A Beneš-Based NoC Switching Architecture for Mixed Criticality Embedded Systems

Problem description

A *predictable processor* attached to an *unpredictable interconnect* is no longer a predictable processor.

Existing NoC structures have the following problems:

- Topologies produce variable timing bounds.
- Congestion can introduce further delays.
- Task placement highly sensitive to topology.

Contributions

- Specification of a **NoC switching architecture** that meets multi-core MCE needs.
- A novel, **non-blocking and timing-predictable** implementation addressing these requirements.
- Demonstration of **scaling** and performance.
- **Formal verification** [1] to prove correct behaviour.

Beneš networks

- Multi-stage networks: Clos [2], Beneš [3].
- Low switching costs (telephony, VLSI [4]).
- All nodes equidistant (num. stages).

Figure: Eight port, three stage Clos / Five stage Beneš conceptual comparison.

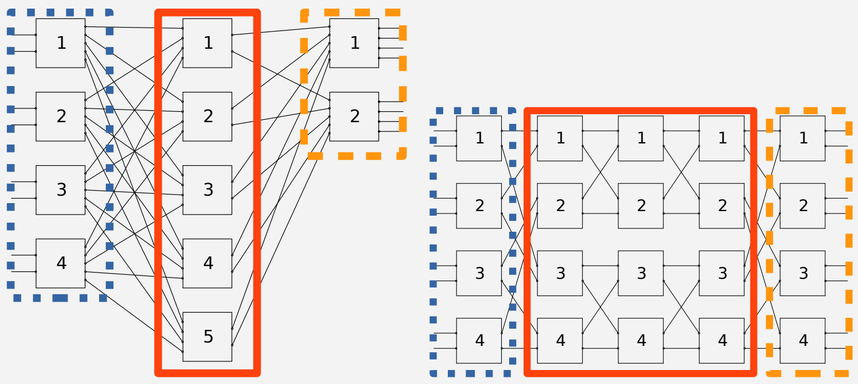

(a) Clos (b) Beneš

Figure: Folding a Beneš network to create node connections.

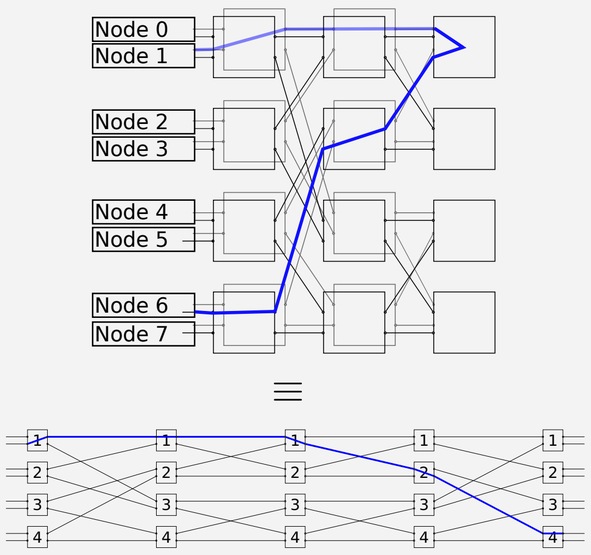

Figure: An eight-node network using two-port switching elements, eight routes established.

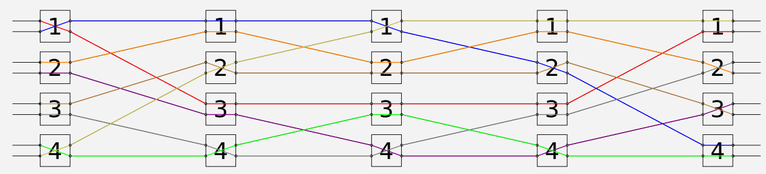

Figure: Two sub-networks of four nodes each, created from an eight node system.

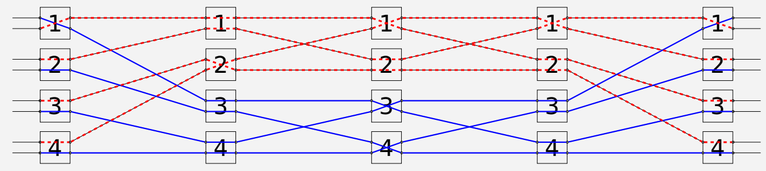

Verification goals

Set of high-level requirements, e.g.:

- Communication between two nodes is non-blocking.
- Invalid or erroneous communications cannot interfere with more critical communications.
- Routing decisions must be **statically resolvable** with **minimal overhead** at large scales.

Refined into *system-*, *network-* and *switch-level* specs:

ID	Description	Property
C1	No two active inputs can share the same output channel.	`no_shared_direction`
C2	Incoming data to an active switch propagates to the correct output on the next clock cycle.	`data_fwd_timing, ...`
C15	A port rejects further data in the event of an inbound error signal and propagates the error signal.	`reject_on_err`
N4	All routes through the network connect to the expected destination if no routing conflict arises and the target port does not assert an error.	`route_correct`
S4	Higher priority tasks must create their routes before lower priority tasks.	N/A

These can then be formalised during implementation.

