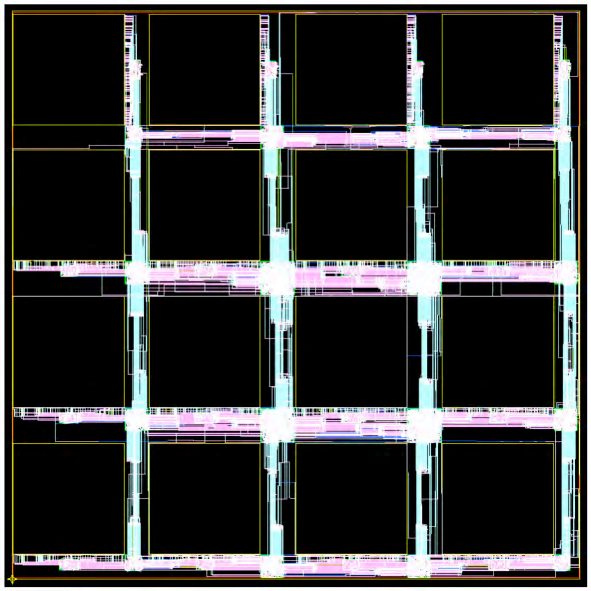

65 nm, 1 mm^2 cores

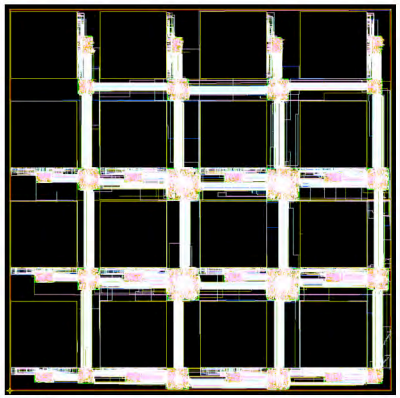

65 nm, 0.4 mm^2 cores

Technology Scaling Results

Max Frequency 38 bit	Max Frequency	NoC Cell Area	Power	Bandwidth
90 nm, 1 mm^2	1 GHz	1.31 mm^2	784.6 mW	228 GB/s
65 nm, 1 mm^2	1.25 GHz	0.64 mm^2	520.1 mW	285 GB/s
65 nm, .4 mm^2	1.25 GHz	0.63 mm^2	495.3 mW	285 GB/s

- Links
  - Always short (<1.2 mm) → non-pipelined
  - However
    - 90 nm 1 mm^2: 3.1 mW
    - 65 nm 1 mm^2: 3.6 mW (tightest fit → more buffering)
    - 65 nm 0.4 mm^2: 2.2 mW
  - Power shifting from switches/NIs to links (buffering)

"High-level" Design Flow

- Automatically find topology, architecture
  - Minimize area, power, latency
  - Satisfying design constraints

