FlooNoC-QoS: Enhancing a RISC-V-based FPGA Overlay with an AXI4-compliant NoC featuring QoS Control Capability

Gianluca Bellocchi[1], Alessandro Capotondi[1], Luca Benini[2,3] and Andrea Marongiu[1]

[1]University of Modena and Reggio Emilia, [2]ETH Zürich, [3]University of Bologna

Introduction and Motivation

Heterogeneous Systems-on-Chip (HeSoCs) [1] combine:

- General-purpose processors (Host): programmability, legacy SW;
- Domain-specific accelerators (DSAs): high performance, energy efficiency;
- Memory subsystem (Mem): shared, distributed;
- **Network-on-Chip (NoC)**: system scalability, high bandwidth.

Real-time applications impose strict requirements for **timing predictability**, thus various challenges must be overcome for HeSoCs adoption:

- **Memory sharing interference**: *unpredictable* timing behaviour;
- **Memory-intensive workloads**: *high latency* toward the off-chip memory;
- **Heterogeneous traffic**: Host (cache-line sized) vs. DSA (bulk transfers).

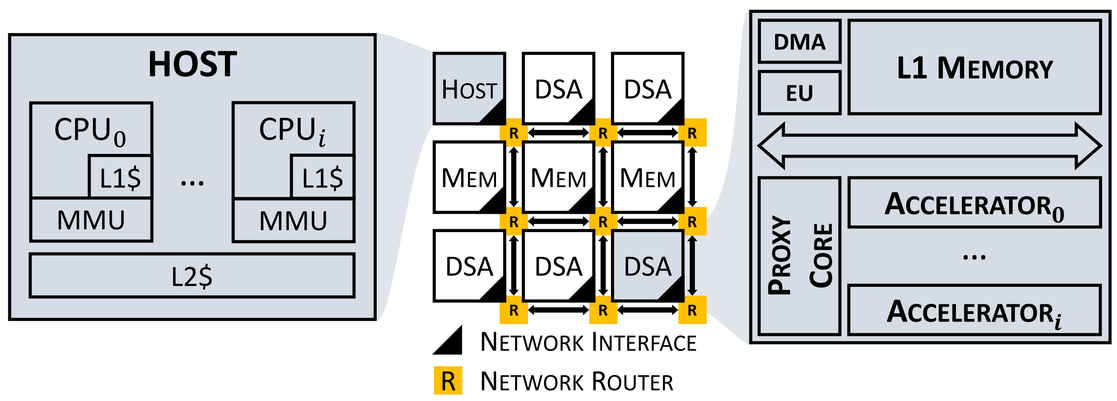

Figure 1: HeSoC architecture

About FlooNoC-QoS

FlooNoC-QoS is an ongoing research project:

- Based on FlooNoC [2], part of the open-source PULP Platform [3];
- *Wormhole NoC fabric* with multiple physical channels (w/o virtual links);
- *Full AXI4 compliance*:
  - Decoupled link-level protocol (flit-based) for better scalability;
  - Support of **multiple burst-based outstanding transactions**;
  - Streamlined on-chip integration of heterogeneous components.
- **QoS control**: *memory bandwidth regulation (monitoring + throttling)* [4];
- Implementation: **FPGA** (focus of this work), ASIC.

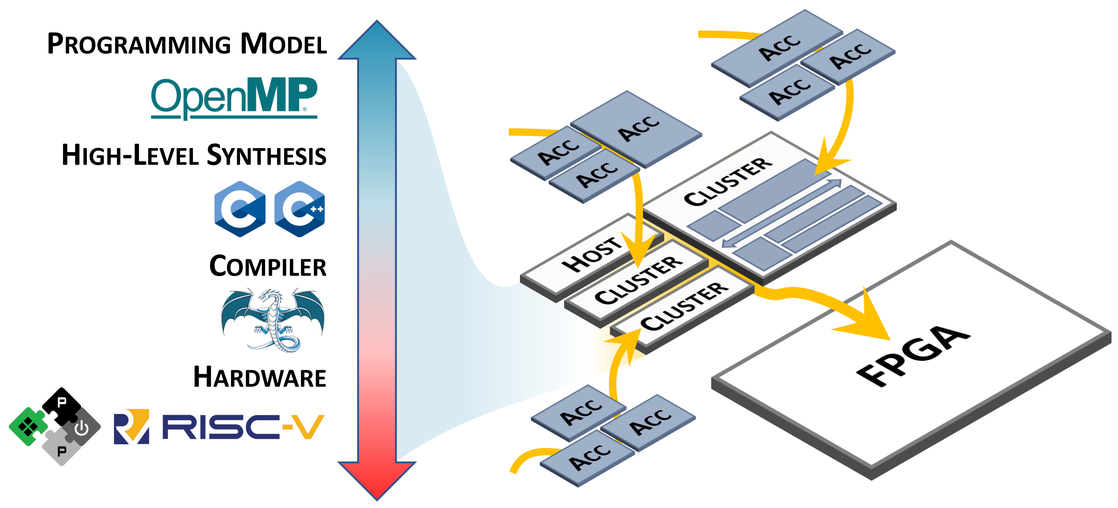

Figure 2: Richie SLD framework (under review)

Results: FPGA Porting and Deployment

1. Porting of *2D mesh* to COTS FPGA-based HeSoC (**ZU9EG, f_{ck}=100MHz**);
2. Characterization of *FPGA resource utilization* (LUT, FF, see Fig. 3):
   - NoC more beneficial over FPGA-optimized XBAR when $N_{tiles} \geq 7$.
3. Full-stack emulation of a **RISC-V-based FPGA overlay** [5]:
   - Integration and deployment exploiting the **Richie SLD framework**;
   - Testing of Host-DSA interaction with OpenMP-based applications.

