

Physically-Aware Preemptive Virtual Channels for Deadlock-Free AXI Networks-on-Chip

Lorenzo Leone¹, Luca Colagrande¹, Luca Benini^{1,2}

¹Integrated Systems Laboratory, ETH Zurich

²Department of Electrical, Electronic, and Information Engineering, University of Bologna




1 Motivation and Challenges


- As many-core SoC scale, Network-on-Chips must sustain higher memory bandwidth adding pressure on the NoC routing.
- Providing industry-standard AXI4 interfaces at the network boundaries is critical for seamless IPs integration.
- Deadlock freedom remains central: endpoints can create protocol-level deadlocks invisible to the routing algorithm.

How do we map AXI channels preserving deadlock freedom?

Multiplane




Lightweight Virtual Channels



2 Implementation


We explore three lightweight VC implementations, progressively removing the timing and buffering overheads of conventional flow control while keeping the wiring efficiency of Virtual Channels.

Naïve




- Single physical data link (512 bit)
- Replicate valid-ready per VC
- Valid-ready combinational dependency
- Tile-to-tile delay on the critical path

Preemptive



- Single physical data link (512 bit)
- Replicate valid-ready per VC
- Cut valid-ready dependency
- Full bandwidth with no extra buffering



Credit Based



- Single physical data link (512 bit)
- Credit signal for downstream flit consumption
- Extra buffering for full bandwidth

3 Results and Discussion

We evaluate the four deadlock-free AXI4 traffic-separation schemes in a representative 4x4 multicluster SoC implemented in TSMC 7 nm, comparing timing, area, routing-resource usage, and bandwidth utilization. The naïve VC design incurs up to a **15% frequency penalty** due to its long combinational backpressure path, while the credit-based implementation removes this path at the cost of up to **28% more area** at 1 GHz. In contrast, **Preemptive VCs** sustain full bandwidth while reaching **1.7 GHz**, matching the multiplane design with only **3% router-area overhead**. By sharing the wide physical link, they also **reduce link-resource requirements by up to 76%** relative to the multiplane baseline.



Configuration	BW util.	Area [kGE]	Max. Freq. [GHz]	#Metal tracks
Multiplane	100%	270	1.68 (-1.5%)	2776 (+76%)
Naïve	100%	275 (+2%)	1.42 (-16%)	1578
CreditBased 2 buff.	66%	294 (+9%)	1.62 (-4.9%)	1610 (+2%)
CreditBased 3 buff.	100%	345 (+28%)	1.62 (-4.9%)	1610 (+2%)
Preemptive	100%	279 (+3%)	1.70	1578

Percentages are relative to the best value for each metric, highlighted in bold.

Full bandwidth is preserved by all designs except the 2-buffer Credit-Based VC, whose credit latency causes up to **33% throughput loss**.

Preemptive VCs recover multiplane timing, while Credit-Based VCs pay a large area cost to sustain full bandwidth.

Preemptive VCs provide the best overall trade-off: **100% bandwidth** utilization, **1.7 GHz**, only **+3% area**, while **saving 76%** routing resources.

Preemptive VC, the best of both worlds: VC wiring efficiency with multiplane-like simplicity.