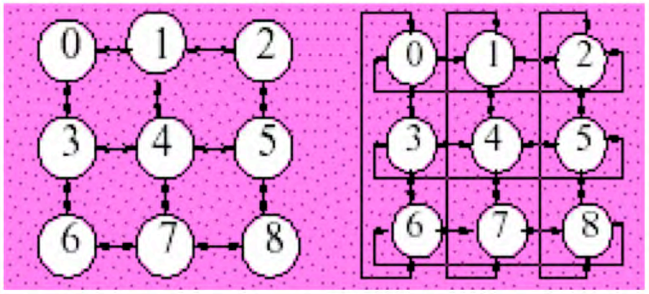

a. Mesh b. Torus

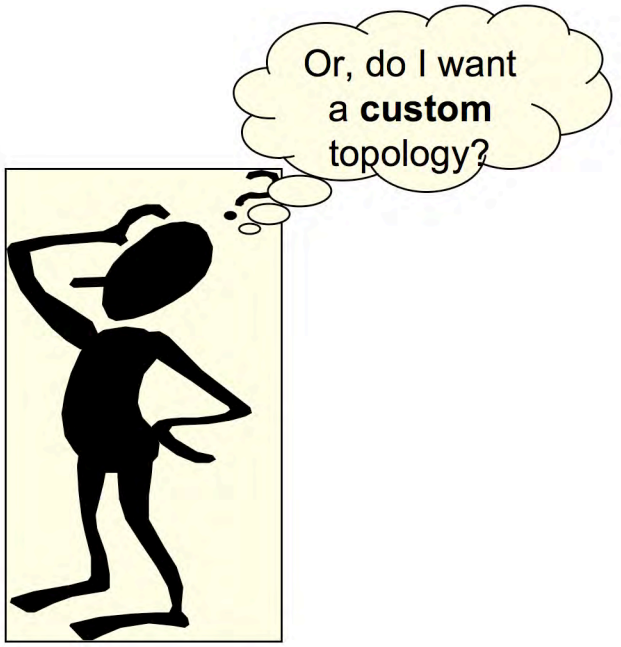

SUNFLOOR & Xpipes Flow

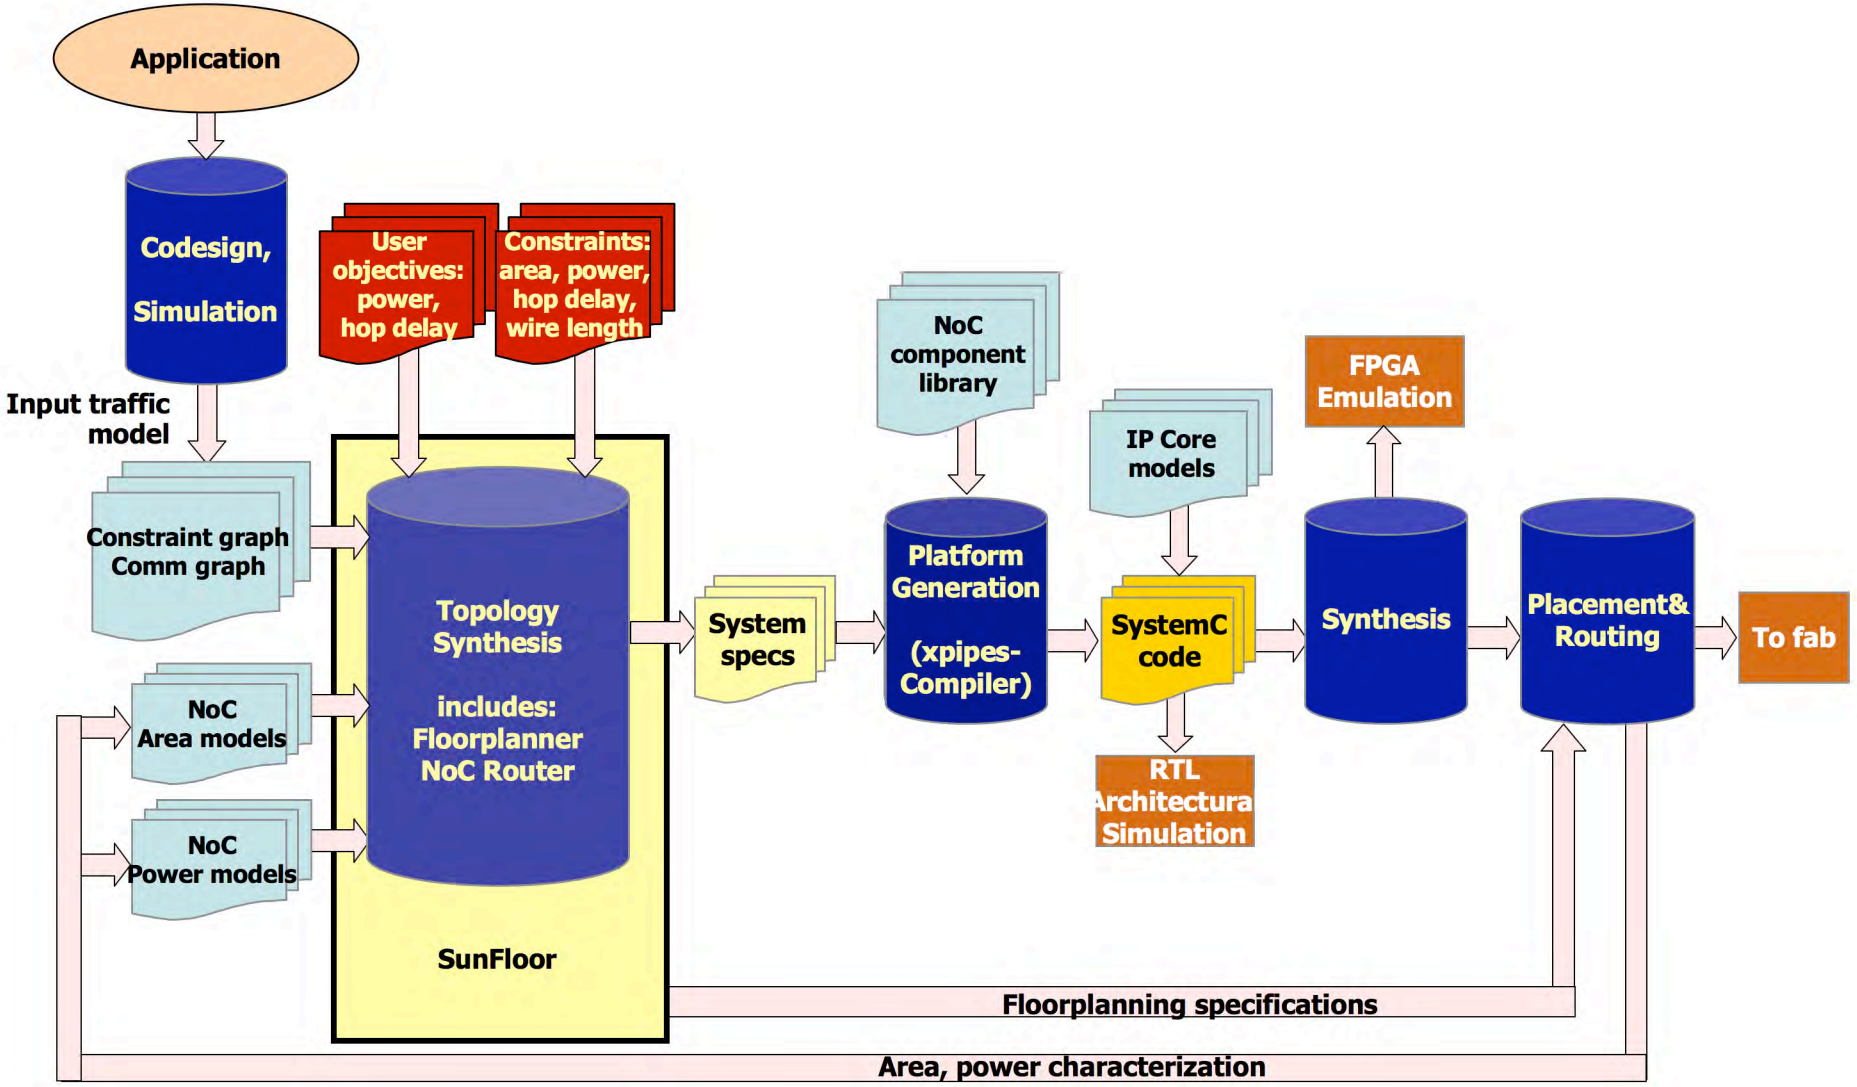

SUNFLOOR & Xpipes Flow

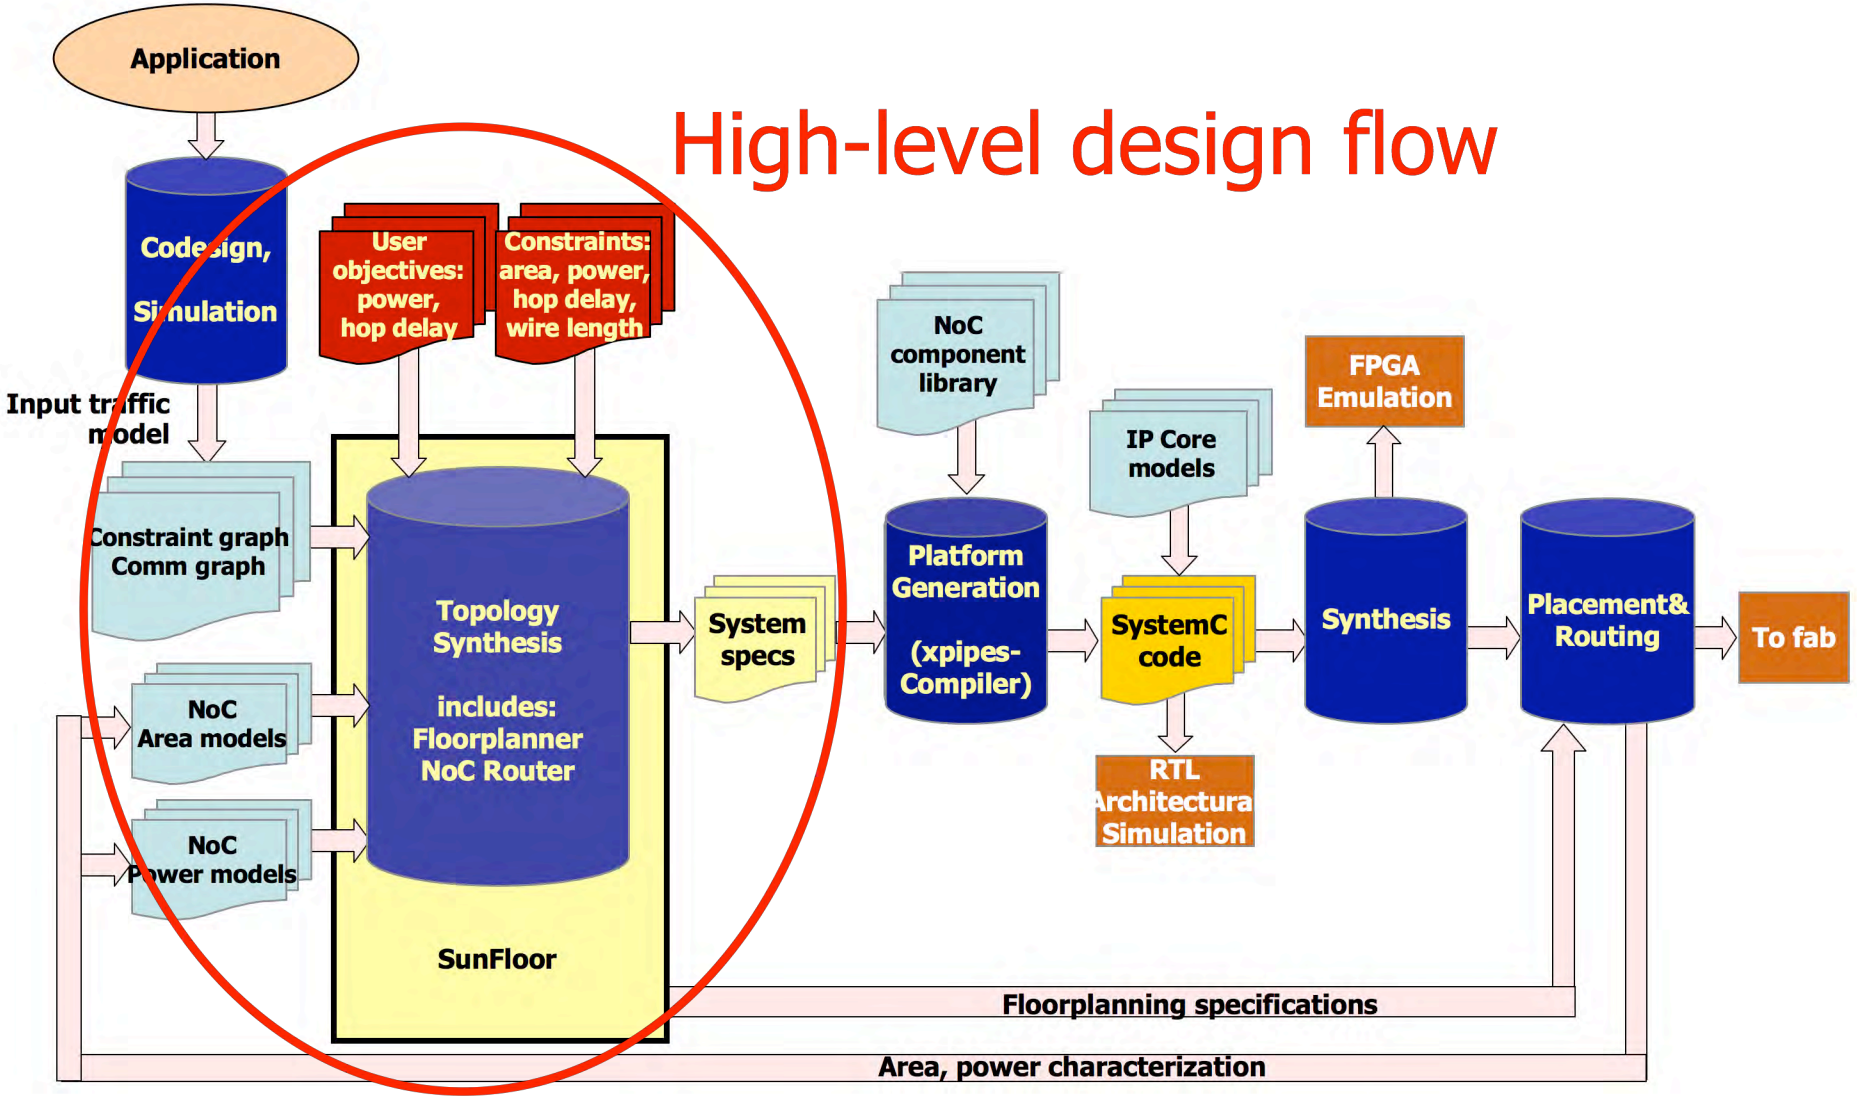

Custom Topology & Mapping

- Objectives
  - Design fully application-specific custom topologies
  - Generate deadlock-free networks: both routing and message-level deadlocks are removed
  - Optimize architectural parameters of the NoC (frequency, flit size), tuning based upon application requirements

Custom Topology & Mapping

- Objectives
  - Design fully application-specific custom topologies
  - Generate deadlock-free networks: both routing and message-level deadlocks are removed
  - Optimize architectural parameters of the NoC (frequency, flit size), tuning based upon application requirements

Physical design awareness

  - Leverage accurate analytical models for area and power, back-annotated from layouts
  - Integrated floorplanner to achieve design closure while also considering wiring complexity

