




| The European Synchrotron

Development of an advanced data acquisition framework for 2D X-Ray Detectors

27/10/2020

Aurélien BIDEAUD and **Samuel DEBIONNE**

on behalf of many ESRF staff



Introduction

A. RASHPA: RDMA-based Data Acquisition framework

- I. Introduction
- II. Architecture
- III. Implementation
- IV. Example

B. LIMA2: Distributed Acquisition for High Performance Detectors

- I. Motivations
- II. Software Stack
- III. Data transfers
- IV. Processing
- V. Plugins

Conclusion & perspectives

Increase X-Ray source brilliance

High frame rate: 1 - 10 KHz


Large detectors

Pixel detectors

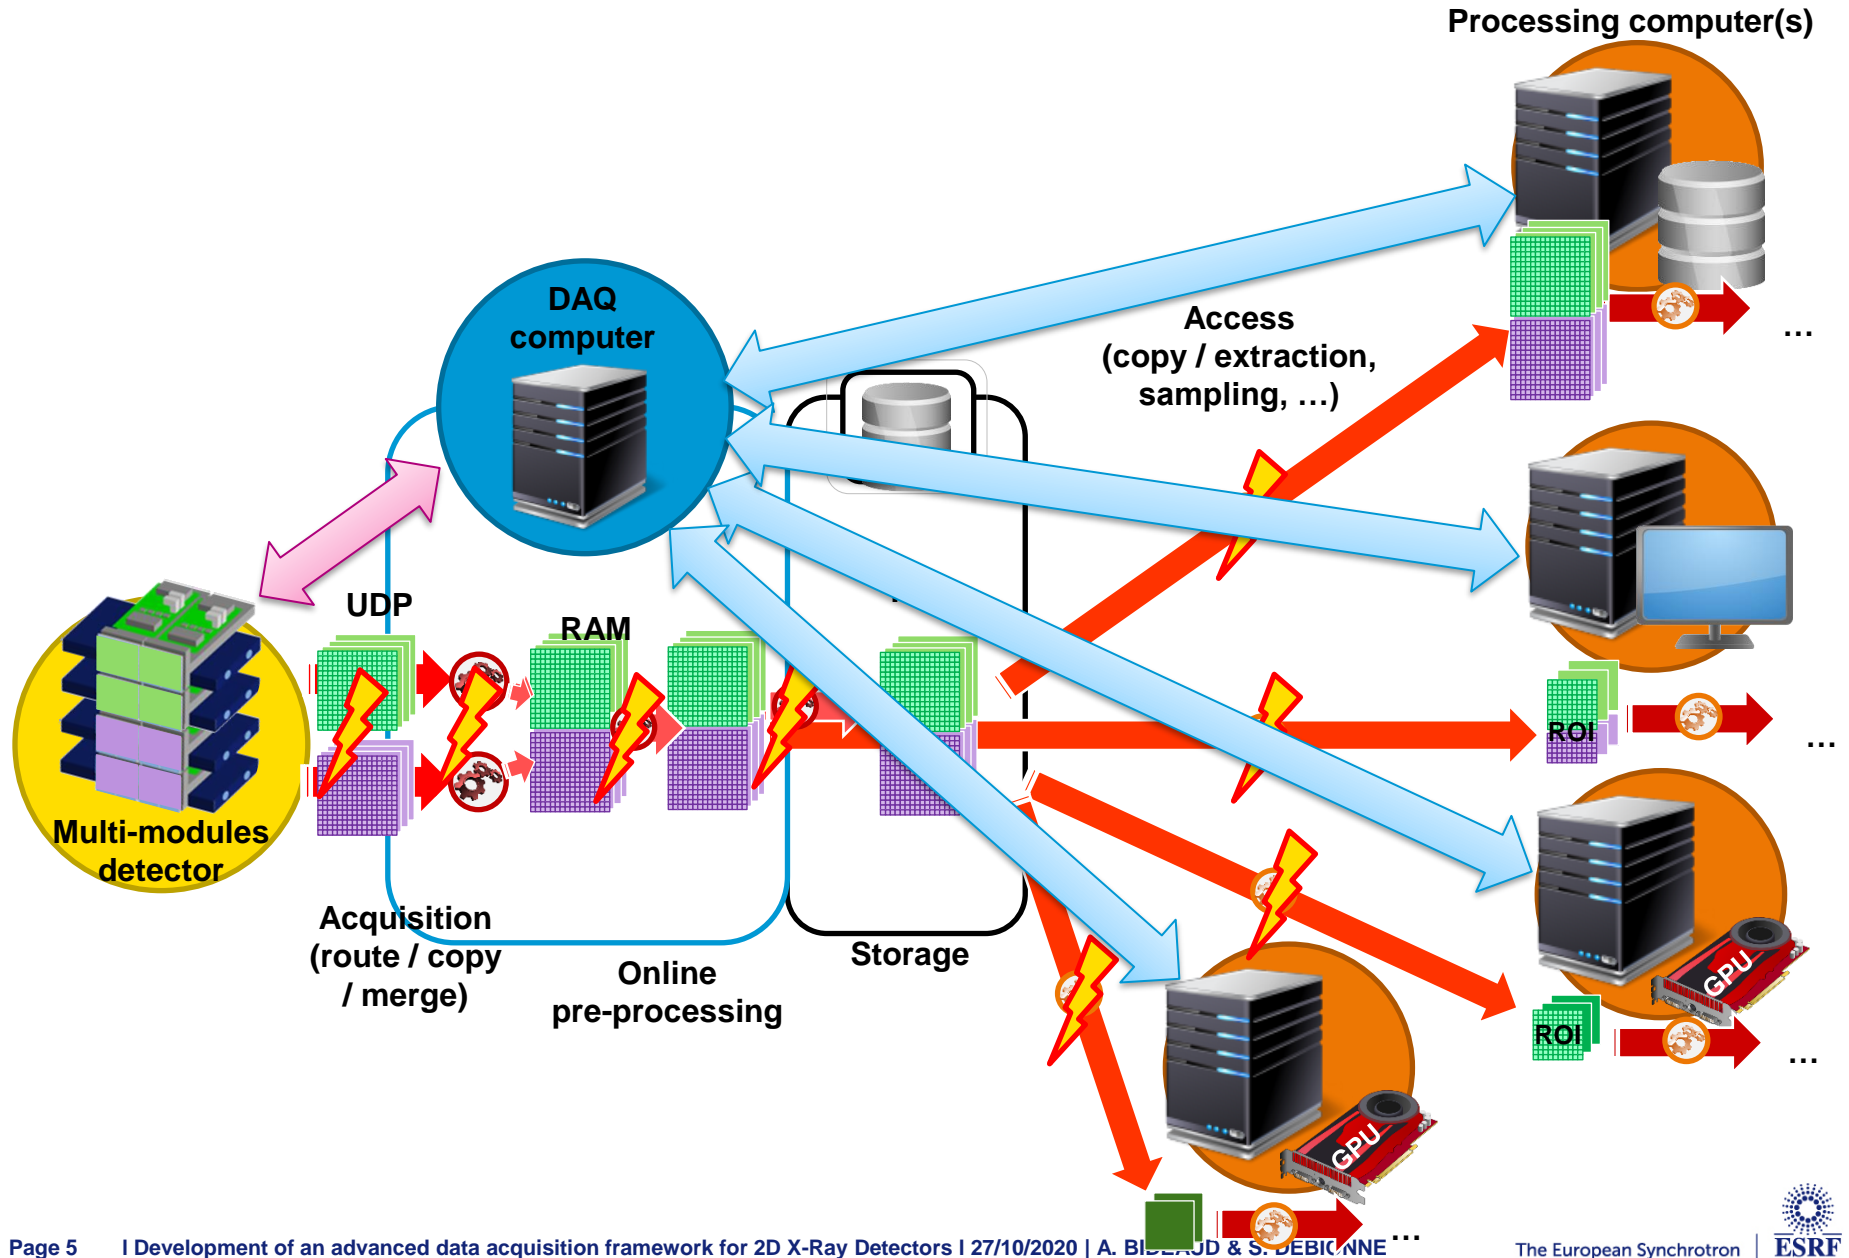
- Single chips: **100 MB/s**
- Basic modules: **0.5 – 1 GB/s**
- Multi-modules: **10+ GB/s**