MCENoC implementation & verification [5]

- SystemVerilog & Verilog HDL.
- SystemVerilog Assertions (SVA). E.g. **C2**:

```
property data_fwd_timing(i);
  ( port_is_active(i) and act_in[i] ) |=>
  ( dat_out[direction[i]] == $past(dat_in[i], 1)
    and clm_out[direction[i]] and act_out[direction[i]] );
endproperty
```

- Configurable switch size to reduce stage-count.
- Verified formally using Cadence Jasper Gold.

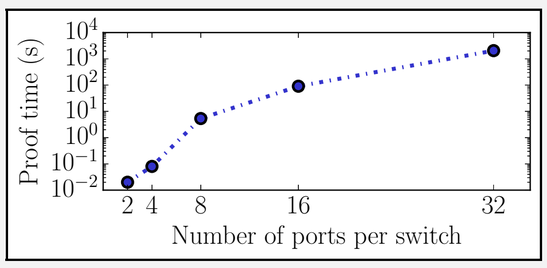

(a) Proof time for a single switching element

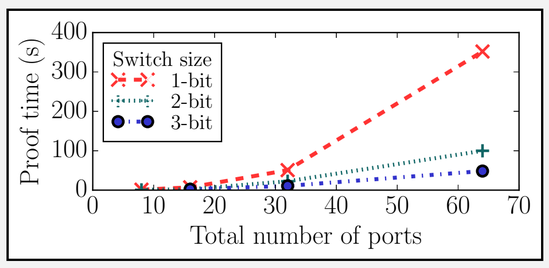

(b) Proof time for networks of varying stage-count & switching element size

- Routing setup/teardown scales into thousands of nodes, retaining **real-time response** capabilities.

Figure: Time to cycle through N communications in an N-node network at 364 Mbit/s/port with 2-port switches.

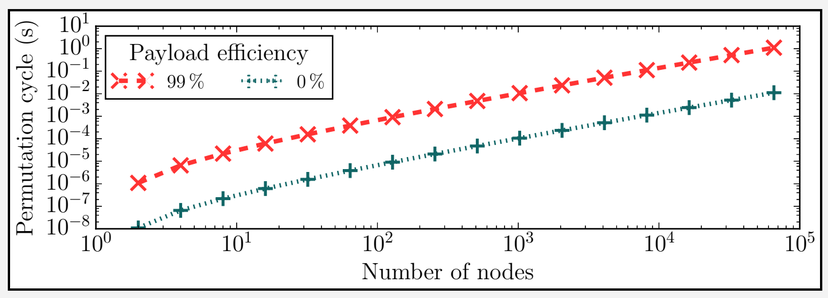

Ongoing & Future work

- Integrating with RISC-V cores in FPGA: 32-way system is now synthesized and implemented at 100 MHz, 25.6 Gbit/s bisection bandwidth.
- Gathering use cases for application-level benchmarks. Interested in:
  - Different communication patterns (mesh, graph, . . .).
  - Mixed use cases.
  - **Talk to us!**
- Combining static and dynamic communication scheduling.
- Formal verification of the software (kernel).
- Fault injection & handling.

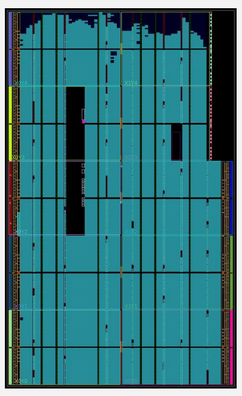

References

[1] Erik Seligman, Tom Schubert, and MV Achutha Kiran Kumar. *Formal Verification: An Essential Toolkit for Modern VLSI Design*. Morgan Kaufmann, 2015.

[2] Charles Clos. "A Study of Non-Blocking Switching Networks". In: *Bell System Technical Journal* (1952), pp. 406–424. DOI: 10.1002/j.1538-7305.1953.tb01433.x.

[3] V. E. Beneš. "On Rearrangeable Three-Stage Connecting Networks". In: *Bell System Technical Journal* 41.5 (Sept. 1962), pp. 1481–1492. DOI: 10.1002/j.1538-7305.1962.tb03990.x.

[4] Yikun Jiang and Mei Yang. "On circuit design of on-chip non-blocking interconnection networks". In: *2014 27th IEEE International System-on-Chip Conference (SOCC)*. Vol. 40. 8. IEEE, Sept. 2014, pp. 192–197. DOI: 10.1109/SOCC.2014.6948925.

[5] Steve Kerrison, David May, and Kerstin Eder. "A Benes Based NoC Switching Architecture for Mixed Criticality Embedded Systems". In: *IEEE 10th International Symposium on Embedded Multicore/Many-core Systems-on-Chip (MCSoC 2016) [to appear]*. Lyon, France, 2016, pp. 1–8.

Steve Kerrison, David May, Kerstin Eder

firstname.lastname@bristol.ac.uk

University of Bristol, United Kingdom

The research leading to these results has received funding from the ARTEMIS Joint Undertaking under grant agreement number 621429 (project EMC2).

University of BRISTOL

EMC²