



exploring Manycore Architectures for Next-GeneratiOn HPC systems

# The MANGO Process for Designing and Programming Multi-Accelerator Multi-FPGA Systems


H2RC'18

Sunday, November 11th, 2018  
Dallas, Texas, USA

Rafael Tornero, José Flich, José María Martínez, Tomás Picornell, Vincenzo Scotti  
Email: [ratorga@disca.upv.es](mailto:ratorga@disca.upv.es)

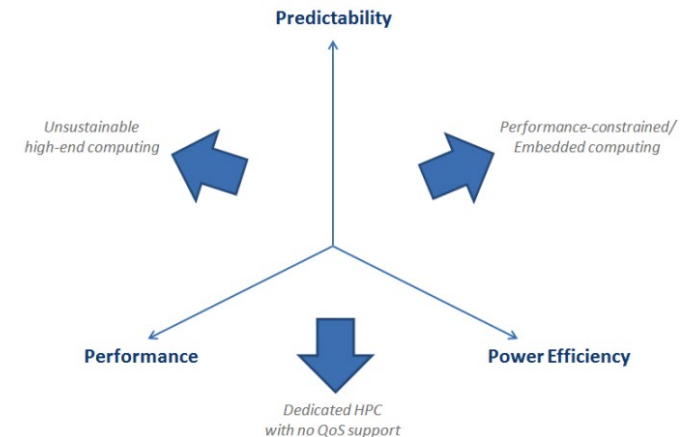


This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 671668



# MANGO Context

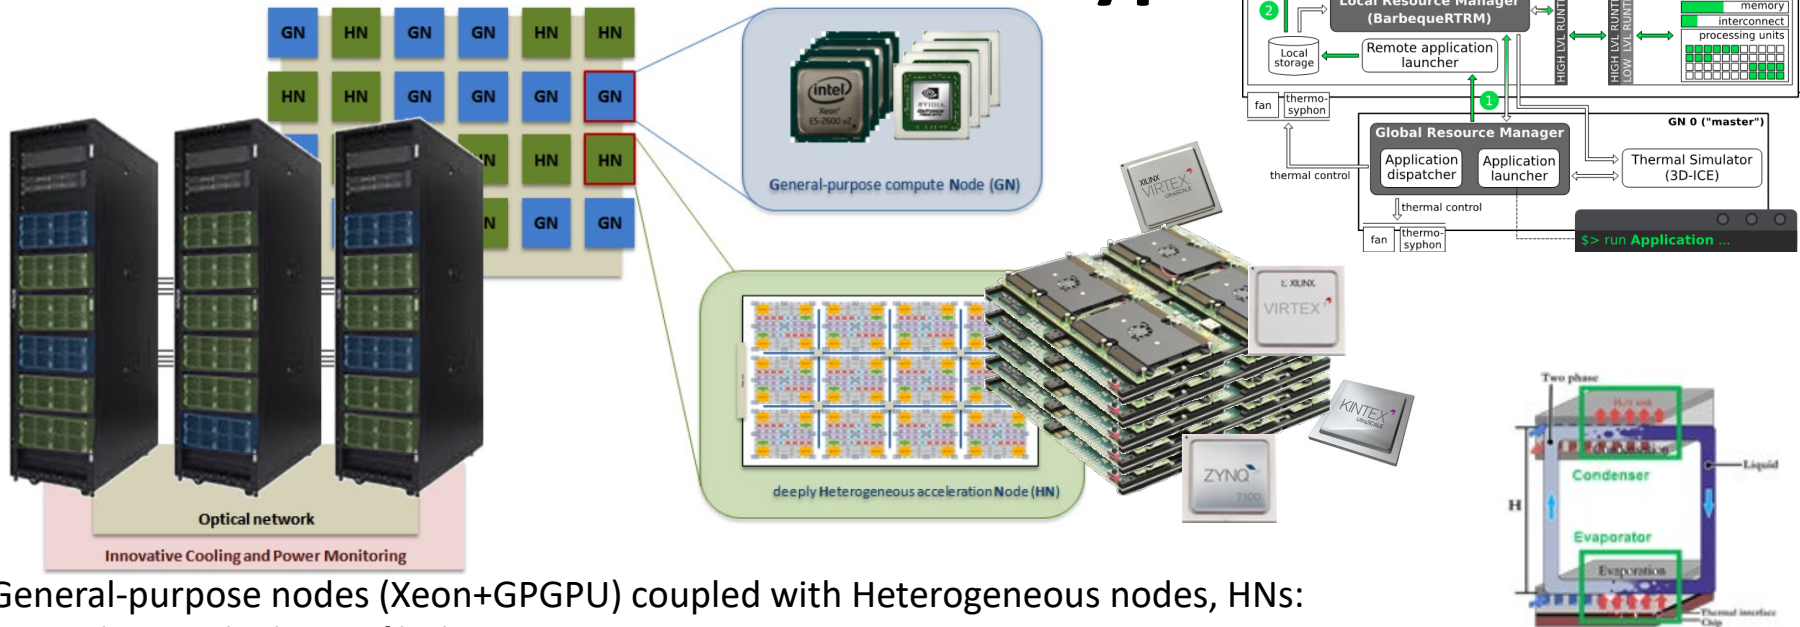
- MANGO FETHPC-2014 project:
  - is about *manycore architecture exploration in HPC*
- HPC quest for performance/power improvement
  - Trend in using heterogeneous components
    - GPUs, manycores, and even FPGAs
    - Goal is to get closer to the Intrinsic Computational Efficiency (ICE)
  - MANGO focuses on **heterogeneity**
    - How we combine heterogeneous components for the best achievement of computational efficiency
    - How to program/manage them for the best achievement of computational efficiency
- Emerging requirements on HPC systems:
  - Predictability (QoS; time sensitivity)
    - Due to the merging of HPC with Big Data
  - Capacity computing
    - Run as many application instances as possible
  - MANGO addresses **predictability** and **capacity computing**
    - **3P model** (Performance/Power/Predictability)




# MANGO Context

- MANGO builds a **prototyping system** for 3P space exploration
- Goals:
  - Hardware
    - Develop a flexible prototype for rapid exploration of architectures
    - Explore new deeply heterogeneous manycore architectures
    - Real-time support exploring the PPP design space
    - Provide a unified and simple (homogeneous) access to the system via a smart interconnect
  - Software
    - Adapt programming models and compiler support to the new architectures
    - Develop the right resource manager to deal with the system
  - Infrastructure
    - Provide new monitoring tools to the system
    - Provide new cooling techniques to the system
  - Applications
    - Analyze impact of on a set of real applications
    - Support of video transcoding, medical imaging, security and surveillance applications

# MANGO Prototype



- General-purpose nodes (Xeon+GPGPU) coupled with Heterogeneous nodes, HNs:
  - A large-scale cluster of high-capacity FPGAs
  - A robust, scalable interconnect for a **multi-FPGA manycore** system
  - Will enable FPGA acceleration *at scale*:
    - a key ingredient for the EsD roadmap
  - A continuum from FPGA emulation to the final physical platform (might be an ASIC manycore, FPGA, mixed...)
    - **under a stable software environment**
  - Native isolation and partitioning mechanisms for **QoS-aware capacity computing** HPC applications
- Two-phase passive **energy-efficient cooling**
- Demonstrated applications with stringent high-performance and QoS requirements