Data transfer can limit sensor speed

Proprietary solutions



ARCHITECTURE OF THE CURRENT DAQ FRAMEWORK @ ESRF



Introduction

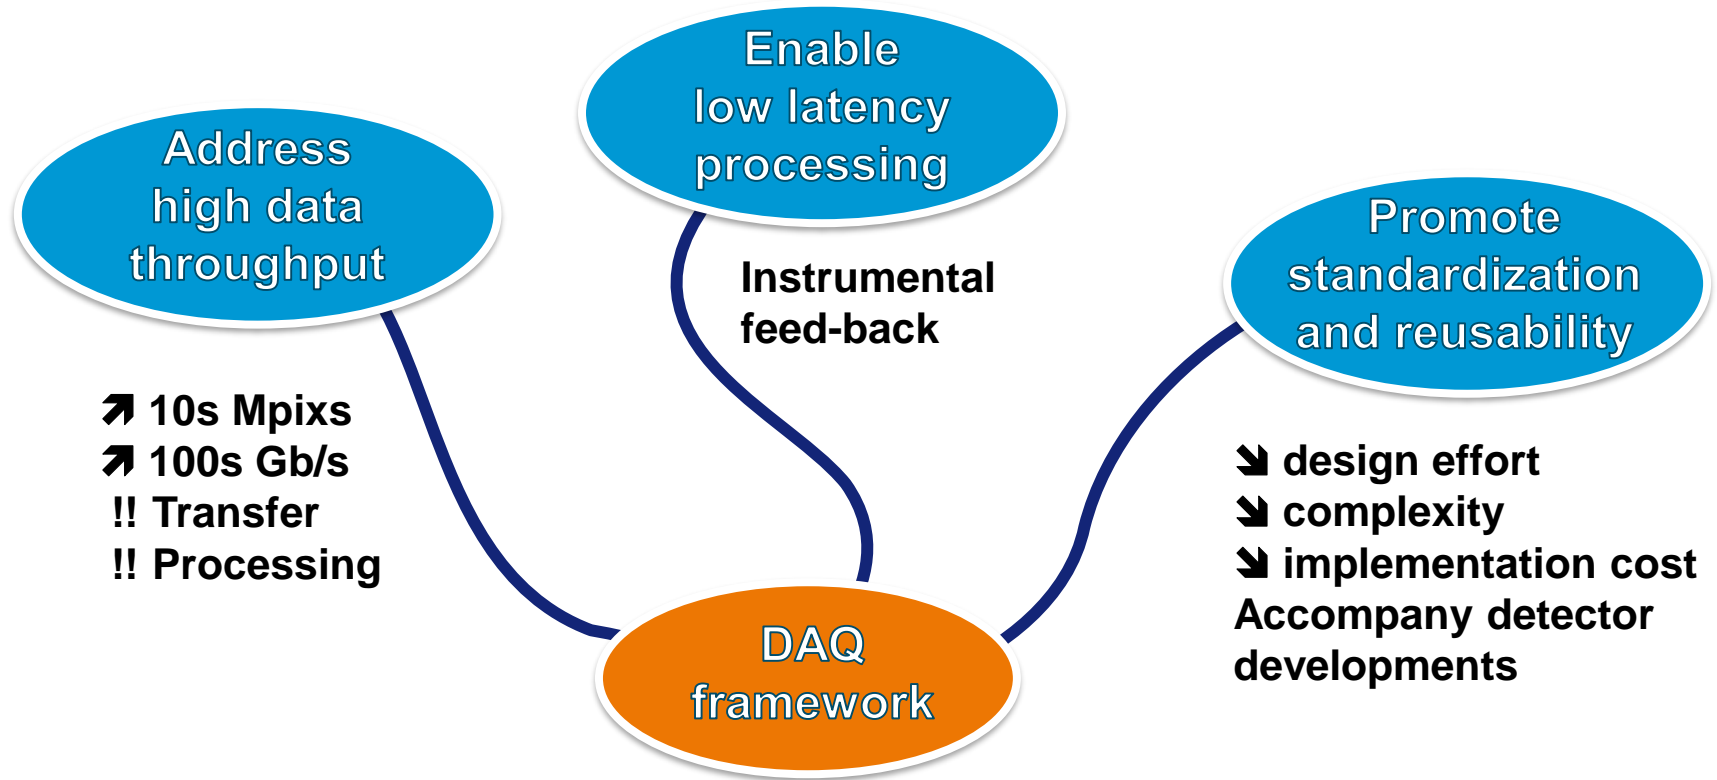
A. RASHPA: RDMA-based Data Acquisition framework