SUNFLOOR Steps

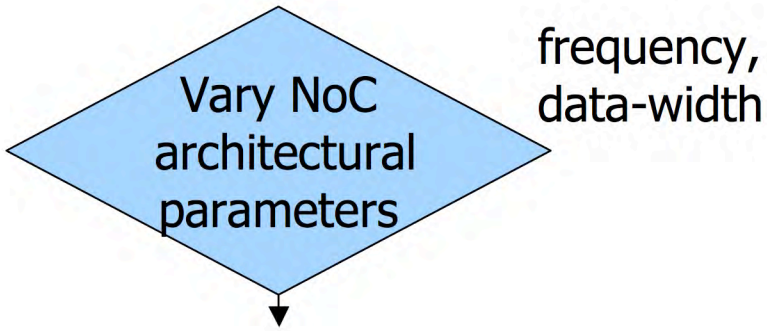

SUNFLOOR Steps

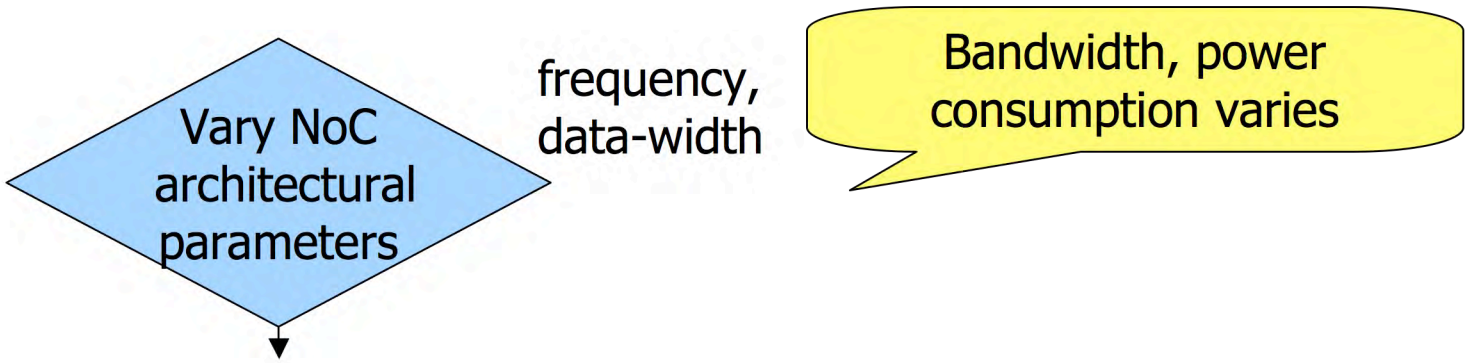

SUNFLOOR Steps

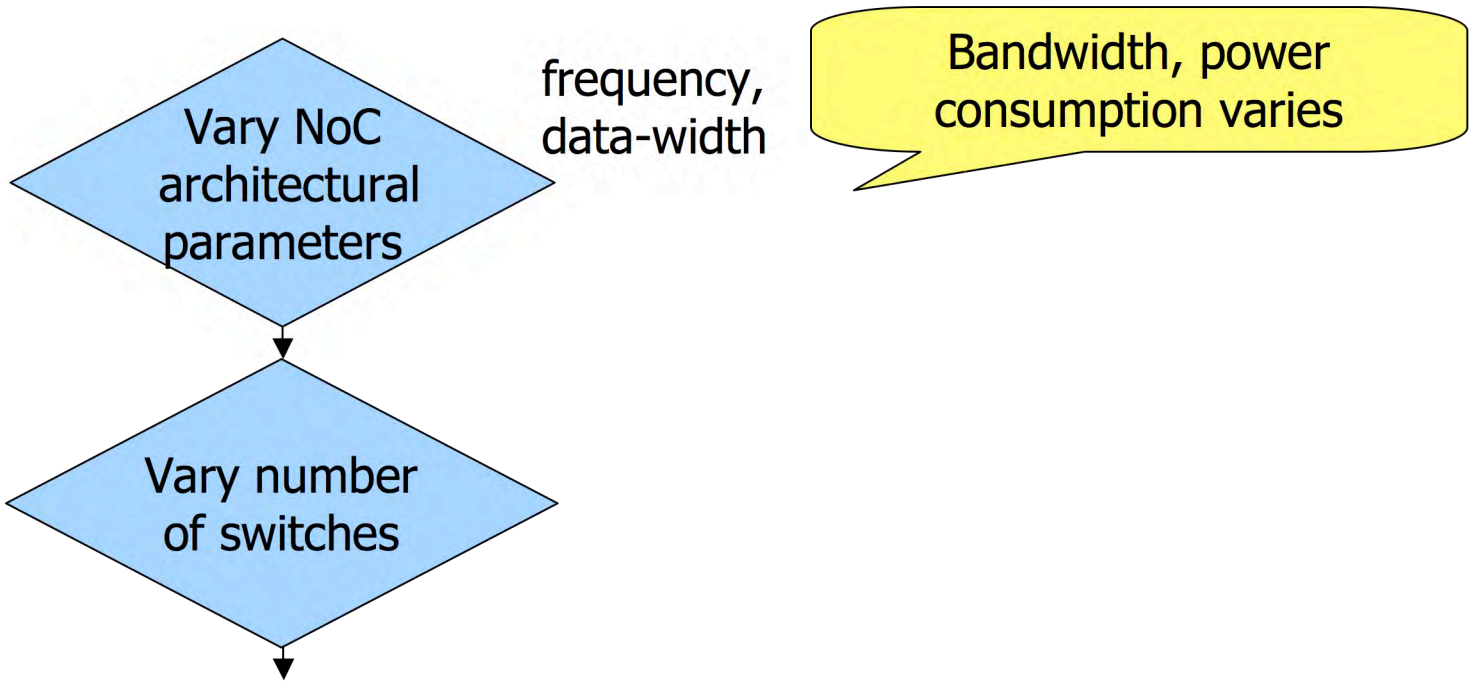

SUNFLOOR Steps

Vary Nc
architect
paramet

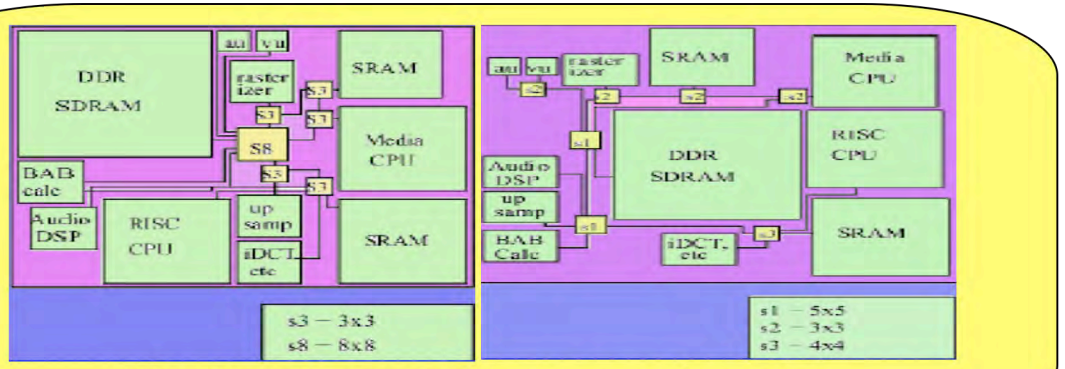

Which is better ? – Do not know !!

Vary number of switches

SUNFLOOR Steps

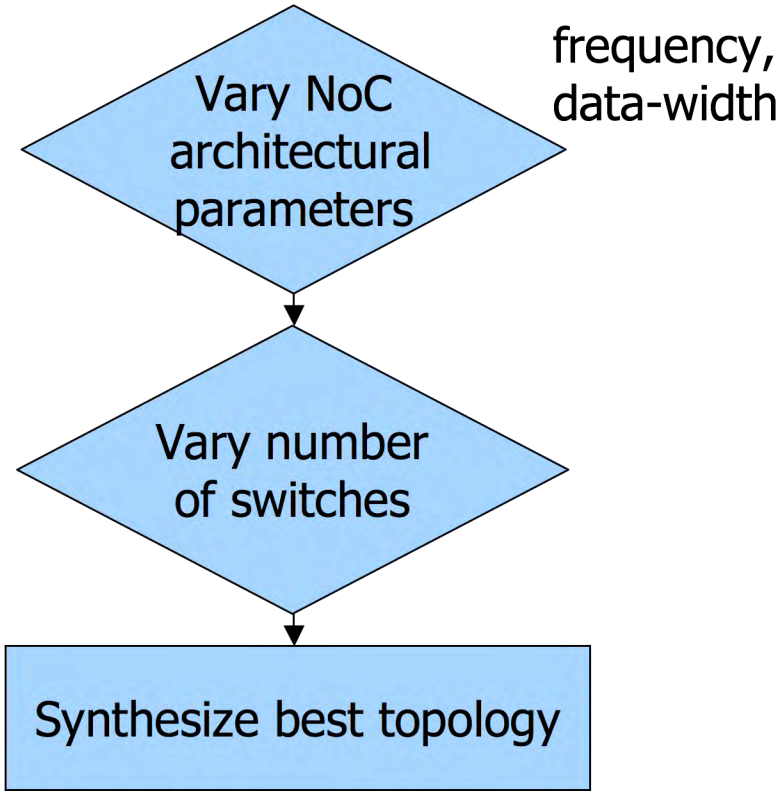

SUNFLOOR Steps

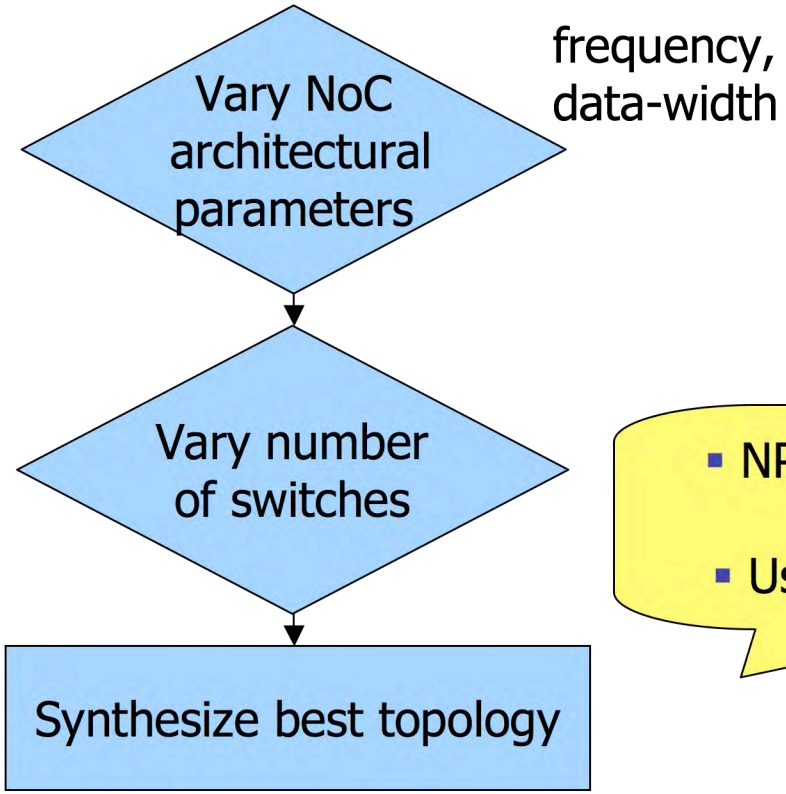

- NP-Hard problem (single path multi-commodity flow)
- Use fast & efficient heuristics

