Reconfigurable Virtual Instrumentation

From Advanced Instrumentation Towards Supercomputing

Andres Cicuttin

ICTP – MLAB

Multidisciplinary Laboratory of The Abdus Salam International Centre for Theoretical Physics

Trieste, Italy

Outline

1. Supercomputing and Custom Computing
   - Definitions
   - Time Computation vs. Space Computation
   - Problems and different approaches
2. Scientific Instrumentation based on FPGA
   - Based on Single FPGA (RVI and SoC FPGA)
   - Based on Multiple FPGAs (Distributed and massively parallel)
3. Abstract model for reconfigurable systems
   - Three-Dimensional extension of FPGA (hyperFPGA)
   - Extended Memory mapping
   - Universal Direct Memory Access (UDMA) Instructions
   - Architecture and Implementation
   - Physical and logical topology of clusters
   - Data packets and routing
4. Opportunities for open collaboration
   - Experimental hardware platforms
   - Software support, Operating systems
   - Brief description of ICTP and its main programs

Supercomputing

Reconfigurable Computing	Custom Computing

The *reconfigurable* hardware infrastructure for *custom* supercomputing should ideally be:

1) ***Versatile***

Must allow the implementation of many different computing architectures and strategies

2) ***Homogeneous***

Any logical subsystem should behave in the same way independently of where it is implemented

3) ***Scalable***

It should be possible to be implemented at different sizes preserving its basic logic and physical structure. It should also be conceived to be compatible with different types of FPGA within a wide range of cost-performance trade-offs

4) ***Efficient***

Must achieve a large number of arithmetic/logic operations per units of time, money and energy.

5) ***Portable***

Must be, as much as possible, FPGA vendor independent

6) ***Updateable***

Can be updated with newer devices without changing the basic structure and preserving as much as possible code compatibility

7) ***Upgradable***

Can be easily upgraded by adding more RAM or storage memory, or by replacing the main devices with more powerful ones

The Custom Computing Problem

- Which is the best reconfigurable hardware infrastructure?
- Which language should be used to capture a computational problem and express its solution?
- Which tools should be developed to configure the hardware to implement the best custom computer?
- Which tools should be developed to compile the code for its efficient execution in the configured custom computer?

None of these questions can be separately solved

It needs solid experimental knowledge and multidisciplinary contribution

Two Main Computational Paradigms

Scarcity of area & low circuit integration =>

***The uProcessor paradigm**:*

- Intensive reutilization of limited HW resources

- Computation along time (time computation)

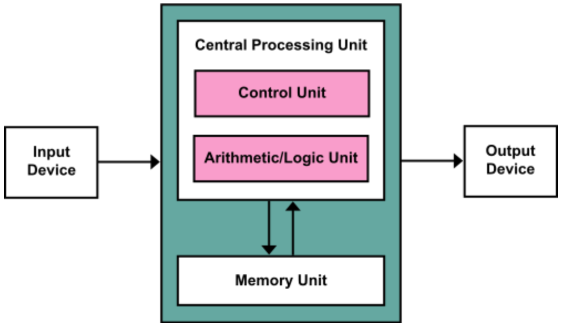

Abundance of area & high circuit integration =>

***The FPGA paradigm**:*

- Allocation of HW resources as needed

- Computation along space (space computation)

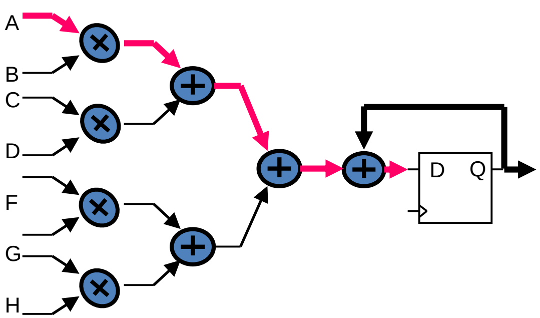

Desirable features of Advanced Instrumentation

	Scientific	Industrial	Commercial	Academic	Military
Performance	max				max
Accuracy, Precision	max	high			high
Reconfigurability	high		sometimes		
Massively parallel	sometimes	sometimes			sometimes
Physically Distributed	sometimes				sometimes
Cost			low	low	
Design time	sometimes		low	low	
Reliability		high			high

Advanced Instrumentation based on FPGA

Reconfigurable Virtual Instrumentation based on FPGA and SoC FPGA

Massively parallel and distributed instrumentation in large high energy physics experiments (Multiple units)