# Consortium



UNIVERSITAT  
POLITÈCNICA  
DE VALÈNCIA



**prodesign**  
electronic gmbh

**THALES**

**EAT•N**  
Powering Business Worldwide

**PHILIPS**

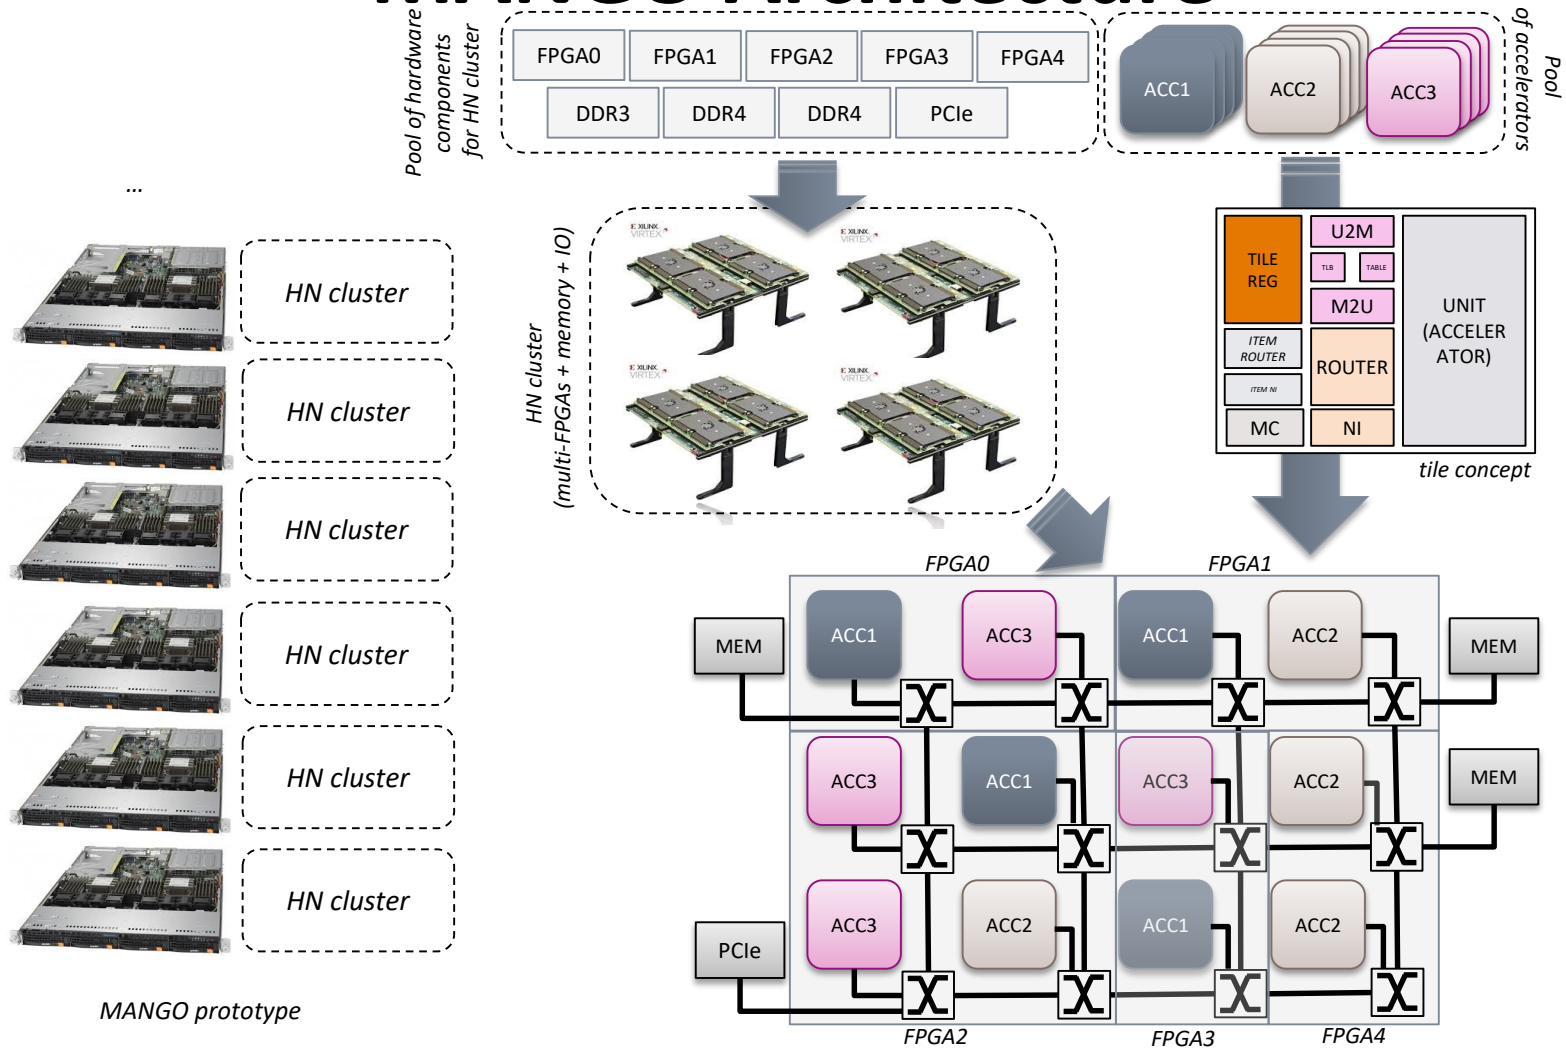
***MANGO: exploring Manycore Architectures  
for Next-GeneratiOn HPC systems***



# Agenda

- MANGO Architecture
  - HN Hardware and assembly
  - Heterogeneity
  - Network
  - Accelerator Interface
  - MANGO Design Flow
- FPGA resource utilization
- Conclusions

# MANGO Architecture







# HN Hardware and Assembly


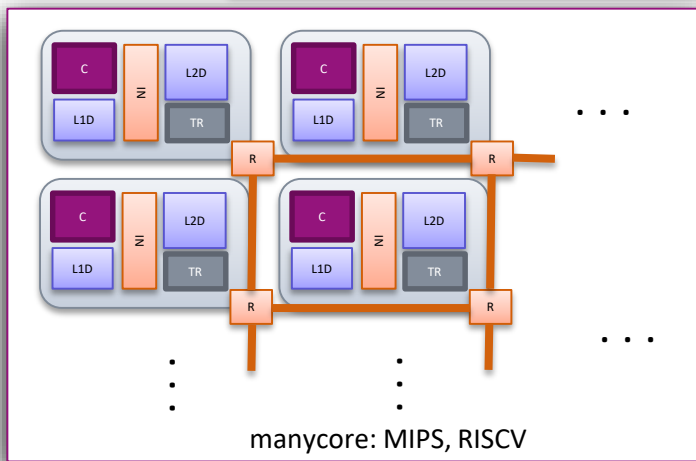
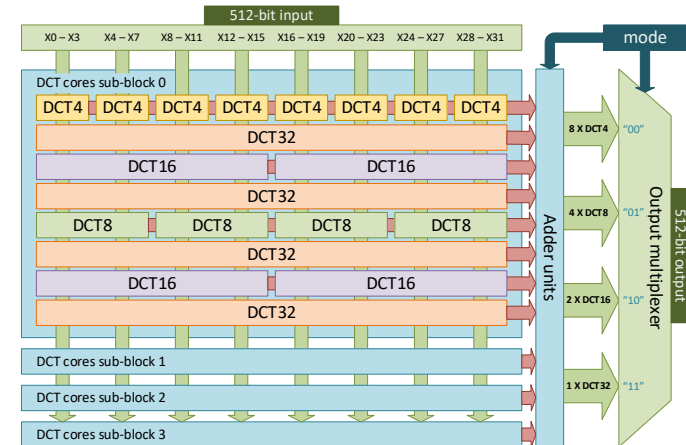
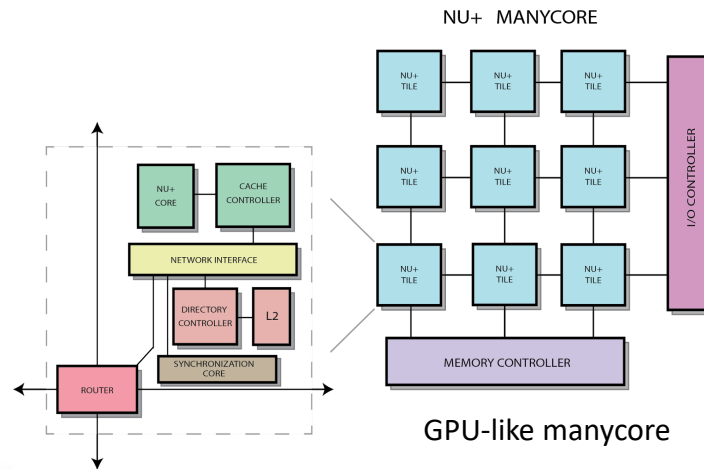


