

On Program Restructuring for Cluster-based Highly Parallel

Embedded Architectures

Motivation

- **Embedded systems** more and more require to process large amounts of data coming from multiple sensors
- **Many-core processors** are starting to appear in the embedded domain, e.g. the Kalray Multi-Purpose Processor Array (**MPPA**) many-core
- The transformation of current single-core applications to parallel multi-threaded applications is a challenge
- This work evaluates strategies to **parallelize real-time applications** into many-cores
- **Goal:** To take an embedded application and parallelize it into the MPPA many-core to compare and evaluate different approaches for parallelization

Application Use Case

- A parallel and distributed **traffic simulator**: It computes the vehicle movements across lanes during T time steps
- 3 application levels: **External Linux-based, I/O and Compute Cluster applications**
- **Master/Slave** scheme: The master process running on the I/O spawns slave processes on the compute clusters that run the simulation
- **I/O application** starts-up the compute clusters, sends a dataset to work on and gathers the results
- **Cluster application** runs the simulation with the received data and sends back the results to the I/O

Algorithm 1 IO-level application

- 1: receive the input data file from Linux-based application
- 2: open and prepare the communication and synchronization connectors
- 3: starts up the compute clusters
- 4: wait until the clusters are ready (global sync)
- 5: send data to compute clusters (portal transfers)
- 6: wait until the clusters open the channels (global sync)
- 7: unblock each single cluster (individual sync)
- 8: wait for the results
- 9: show results