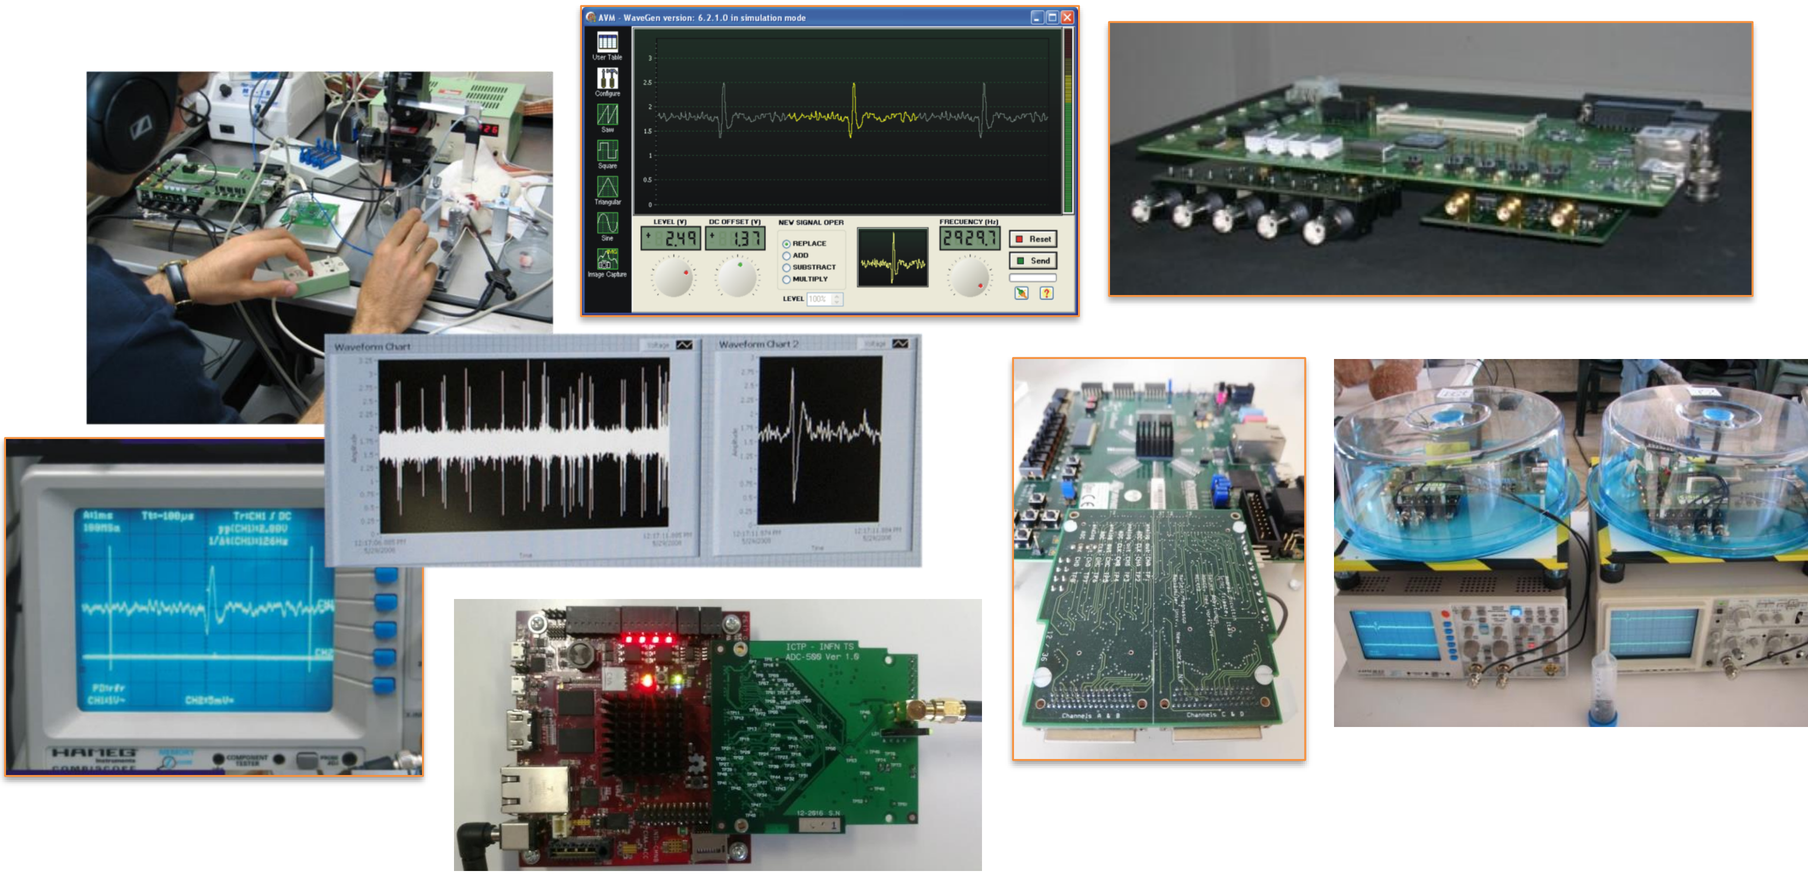

Reconfigurable Virtual Instrumentation

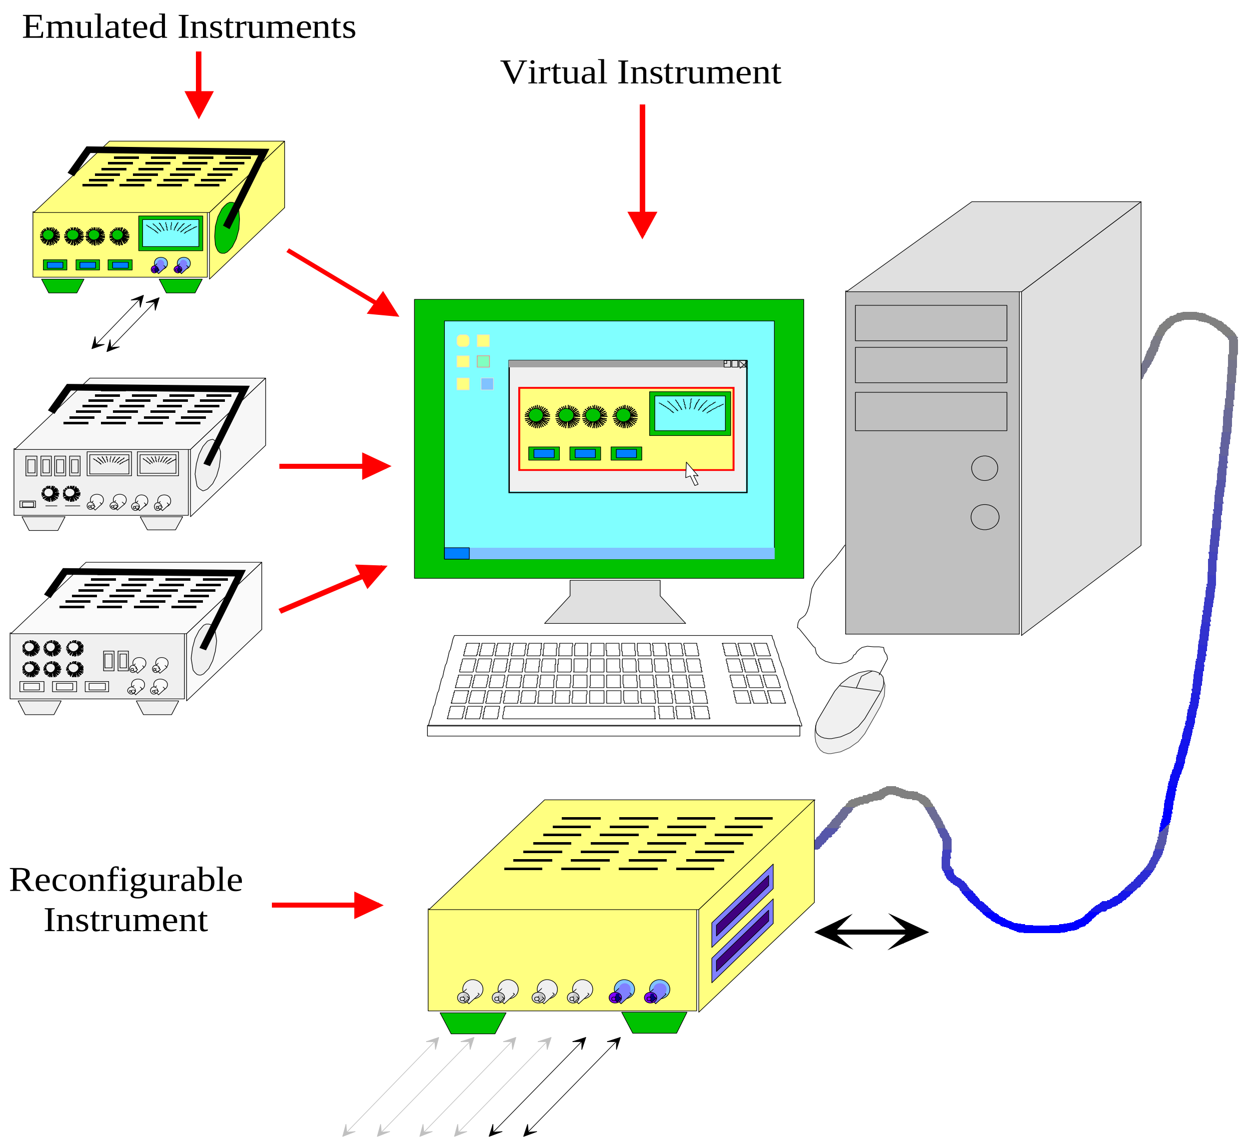

Reconfigurable Virtual Instrumentation

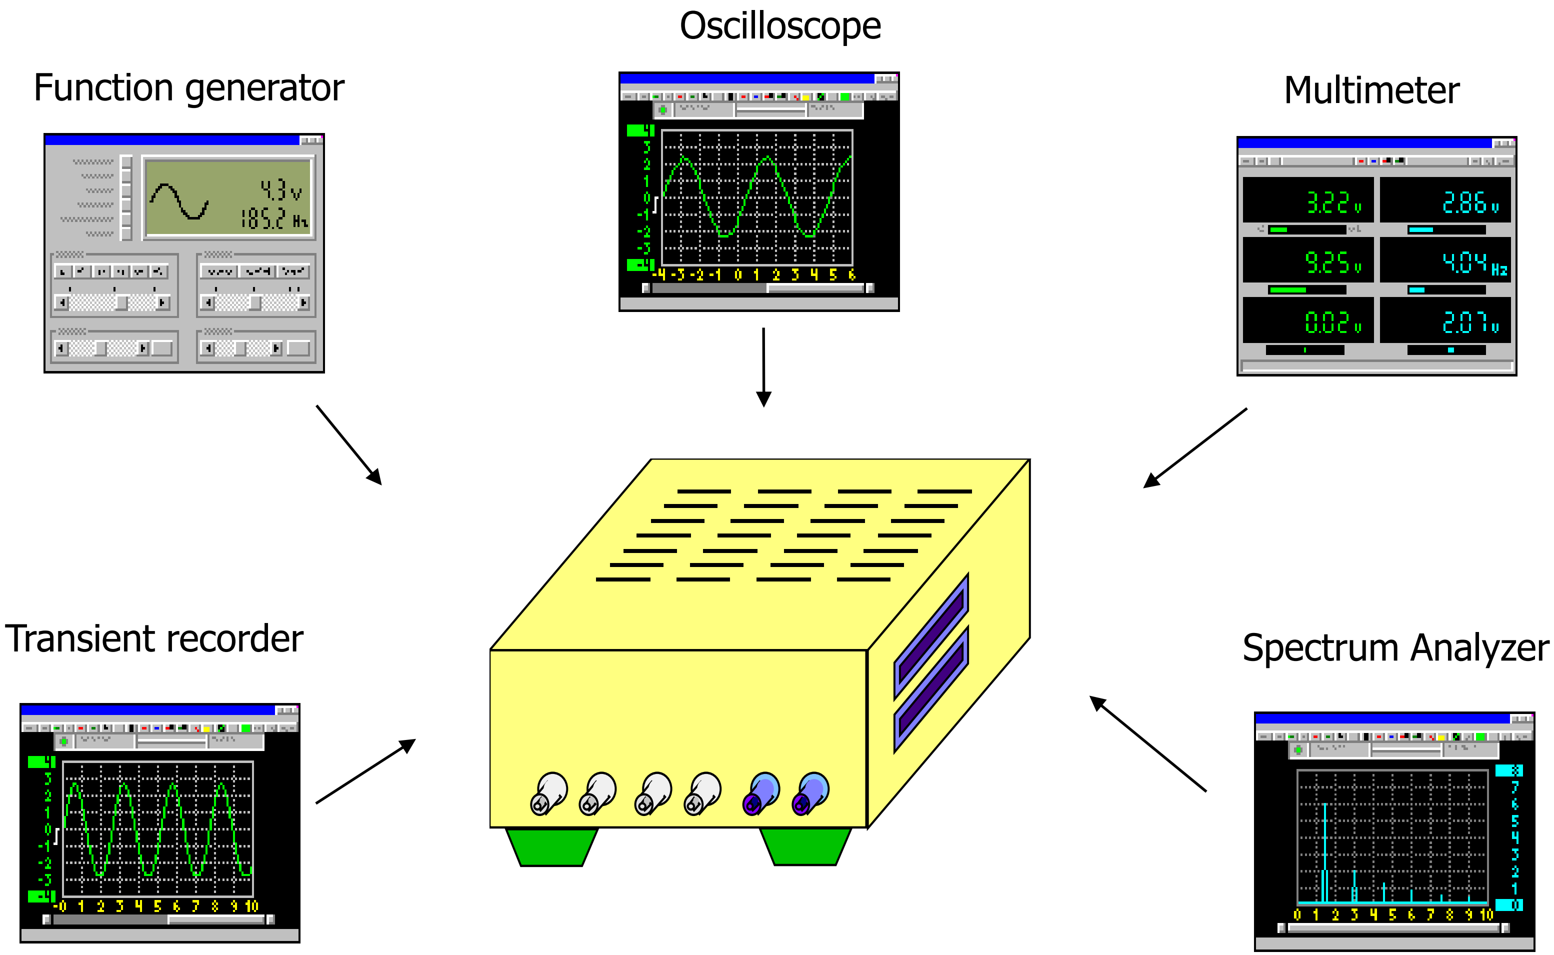

Reconfigurable **V**irtual **I**nstrumentation based on FPGA Global Architecture

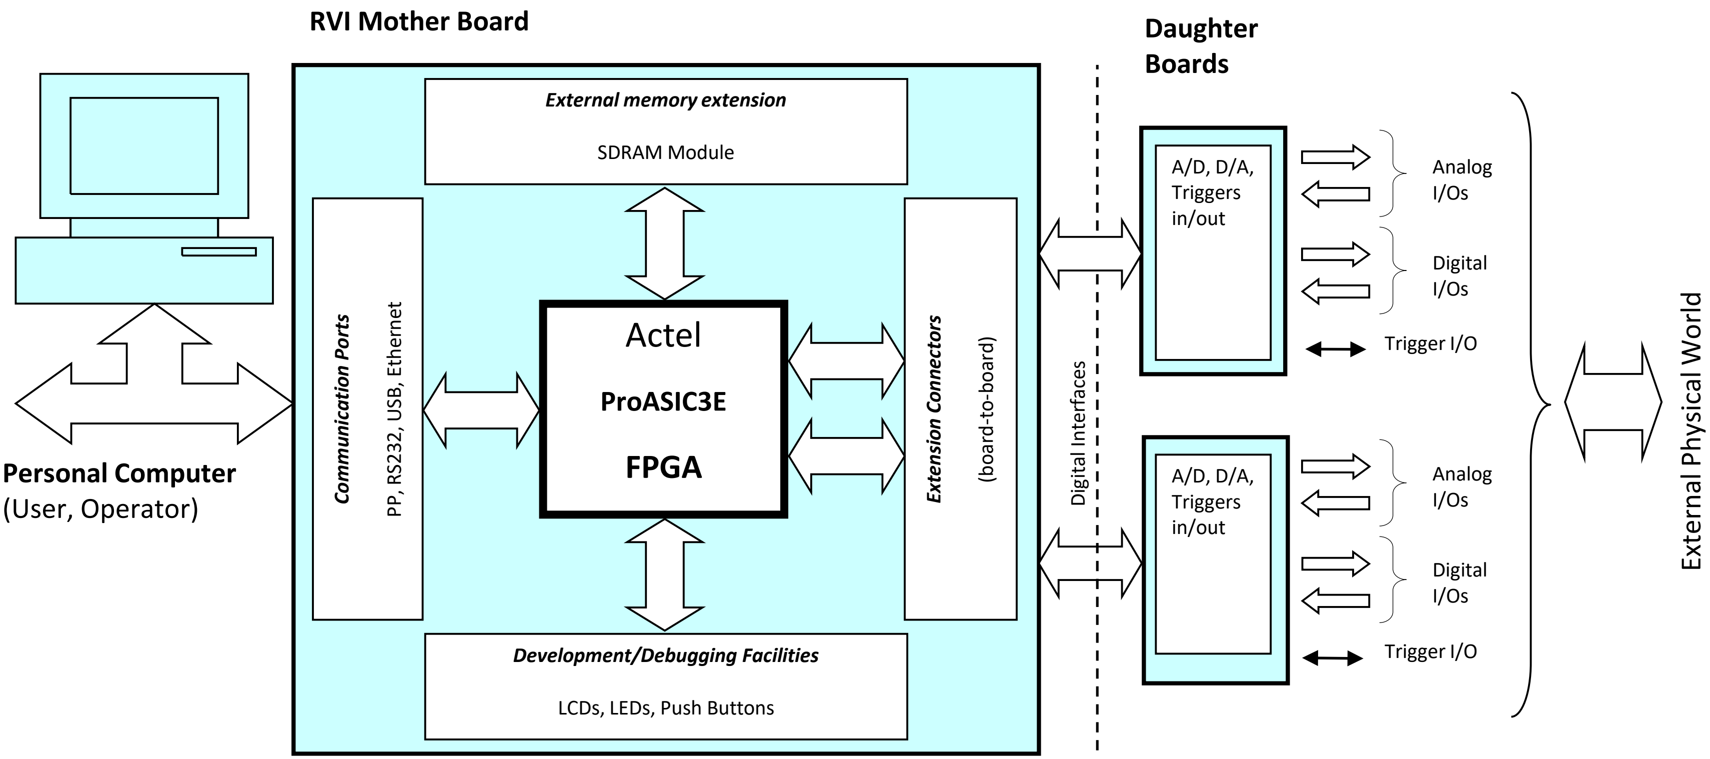

Reconfigurable Instrumentation: Architectural approach and modular structure

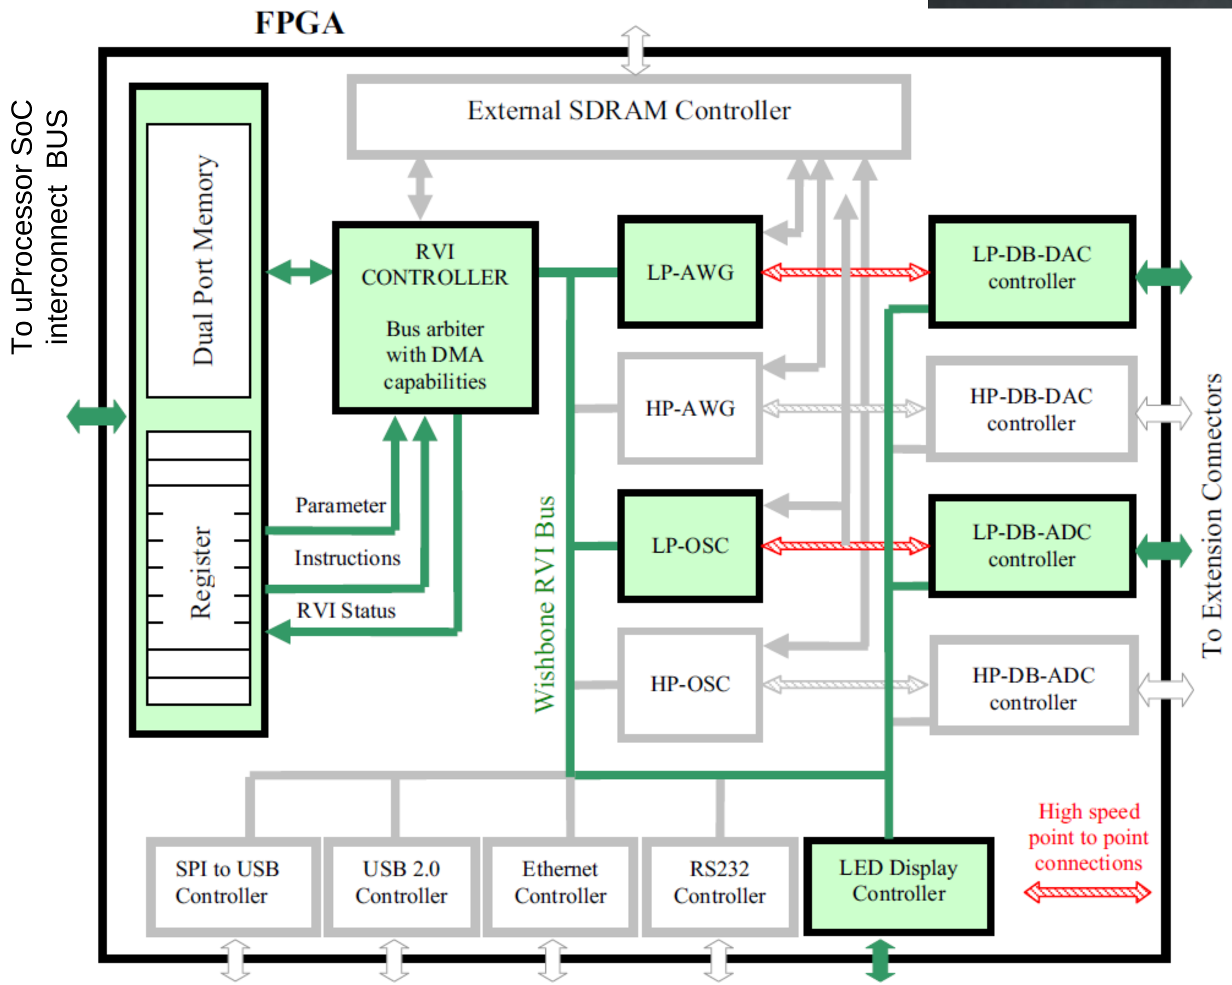

Reconfigurable Virtual Instrumentation based on SoC FPGA Global Architecture

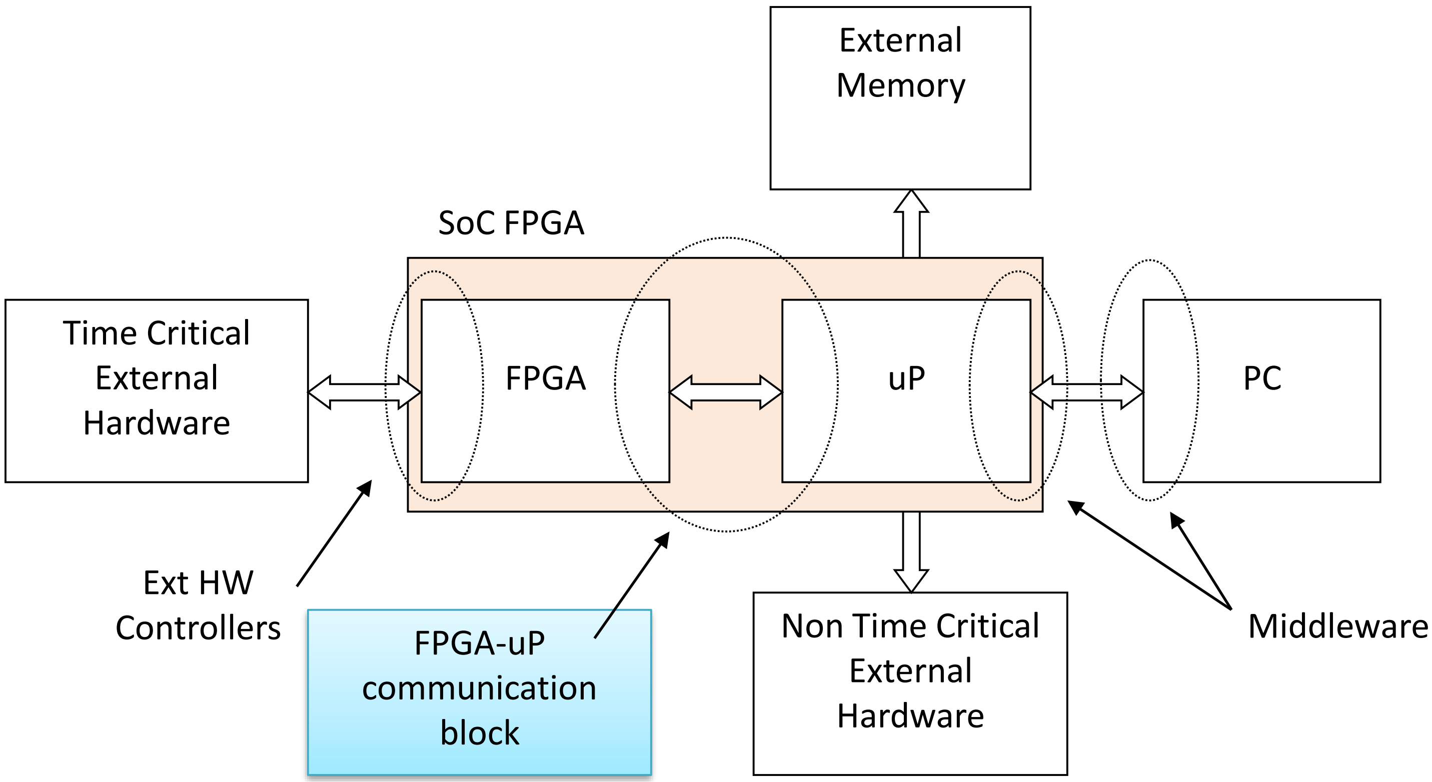

SoC FPGA Based Reconfigurable Virtual Instrumentation Typical Global Architecture