- I. **Introduction**
- II. Architecture
- III. Implementation
- IV. Example

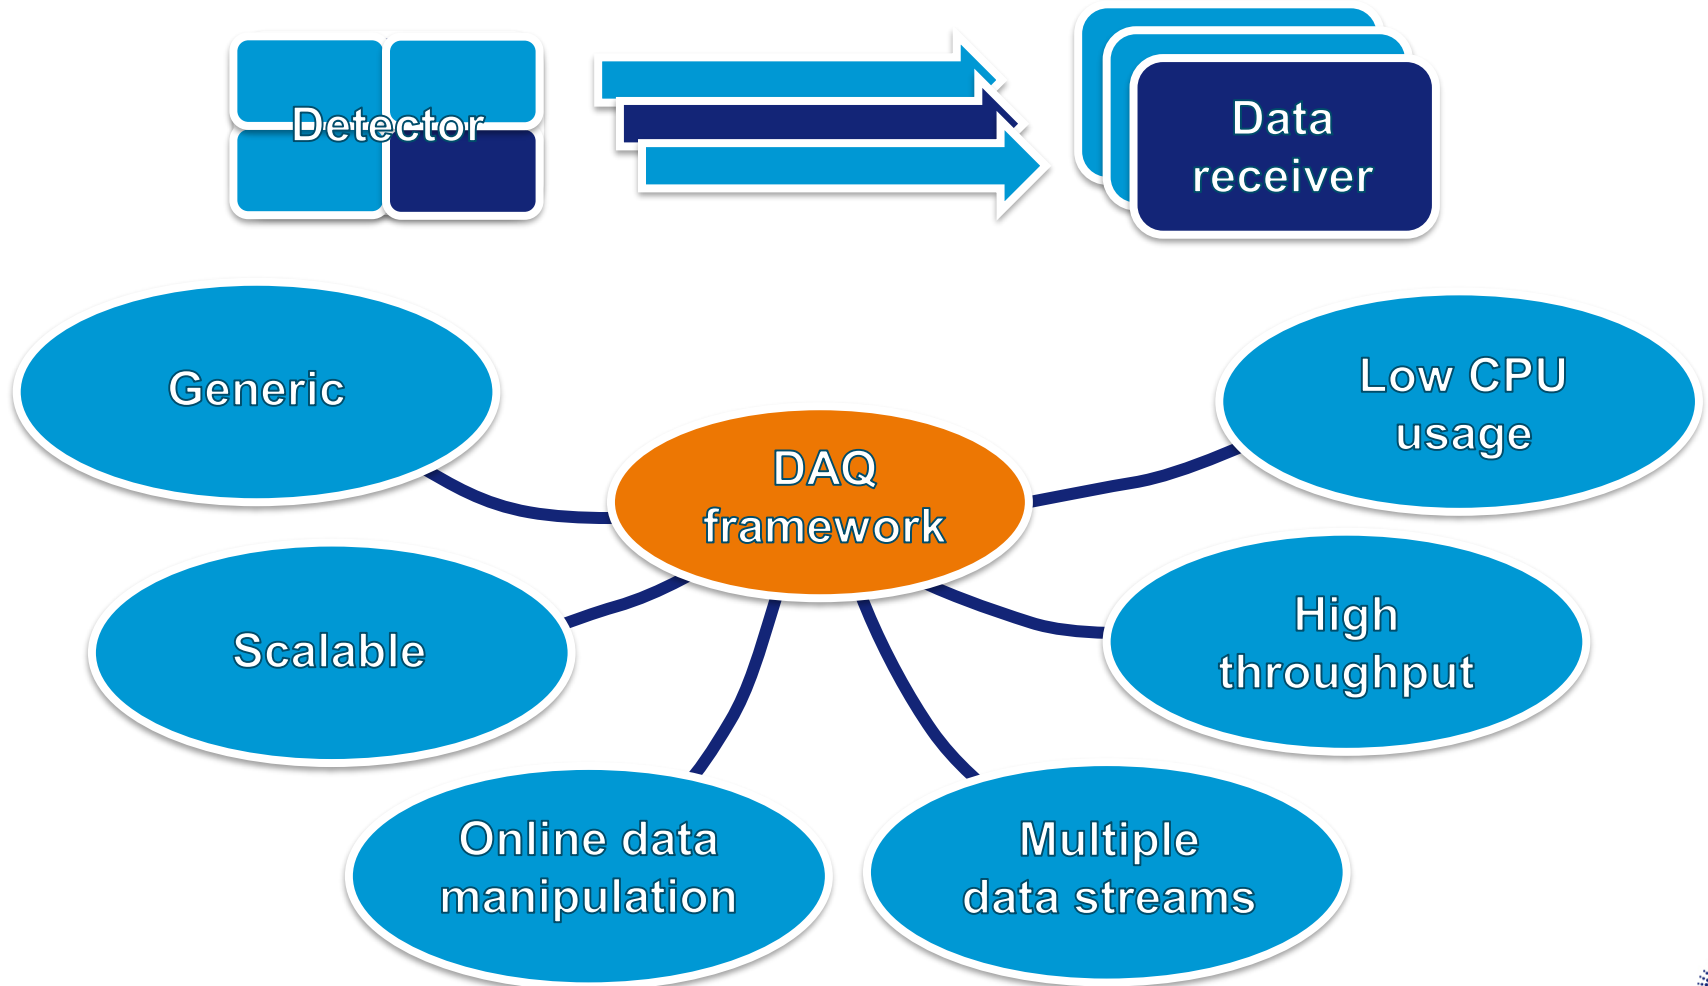
B. LIMA2: Distributed Acquisition for High Performance Detectors

- I. Motivations
- II. Software Stack
- III. Data transfers
- IV. Processing
- V. Plugins

Conclusion & perspectives



RDMA-based Acquisition System for High Performance Applications Remote Direct Memory Access



Introduction

A. RASHPA: RDMA-based Data Acquisition framework

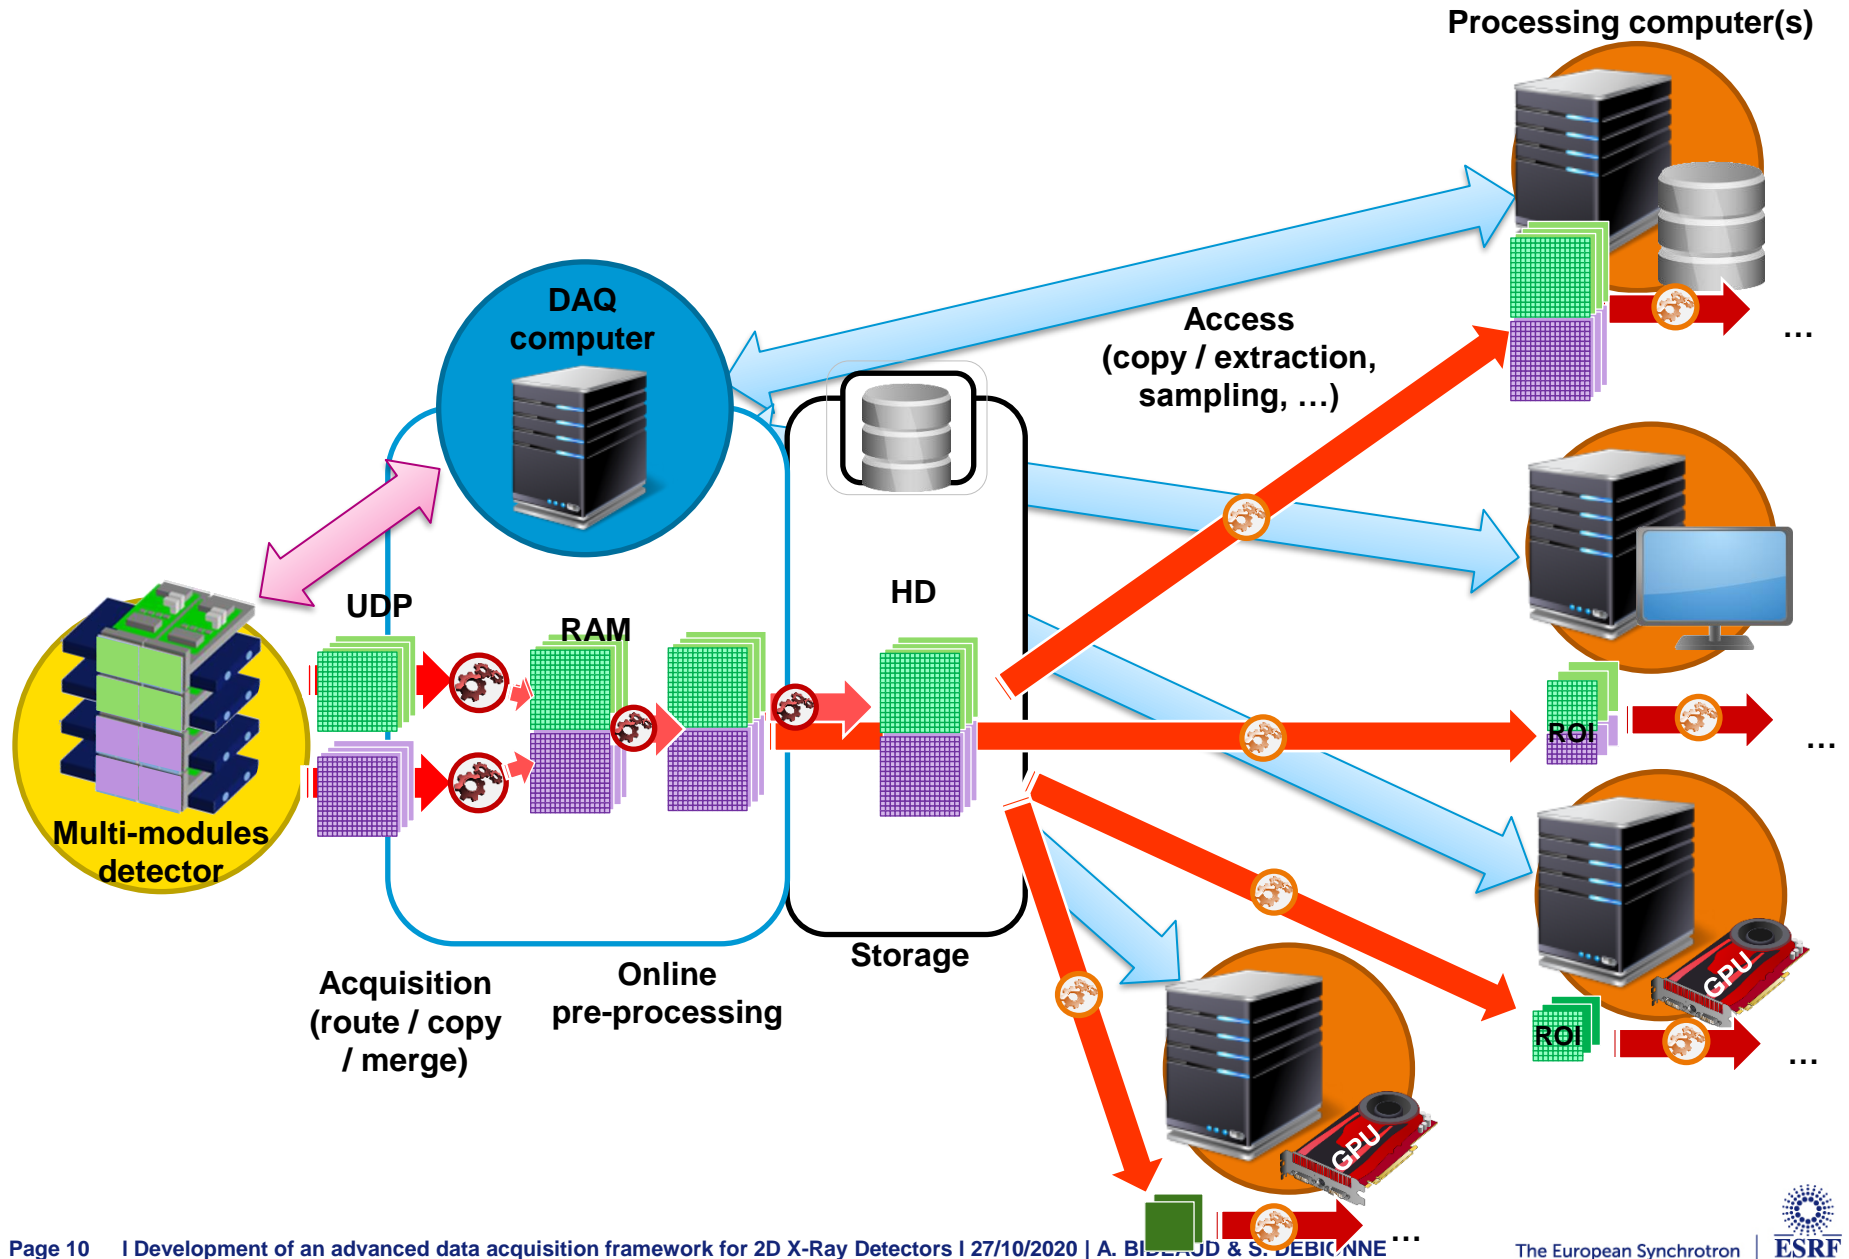
- I. Introduction
- II. Architecture**
- III. Implementation
- IV. Example