Lego-like exercise





# Heterogeneity: Pool of Accelerators

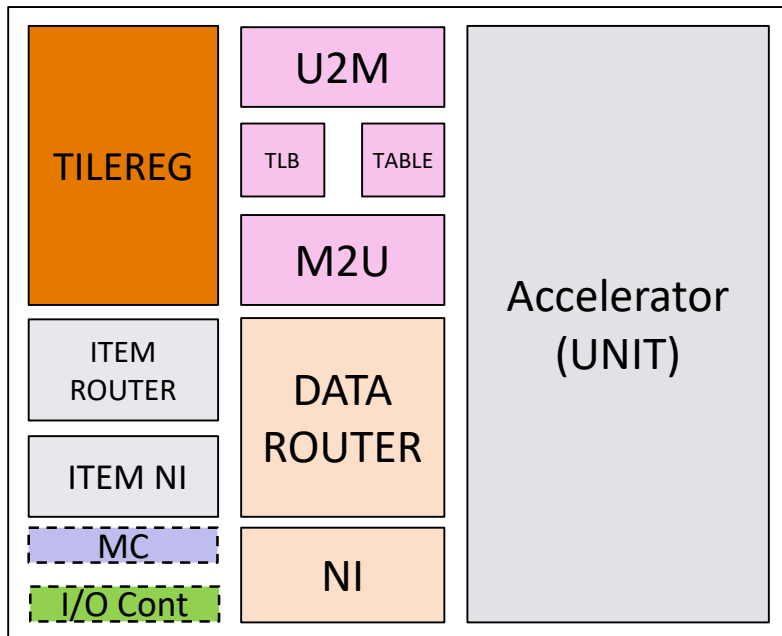


General purpose

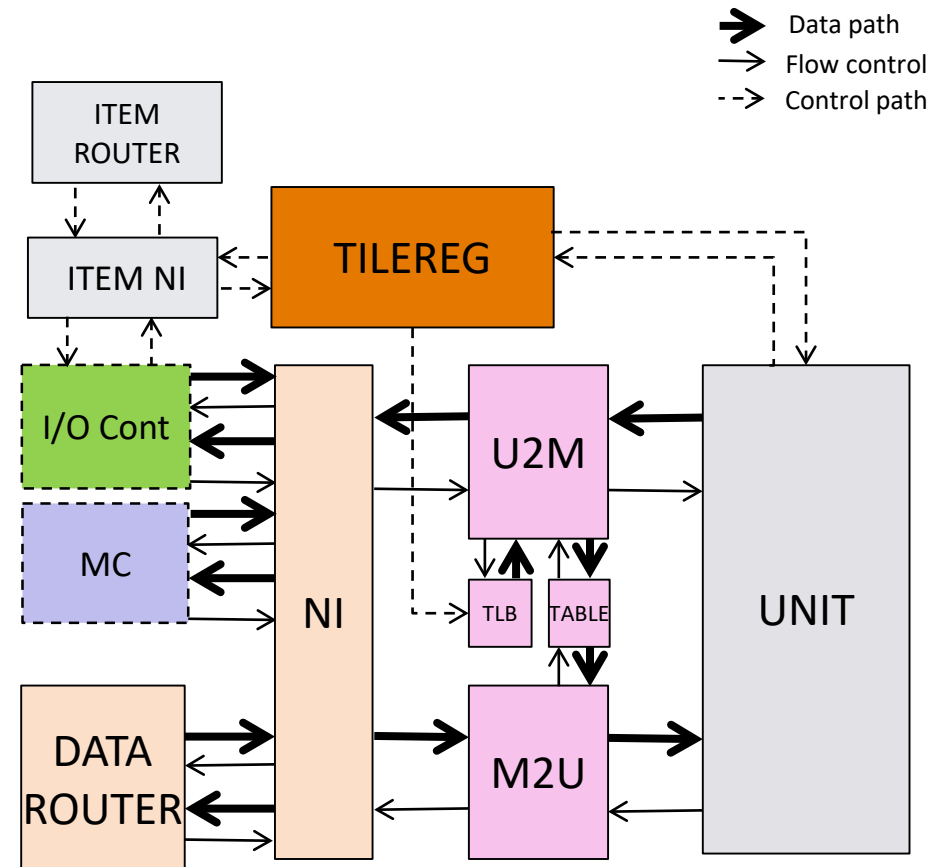
Custom made

# Heterogeneity


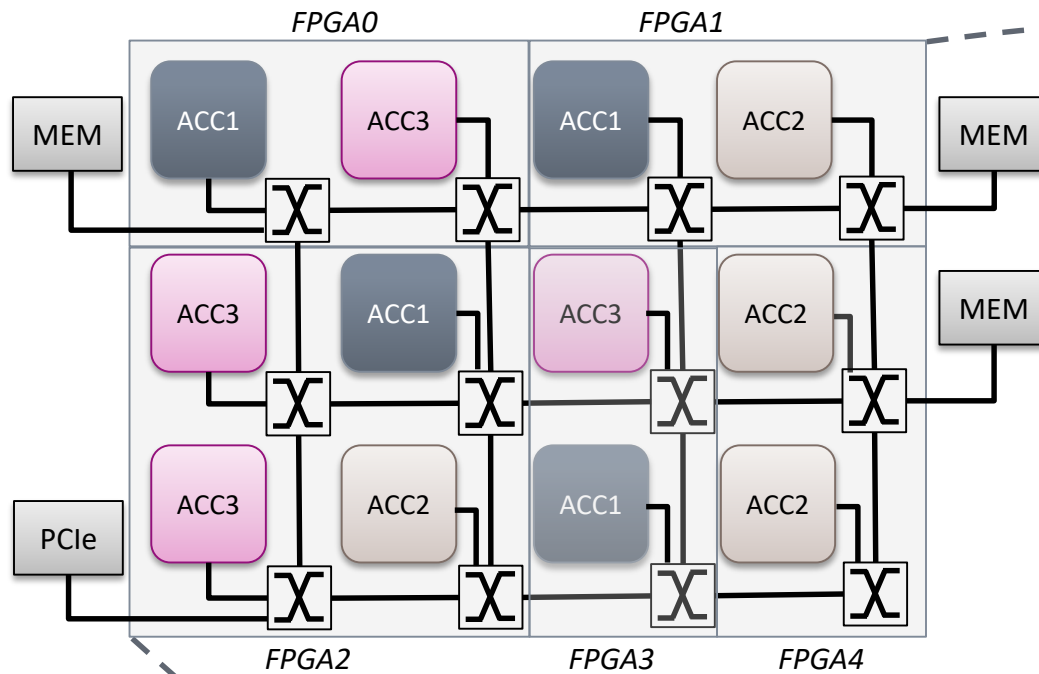
## Tile concept



## Element connection


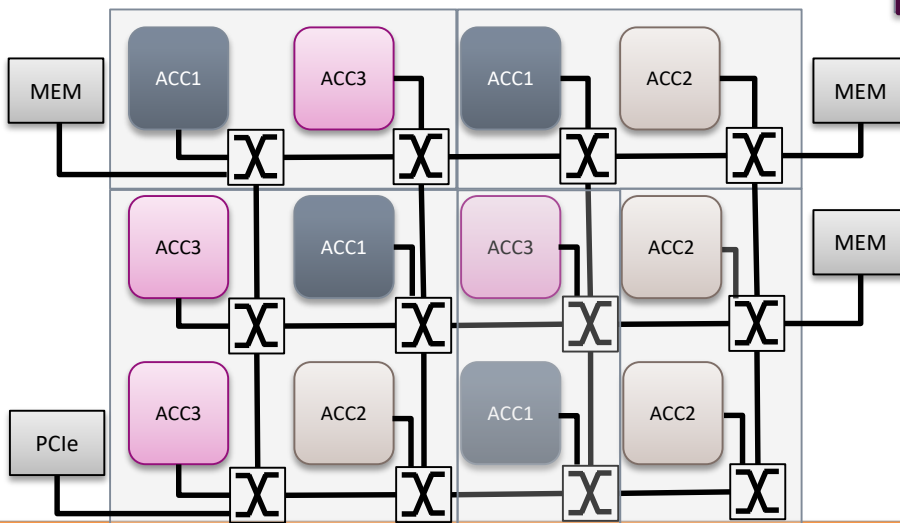




# Heterogeneity, but regular layout



# Data Network


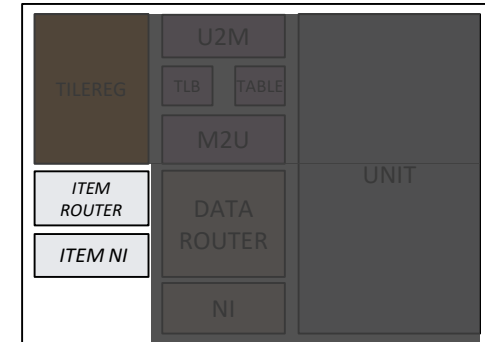
- DATA Routers connected among them in a 2D mesh layout.
- NI decouples network from tile components.
- Support for different Virtual networks (VN) with different number of Virtual channels (VC).
- Support for dynamic assignment of bandwidth per VN.
- Support for capacity computing







# Control Network

- Used for configuration and monitoring
  - MANGO Infrastructure
  - Accelerators or units
- Flexible and generic to let the accelerators be configured based on their complexities
  - Ad-hoc protocols