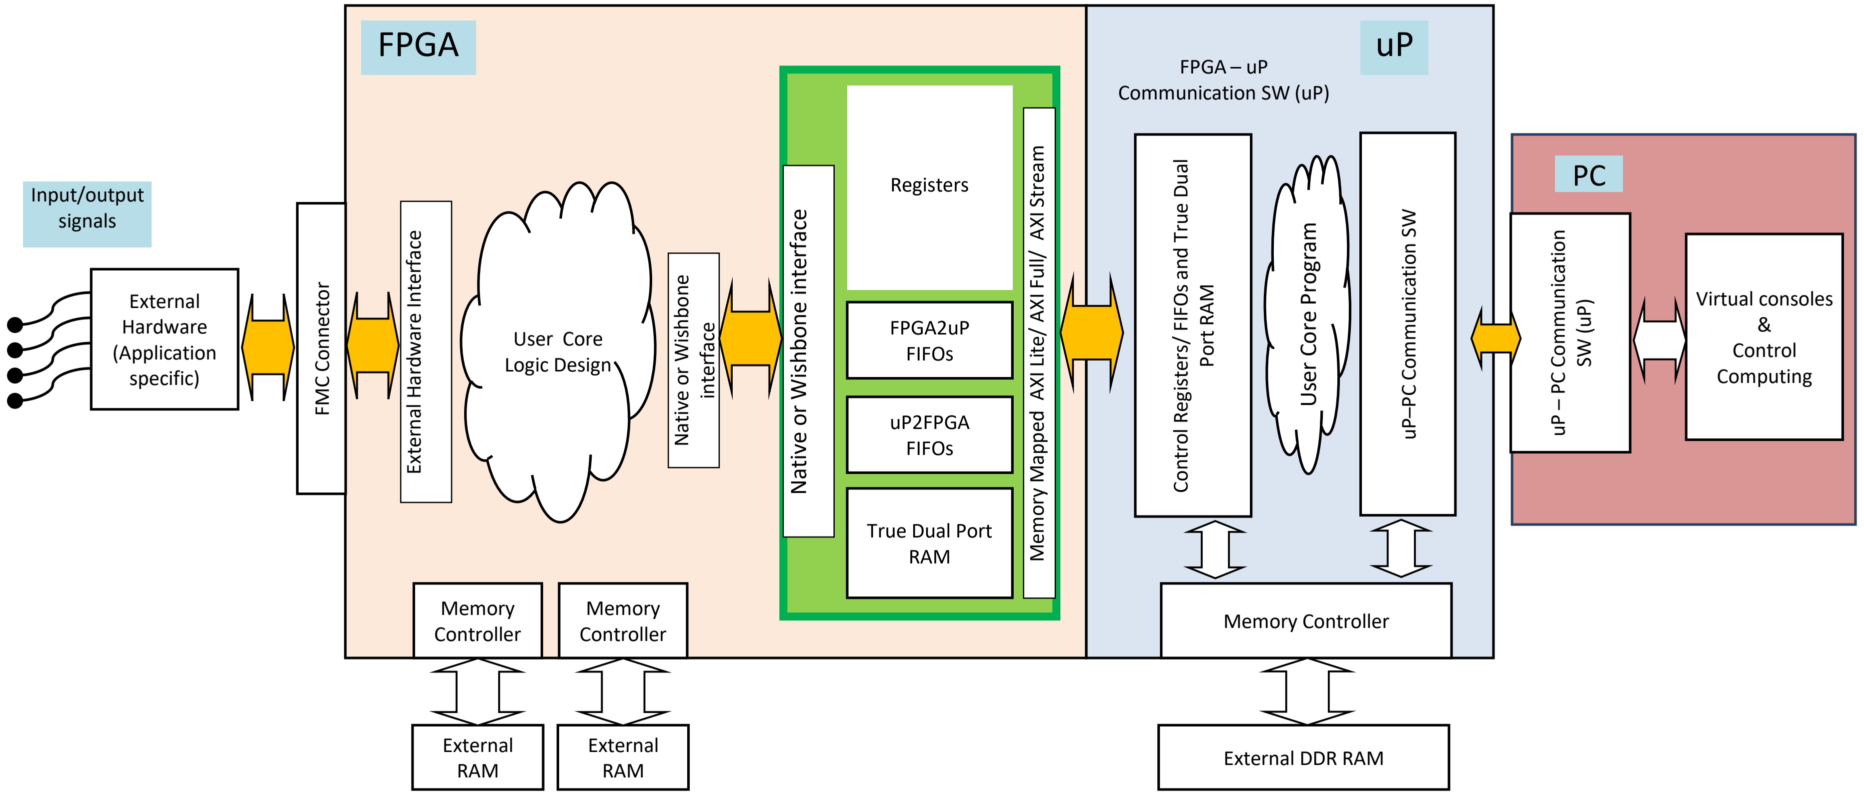

Advanced Instrumentation based on FPGA

Reconfigurable Virtual
 Instrumentation based on
FPGA and SoC FPGA

Massively parallel and distributed
instrumentation in large high-energy
physics experiments

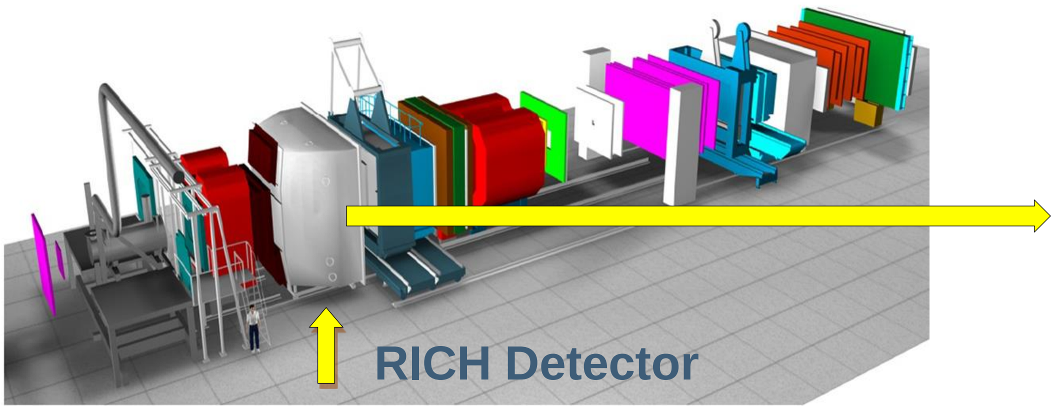

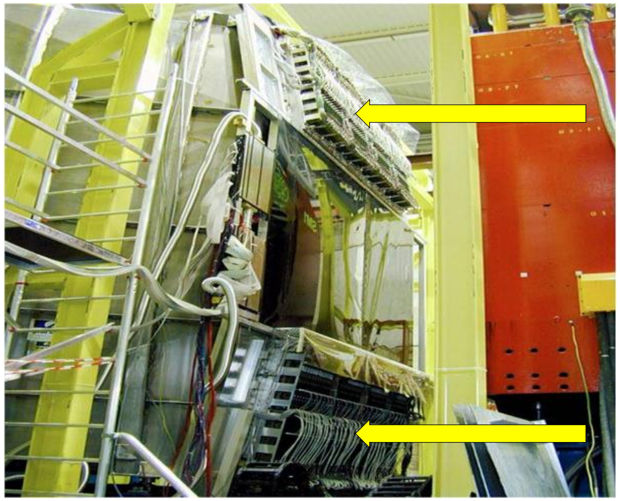

Artistic view of the 60 m long
COMPASS two-stage spectrometer.
The large gray box is the RICH-1
detector.
Approximate size:: 4 m x 4 m x 2 m

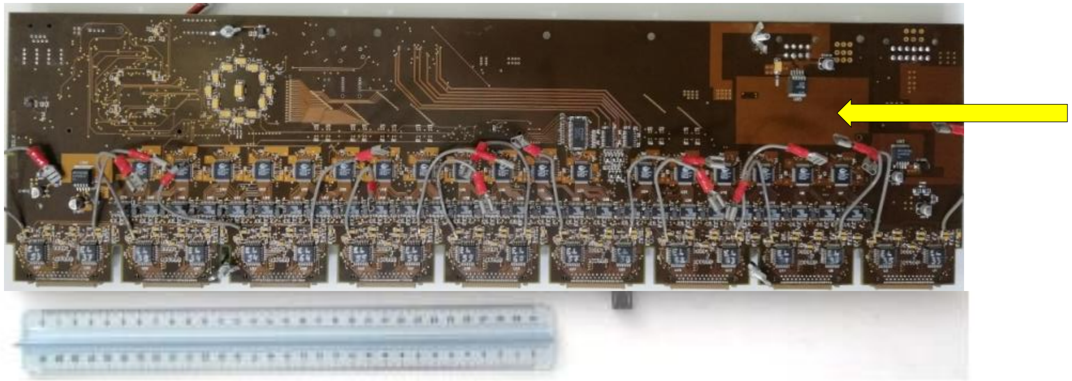

Global Architecture

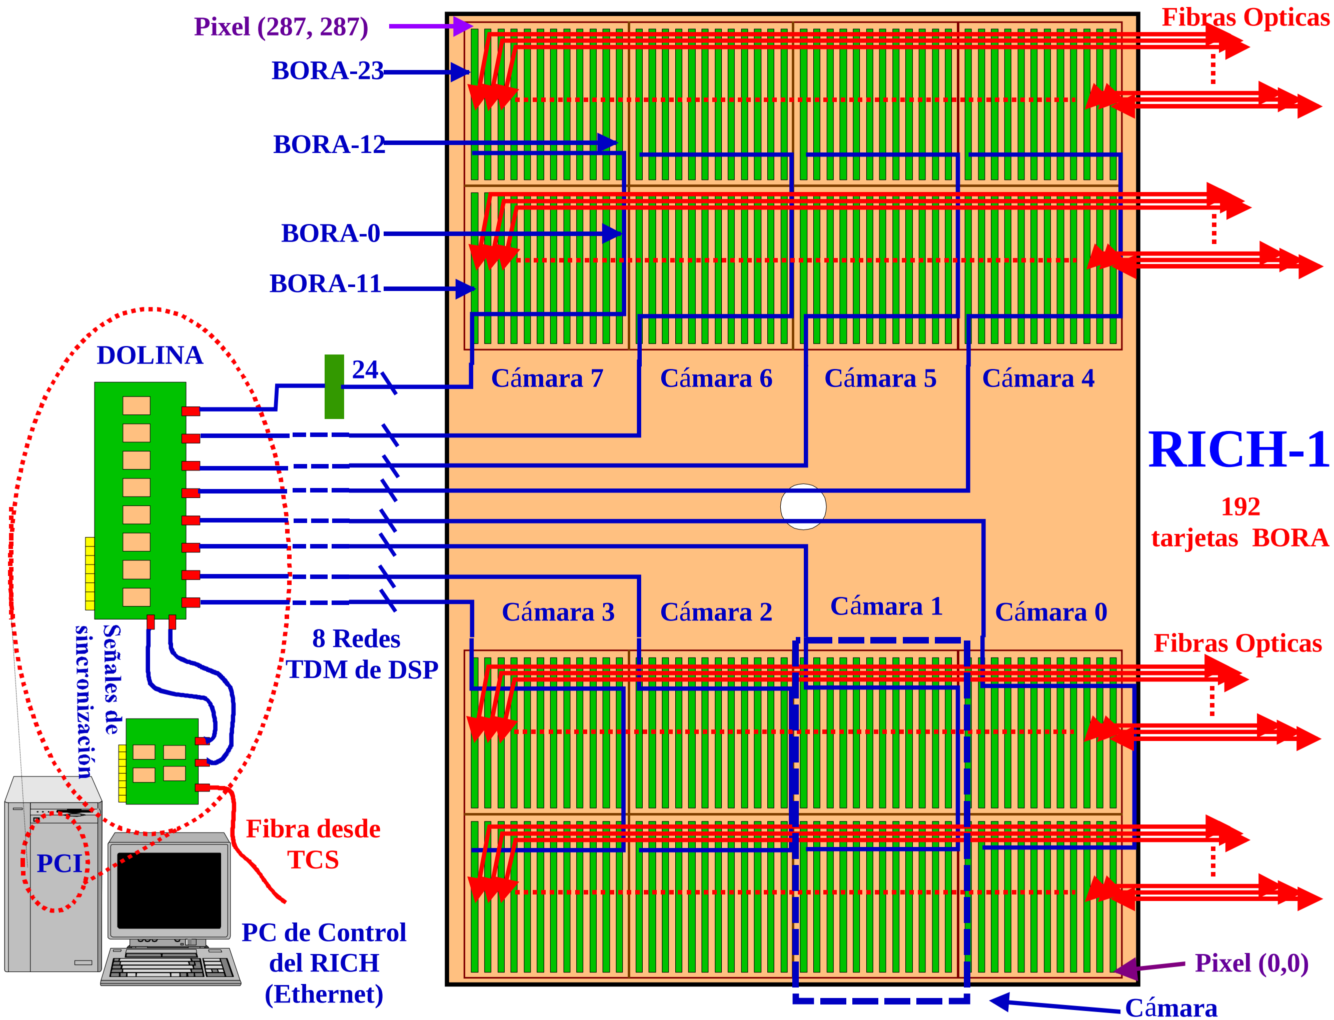

Advanced Instrumentation based on FPGA

Reconfigurable Virtual Instrumentation based on FPGA and SoC FPGA

Massively parallel and distributed instrumentation in large high-energy physics experiments

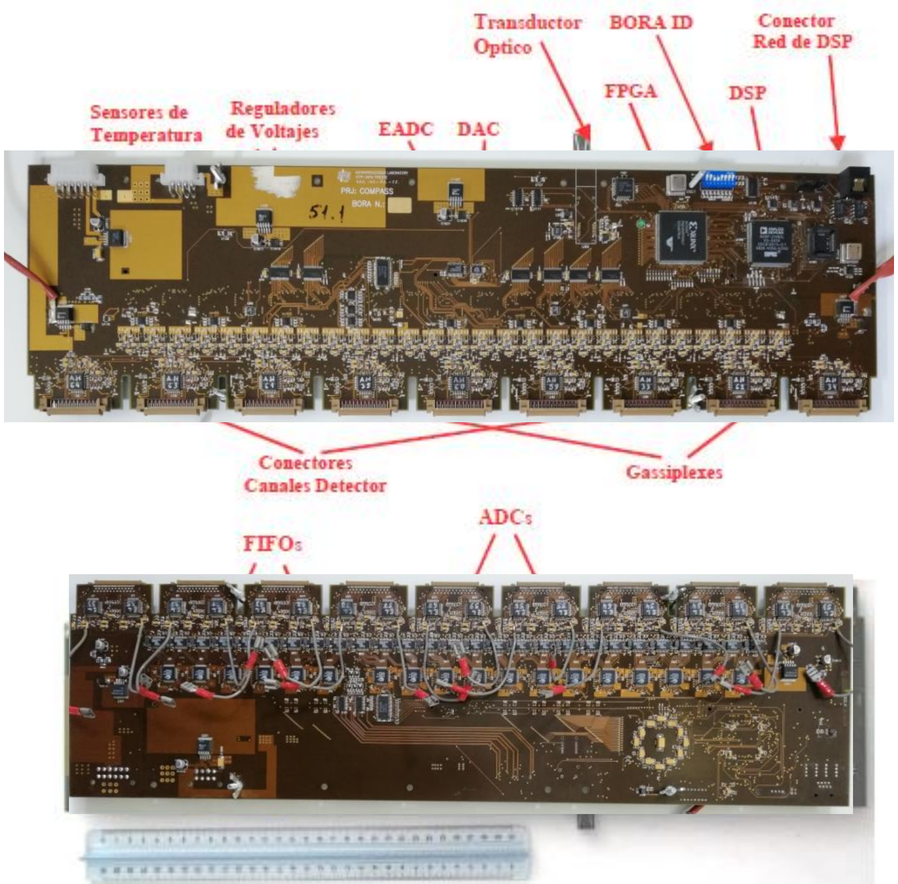

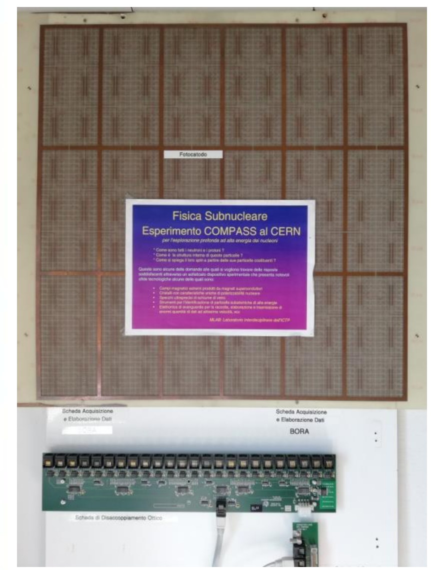

Reconfigurable Virtual Instrumentation based on FPGA

Reconfigurable Virtual Instrumentation based on FPGA and SoC FPGA

Massively parallel and distributed instrumentation in large high-energy physics experiments