B. LIMA2: Distributed Acquisition for High Performance Detectors

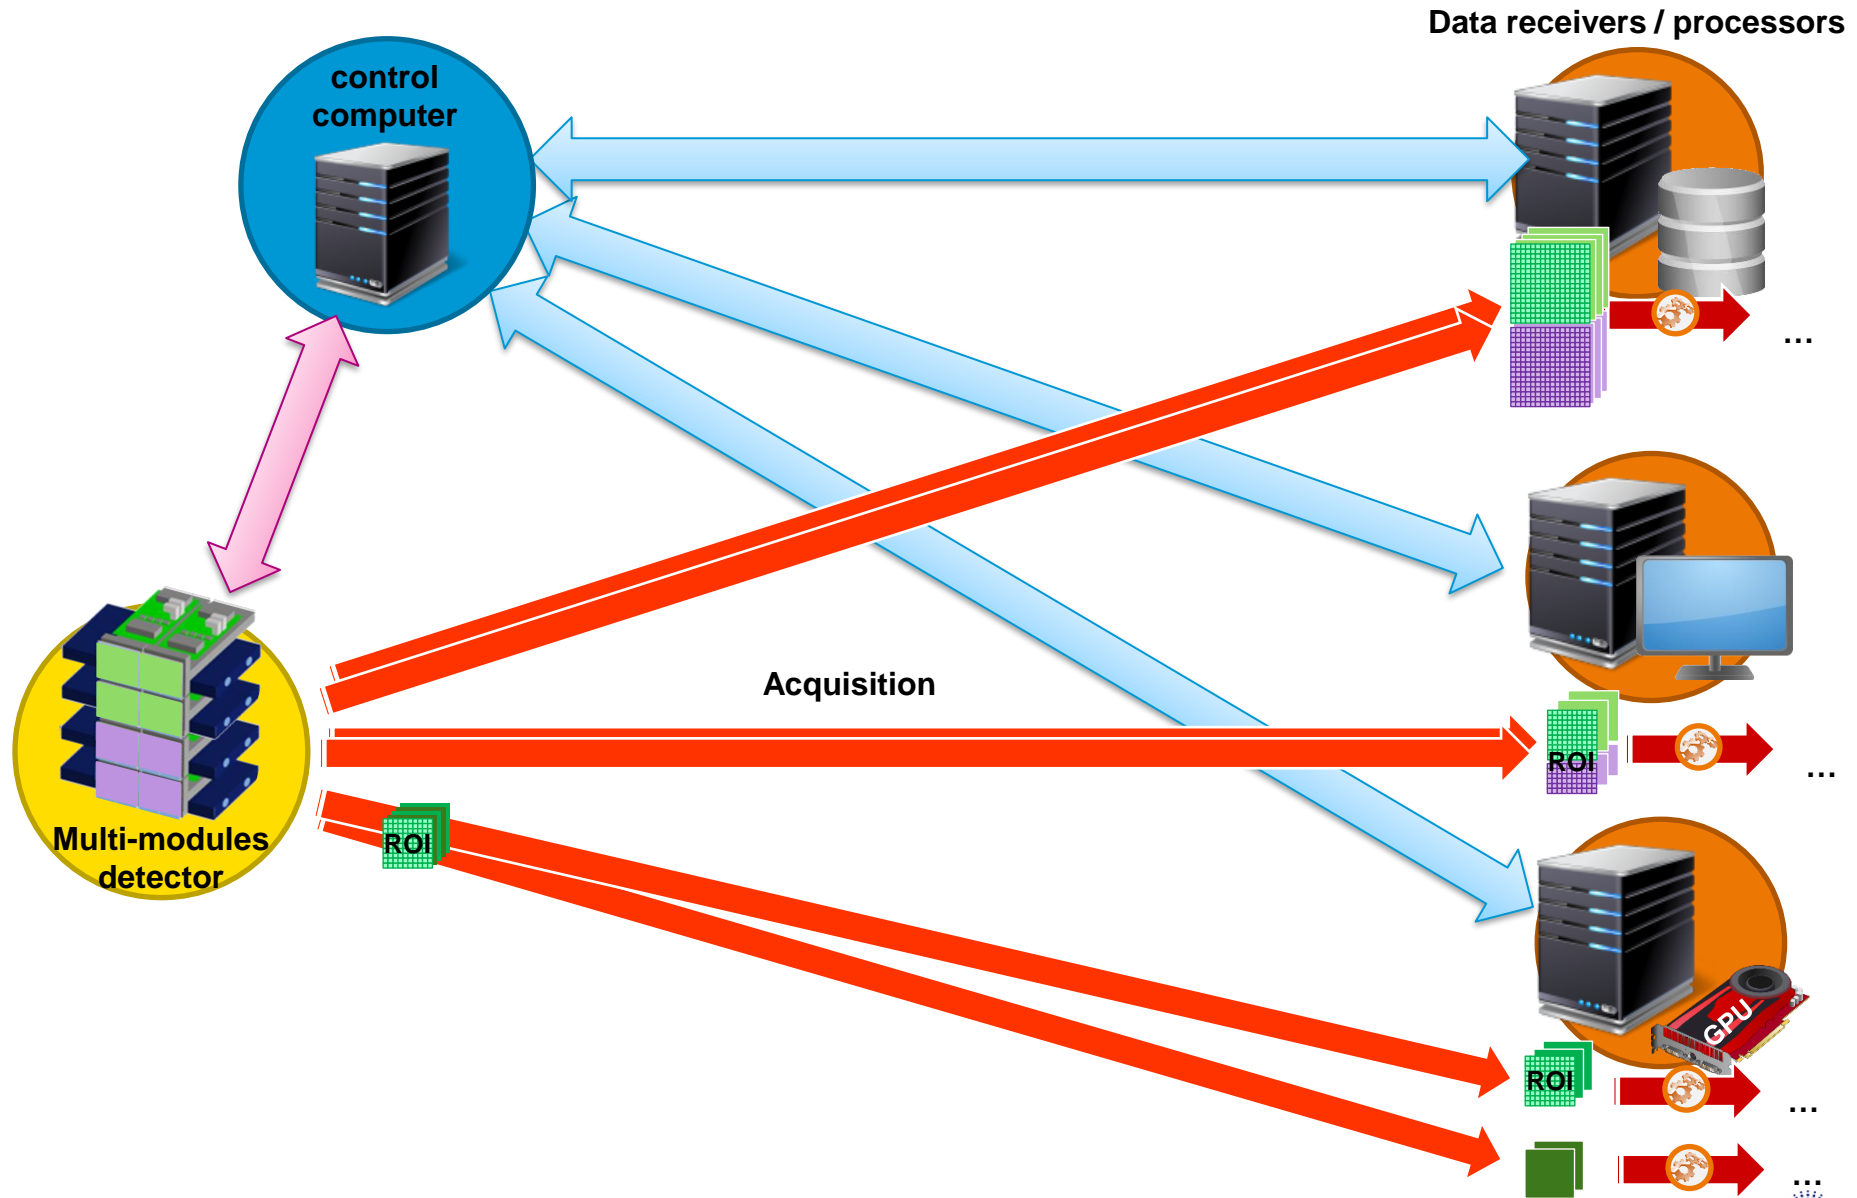
- I. Motivations
- II. Software Stack
- III. Data transfers
- IV. Processing
- V. Plugins