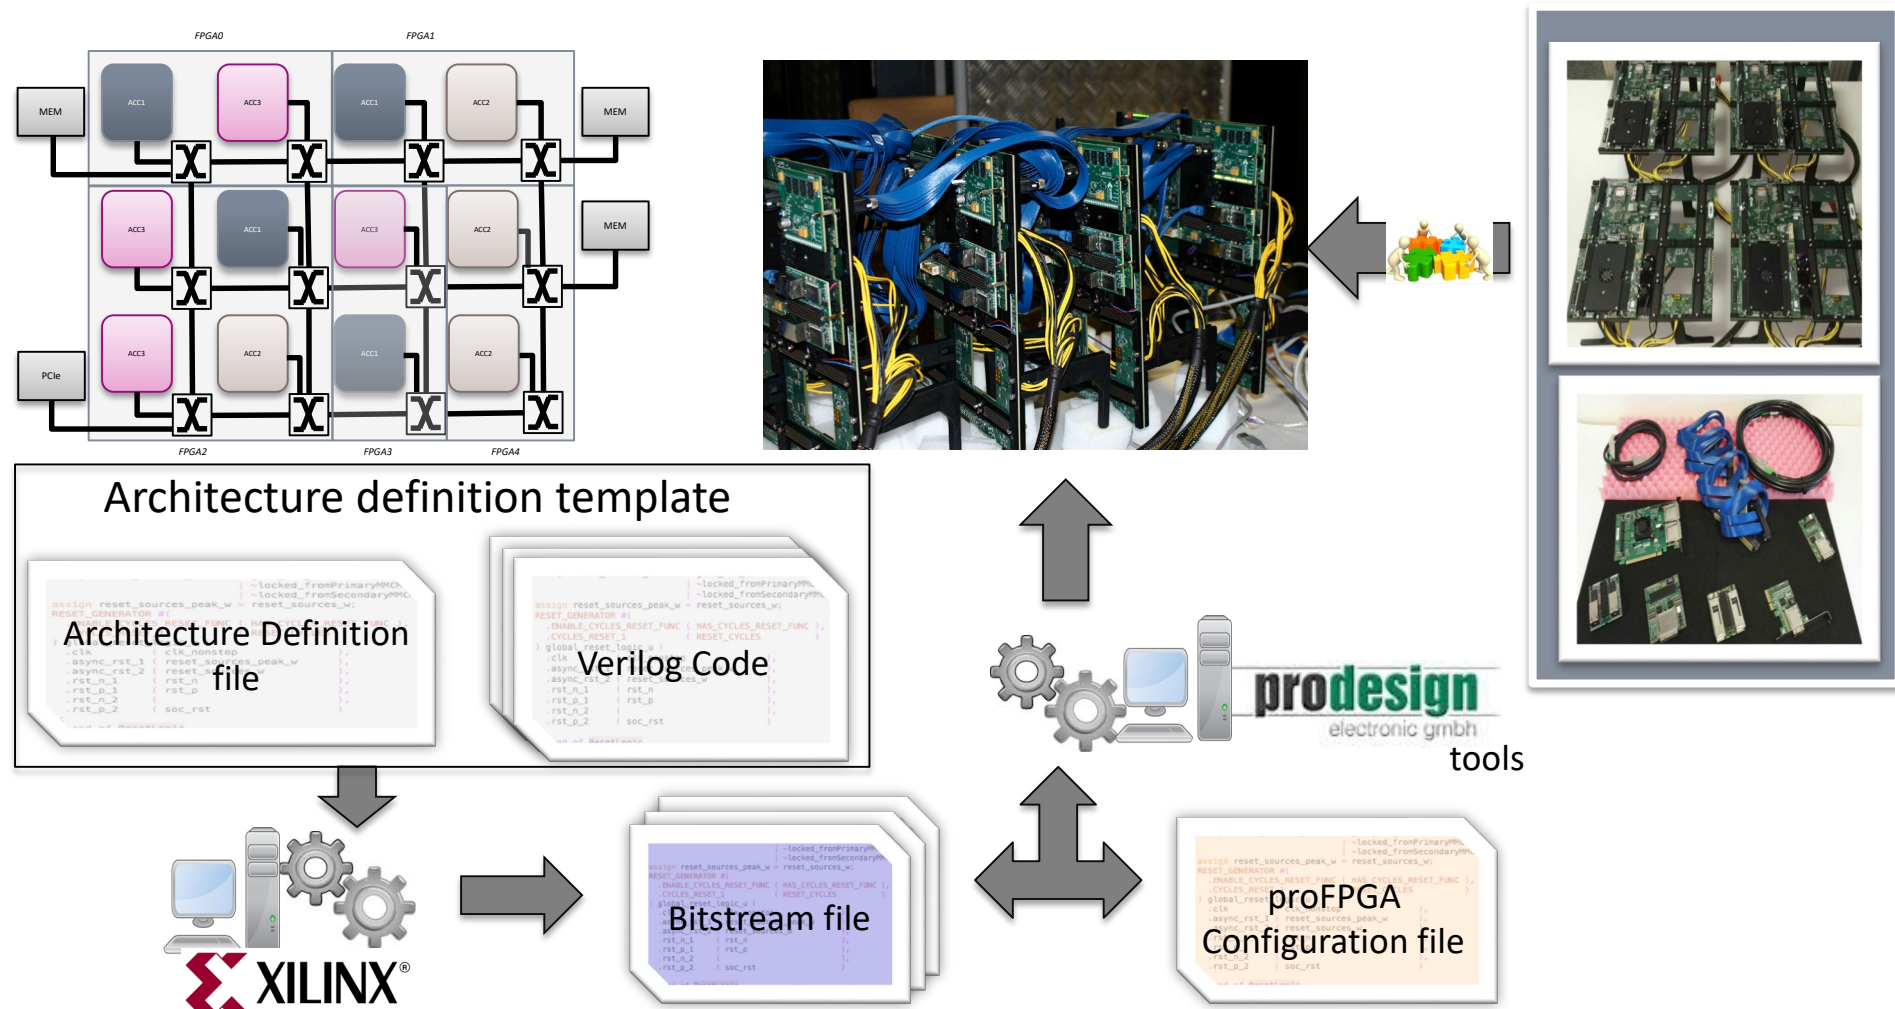


# Accelerator interface

- Decouple the UNIT from the rest of the MANGO platform
- Allow the implementation of a unique interface for every UNIT
- Unify memory access
  - Byte, Half (16 bits), word (32 bits) & block (512 bits) memory access types
- Allow to map synchronization registers in the virtual memory address space



# MANGO Design Flow



# Architecture Definition File

## ARCHITECTURE ID CONCEPT

- Sections
  - General
  - FPGA/multi-FPGA
  - Memory devices
  - I/O devices
  - Network

```
// * MULTI_FPGA_SECTION,
define N_FPGAS 3 // How many FPGAs are used for this architecture de.
define TILE_IDS_PER_FPGA_VECTOR /*FPGA_2*/10'd9,10'd8,/*FPGA_1*/10'd7,10'd6,10'd3,10'd2,/*FPGA_0*/10'd1
define TILE_TYPE_VECTOR 'TILE_NONE_TYPE','TILE_NONE_TYPE','TILE_NONE_TYPE','TILE_NONE_TYPE','TILE_N

define TILE_IDS_PER_FPGA_OFFSET_VECTOR 10'd80,10'd40,10'd0
define MULTI_FPGA
define USE_DAISSY_CHAIN
define DAISY_CHAIN_VECTOR 'MB0_A3_ID','MB0_A1_ID','MB0_C1_ID
define DAISY_CHAIN_DIR_FROM_FIRST_VECTOR 'NORTH_DIR','WEST_DIR','SOUTH_DIR
define DAISY_CHAIN_DIR_FROM_LAST_VECTOR 'NORTH_DIR','WEST_DIR','EAST_DIR
define MB0_A1_MB0_A3_PINES 10'd124 // number of pines available to route VN wires between two fpgas i
define MB0_A3_MB0_A1_PINES 10'd124 // number of pines available to route VN wires between two fpgas i
define MB0_A1_MB0_C1_PINES 10'd90 // number of pines available to route VN wires between two fpgas i
define MB0_C1_MB0_A1_PINES 10'd90 // number of pines available to route VN wires between two fpgas i
define CONNECTIVITY_PORTS 10'd4
define FPGA_CONNECTIVITY_OUT_VECTOR /*FPGA_2*/1'b1,1'b0,1'b0,1'b0,/*FPGA_1*/1'b0,1'b0,1'b1,1'b0,/*FPGA_0*/1'b0,1'b1,
define FPGA_CONNECTIVITY_IN_VECTOR /*FPGA_2*/1'b1,1'b0,1'b0,1'b0,/*FPGA_1*/1'b0,1'b0,1'b1,1'b0,/*FPGA_0*/1'b0,1'b1,
define FPGA_NEIGHBOURS_VECTOR /*FPGA_2*/'MB0_A1_ID,10'd0,10'd0,10'd0,/*FPGA_1*/10'd0,10'd0,'MB0_A1_ID,10'd0,
define FPGA_PINOUT_VECTOR /*FPGA_2*/'MB0_A3_MB0_A1_PINES,10'd0,10'd0,10'd0,/*FPGA_1*/10'd0,10'd0,'MB0_C1_
define FPGA_PININ_VECTOR /*FPGA_2*/'MB0_A1_MB0_A3_PINES,10'd0,10'd0,10'd0,/*FPGA_1*/10'd0,10'd0,'MB0_A1_

// * CLOCK_SECTION
define MMCM_CLKIN_FREQ 100.0 // Frequency of the input oscillator for the Primary MMCM
define ACCELERATOR_FREQ 5.0 // Accelerator frequency (Mhz)

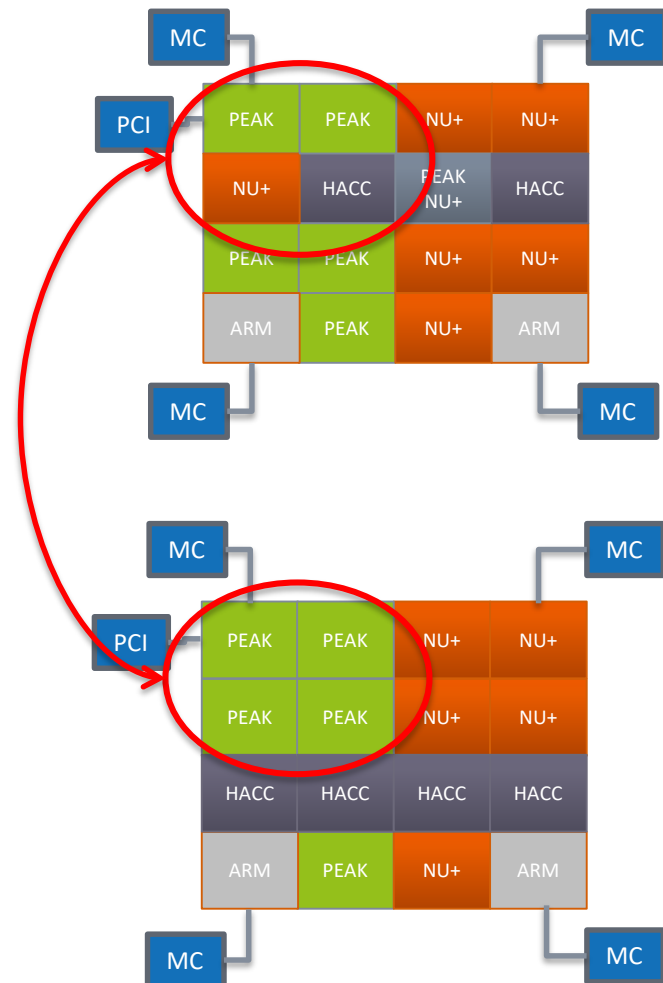
// * TOPOLOGY_SECTION
define MB0_A1_TOPOLOGY 'TOPOLOGY_2x2_1_4 // Mother board 0, FPGA A1 implements a 2x2_1_4 topology cc
define MB0_A3_TOPOLOGY 'TOPOLOGY_2x1_1_2 // Mother board 0, FPGA A3 implements a 2x1_1_2 topology cc
define MB0_C1_TOPOLOGY 'TOPOLOGY_2x2_1_4 // Mother board 0, FPGA C1 implements a 3x2_1_6 topology cc
define GLOBAL_TOPOLOGY 'TOPOLOGY_4x3_1_12
define TOPOLOGY_PER_FPGA_VECTOR 'GLOBAL_TOPOLOGY','MB0_A3_TOPOLOGY','MB0_C1_TOPOLOGY','MB0_A1_TOPOLOGY
define TOPOLOGY_PER_FPGA_VECTOR_w ('MESH_2D_w*4)
define TOPOLOGY_PER_FPGA_OFFSET_VECTOR 10'd120,10'd80,10'd40,10'd0
define TOPOLOGY_WIDTH_VECTOR 'MESH_2D_w','MESH_2D_w','MESH_2D_w','MESH_2D_w'
```