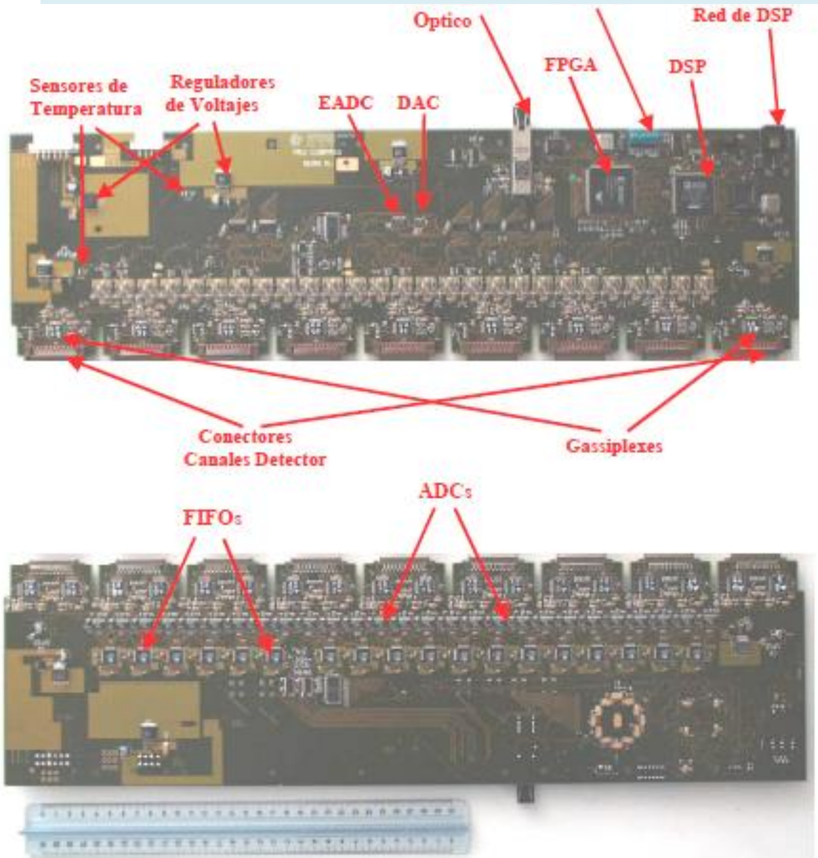

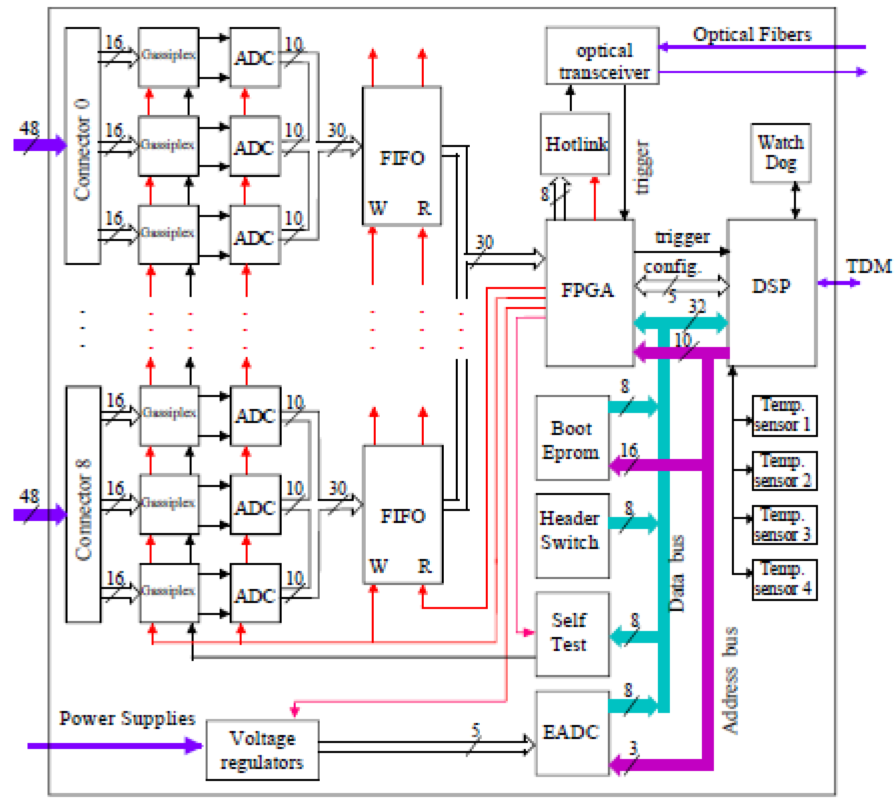

Data movement through distributed instrumentation

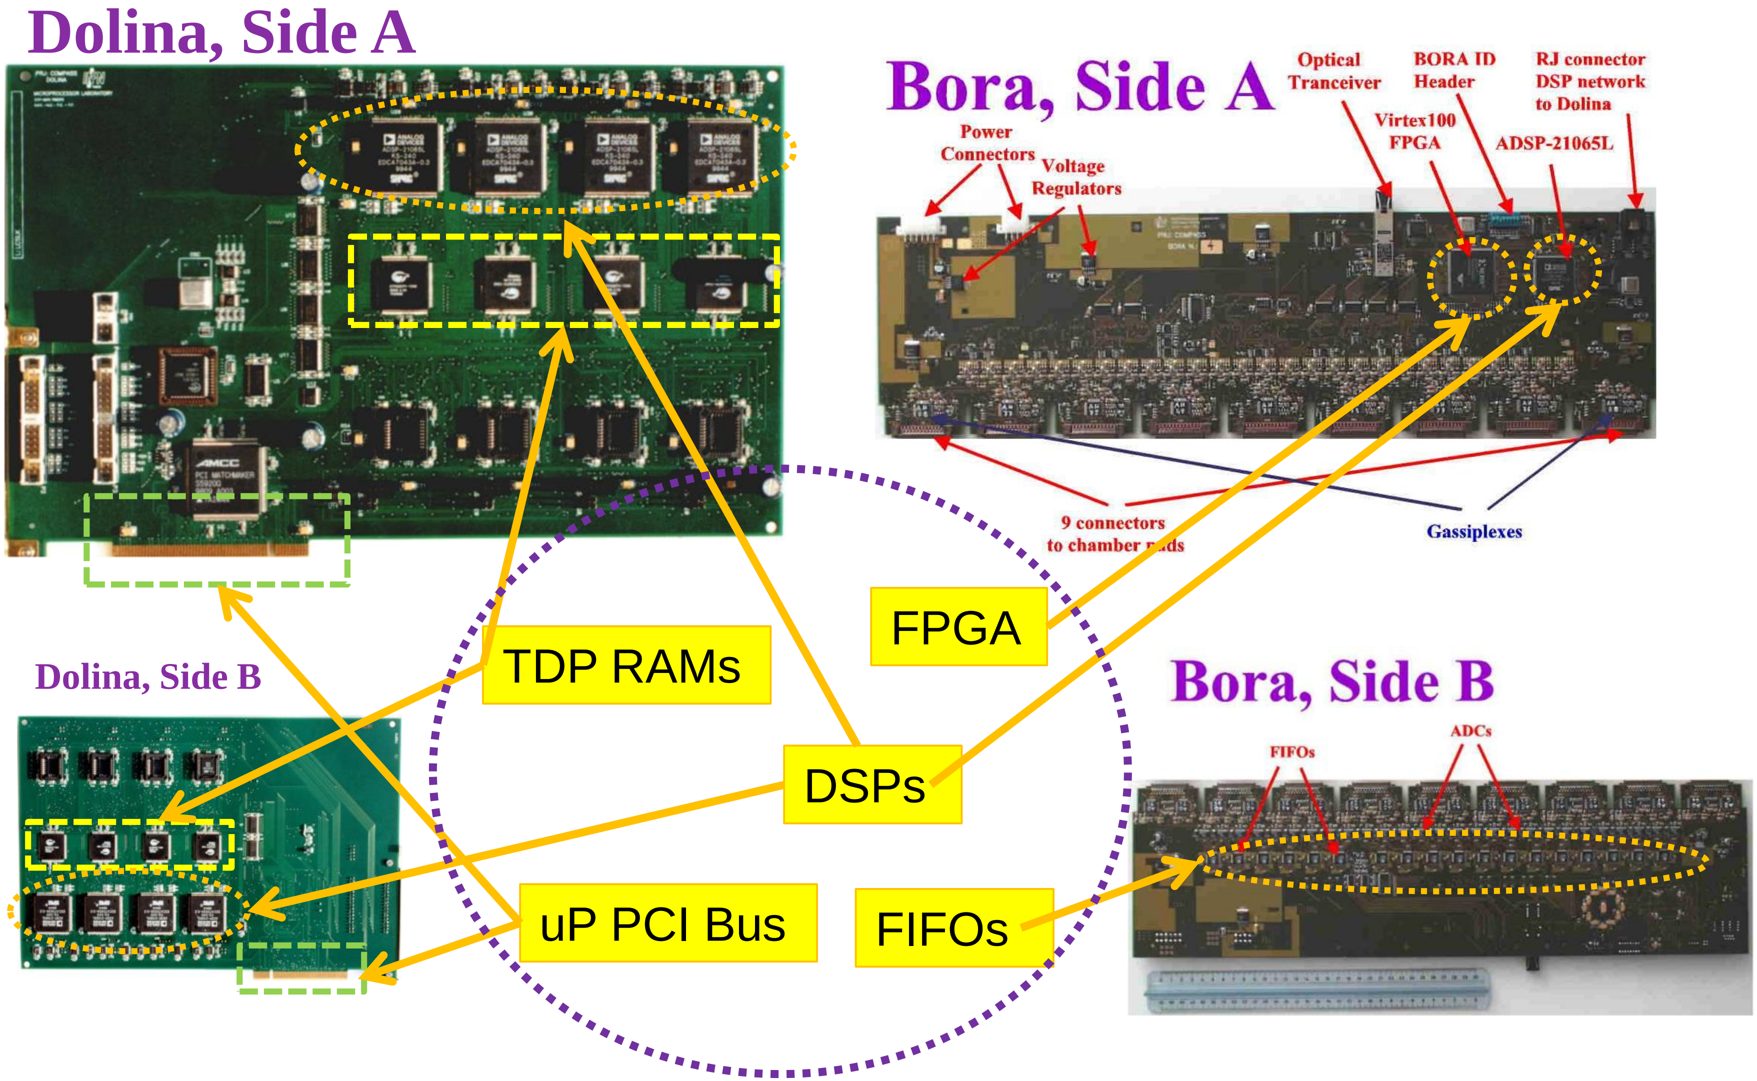

Reconfigurable Virtual Instrumentation based on FPGA

Reconfigurable Virtual Instrumentation based on FPGA and SoC FPGA

Massively parallel and distributed instrumentation in large high-energy physics experiments

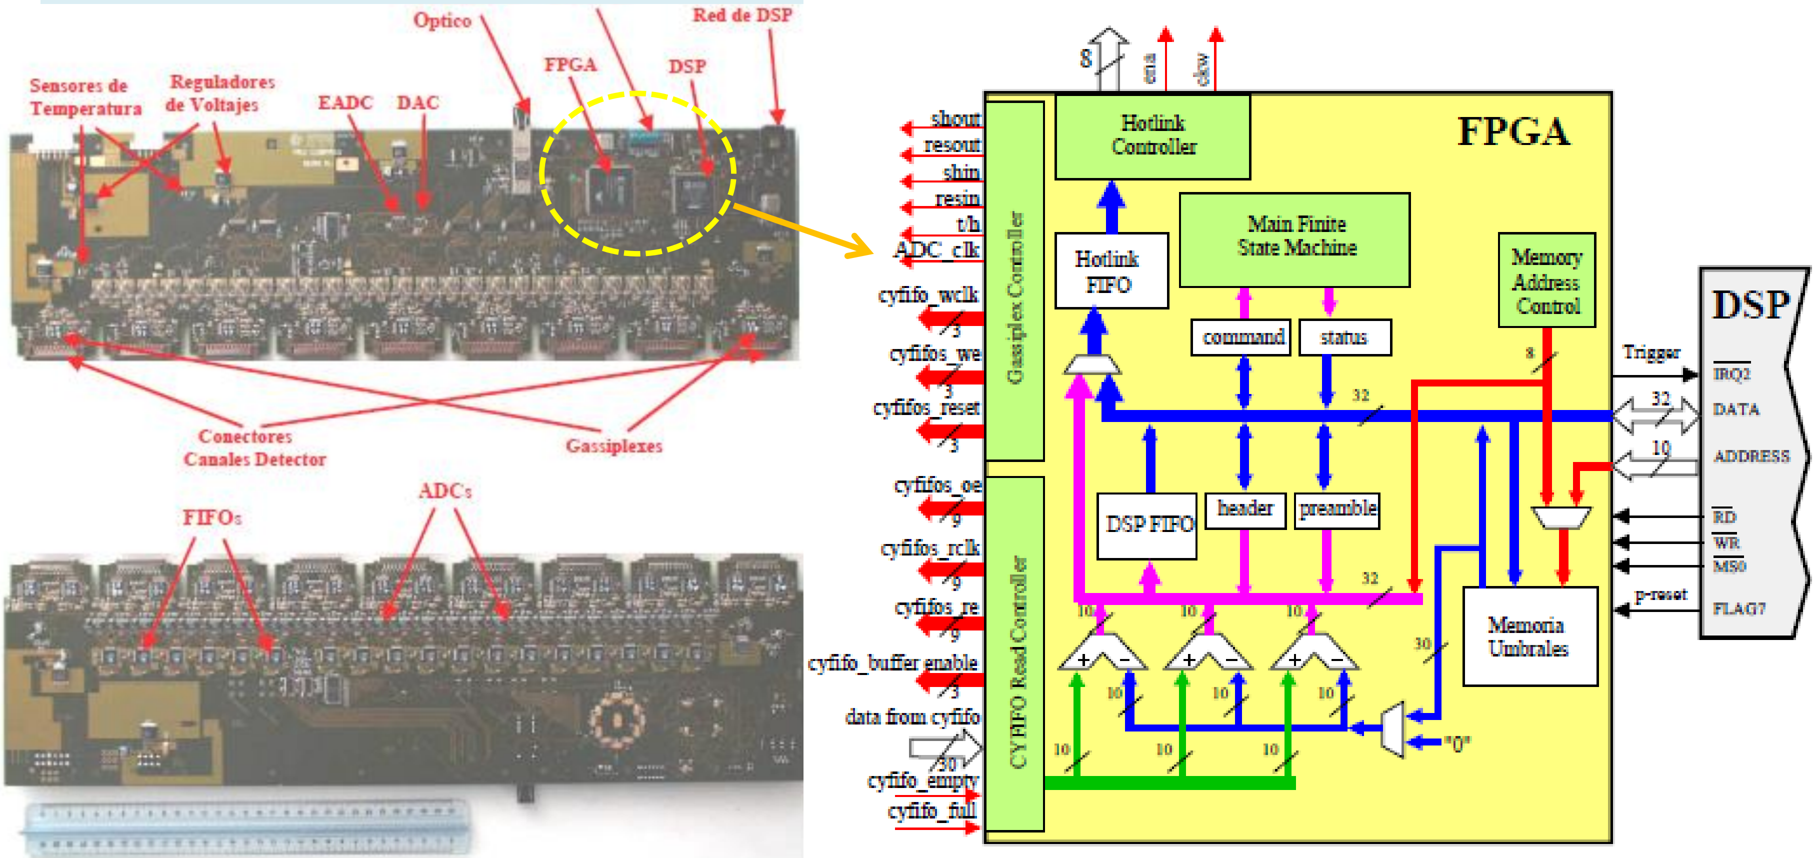

Description of Complex Systems

Modularity

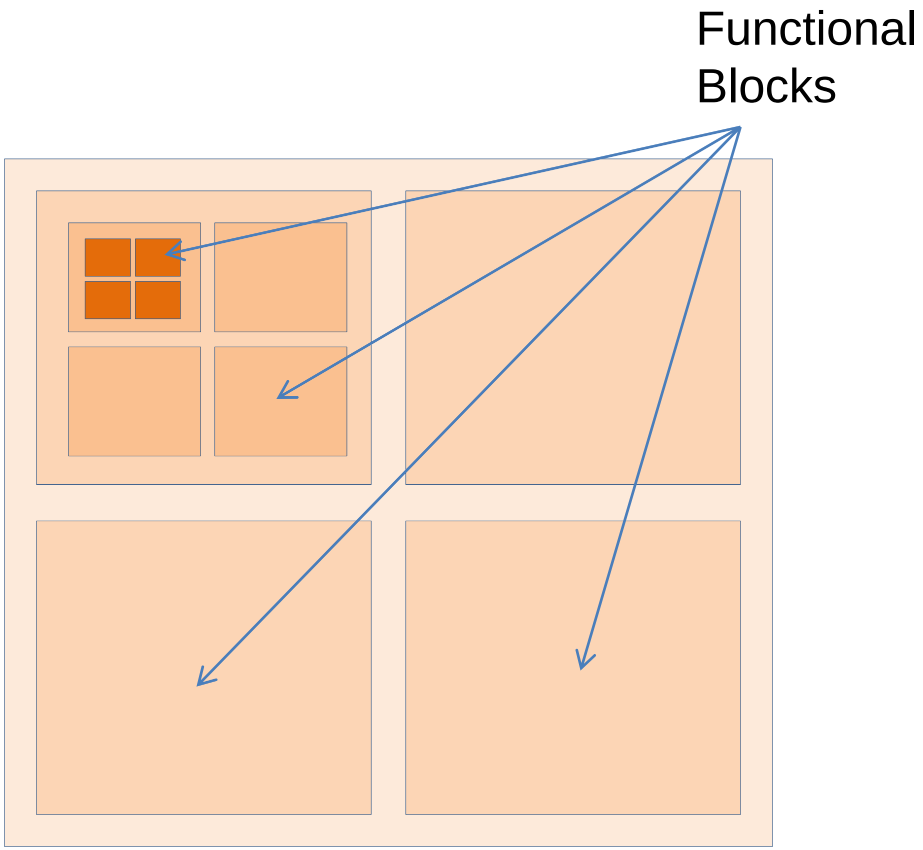

Hierarchy

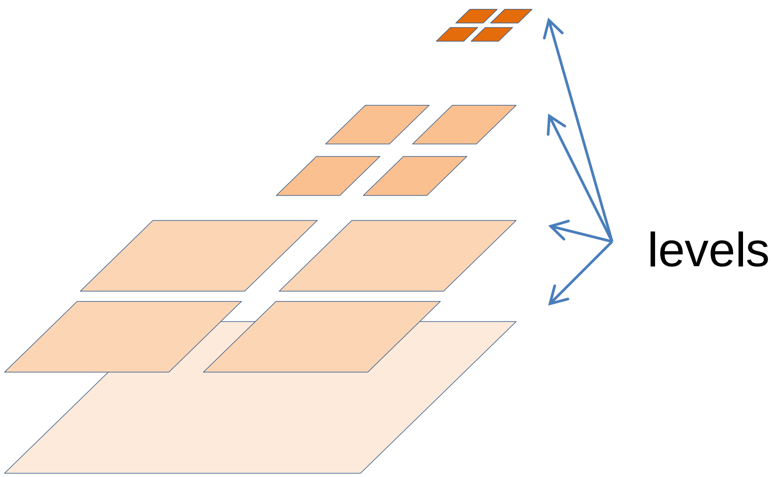

What activity at given hierarchical level?