SUNFLOOR Steps

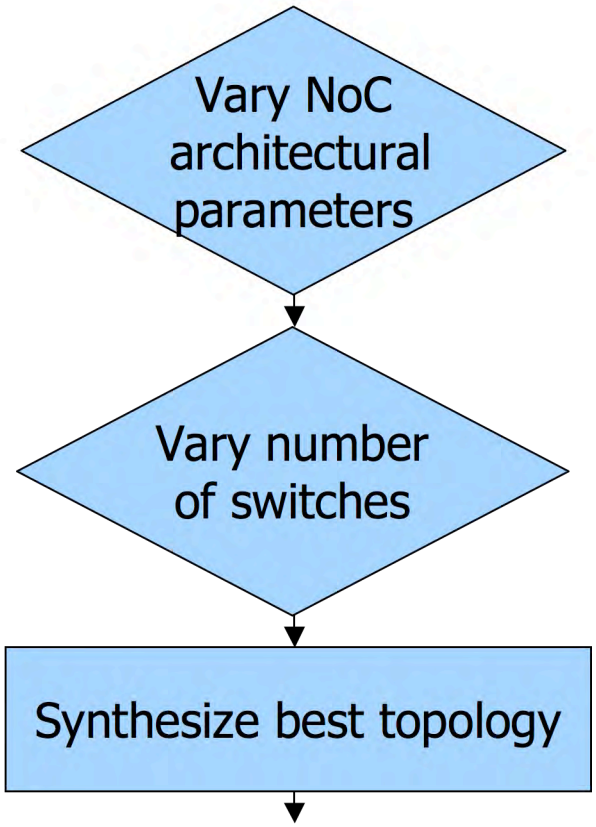

SUNFLOOR Steps

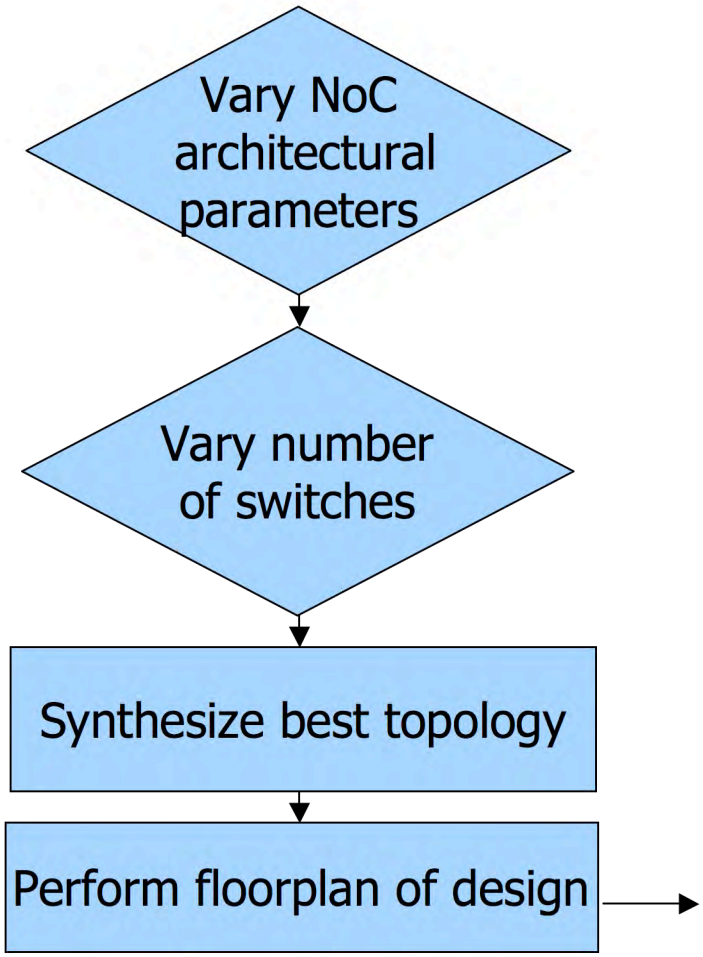

SUNFLOOR Steps

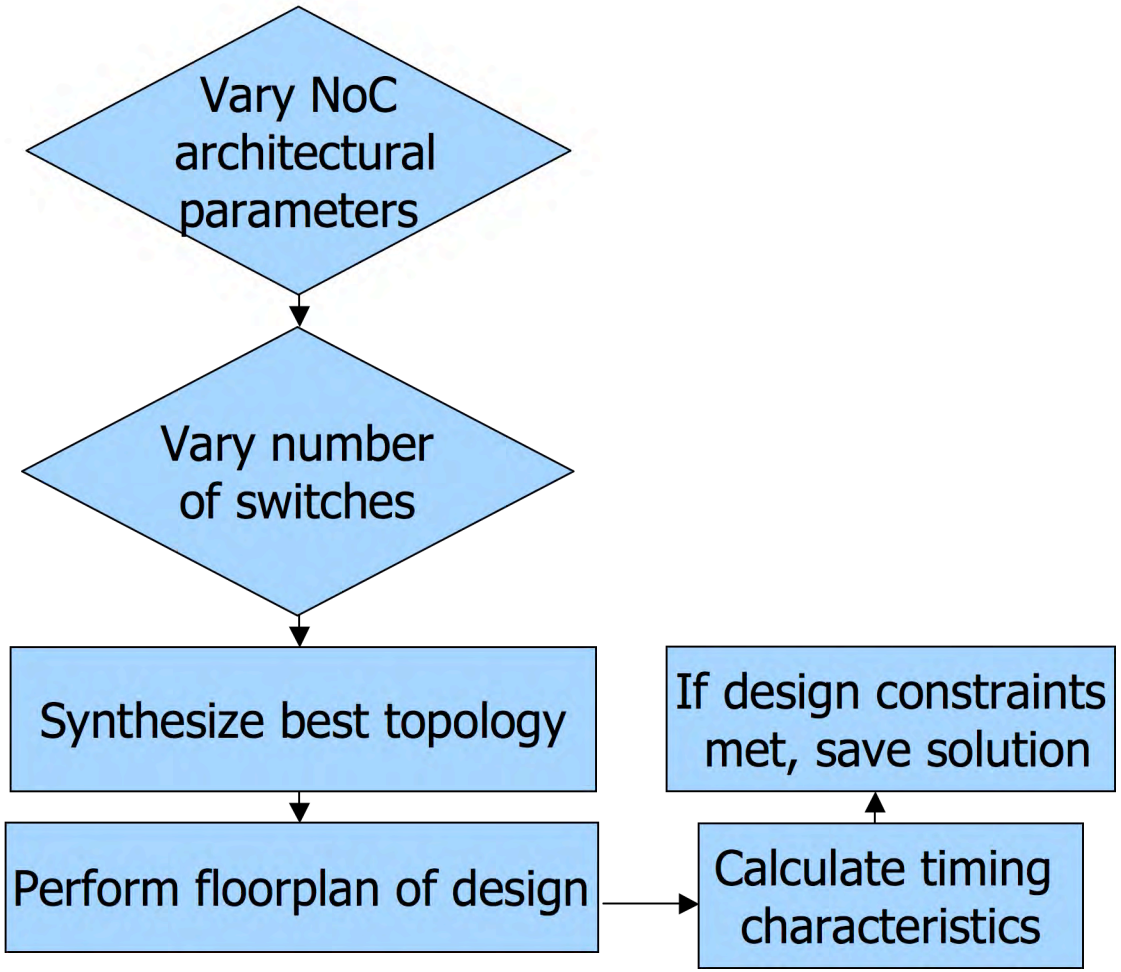

SUNFLOOR Steps

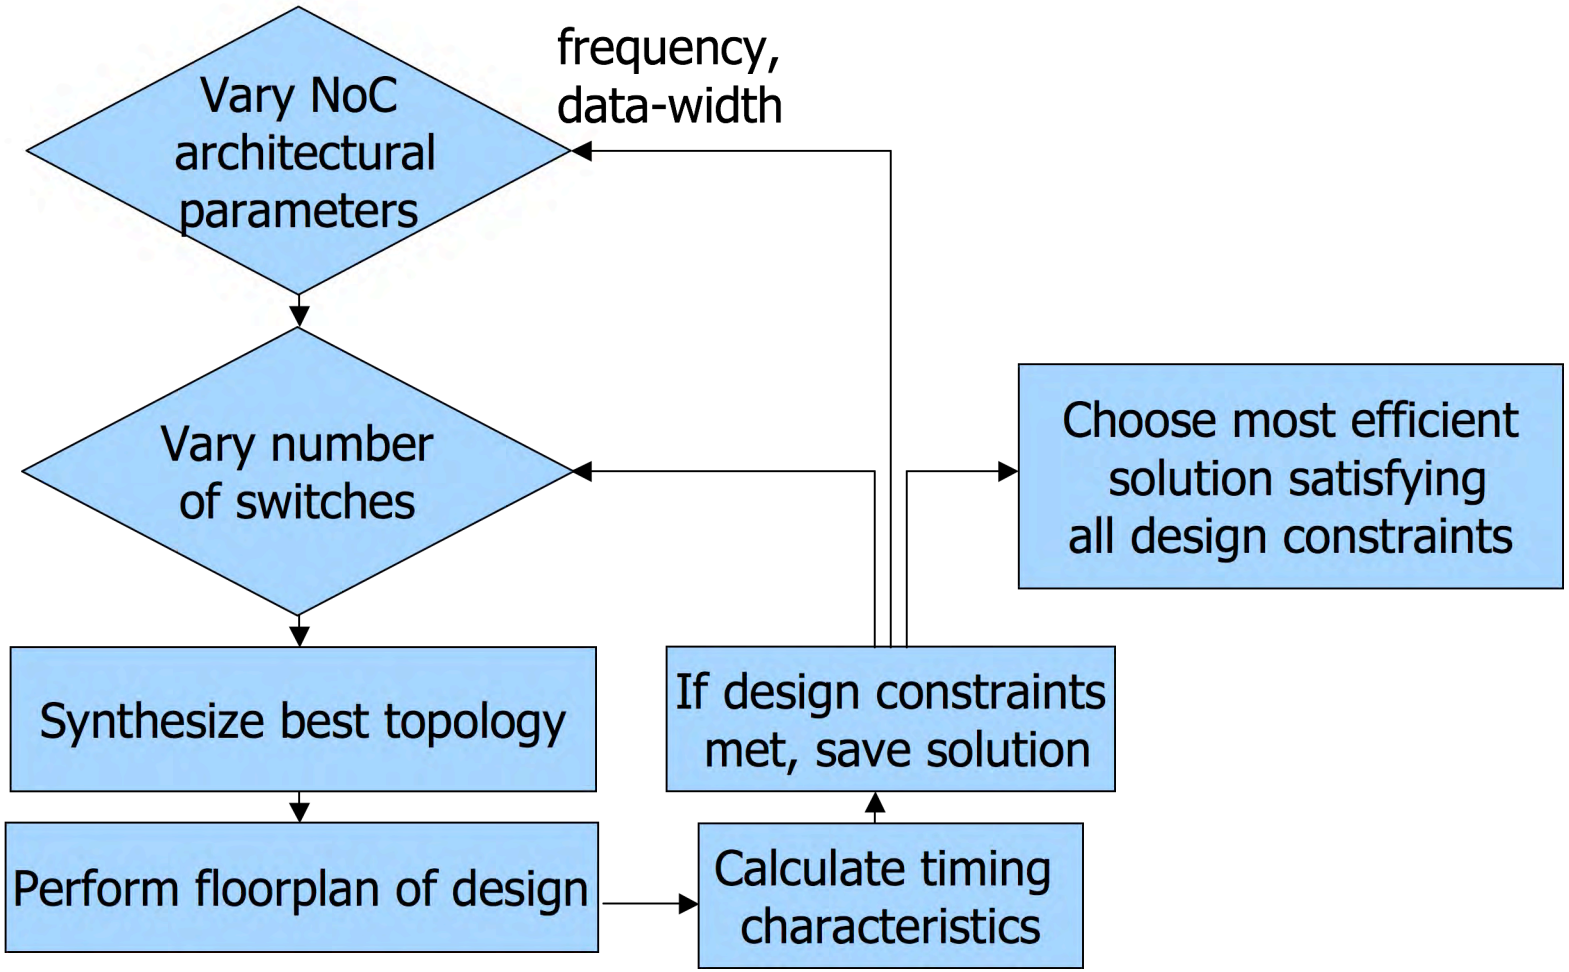

Synthesis Algorithm

Refer to Murali et al. ICCAD06 for full details

Synthesis Algorithm

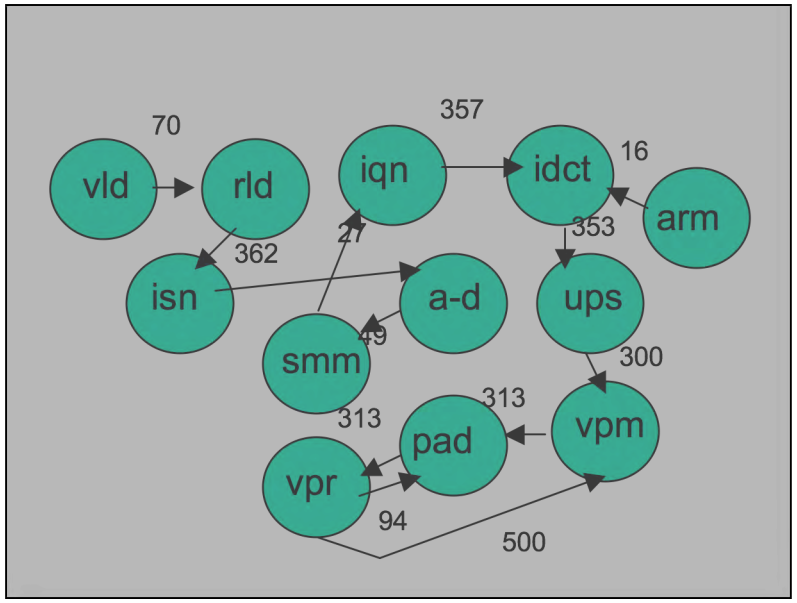

Core graph

Refer to Murali et al. ICCAD06 for full details

Synthesis Algorithm

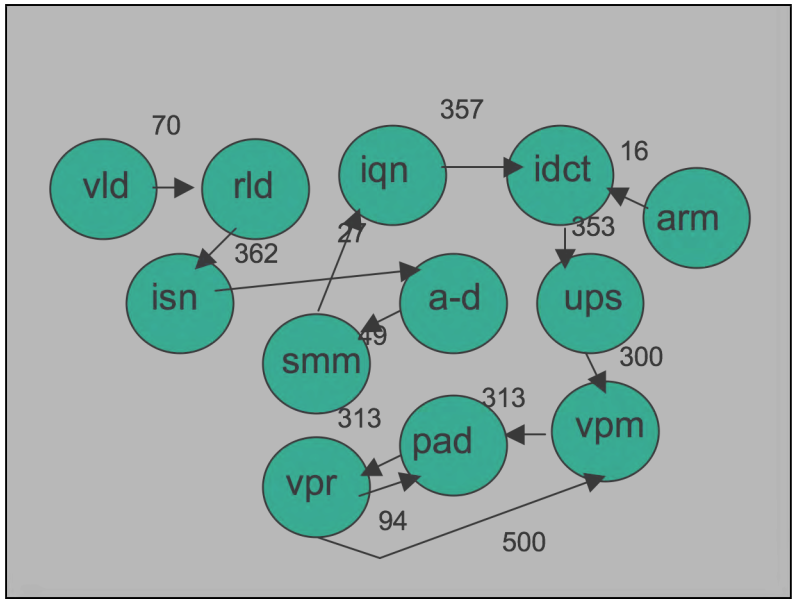

Core graph

- Obtain min-cut partitions of core-graph
- Cores in a partition share a switch