# Architecture Definition Template

## Key aspects

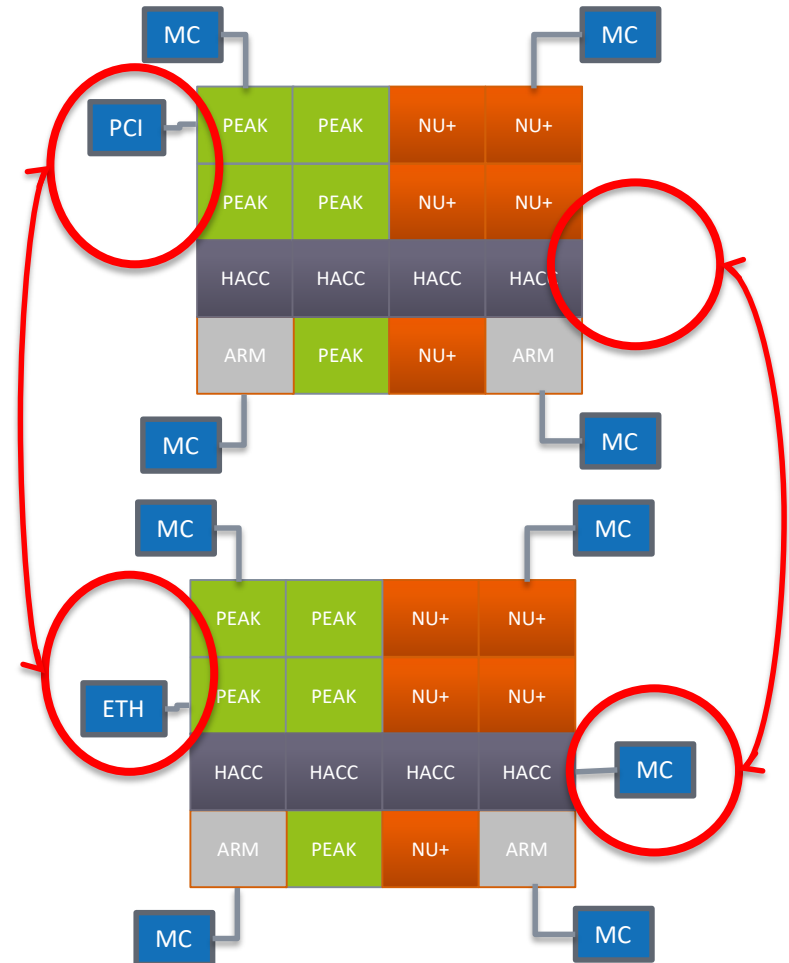
- Ease to configure the Unit type for every tile



# Architecture Definition Template

## Key aspects

- Ease to configure the Unit type for every tile
- Ease to attach Memory Controllers (MC) and I/O devices (PCIe, Ethernet) to any tile



# Architecture Definition Template

## Key aspects


- Ease to configure the Unit type for every tile
- Ease to attach Memory Controllers (MC) and I/O devices (PCIe, Ethernet) to any tile
- Ease to configure the number of tiles



# Architecture Definition Template

## Key aspects

- Ease to configure the Unit type for every tile
- Ease to attach Memory Controllers (MC) and I/O devices (PCIe, Ethernet) to any tile
- Ease to configure the number of tiles
- Ease to implement multi-FPGA designs

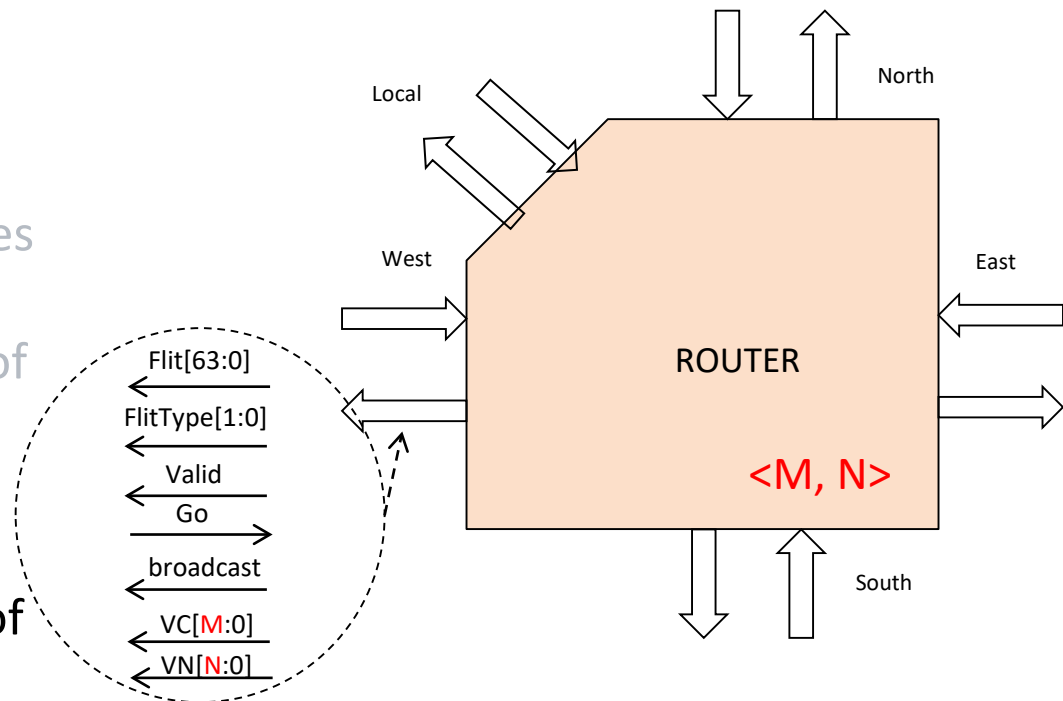




# Architecture Definition Template

## Key aspects

- Ease to configure the Unit type for every tile
- Ease to attach Memory Controllers (MC) and I/O devices (PCIe, Ethernet) to any tile
- Ease to configure the number of tiles
- Ease to implement multi-FPGA designs
- Easy to configure the number of virtual networks and virtual channels to achieve QoS guarantee