Functional blocks Ports Data exchange

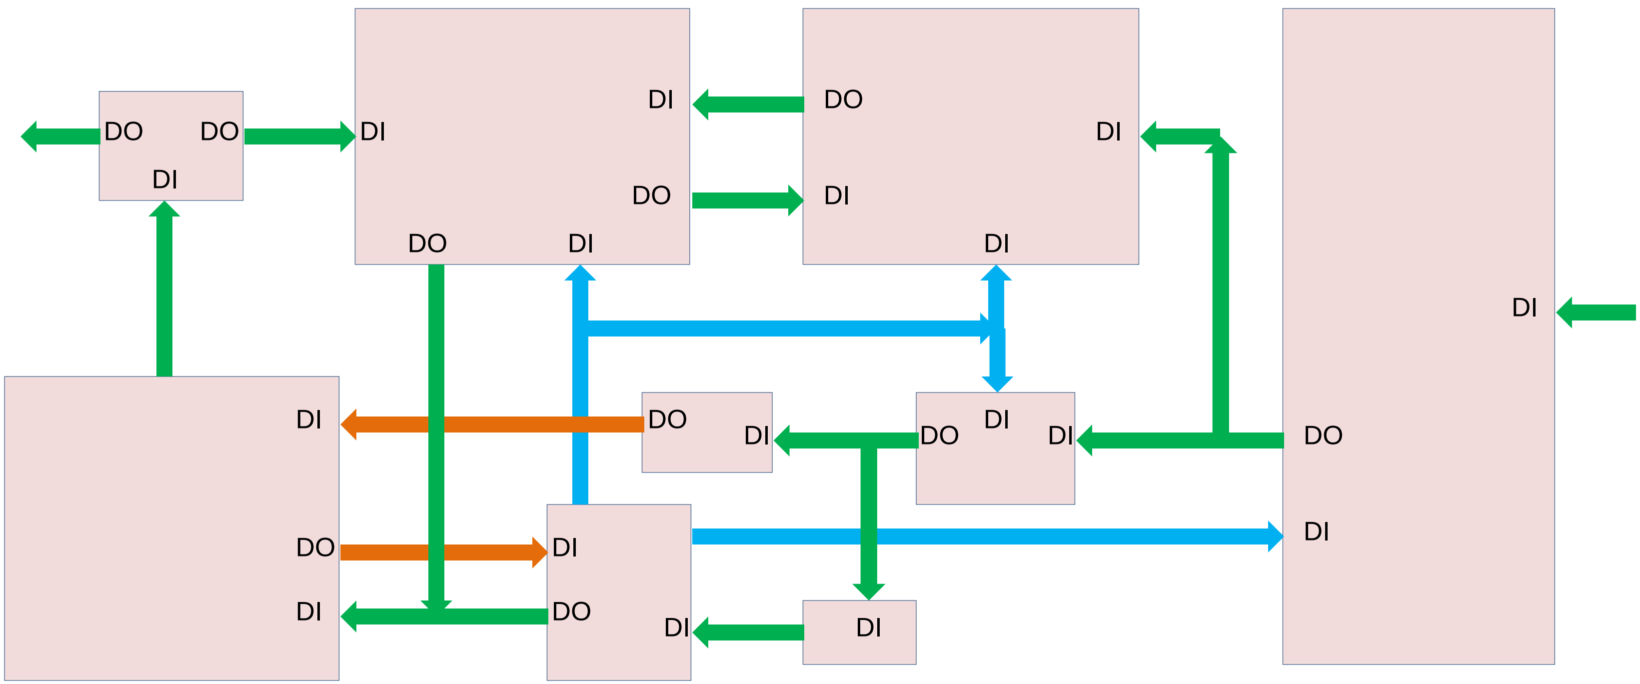

FPGA-Based Reconfigurable Instrument: Abstract Model

Hardware
Configuration

Instantiation of
functional blocks

Memory mapping
of registered ports

All HW Resources

Address	Ports
0x00000001 0x0000FFFF	Ext_RAM
0x000A0000 0x000AEEEE	Ext_ROM
0x000AEEEE	FIFO_a_in
0x000AEEF0	FIFO_b_out
0x001A0000 0x001AEEEE	RAM_block_p
0x002A0000	Register_h
0x002A000A	Operand _i
0x002A000B	Operand _j
0x003A0001	Operator_m_Out_k
0x003A0001	Ext_HW_in_port_x
0x003A0001	Ext_HW_out_port_y
	Register_k

Software
Programming

Description of the
HW actvity

Concurrent execution
of ***U**niversal **D**irect*
***M**emory **A**ccess*
*(**UDMA**)* instructions

Universal Direct Memory Access Instruction

UDMA SA DA SAinc DAinc N <BC> <activate, suspend, abort>

- SA: Source Address
- DA: Destination Address
- SAinc: Increment of Source Address
- DAinc: Increment of Destination Address
- N: Number of Words
- <BC>: Boolean condition
- <activate, suspend, abort>: Reaction

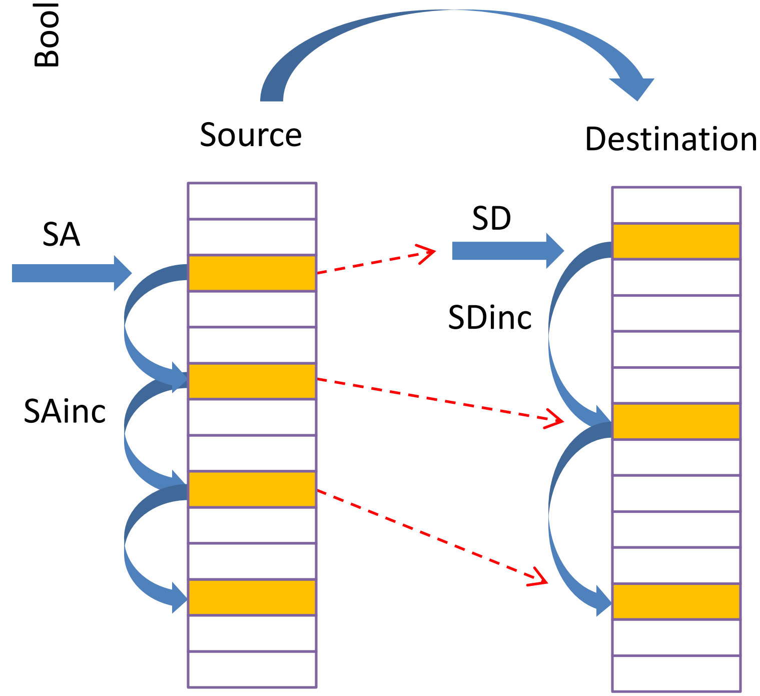

Universal Direct Memory Access Instruction

Some examples

UDMA 0x0000F001 0x0000F00A 1 1 256	*RAM to RAM*
UDMA 0x0000F002 0x0002F00B 1 0 1024	*RAM to FIFO*
UDMA 0x0000F003 0x0004F00C 0 1 1024	*FIFO to RAM*
UDMA 0xAAAAF003 0x008FAA80 4 1 2000	*RAM to RAM*
UDMA 0xAAAA4004 0x000FAA40 0 0 0	*Permanent link*
UDMA 0xFFFF4004 0x000FAA00 4 1 1024 “timer > countmax“ Abort	*Conditional data transfer*
UDMA 0xFFFF4004 0x000FAA00 4 1 1024 “counter1 == 31“ Suspend	*Conditional data transfer*

Universal Direct Memory Access Instruction for **Distribution**

UDMA_BC SA $\{DA_1,\ldots,DA_k\}$ N NC <BC> <R>

- SA: Source Address
- $\{DA_1,\ldots,DA_k\}$: Destination Addresses
- N: Number of Words per cycle
- NC: Number of cycles
- <BC>: Boolean condition
- <R>: Reaction

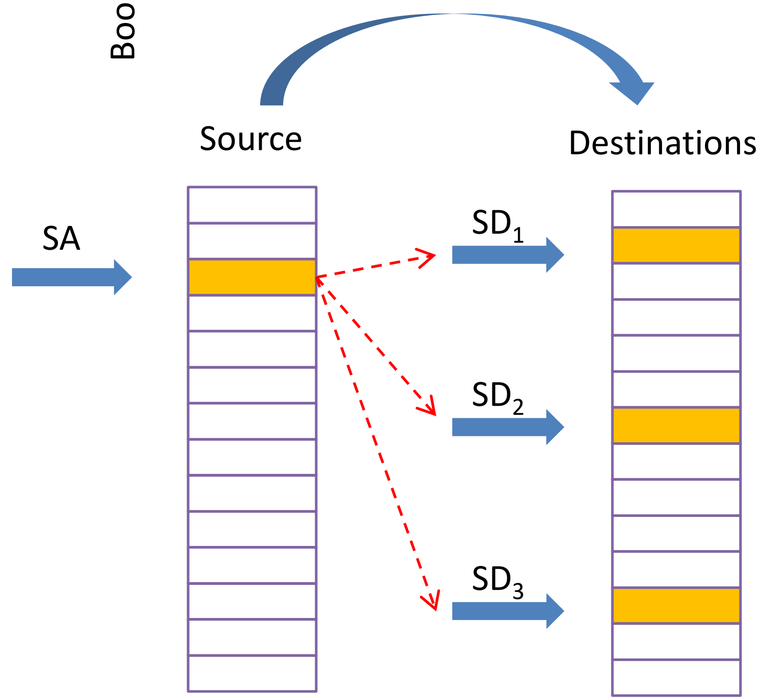

UDMA_BC (BroadCast)

UDMA_DRR (Distribute Round Robin)

UDMA_DTFF (Distribute Till FIFO Full)

Universal Direct Memory Access Instruction for **Gathering**

UDMA_GRR {SA_1,…,SA_k} DA N NC <BC> <R>

- {SA_1,…,SA_k}: Source Addresses
- DA: Destination Address
- N: Number of Words per cycle
- NC: Number of cycles
- <BC>: Boolean condition
- <R>: Reaction

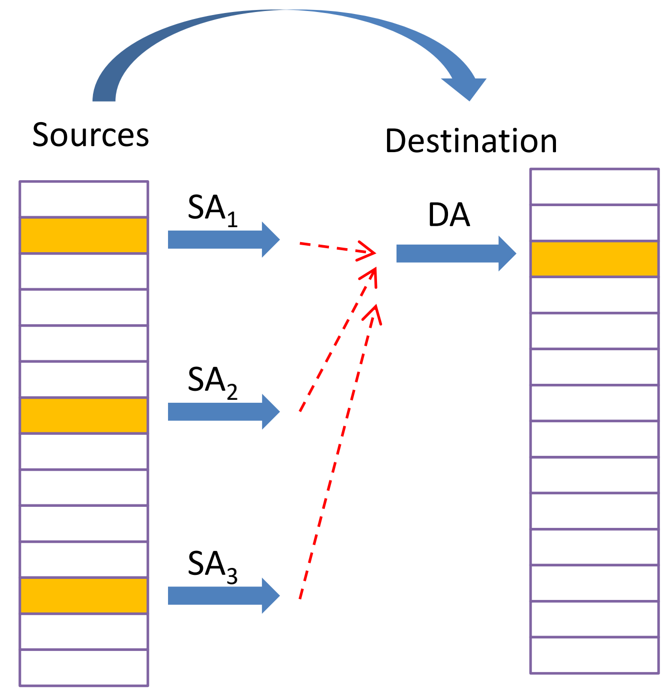

UDMA_GRR (Gathering Round Robin)

UDMA_GTFE (Gathering Till FIFO Empty)

System on Chip: The Wishbone Bus-Interface Standard Definitions

The four main components of the Wishbone system:

Master and *Slave interfaces*, *Syscon* and *Intercon*.

SYSCON: drives the system clock and reset signals.

MASTER: IP Core interface that generates bus cycles.

SLAVE: IP Core interface that receives bus cycles.

INTERCON: an IP Core that connects all of the MASTER and SLAVE interfaces together.