Conclusion & perspectives

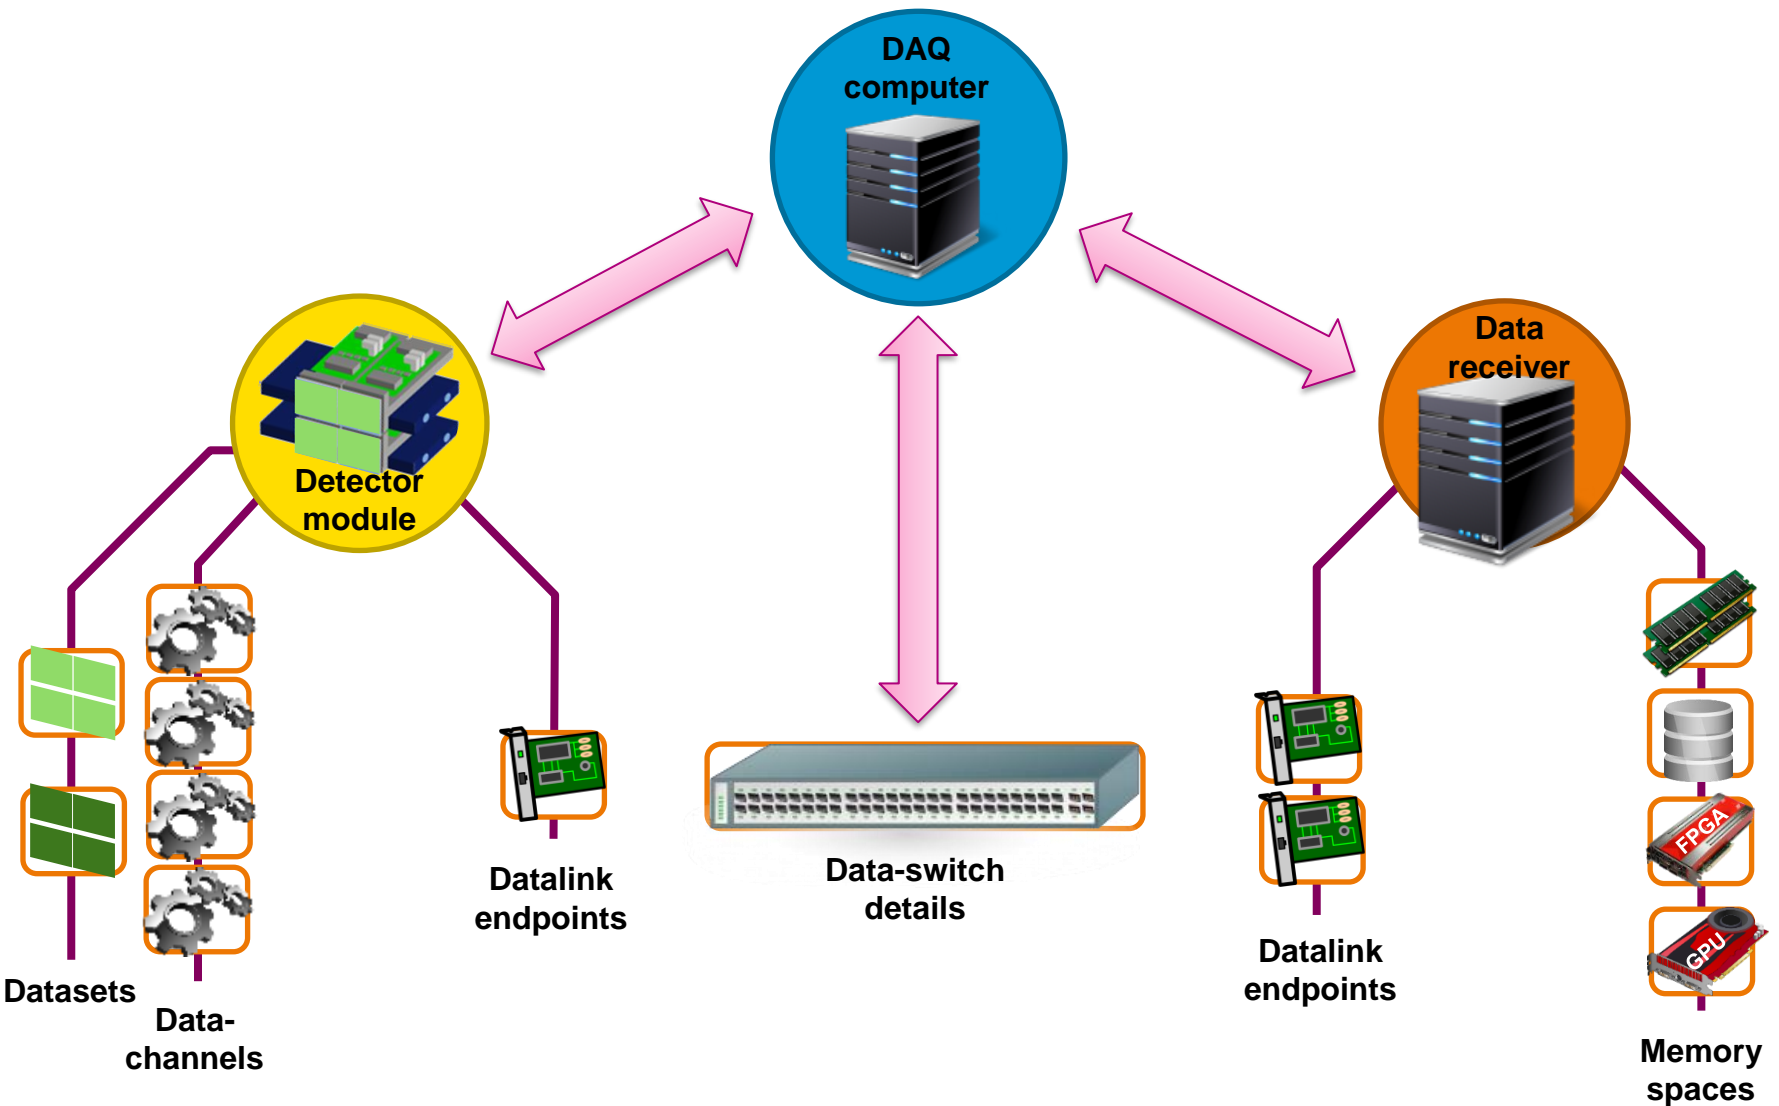
RASHPA ARCHITECTURE – CURRENT DAQ FRAMEWORK @ ESRF