# Architecture Definition Template

## Key aspects

- Ease to configure the Unit type for every tile
- Ease to attach Memory Controllers (MC) and I/O devices (PCIe, Ethernet) to any tile
- Ease to configure the number of tiles
- Ease to implement multi-FPGA designs
- Easy to configure the number of virtual networks and virtual channels to achieve QoS guarantee

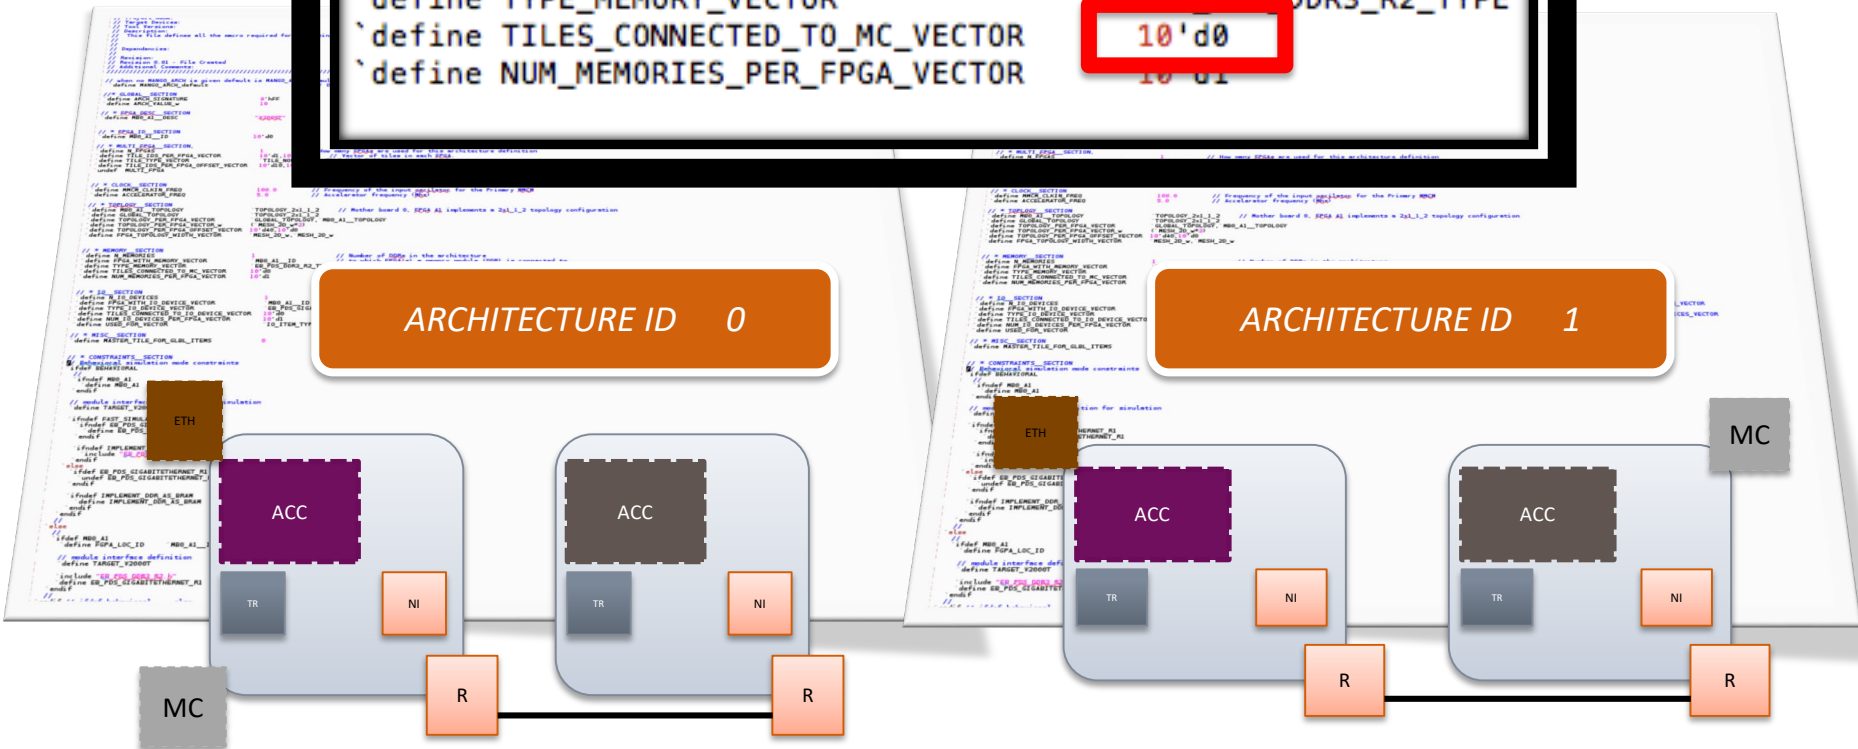


# Architecture Definition File: Example

```
// * MEMORY_SECTION
#define N_MEMORIES 1
#define FPGA_WITH_MEMORY_VECTOR MB0_A1_ID
#define TYPE_MEMORY_VECTOR FB_DDR3_DDR3_R2_TYPE
#define TILES_CONNECTED_TO_MC_VECTOR 10'd0
#define NUM_MEMORIES_PER_FPGA_VECTOR 10'd1
```

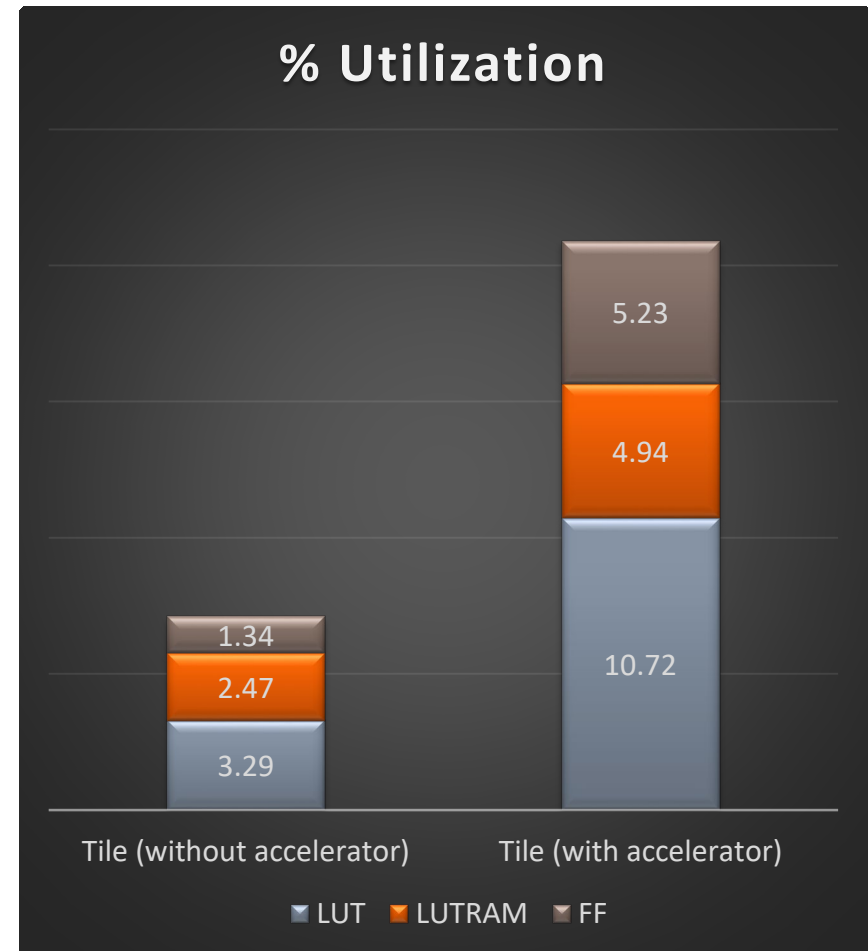
ARCHITECTURE ID 0

ARCHITECTURE ID 1



# FPGA Resource utilization

- Xilinx Virtex V2000T speedgrade -1
- Tile with Accelerator:
  - MIPS-based cache coherent 2-core accelerator
    - 32K L1I
    - 512K L1D
    - 1MB L2D



# Conclusions

- The MANGO approach for supporting the implementation of multiple accelerators on a multi-FPGA platform
  - High customization of the cluster
    - Flexibility for architecture exploration
  - Flexibility in configuring an architecture
    - Rapid architecture exploration
  - Effectivity in the system communications
    - QoS guarantee
  - Effectivity in monitoring the system
- Percentage of resources needed per tile quite slow



# Thank you for your attention

- Contact us at
  - [www.mango-project.eu](http://www.mango-project.eu)
- Other directly related EU project
  - [www.recipe-project.eu](http://www.recipe-project.eu)
- The MANGO project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 671668.
- The content of this presentation reflects only the authors' views and the European Commission is not responsible for any use that may be made of the information it contains.