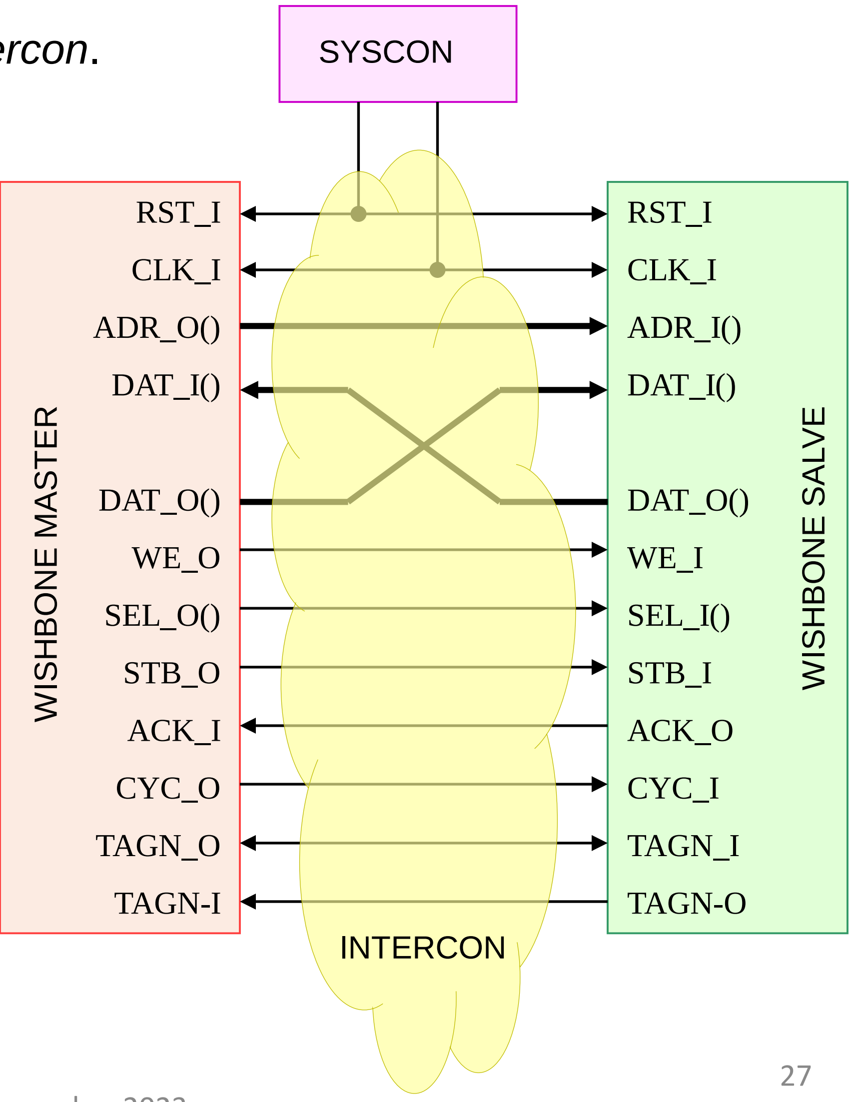

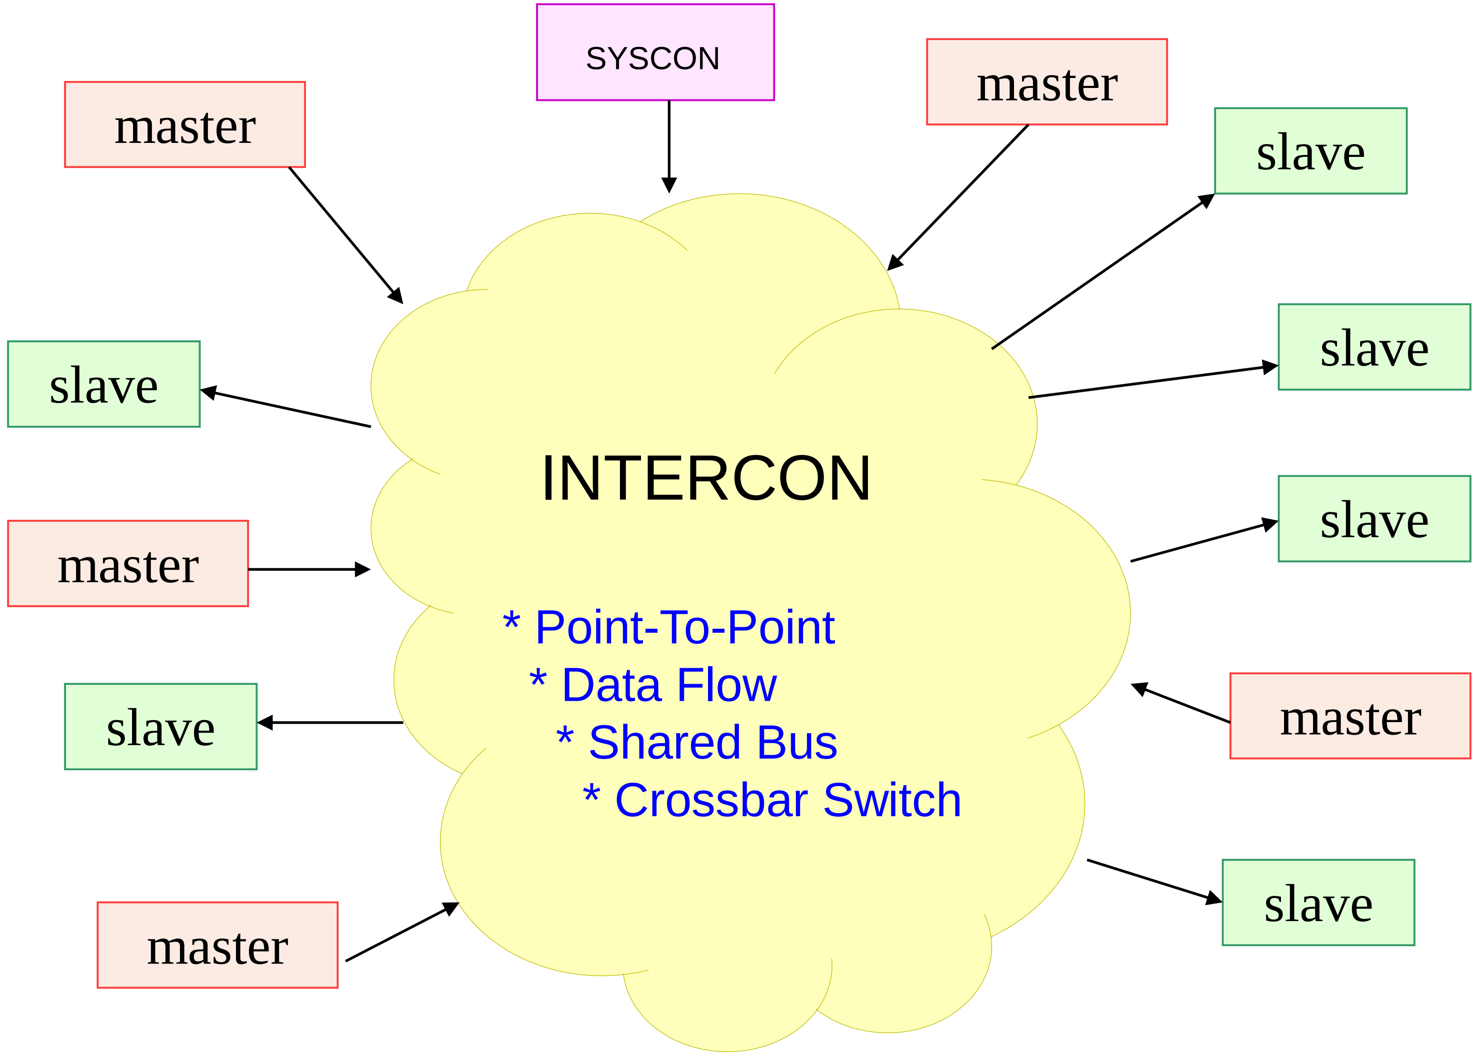

The Wishbone Interconnection is created by the SYSTEM INTEGRATOR, who has total control of its design