Refer to Murali et al. ICCAD06 for full details

Synthesis Algorithm

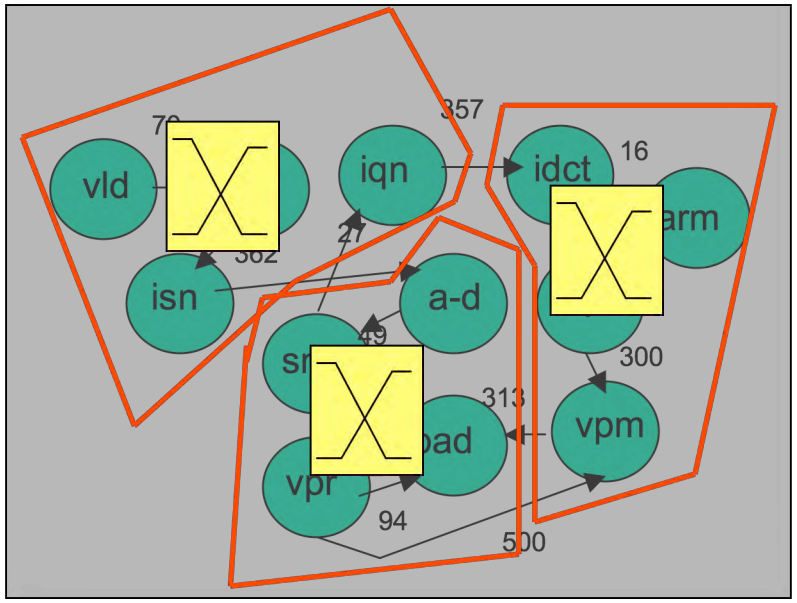

Core graph

- Obtain min-cut partitions of core-graph
- Cores in a partition share a switch

Refer to Murali et al. ICCAD06 for full details

Synthesis Algorithm

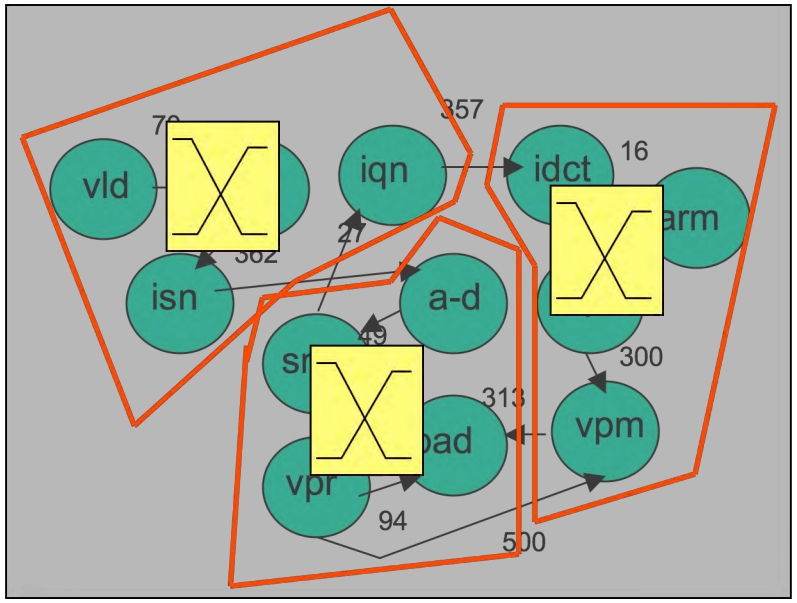

Core graph

- Obtain min-cut partitions of core-graph
- Cores in a partition share a switch
- Find lower bound on switch sizes, switch power
- Establish switch connectivity by routing flows
  - Account for constraints on # hops, deadlock avoidance, switch size
  - Minimize power