# Architecture Definition File

## ○Sections

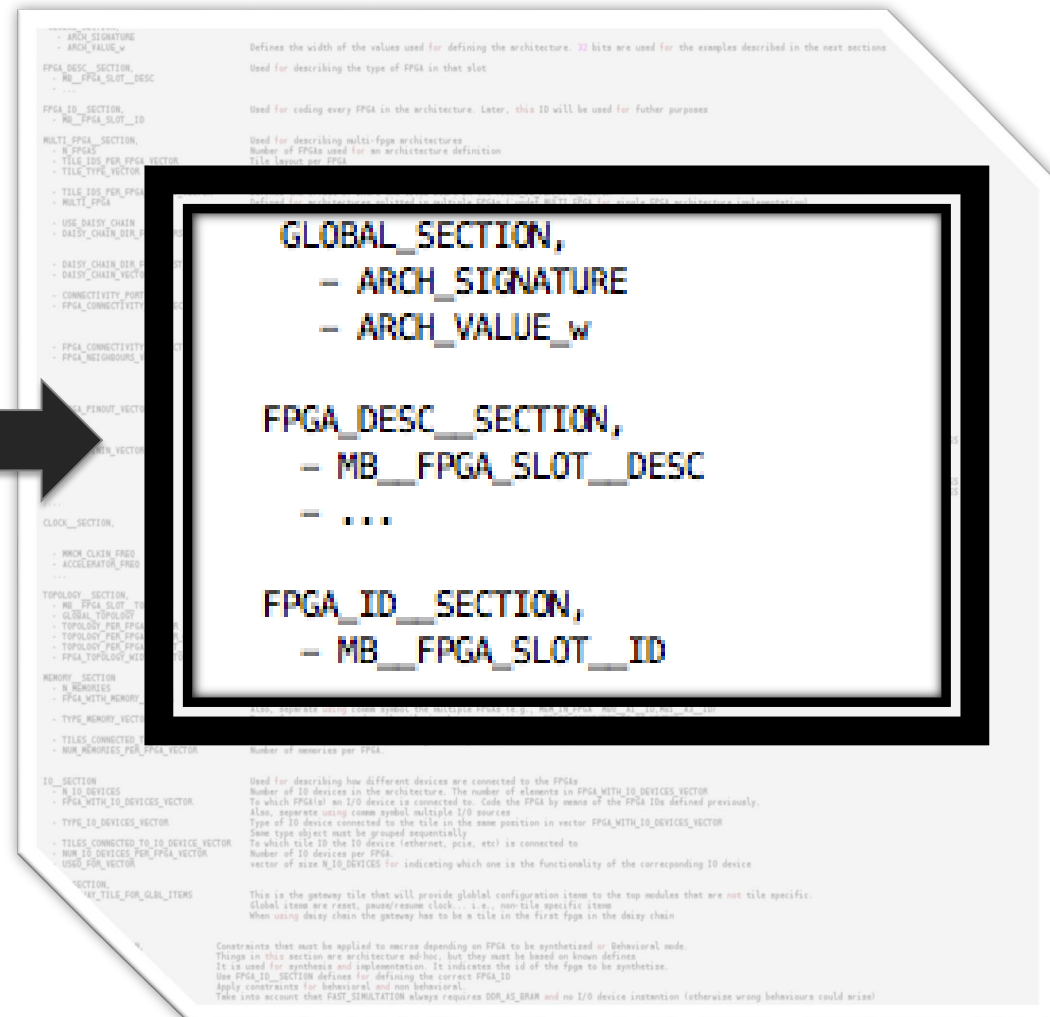
### ○General

### ○FPGA/multi-FPGA

### ○Memory devices

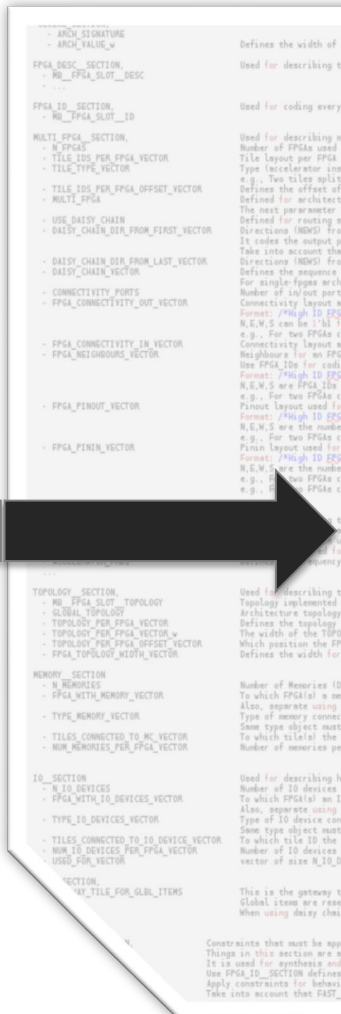
### ○I/O devices

### ○Network



# Architecture Definition

- Sections
  - General
  - **FPGA/multi-FPGA**
  - Memory devices
  - I/O devices
  - Network



```
MULTI_FPGA_SECTION,  
- N_FPGAS  
- TILE_IDS_PER_FPGA_VECTOR  
- TILE_TYPE_VECTOR  
  
- TILE_IDS_PER_FPGA_OFFSET_VECTOR  
- MULTI_FPGA  
  
- USE_DAIY_CHAIN  
- DAIY_CHAIN_DIR_FROM_FIRST_VECTOR  
  
- DAIY_CHAIN_DIR_FROM_LAST_VECTOR  
- DAIY_CHAIN_VECTOR  
  
- CONNECTIVITY_PORTS  
- FPGA_CONNECTIVITY_OUT_VECTOR  
  
- FPGA_CONNECTIVITY_IN_VECTOR  
- FPGA_NEIGHBOURS_VECTOR  
  
- FPGA_PINOUT_VECTOR  
  
- FPGA_PININ_VECTOR  
  
....  
  
CLOCK_SECTION,  
  
- MMCM_CLKIN_FREQ  
- ACCELERATOR_FREQ  
...  
  
TOPOLOGY_SECTION,  
- MB_FPGA_SLOT_TOPOLOGY  
- GLOBAL_TOPOLOGY  
- TOPOLOGY_PER_FPGA_VECTOR  
- TOPOLOGY_PER_FPGA_VECTOR_W  
- TOPOLOGY_PER_FPGA_OFFSET_VECTOR  
- FPGA_TOPOLOGY_WIDTH_VECTOR
```

# Architecture Definition File

- Sections
  - General
  - FPGA/multi-FPGA
  - **Memory devices**
  - I/O devices
  - Network

```

- ARCH_SIGNATURE
- ARCH_VALUE_w

FPGA_DESC_SECTION
- NO_FPGA_SLOT_DESC
...

FPGA_ID_SECTION
- NO_FPGA_SLOT_ID

MULTI_FPGA_SECTION
- N_FPGAS
- TILES_ID_PER_FPGA_VECTOR
- TILES_TYPE_VECTOR
- TILES_ID_PER_FPGA_OFFSET_VECTOR
- MULTI_FPGA
- USE_DAIY_CHAIN
- DAIY_CHAIN_DIR_FROM_FIRST_VECTOR
...
- DAIY_CHAIN_DIR_FROM_LAST_VECTOR
- DAIY_CHAIN_VECTOR
- CONNECTIVITY_PORTS
- FPGA_CONNECTIVITY_OUT_VECTOR

- FPGA_CONNECTIVITY_IN_VECTOR
- FPGA_NEIGHBOUR_VECTOR

- FPGA_PINOUT_VECTOR

- FPGA_PININ_VECTOR

...

CLOCK_SECTION
- WNOCK_CLKIN_FREQ
- ACCELERATOR_FREQ
...

TOPOLOGY_SECTION
- ...

MEMORY_SECTION
- N_MEMORIES
- FPGA_WITH_MEMORY_VECTOR
- TYPE_MEMORY_VECTOR
- TILES_CONNECTED_TO_MC_VECTOR
- NUM_MEMORIES_PER_FPGA_VECTOR

I/O_SECTION
- N_I/O_DEVICES
- FPGA_WITH_I/O_DEVICES_VECTOR
- TYPE_I/O_DEVICES_VECTOR
- TILES_CONNECTED_TO_I/O_DEVICES
- NUM_I/O_DEVICES_PER_FPGA_VECTOR
- USAGE_FOR_VECTOR

...
                    
```

Defines the width of the values used for defining the architecture. 32 bits are used for the examples described in the next sections

Used for describing the type of FPGA in that slot

Used for coding every FPGA in the architecture. Later, this ID will be used for further purposes

Used for describing multi-fpga architectures

Number of FPGAs used for an architecture definition

Tile layout per FPGA

Type (accelerator/inserter) of each tile

e.g., Two tiles splitted in two FPGAs with one tile per FPGA /FPGA\_0/32'd0, /FPGA\_0/32'd0 // Vector of tiles in each FPGA

Defines the offset of where the tiles start in the TILES\_ID\_PER\_FPGA\_VECTOR

Defined for architectures splitted in multiple FPGAs (under MULTI\_FPGA for single FPGA architecture implementation)

The next parameter in this section are required for MULTI\_FPGA designs only.