Algorithm 2 Cluster-level application

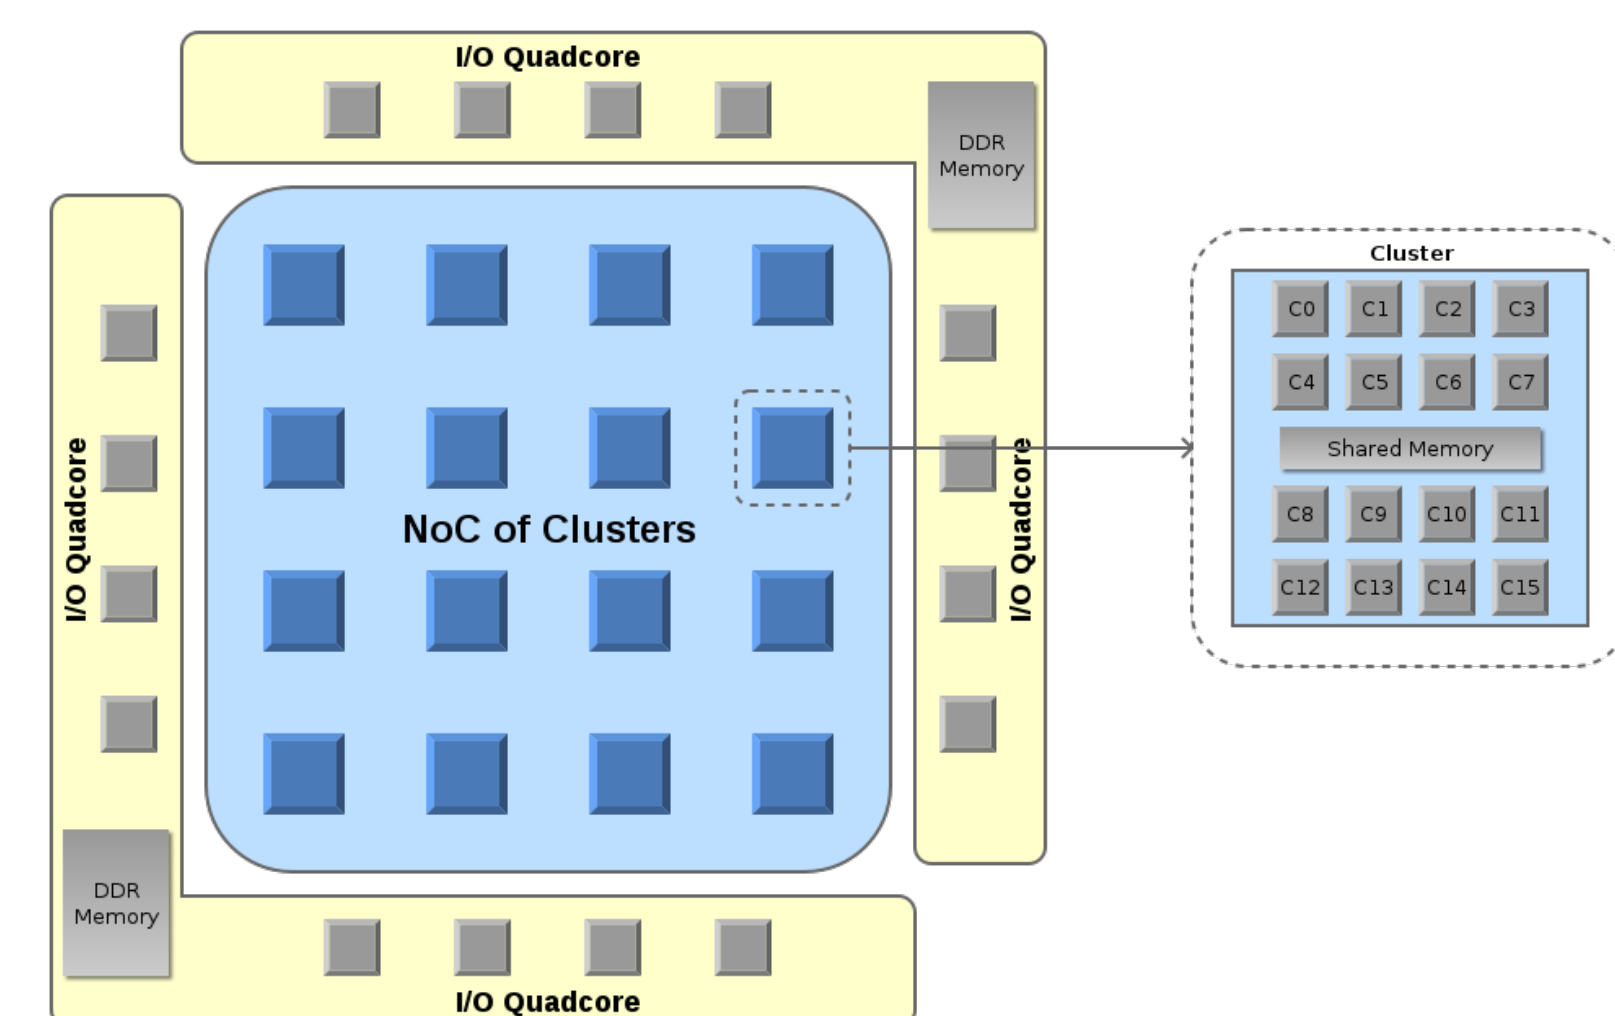
- 1: open and prepare the communication and synchronization connectors
- 2: contribute to unblock the master process (global sync)
- 3: wait for data
- 4: open the R/W channels
- 5: notify to the I/O that it has already opened its channels (global sync)
- 6: wait for the others (individual sync)
- 7: **for** $i = 0 \rightarrow SIMULATION_TIME$ **do**
- 8: run one-step simulation
- 9: exchange data on frontiers with other clusters (channels)
- 10: **end for**
- 11: send partial results (portal transfer)

- They use **portals** (multipoint transfers), **channels** (point-to-point) and **sync** (synchronization) connectors for communication and synchronization

- **OpenMP** is used to exploit the thread-level parallelism

MPPA Architecture


- **16 Compute Clusters** connected through a high-throughput Network-on-Chip with 16 cores each one
- **4 I/O subsystems** enable the access to the many-core processor



- It uses message passing (**MPPA IPC API**) and shared memory programming models (OpenMP and Pthreads)
- The MPPA IPC API provides connectors for communication and synchronization of processes

Performance Evaluation

- **Inter-cluster** parallelization (N clusters with 1 core) vs. **Intra-cluster** parallelization (N cores at 1 cluster)



- **Inter- and intra-cluster parallelization** are jointly used (N cores at M clusters)