Refer to Murali et al. ICCAD06 for full details

Case Study 1: Comparison Against Hand-Mapped Topology

- On our 30-core multimedia benchmark

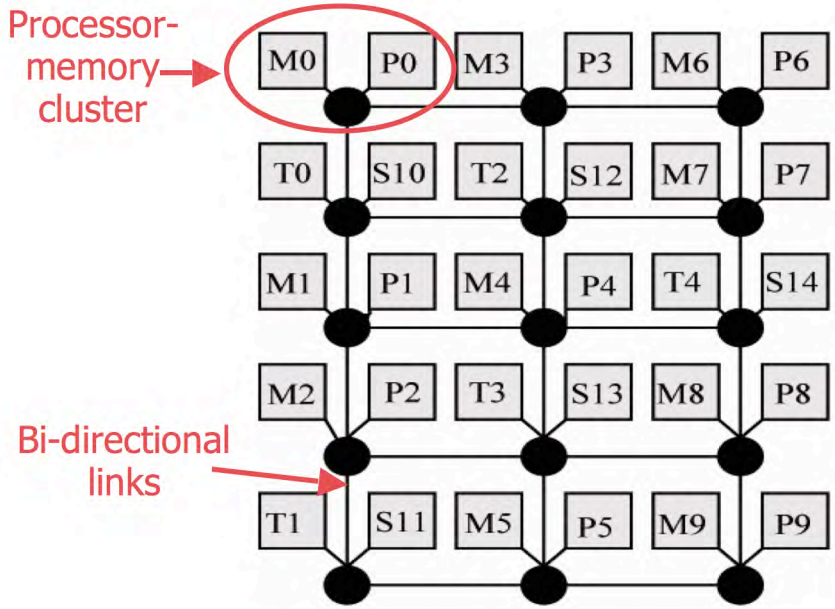

Hand-mapped topology

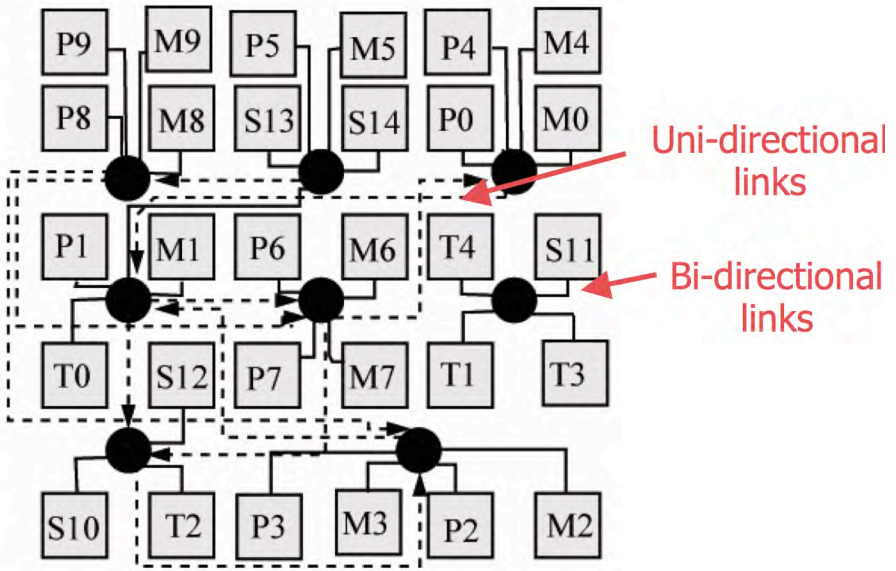

SUNFLOOR custom topology

P-processors, M-private memories,
T-traffic generators, S-shared slaves

Case Study 1: Results vs Hand-Mapped

Hand-mapped design:

- Topology: 5x3 mesh (15 switches)
- Operating frequency: 793 MHz (post-layout)
- Power consumption: 368 mW
- Floorplan area: 35.4 mm^2
- Design time: weeks
- 0.13 µm technology

SunFloor:

- Topology: custom (8 switches)
- Operating frequency: 793 MHz (post-layout)
- Power consumption: 277 mW **(-25%)**
- Cell area: 37 mm^2 **(+4%)**
- Design time: **4 hours design to layout**
- 0.13 µm technology

Benchmark execution time comply with application requirements and are even 10% better on SunFloor topology.

Case Study 2: SUNFLOOR Vs Regular Topologies

Application	Topology	Power(mW)	Avg. nr. hops
VPROC (42 cores)	Custom	79.64	1.67
	Mesh	301.8	2.58
	Opt-mesh	136.1	2.58
MPEG4 (12 cores)	Custom	27.24	1.50
	Mesh	96.82	2.17
	Opt-mesh	60.97	2.17
VOPD (12 cores)	Custom	30.00	1.33
	Mesh	95.94	2.00
	Opt-mesh	46.48	2.00
MWD (12 cores)	Custom	20.53	1.15
	Mesh	90.17	2.00
	Opt-mesh	38.60	2.00

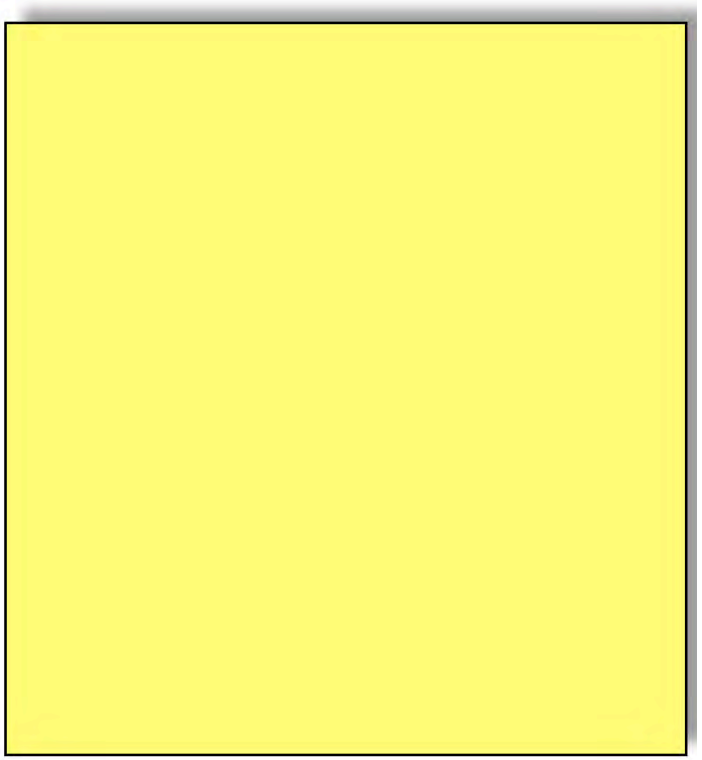

Case Study 2: SUNFLOOR Vs Regular Topologies

Application	Topology	Power(mW)	Avg. nr. hops
VPROC (42 cores)	Custom	79.64	1.67
	Mesh	301.8	2.58
	Opt-mesh	136.1	2.58
MPEG4 (12 cores)	Custom	27.24	1.50
	Mesh	96.82	2.17
	Opt-mesh	60.97	2.17
VOPD (12 cores)	Custom	30.00	1.33
	Mesh	95.94	2.00
	Opt-mesh	46.48	2.00
MWD (12 cores)	Custom	20.53	1.15
	Mesh	90.17	2.00
	Opt-mesh	38.60	2.00

- On average, SunFloor custom topologies:

Case Study 2: SUNFLOOR Vs Regular Topologies

Application	Topology	Power(mW)	Avg. nr. hops
VPROC (42 cores)	Custom	79.64	1.67
	Mesh	301.8	2.58
	Opt-mesh	136.1	2.58
MPEG4 (12 cores)	Custom	27.24	1.50
	Mesh	96.82	2.17
	Opt-mesh	60.97	2.17
VOPD (12 cores)	Custom	30.00	1.33
	Mesh	95.94	2.00
	Opt-mesh	46.48	2.00
MWD (12 cores)	Custom	20.53	1.15
	Mesh	90.17	2.00
	Opt-mesh	38.60	2.00

- On average, SunFloor custom topologies:
  - **2.75x** less power consumption

Case Study 2: SUNFLOOR Vs Regular Topologies

Application	Topology	Power(mW)	Avg. nr. hops
VPROC (42 cores)	Custom	79.64	1.67
	Mesh	301.8	2.58
	Opt-mesh	136.1	2.58
MPEG4 (12 cores)	Custom	27.24	1.50
	Mesh	96.82	2.17
	Opt-mesh	60.97	2.17
VOPD (12 cores)	Custom	30.00	1.33
	Mesh	95.94	2.00
	Opt-mesh	46.48	2.00
MWD (12 cores)	Custom	20.53	1.15
	Mesh	90.17	2.00
	Opt-mesh	38.60	2.00

- On average, SunFloor custom topologies:
  - **2.75x** less power consumption

Case Study 2: SUNFLOOR Vs Regular Topologies

Application	Topology	Power(mW)	Avg. nr. hops
VPROC (42 cores)	Custom	79.64	1.67
	Mesh	301.8	2.58
	Opt-mesh	136.1	2.58
MPEG4 (12 cores)	Custom	27.24	1.50
	Mesh	96.82	2.17
	Opt-mesh	60.97	2.17
VOPD (12 cores)	Custom	30.00	1.33
	Mesh	95.94	2.00
	Opt-mesh	46.48	2.00
MWD (12 cores)	Custom	20.53	1.15
	Mesh	90.17	2.00
	Opt-mesh	38.60	2.00

- On average, SunFloor custom topologies:
  - **2.75x** less power consumption
  - **1.55x** less hop delay

Case Study 2: SUNFLOOR Vs Regular Topologies

Application	Topology	Power(mW)	Avg. nr. hops
VPROC (42 cores)	Custom	79.64	1.67
	Mesh	301.8	2.58
	Opt-mesh	136.1	2.58
MPEG4 (12 cores)	Custom	27.24	1.50
	Mesh	96.82	2.17
	Opt-mesh	60.97	2.17
VOPD (12 cores)	Custom	30.00	1.33
	Mesh	95.94	2.00
	Opt-mesh	46.48	2.00
MWD (12 cores)	Custom	20.53	1.15
	Mesh	90.17	2.00
	Opt-mesh	38.60	2.00

- On average, SunFloor custom topologies:
  - **2.75x** less power consumption
  - **1.55x** less hop delay

Case Study 2: SUNFLOOR Vs Regular Topologies

Application	Topology	Power(mW)	Avg. nr. hops
VPROC (42 cores)	Custom	79.64	1.67
	Mesh	301.8	2.58
	Opt-mesh	136.1	2.58
MPEG4 (12 cores)	Custom	27.24	1.50
	Mesh	96.82	2.17
	Opt-mesh	60.97	2.17
VOPD (12 cores)	Custom	30.00	1.33
	Mesh	95.94	2.00
	Opt-mesh	46.48	2.00
MWD (12 cores)	Custom	20.53	1.15
	Mesh	90.17	2.00
	Opt-mesh	38.60	2.00

- On average, SunFloor custom topologies:
  - **2.75x** less power consumption
  - **1.55x** less hop delay
- Despite large design space, maximum run time of 1 hour for VPROC