Defined for routing specific (reset, congested, etc.) signals among FPGA in multi-fpga designs using daisy chain

Directions (INDXS) from the first fpga in daisy chain point of view

It codes the output ports (INDXS) that each FPGA provides for connecting to the next one in the daisy chain vector

Take into account that the indexes of this vector are relative to the FPGA\_IDs and not to the order of the DAIY\_CHAIN\_VECTOR

Directions (INDXS) from the last fpga in daisy chain point of view. Generate this vector as the previous one

Defines the sequence tiles vector in the daisy chain for specific signals connection (from right to left)

For single-fpga architectures the next parameters of this section are not required

Number of input ports for FPGA connections

Connectivity layout among FPGAs for output ports

Format: /High ID (FPGA) N.E.W.S. ... /Low ID (FPGA) N.E.W.S

N.E.W.S can be 1'b1 for indicating connectivity across the stated direction. Use 1'b0 otherwise

e.g., For two FPGAs connected using N-to-S Links: /FPGA\_1/1'b1,1'b0,1'b0,1'b0, /FPGA\_0/1'b0,1'b0,1'b0,1'b1

Connectivity layout among FPGAs for inputs ports. Same format as FPGA\_CONNECTIVITY\_OUT\_VECTOR

Neighbour for an FPGA. It is a full matrix vector

Use FPGA\_IDs for coding the vector

Format: /High ID (FPGA) N.E.W.S. ... /Low ID (FPGA) N.E.W.S

N.E.W.S are FPGA\_IDs or 32'd0 if no neighbour in the given direction

e.g., For two FPGAs connected using N-to-S Links: /FPGA\_1/ NO\_A1\_ID,32'd0,32'd0,32'd0, /FPGA\_0/32'd0,32'd0,32'd0, NO\_A3\_ID

Pinout layout used for connecting different FPGAs

Format: /High ID (FPGA) N.E.W.S. ... /Low ID (FPGA) N.E.W.S

N.E.W.S are the number of output pines required or 32'd0 if no pines required in the given direction

e.g., For two FPGAs connected using N-to-S Links: /FPGA\_1/ NO\_A3\_NUM\_A3\_PINES,32'd0,32'd0,32'd0, /FPGA\_0/32'd0,32'd0,32'd0, NO\_A3\_NUM\_A3\_PINES

Pinin layout used for connecting different FPGAs

Format: /High ID (FPGA) N.E.W.S. ... /Low ID (FPGA) N.E.W.S

N.E.W.S are the number of input pines required or 32'd0 if no pines required in the given direction

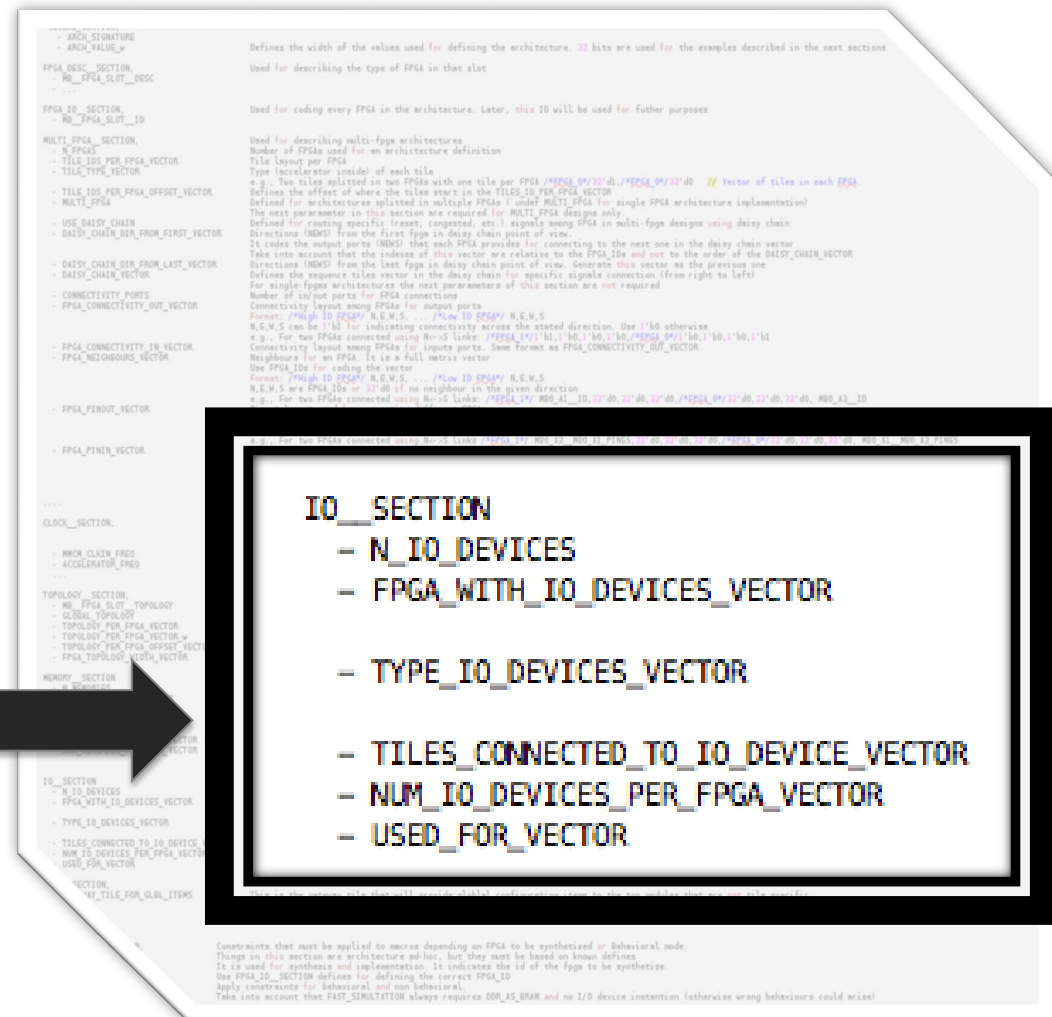
**MEMORY\_SECTION**

- N\_MEMORIES
- FPGA\_WITH\_MEMORY\_VECTOR
- TYPE\_MEMORY\_VECTOR
- TILES\_CONNECTED\_TO\_MC\_VECTOR
- NUM\_MEMORIES\_PER\_FPGA\_VECTOR

The FPGA\_ID\_SECTION defines for defining the correct FPGA\_ID apply constraints for behavioral and non behavioral. Take into account that FAST\_SIMULATION always requires DOR\_A5\_BRAR and no I/O device instantiation (otherwise wrong behaviours could arise)

# Architecture Definition File

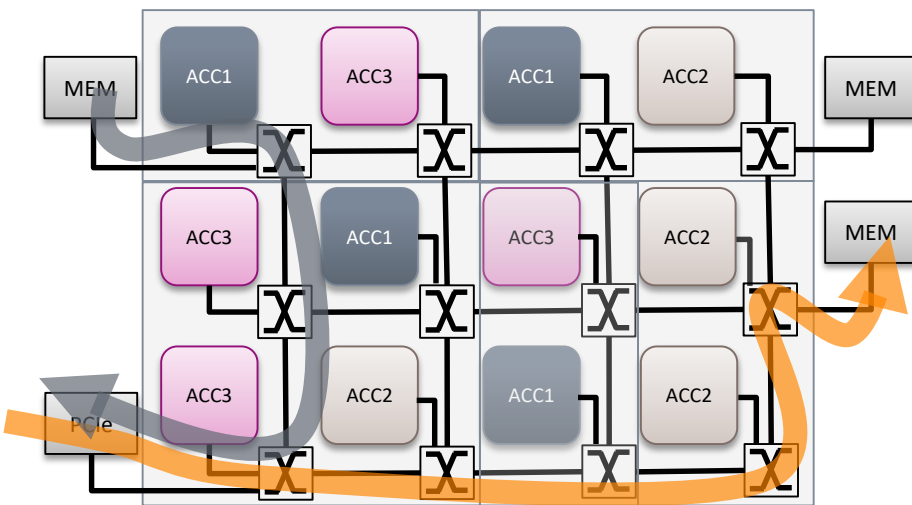
- Sections
  - General
  - FPGA/multi-FPGA
  - Memory devices
  - I/O devices
  - Network






# DMA Transfers

- DMA controller attached to every memory controller.
- Programming parameters: base address, source tile, size of the transfer, destination of the transfer, target address.
- Possibilities:







# Data Network

- DATA Routers connected among them in a 2D mesh layout.
- NI decouples network from tile components.
- Supported different Virtual networks (VN) with different number of Virtual channels (VC).
- Supported dynamic assignment of bandwidth per VN.