RASHPA ARCHITECTURE – QUICK VIEW




RASHPA CONFIGURATION – CAPABILITIES RETRIEVAL




RASHPA CONFIGURATION – DATA FLOWS DEFINITION



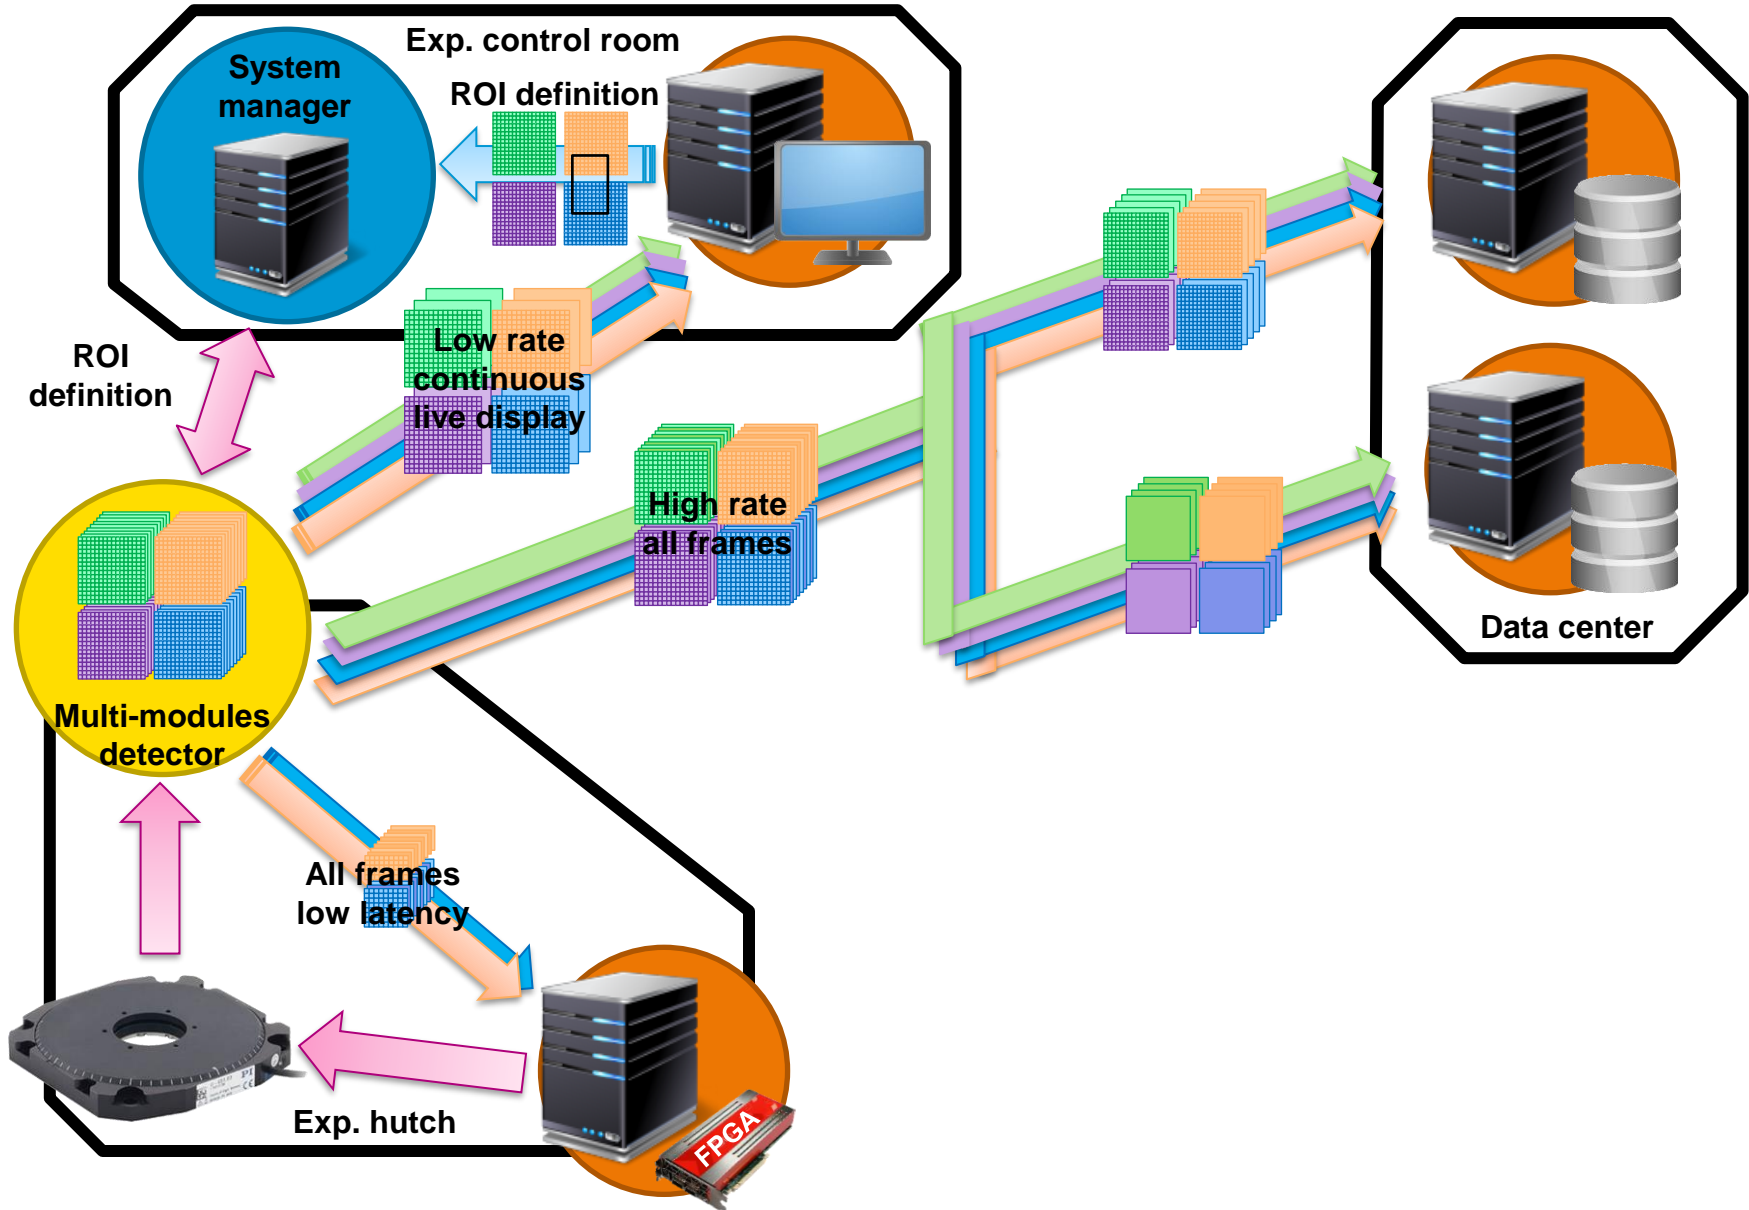
RASHPA CONFIGURATION – DATA-FLOW ROUTING