Case Study 3: High-End Application

Case Study 3: High-End Application

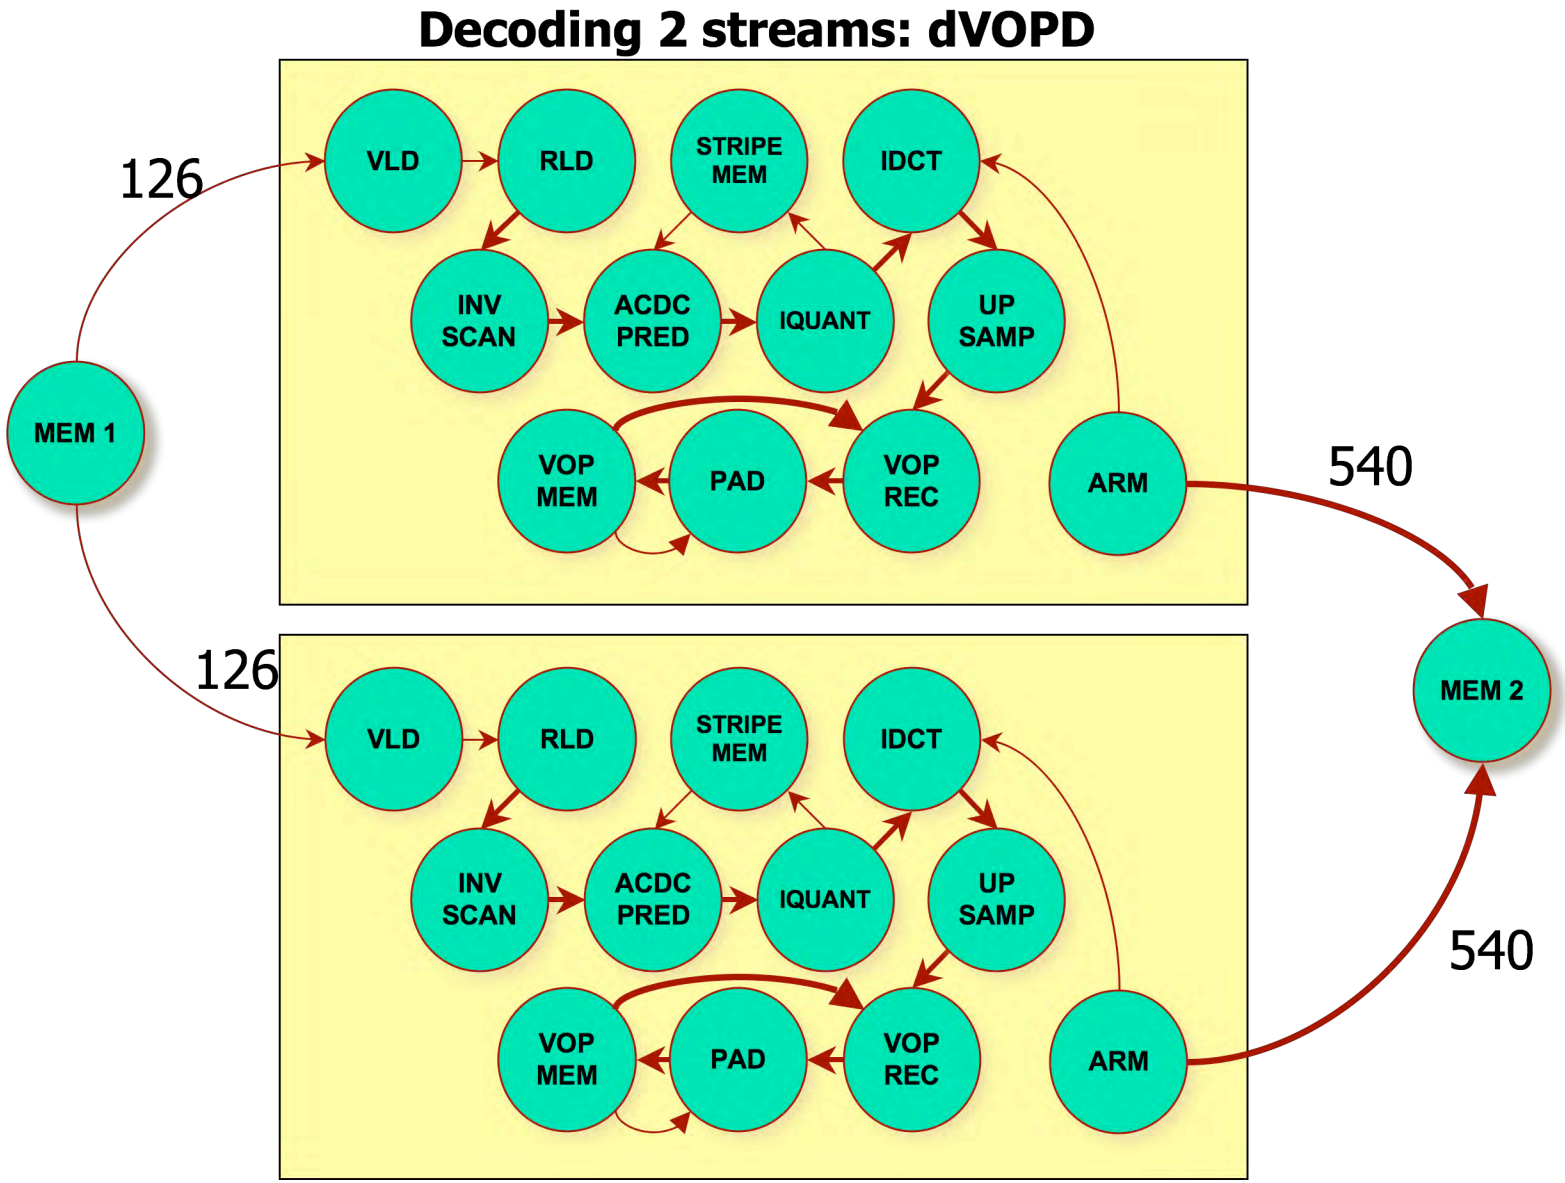

Case Study 3: High-End Application

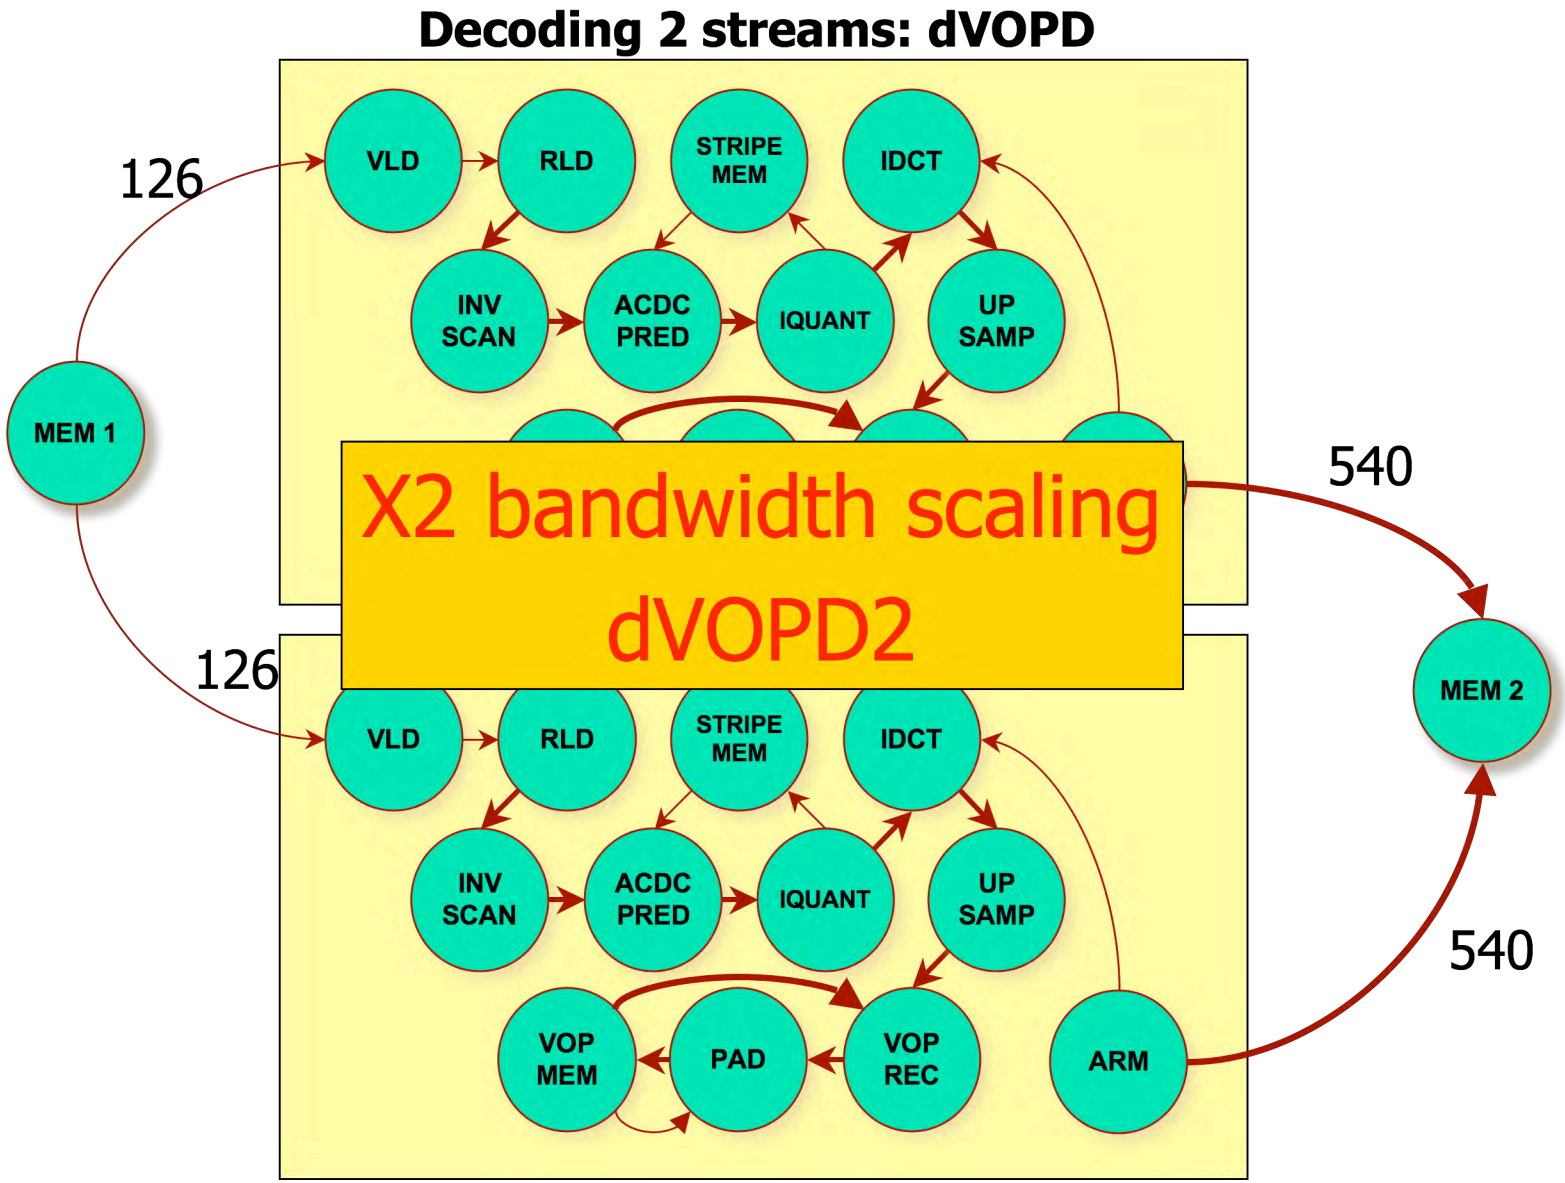

Case Study 3: High-End Application

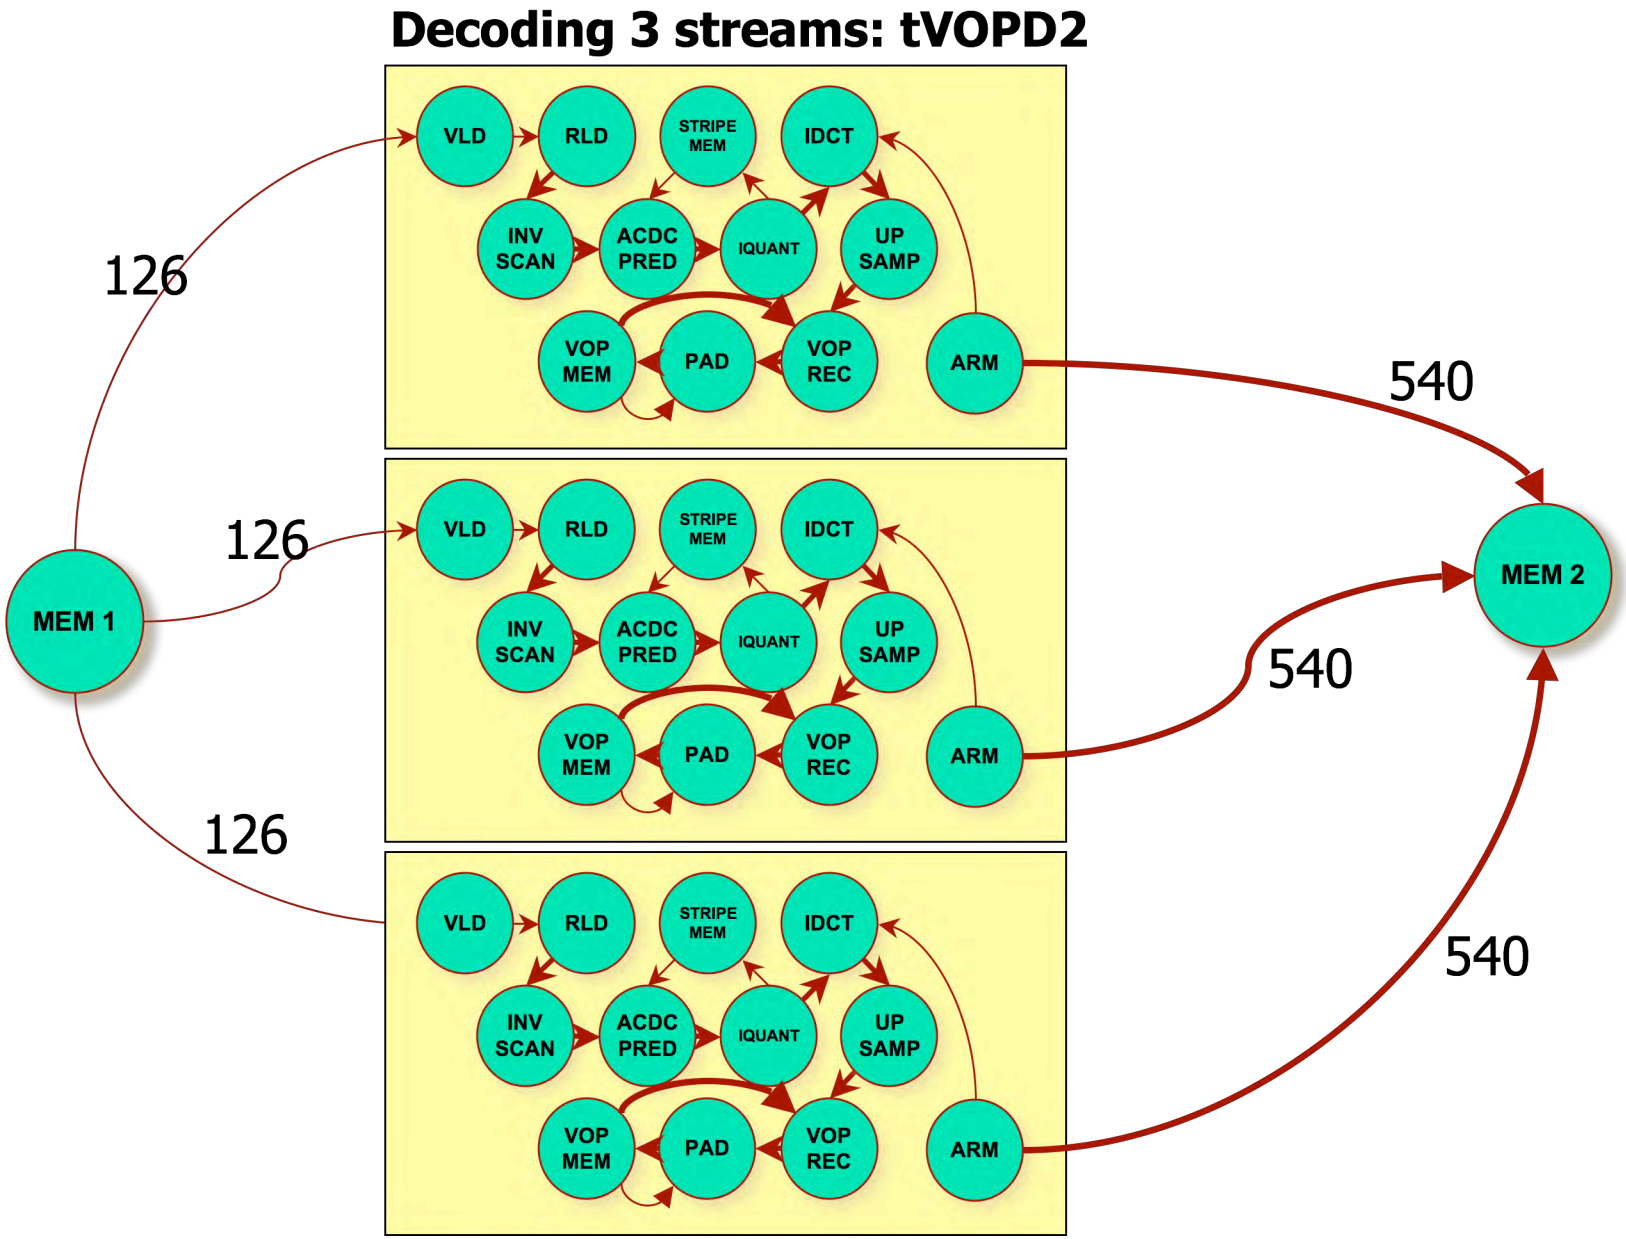

Case Study 3: Network Synthesis Results

	Library	Frequency	Switch Count	Largest Switch	Total NoC Power	Avg. head flit latency
	90nm LVT	400 MHz	4	10x9	175.88 mW	3.42 cycles [3,5]
	90nm SVT	-	-	-	-	-
	65nm LVT	400 MHz	4	10x9	81.96 mW	3.42 cycles [3,5]
	65nm HVT	-	-	-	-	-
dVOPD2	65nm LVT	800 MHz	6	7x6	129.36 mW	4.24 cycles [3,7]
tVOPD2	65nm LVT	800 MHz	10	7x7	196.40 mW	4.35 cycles [3,9]

Case Study 3: Network Synthesis Results

	Library	Frequency	Switch Count	Largest Switch	Total NoC Power	Avg. head flit latency
	90nm LVT	400 MHz	4	10x9	175.88 mW	3.42 cycles [3,5]
	90nm SVT					
	65nm LVT					ycles 5]
	65nm HVT					
dVOPD2	65nm LVT	800 MHz	6	7x6	129.36 mW	ycles [3,7]
tVOPD2	65nm LVT	800 MHz	10	7x7	196.40 mW	4.35 cycles [3,9]

- Observations:
  - Lower power in 65nm for same design
  - 65 nm supports 2x BW, at lower power!
  - NoC for a big design (38 cores) operates at 800 MHz
  - With increasing app BW or number of cores, more switches needed (due to freq limit of switches)

Case Study 4: Low Bandwidth & Power Application

- Parallel encryption engine (18 cores)

Library	Frequency	Switch Count	Largest Switch	Total NoC Power	Avg. head flit latency
90nm LVT	50 MHz	2	11x11	10.4 mW	3.94 cycles [3,5]
90nm SVT	50 MHz	2	11x11	4.1 mW	3.94 cycles [3,5]
65nm LVT	50 MHz	2	11x11	4.72 mW	3.94 cycles [3,5]
65nm HVT	50 MHz	5	9x9	3.1 mW	4.38 cycles [3,7]

Looking Forward

- Quality of service guarantees for critical traffic
- Run-time configurability
- Robustness w.r.t. to static/dynamic variations, errors
- Network interfaces: interoperability, performance