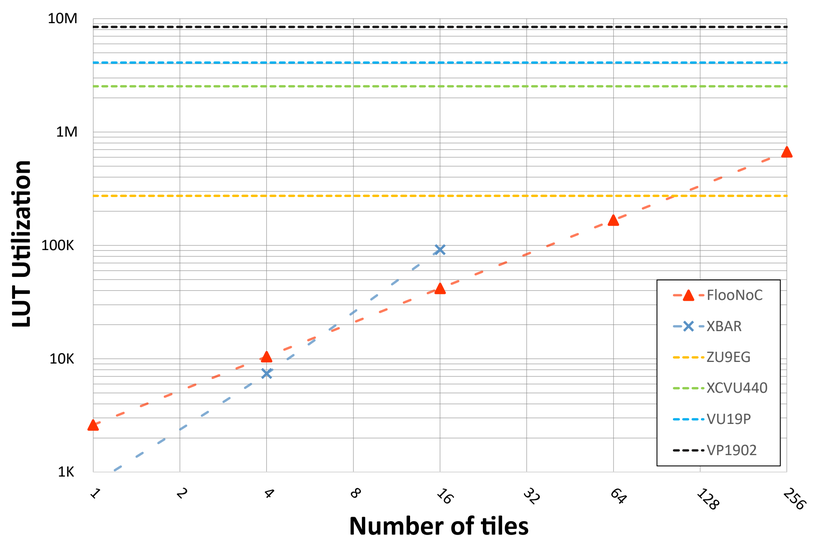

Figure 3: LUT resource utilization

Conclusions and Future Works

- Investigation of *flexible and runtime configurable* **QoS mechanisms**:
  - **Master-controlled**: tightly-coupled monitoring and throttling [4];
  - **Network-controlled**: fragmentation of outstanding transactions [6].
- Design of massively-parallel HeSoCs—accelerator-rich and many-core—for *safety-critical real-time platforms*, e.g., autonomous cars and robots.

References

[1] J. L. Hennessy and D. A. Patterson, "A new golden age for computer architecture," *Communications of the ACM*, vol. 62, no. 2, pp. 48–60, 2019.

[2] T. Fischer, M. Rogenmoser, T. Benz, F. K. Gürkaynak, and L. Benini, "Floonoc: A 645 gbps/link 0.15 pj/b/hop open-source noc with wide physical links and end-to-end axi4 parallel multi-stream support," *arXiv preprint arXiv:2409.17606*, 2024.

[3] D. Rossi, F. Conti, A. Marongiu, A. Pullini, I. Loi, M. Gautschi, G. Tagliavini, A. Capotondi, P. Flatresse, and L. Benini, "Pulp: A parallel ultra low power platform for next generation iot applications," in *2015 IEEE Hot Chips 27 Symposium (HCS)*. IEEE Computer Society, 2015, pp. 1–39.

[4] G. Valente, G. Brilli, T. Di Mascio, A. Capotondi, P. Burgio, P. Valente, and A. Marongiu, "Fine-grained qos control via tightly-coupled bandwidth monitoring and regulation for fpga-based heterogeneous socs," *IEEE Transactions on Parallel and Distributed Systems*, 2024.

[5] G. Bellocchi, A. Capotondi, F. Conti, and A. Marongiu, "A risc-v-based fpga overlay to simplify embedded accelerator deployment," in *2021 24th Euromicro Conference on Digital System Design (DSD)*. IEEE, 2021, pp. 9–17.

[6] T. Benz, A. Ottaviano, R. Balas, A. Garofalo, F. Restuccia, A. Biondi, and L. Benini, "Axi-realm: A lightweight and modular interconnect extension for traffic regulation and monitoring of heterogeneous real-time socs," in *2024 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2024, pp. 1–6.

Links

FlooNoC
Fully AXI4-compliant NoC implementation

Richie
SLD framework for RISC-V-based HeSoCs

PULP Platform
HW/SW ecosystem of RISC-V-based IPs

Acknowledgments

This work was supported by the European Union under the NextGenerationEU Programme within the Plan "PNRR - Missione 4 "Istruzione e Ricerca" - Componente C2 Investimento 1.1 "Fondo per il Programma Nazionale di Ricerca e Progetti di Rilevante Interesse Nazionale (PRIN)" by the Italian Ministry of University and Research (MUR)".

pulp-platform.org pulp-platform @pulp_platform gianluca.bellocchi@unimore.it