RASHPA CONFIGURATION – MULTI-NODES



RASHPA CONFIGURATION – EXAMPLE OF USE



Introduction

A. RASHPA: RDMA-based Data Acquisition framework


- I. Introduction
- II. Architecture
- III. Implementation**
- IV. Example

B. LIMA2: Distributed Acquisition for High Performance Detectors

- I. Motivations
- II. Software Stack
- III. Data transfers
- IV. Processing
- V. Plugins

Conclusion & perspectives

Detector: SMARTPIX



Client application: LIMA2



SMARTPIX RASHPA capabilities:

- 8 simultaneous data streams.
- PCIe data transfer.
- Gbethernet data transfer in progress.
- Extraction, reconstruction.


Specific developments:

- generic plug-in for all RASHPA-compliant detectors

Emulated Detector

RASHPA EMUlator

REMU



Client application: LIMA2




RASHPA Capabilities:

- N simultaneous data streams.
- PCIe data transfer.
- Gbethernet data transfer in progress.
- Extraction, reconstruction and extra data manipulation possible.

Specific developments:

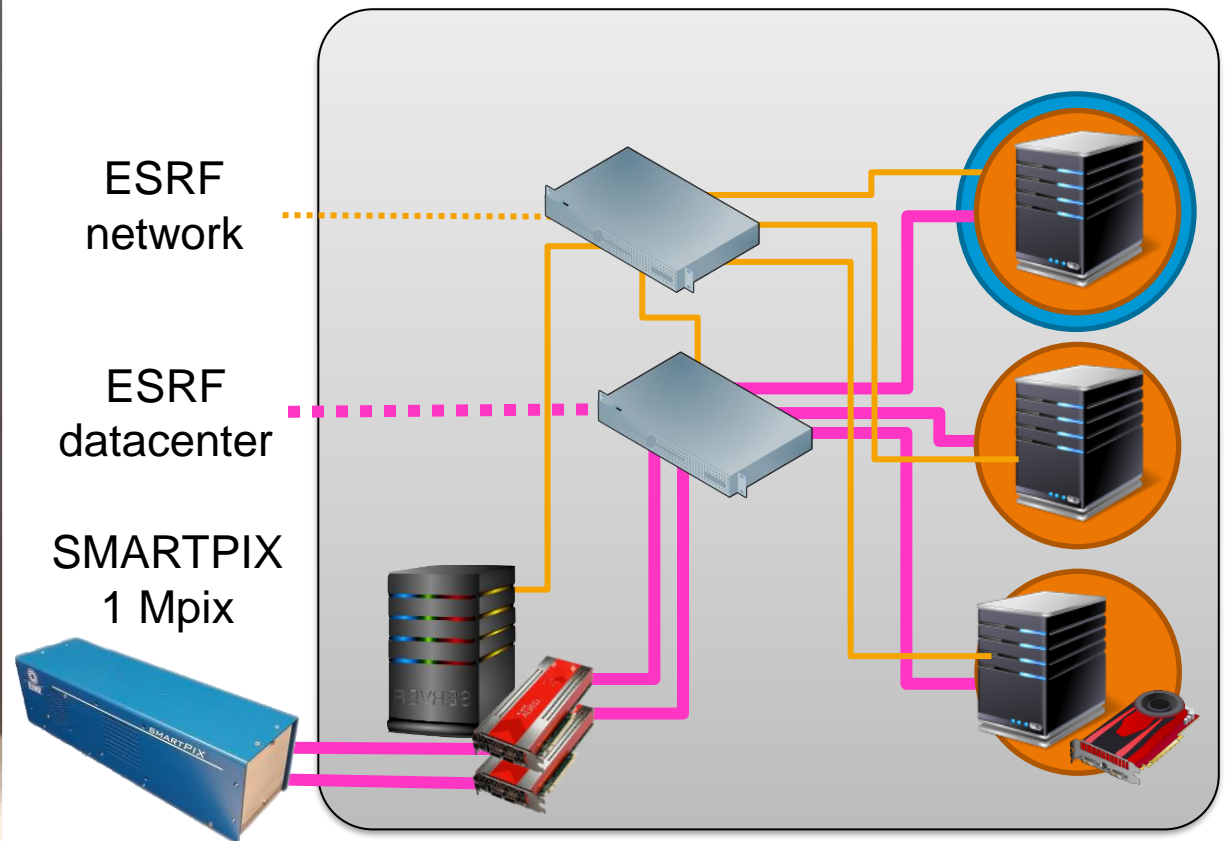
- generic plug-in for all RASHPA-compliant detectors
- REMU plug-in (for control)

RASHPA IMPLEMENTATION – HARDWARE VIEW



Goal:

Demonstrate RASHPA associated to the SMARTPIX 2Gpix-detector on a beamline @ ESRF.



Introduction


A. RASHPA: RDMA-based Data Acquisition framework

- I. Introduction
- II. Architecture
- III. Implementation
- IV. Example**

B. LIMA2: Distributed Acquisition for High Performance Detectors

- I. Motivations
- II. Software Stack
- III. Data transfers
- IV. Processing
- V. Plugins

Conclusion & perspectives



Scene to image




4-modules detector

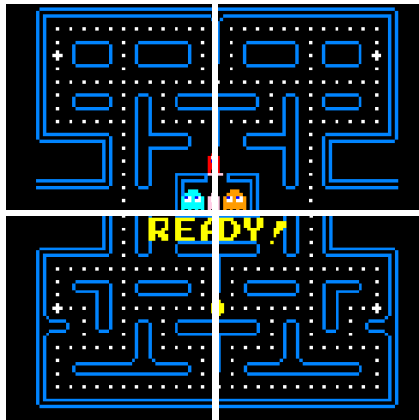

2 data-sets

10 data-channels

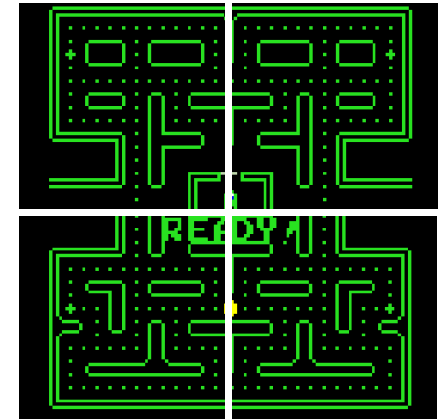


3 data receivers





Module size : 240x240
Total size : 500 x 500
2 datasets : CTR1 & CTR2



RASHPA EXAMPLE – DATA FLOW #1 (3/6)