Interconnections II

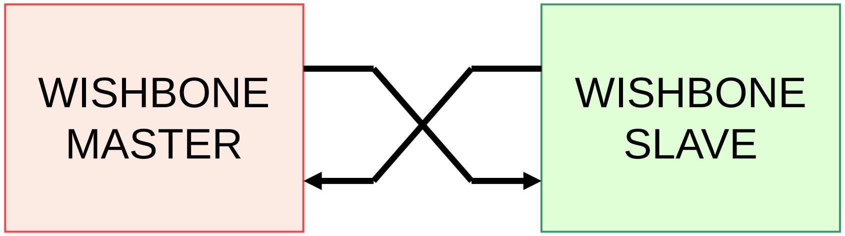

Point-To-Point

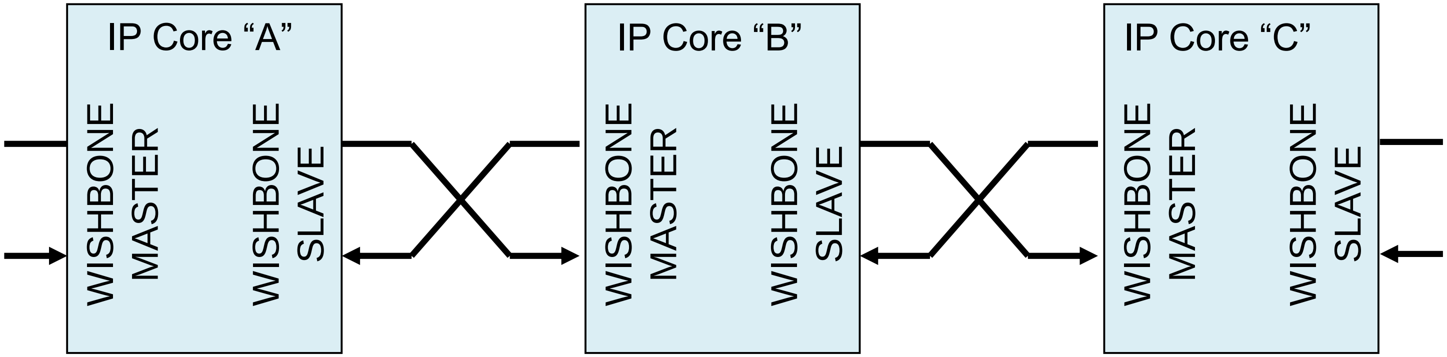

Data Flow

Interconnections III

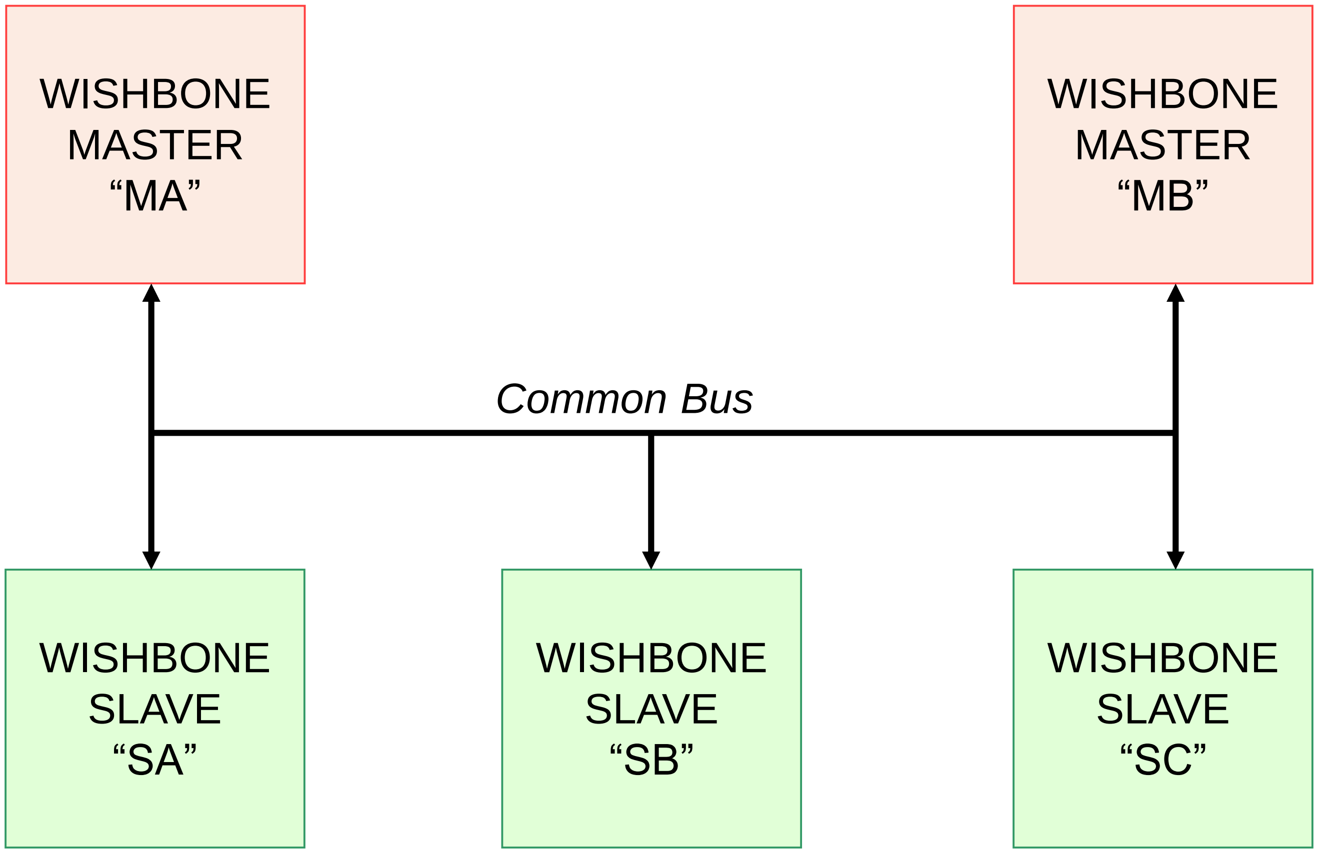

Shared Bus

Interconnections IV

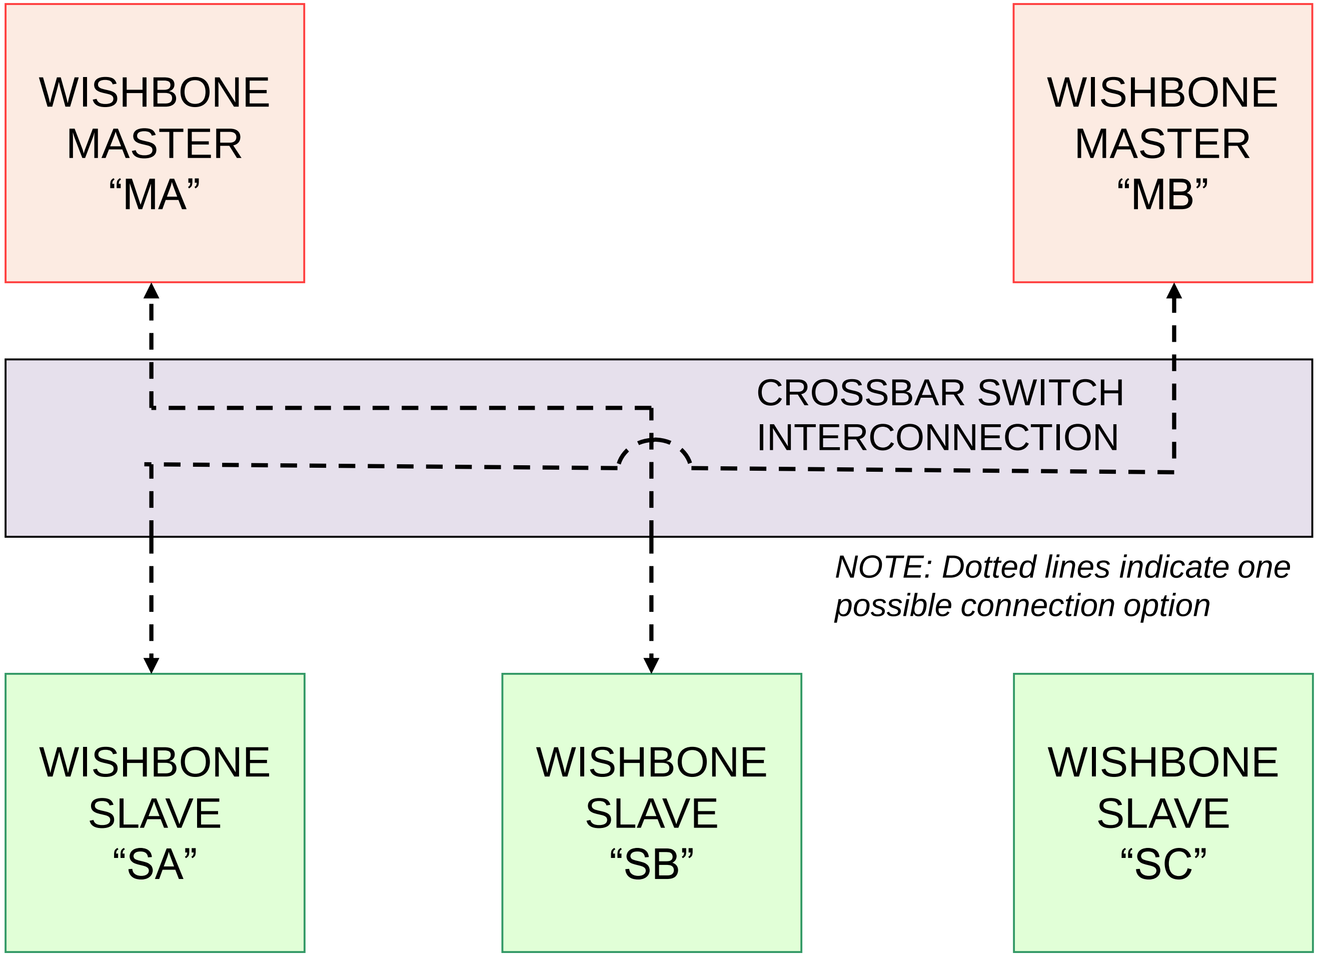

Crossbar Switch

UDMA controller for a system based on Wishbone compliant modules

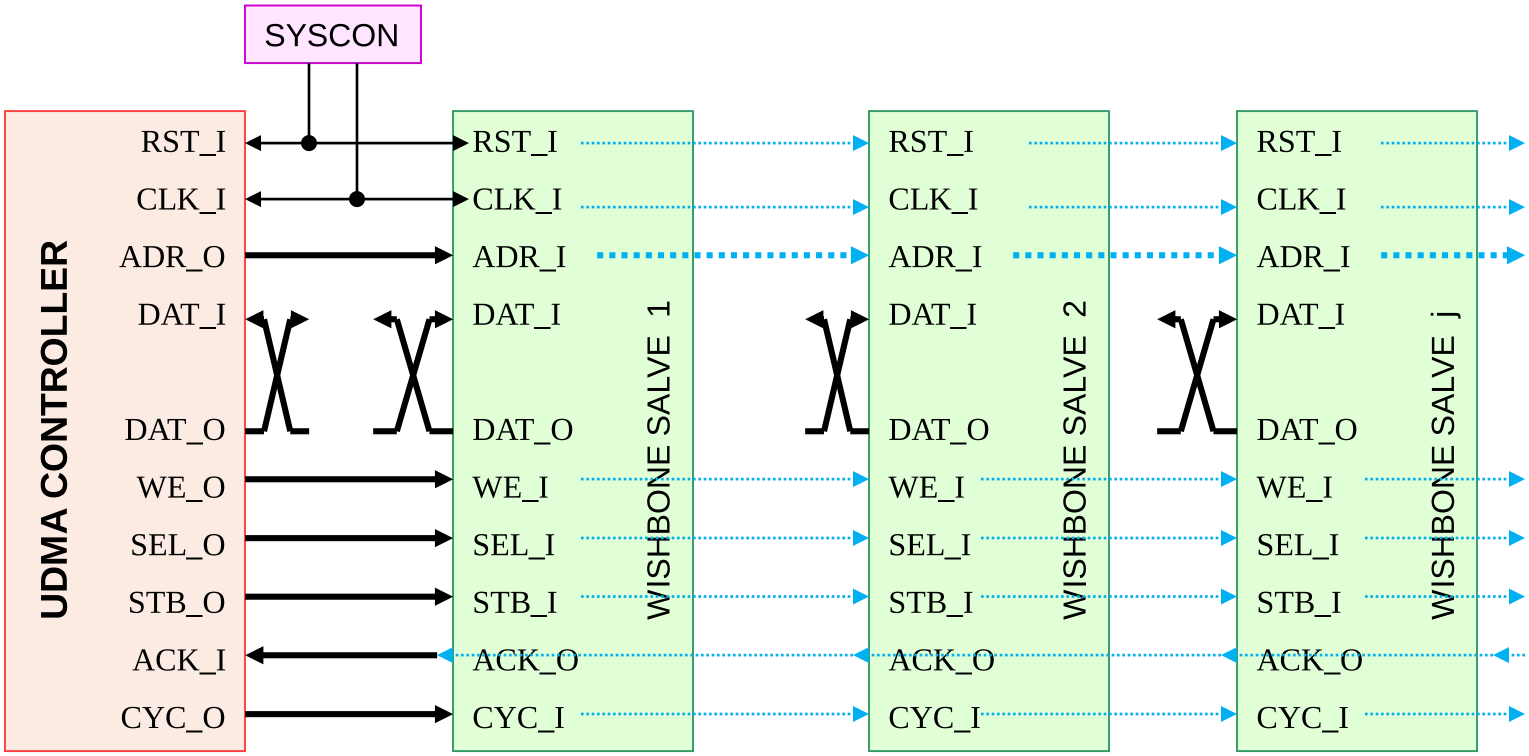

- UDMA instructions could be stored in a WB module
- One WB module must be a ***communication block*** which could also store UDMA Instructions in a reserved area.

Summary of key concepts so far and its relations

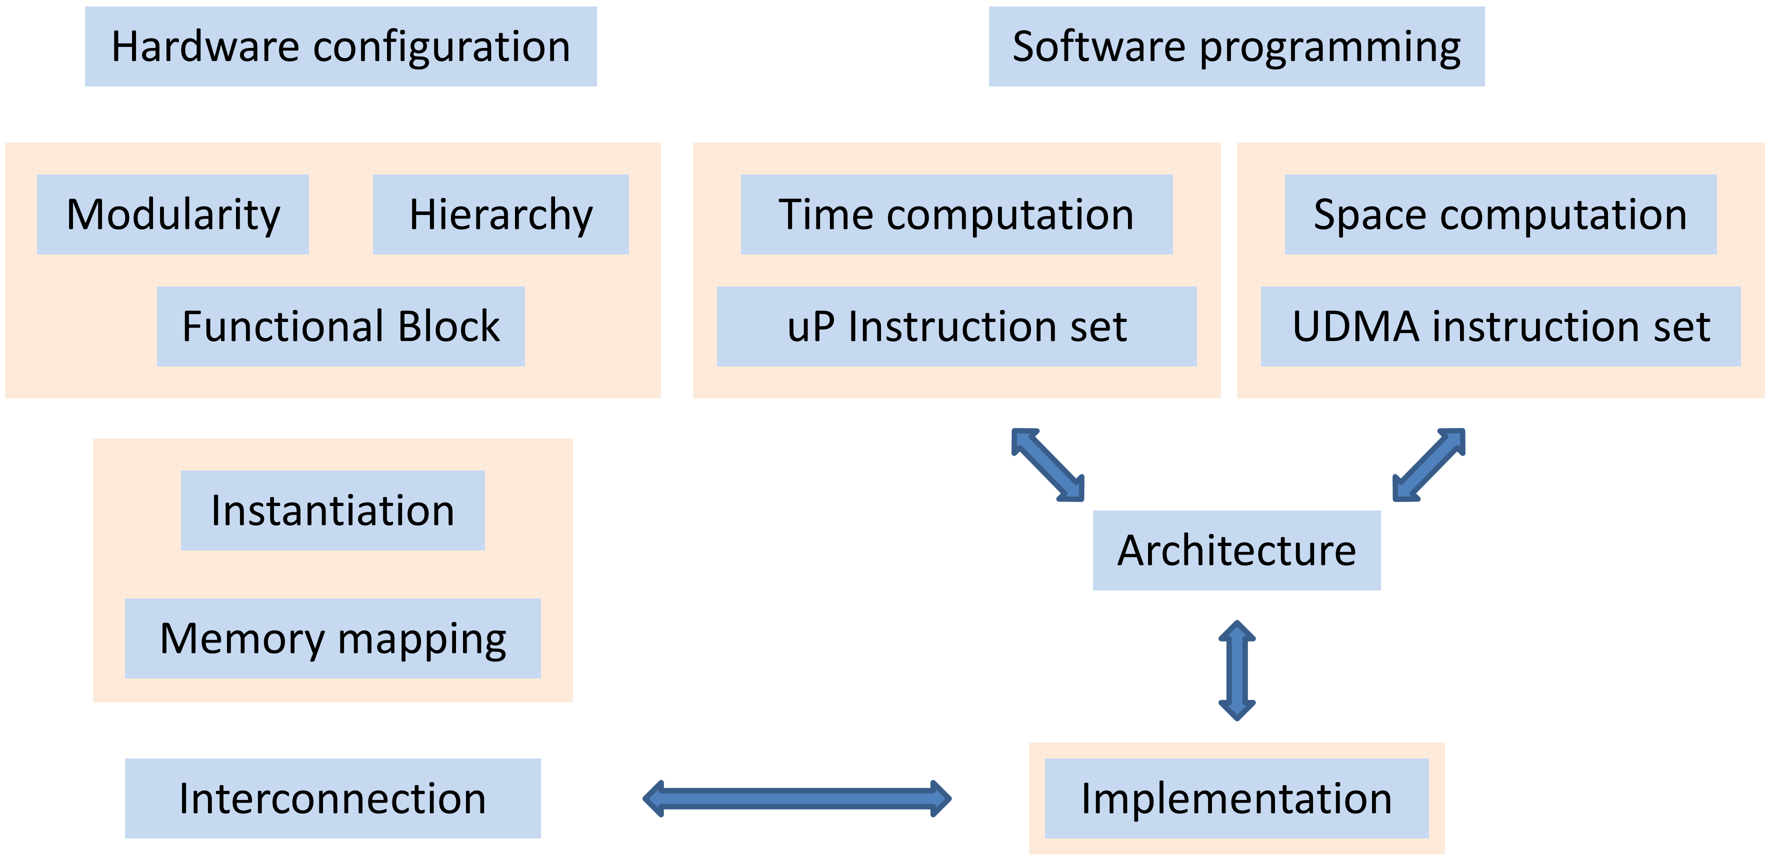

Communication through FPGAs in clusters of reconfigurable computational units

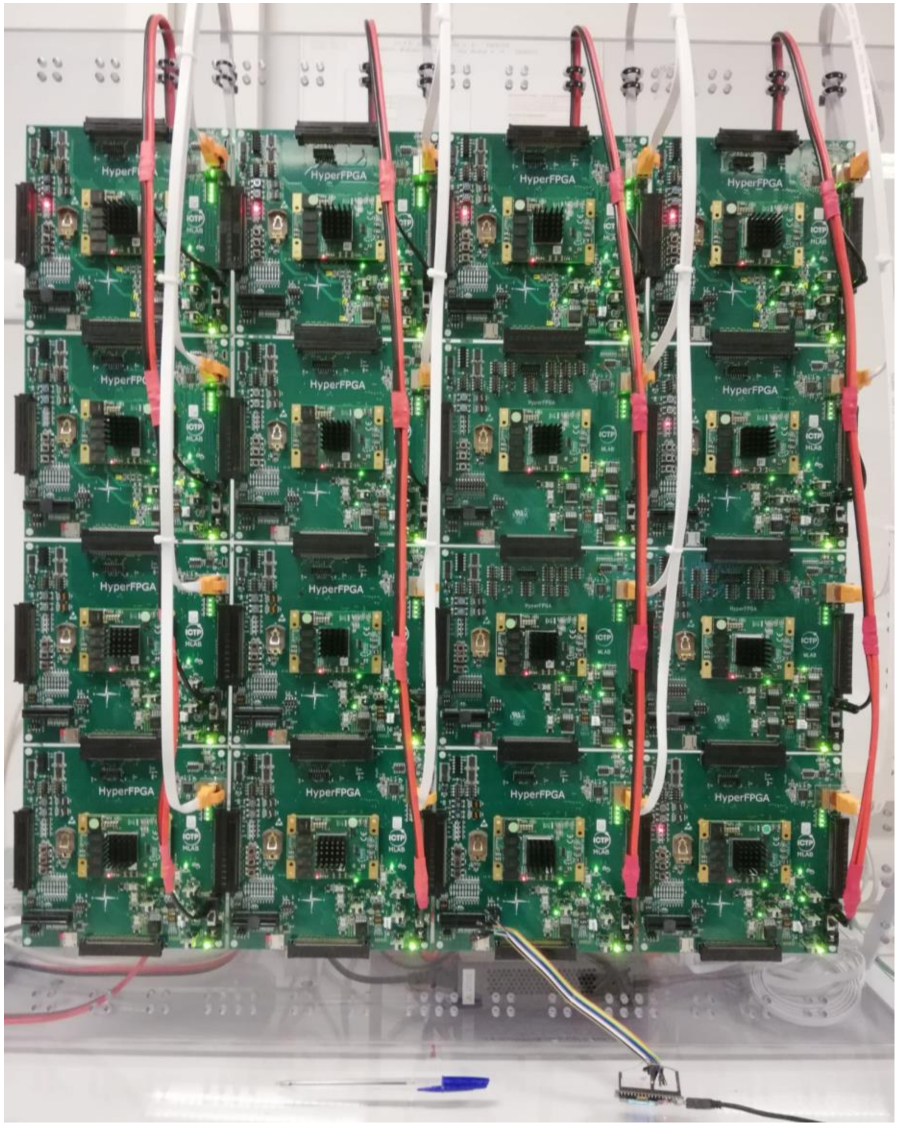

With same physical connections but with different IO configuration and activity programming:

Data packet transmission over

- *On demand point-to-point connections*
- *Buses*
- *Time-Division Multiplexing* on *common signal paths*

Interconnection of Multiple FPGAs

Different Topologies

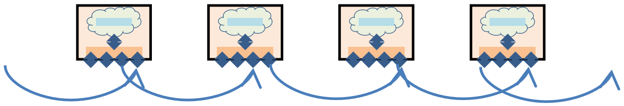

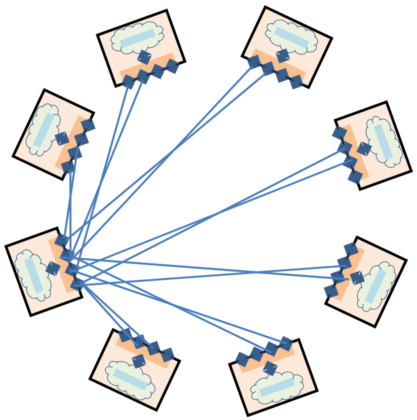

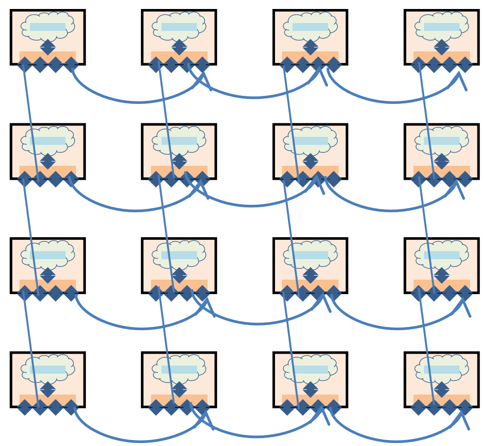

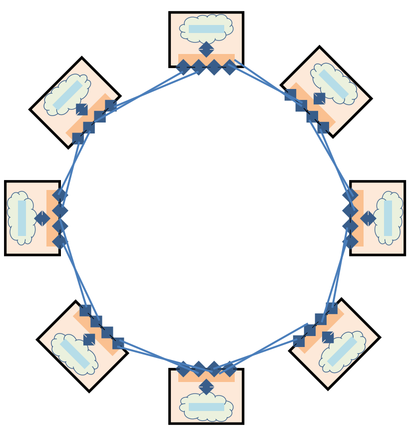

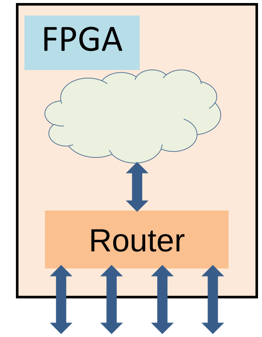

Three main communication layers

- Physical
- Logical
- System

Slave Communication Blocks

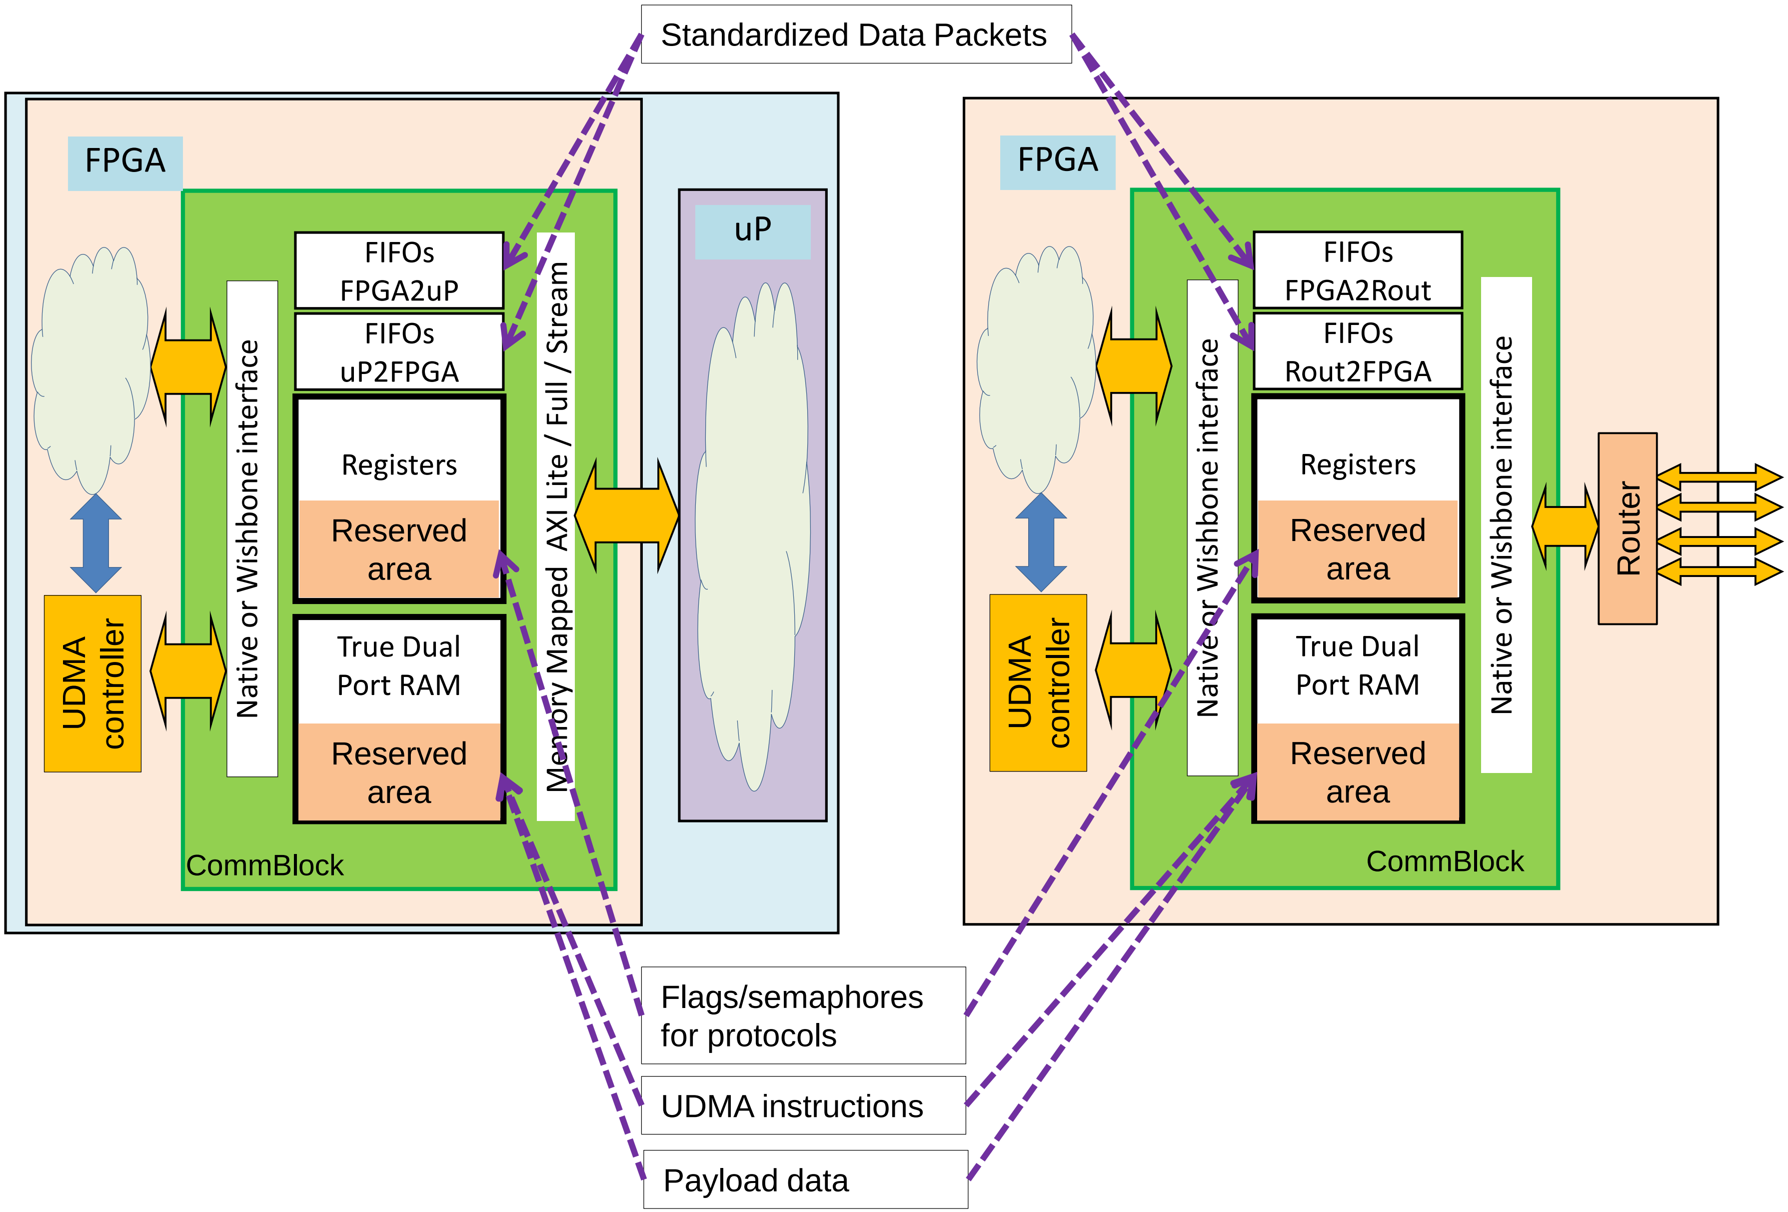

Standardized Data Packets

- Command
- Error Message
- Status report
- Raw Data
- Bit Stream
- Engineering frame
- **UDMA**
- Etc.

Header and Trailer
depend on
Packet type

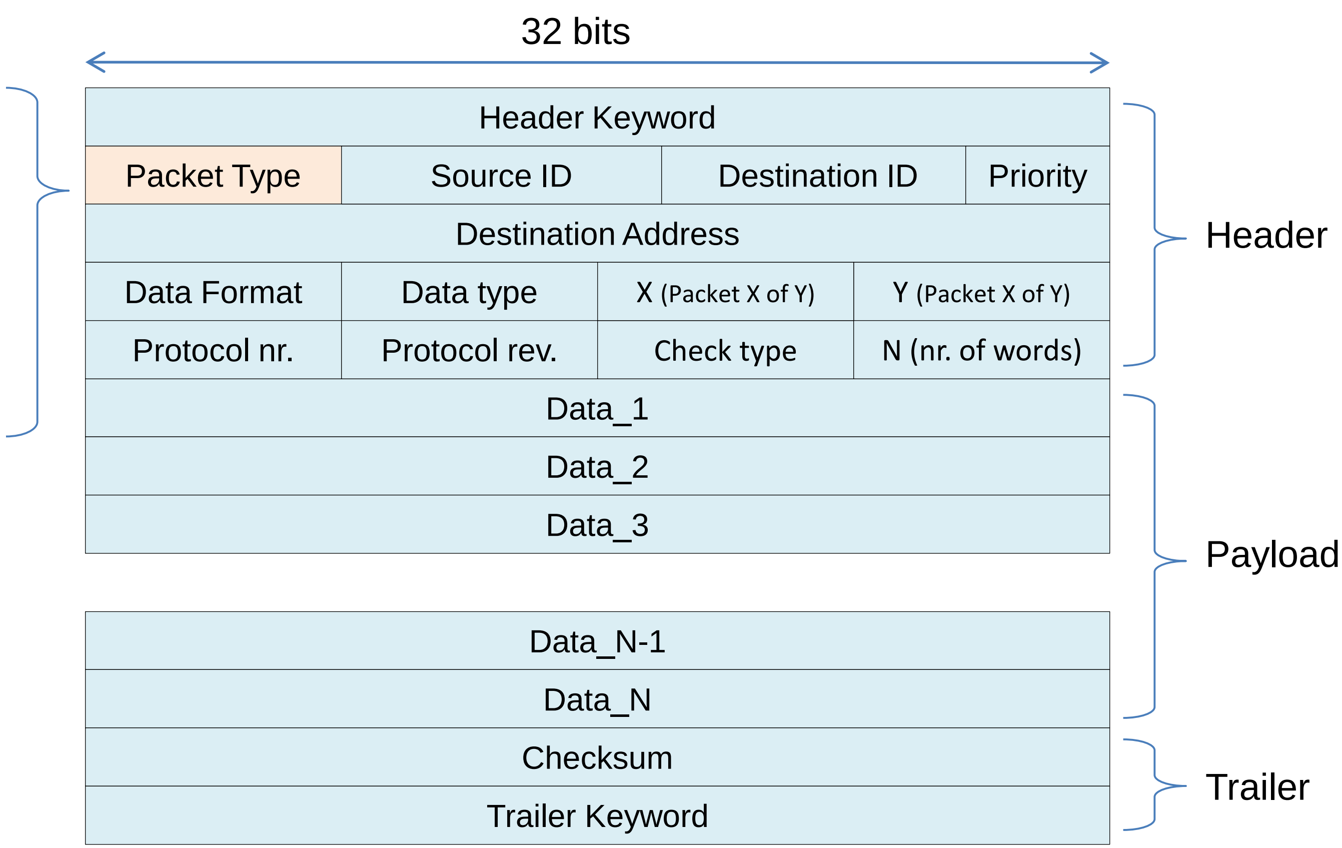

Standardized *data packets* and corresponding *handling mechanisms* for moving data across entire hybrid systems

UDMA SA DA SAinc DAinc N

(1) **UDMA-Packet** is sent from "i" to "j" to move data from data source "j" to destination "k"

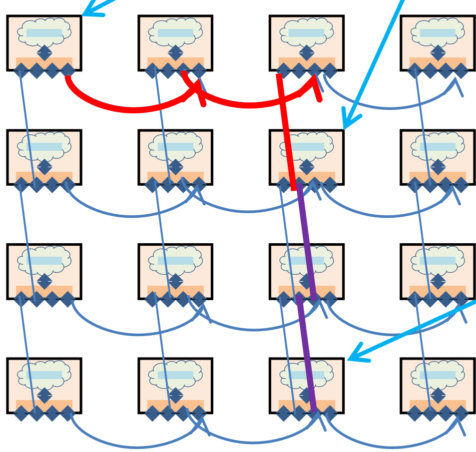

(2) A **Data-Packet** is prepared and sent from data source "j" to destination "k"

At this level of abstraction we don't care about underlying networks and low level communication layers.

Data also include *instructions*, *commands*, *error messages*, etc.

Corresponding *Acknowledge-Packets* can optionally be sent back to conclude transactions

Preliminary conclusions I

- Reconfigurable Hardware abstract models and strategies developed for advanced scientific instrumentation based on FPGA can be adapted for high-performance reconfigurable computing.
- Abundance of reconfigurable hardware resources lead to new computational paradigms inspired on the FPGA model escaping from the limitations of typical von Neumann and similar uP architectures.
- A spatial dimension can be added to the temporal dimension of dominant computing paradigm based on uP instruction set architectures.
- Universal Direct Memory Access (UDMA) instructions appear as a suitable means to describe and program the computational activity of powerful hardware platforms based on modern reconfigurable hybrid devices such as SoC FPGA.

Preliminary conclusions II

Recalling The Custom Computing Problem

- Which is the best reconfigurable hardware infrastructure?
- Which language should be used to capture a computational problem and codify its solution?
- Which tools should be developed to configure the hardware to implement the best custom computer?
- Which tools should be developed to compile the code for its efficient execution in the configured custom computer?

This is still a very complex problem that needs multidisciplinary contributions and positive knowledge experimentally obtained on scalable hardware infrastructures.

Thank you for your attention!

Opportunities for open collaboration on scientific supercomputing based on FPGA technologies

- Synergies between Industry, Universities and Public Research Centers.
- ICTP (UNESCO - IAEA) Programs
  - **TRIL**: Training and Research in Italian Laboratories
  - **Associates** (junior, regular, senior)
  - **Federation Agreements**
  - **Scientific Calendar** of international activities for training and research in Physics, Mathematics and Interdisciplinary areas.

ICTP (UNESCO - IAEA) Programs

TRIL: Training and Research in Italian Laboratories

https://www.ictp.it/tril.aspx

This programme offers scientists from developing countries the opportunity to undertake training and research in an Italian laboratory in different branches of the physical sciences

The ICTP has established agreements of collaboration with more than 400 Italian research institutes, providing young scientists with numerous options. TRIL partners include:

- CNR (Italian National Research Council) institutes
- Elettra-Sincrotrone Trieste (Elettra Synchrotron Light Source)
- ENEA (Italian National Agency for New Technologies, Energy and Sustainable Economic Development)
- INFN (National Institute for Nuclear Physics)
- INGV (Istituto Nazione di Geofisica e Vulcanologia)
- OGS (National Institute of Oceanography and Experimental Geophysics)

ICTP (UNESCO - IAEA) Programs

ICTP Associateship: Junior (<36), Regular (<46), Senior (<63)

https://portal.ictp.it/assoc/associateship-scheme

The Associate Scheme is one of the ICTP's oldest programs, and was established to provide support for distinguished scientists in developing countries in an effort to lessen the brain-drain.

- The Junior Associateship award has a six-year duration throughout which the Junior Associate is entitled to spend up to 180 days (with a maximum duration of 60 days for any single visit) at the Centre, with three fares paid. A fare is granted for visits having a minimum duration of 30 days. For each visit the Centre provides a daily living allowance.

- The Regular Associateships are six-year awards intended exclusively for scientists between the ages of 36 and 45 from and working in developing countries.

- Senior Associateships are intended for scientists from a developing country who have acquired international scientific status. Awards have a six-year duration with a total allocation of 8000 Euro. These funds are made available for visits in the form of a daily living allowance and/or travel expenses. During the six years, Senior Associate Members may apply to visit the Centre as often and for as long as they wish, until the allocation is exhausted, although the maximum foreseen duration of any visit is 60 days.

ICTP Federated Institutes

https://www.ictp.it/programmes/federated-institutes.aspx

The Federated Institutes programme offers young scientific staff, as well as post-doctoral and PhD students from institutes in developing countries, the opportunity to attend meetings at ICTP or to participate in group activities.

Institutes wishing to be considered for the possibility of becoming an ICTP Federated Institute must satisfy the following criteria:

- The institute must be in a developing country;
- The institute must have active research programmes in at least one of the areas of interest to ICTP;
- There should be at least a Masters but preferably a PhD programme in the fields of interest;
- In case the institute is accepted as being Federated, the coordinator (applicant) must be an active member of the institute for the duration of the agreement.
- Former Federated Institutes are eligible to apply again for Federation status. Extensions are not envisaged.

ICTP (UNESCO - IAEA) Programs

ICTP Scientific Calendar

https://www.ictp.it/scientific-calendar.aspx

Each year, ICTP organizes more than 60 international conferences, workshops, and numerous seminars and colloquia for training and research in Physics, Mathematics and Interdisciplinary areas.

- Those interested in attending an activity must complete an online application form.
- To propose a conference, school or workshop check the corresponding guidelines (https://www.ictp.it/call-for-proposals.aspx).
- **The deadline for proposals is typically end of February for** activities to take place in the next year. ICTP announces the call for proposals on its website.
- Travel fellowships and financial support for ICTP conferences and workshops are available.

ICTP (UNESCO - IAEA) Programs

ICTP invites proposals from the international scientific community for any of the following types of activities:

Schools/Colleges: These largely pedagogical events cover a relatively broad scientific field normally through lectures at an expository level, and may include exercise sessions, discussion groups and computer laboratory sessions.

Advanced Schools/Workshops: These events deal with specific or specialized topics. In some cases, particularly when held periodically over time, the main purpose may be to cover developments of the last few years. A fraction of the audience may consist of former participants who should be actively involved in the programme, for instance through poster sessions. Typical length is 2 weeks.

Conferences: These activities last for a few days to a week and consist of presentations of recent results on timely and exciting subjects.

Extended Workshops: These less structured activities last from 2 to 3 months and cover selected research topics.

Outside Activities: Regional activities, to take place in an emerging or developing country, meant for promoting science in the host country and the surrounding region.

Co-sponsored Activities: Proposed activities that typically bring most of their own funding and organization, but seek an international venue and only modest support from ICTP.

Thank you for your attention!