Full frame

dataset #1

1 frame / 10

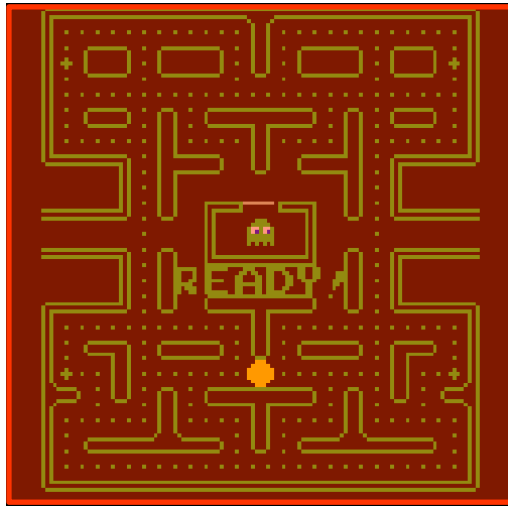
limited buffer size



```
rashpa_gdt_t *GDT = rashpaManagerCreateGDT(manager_pt, "FullDet1", "CTR1"); // GDT named "FullDet1", dataset "CTR1"
rashpaGDTSetSourceFullImage(GDT); // full detector image
rashpaGDTSetAlignment(GDT, 256, 1, 0, 0); // aligned on 256 bits, no header/footer
rashpaGDTSetSourceSlices(GDT, 0, 12, 10); // starting at frame 0, 12 frames, 1 frame / 10
rashpaGDTSetRules(GDT, 1, 0, RASHPA_CONCATENATE | RASHPA_CIRCULAR); // 1 block / group, starting at slot 0
rashpaGDTSetEvents(GDT, 1, RASHPA_EV_GRP); // event every group of data block

// 16 kiB per local buffer
ssize_t lbuff_bytes = 16*1024;
rashpa_buffer_t* buffer = rashpaManagerGetGDTBuffer(manager_pt, GDT, lbuff_bytes);
rashpaBufferCreateLocalBufferInMemSpace(buffer, receiver0_pt, "dataMSpace"); // in receiver 0, in dataMSpace memspace
```

RASHPA EXAMPLE – DATA FLOW #2 (4/6)



dataset #2

1 frame / 5




limited buffer size



```
rashpa_gdt_t *GDT = rashpaManagerCreateGDT(manager_pt, "FullDet2", "CTR2"); // GDT named "FullDet2", dataset "CTR2"
rashpaGDTSetSourceFullImage(GDT); // full detector image
rashpaGDTSetAlignment(GDT, 256, 1, 0, 0); // aligned on 256 bits, no header/footer
rashpaGDTSetSourceSlices(GDT, 0, 22, 5); // starting at frame 0, 22 frames, 1 frame / 5
rashpaGDTSetRules(GDT, 1, 0, RASHPA_CONCATENATE | RASHPA_CIRCULAR); // 1 block / group, starting at slot 0
rashpaGDTSetEvents(GDT, 1, RASHPA_EV_GRP); // event every group of data blocks
```

```
ssize_t lbuff_bytes = 16*1024; // 16 kiB per local buffer
rashpa_buffer_t* buffer = rashpaManagerGetGDTBuffer(manager_pt, GDT, lbuff_bytes);
rashpaBufferCreateLocalBufferInMemSpace(buffer, receiver0_pt, "dataMSpace"); // in receiver 0, in dataMSpace memspace
rashpaBufferCreateLocalBufferInMemSpace(buffer, receiver1_pt, "dataMSpace"); // in receiver 1, in dataMSpace memspace
```



RASHPA EXAMPLE – DATA FLOW #3 (5/6)



dataset #1
1 frame / 5




```
rashpa_gdt_t *GDT = rashpaManagerCreateGDT(manager_pt, "GhostRoom", "CTR1"); // GDT "GhostRoom", dataset "CTR1"  
rashpaGDTSetSourceImage(GDT, 190, 125, 190, 78); // ROI of (125 x 78) with offset (190, 78)  
rashpaGDTSetAlignment(GDT, 1, 1, 0, 0); // no data alignment, no header/footer  
rashpaGDTSetSourceSlices(GDT, 0, 22, 5); // starting at frame 0, 22 frames, 1 frame / 5  
rashpaGDTSetRules(GDT, 1, 0, RASHPA_CONCATENATE | RASHPA_CIRCULAR); // 1 block / group, starting at slot 0  
rashpaGDTSetEvents(GDT, 0, RASHPA_EV_ATEND); // event every data blocks
```

```
ssize_t lbuff_bytes = 10*(1024*1024); // 10 MiB per local buffer  
rashpa_buffer_t* buffer = rashpaManagerGetGDTBuffer(manager_pt, GDT, lbuff_bytes);  
rashpaBufferCreateLocalBufferInMemSpace(buffer, receiver0_pt, "dataMSpace"); // in receiver 0, in dataMSpace memspace  
rashpaBufferCreateLocalBufferInMemSpace(buffer, receiver1_pt, "dataMSpace"); // in receiver 1, in dataMSpace memspace
```

RASHPA EXAMPLE – DATA FLOW #4 (6/6)



dataset #1
all frames



```
rashpa_gdt_t *GDT = rashpaManagerCreateGDT(manager_pt, "LeftCorridor", "CTR1"); // GDT LeftCorridor, dataset CTR1
rashpaGDTSetSourceImage(GDT, 35, 89, 204, 50); // ROI of (89 x 50) with offset (35, 204)
rashpaGDTSetAlignment(GDT, 1, 1, 0, 0); // no data alignment, no header/footer
rashpaGDTSetSourceSlices(GDT, 0, 114, 1); // starting at frame 0, 114 frames, all frames
rashpaGDTSetRules(GDT, 4, 0, RASHPA_OVERWRITE | RASHPA_ONESHOT); // 4 blocks / group, starting at slot 0
rashpaGDTSetEvents(GDT, 1, RASHPA_EV_GRP | RASHPA_EV_ATEND); // event every group of data blocks
```

```
ssize_t lbuff_bytes = 10*(1024*1024); // 10 MiB per local buffer
rashpa_buffer_t* buffer = rashpaManagerGetGDTBuffer(manager_pt, GDT, lbuff_bytes);
rashpaBufferCreateLocalBufferInMemSpace(buffer, receiver0_pt, "GPUMSpace"); // in receiver 0, in dataMSpace memspace
rashpaBufferCreateLocalBufferInMemSpace(buffer, receiver1_pt, "GPUMSpace"); // in receiver 1, in dataMSpace memspace
rashpaBufferCreateLocalBufferInMemSpace(buffer, receiver0_pt, "GPUMSpace"); // in receiver 0, in dataMSpace memspace
rashpaBufferCreateLocalBufferInMemSpace(buffer, receiver1_pt, "GPUMSpace"); // in receiver 1, in dataMSpace memspace
```

Introduction

A. RASHPA: RDMA-based Data Acquisition framework

- I. Introduction
- II. Architecture
- III. Implementation
- IV. Example

B. LIMA2: Distributed Acquisition for High Performance Detectors

- I. Motivations**
- II. Software Stack**
- III. Data transfers**
- IV. Processing**
- V. Plugins**

Conclusion & perspectives

Scalable to support the ever growing data throughput of 2D detector

→ Distributed DAQ and Processing

Flexible and Open-Ended to support the most exotic experiments of our users

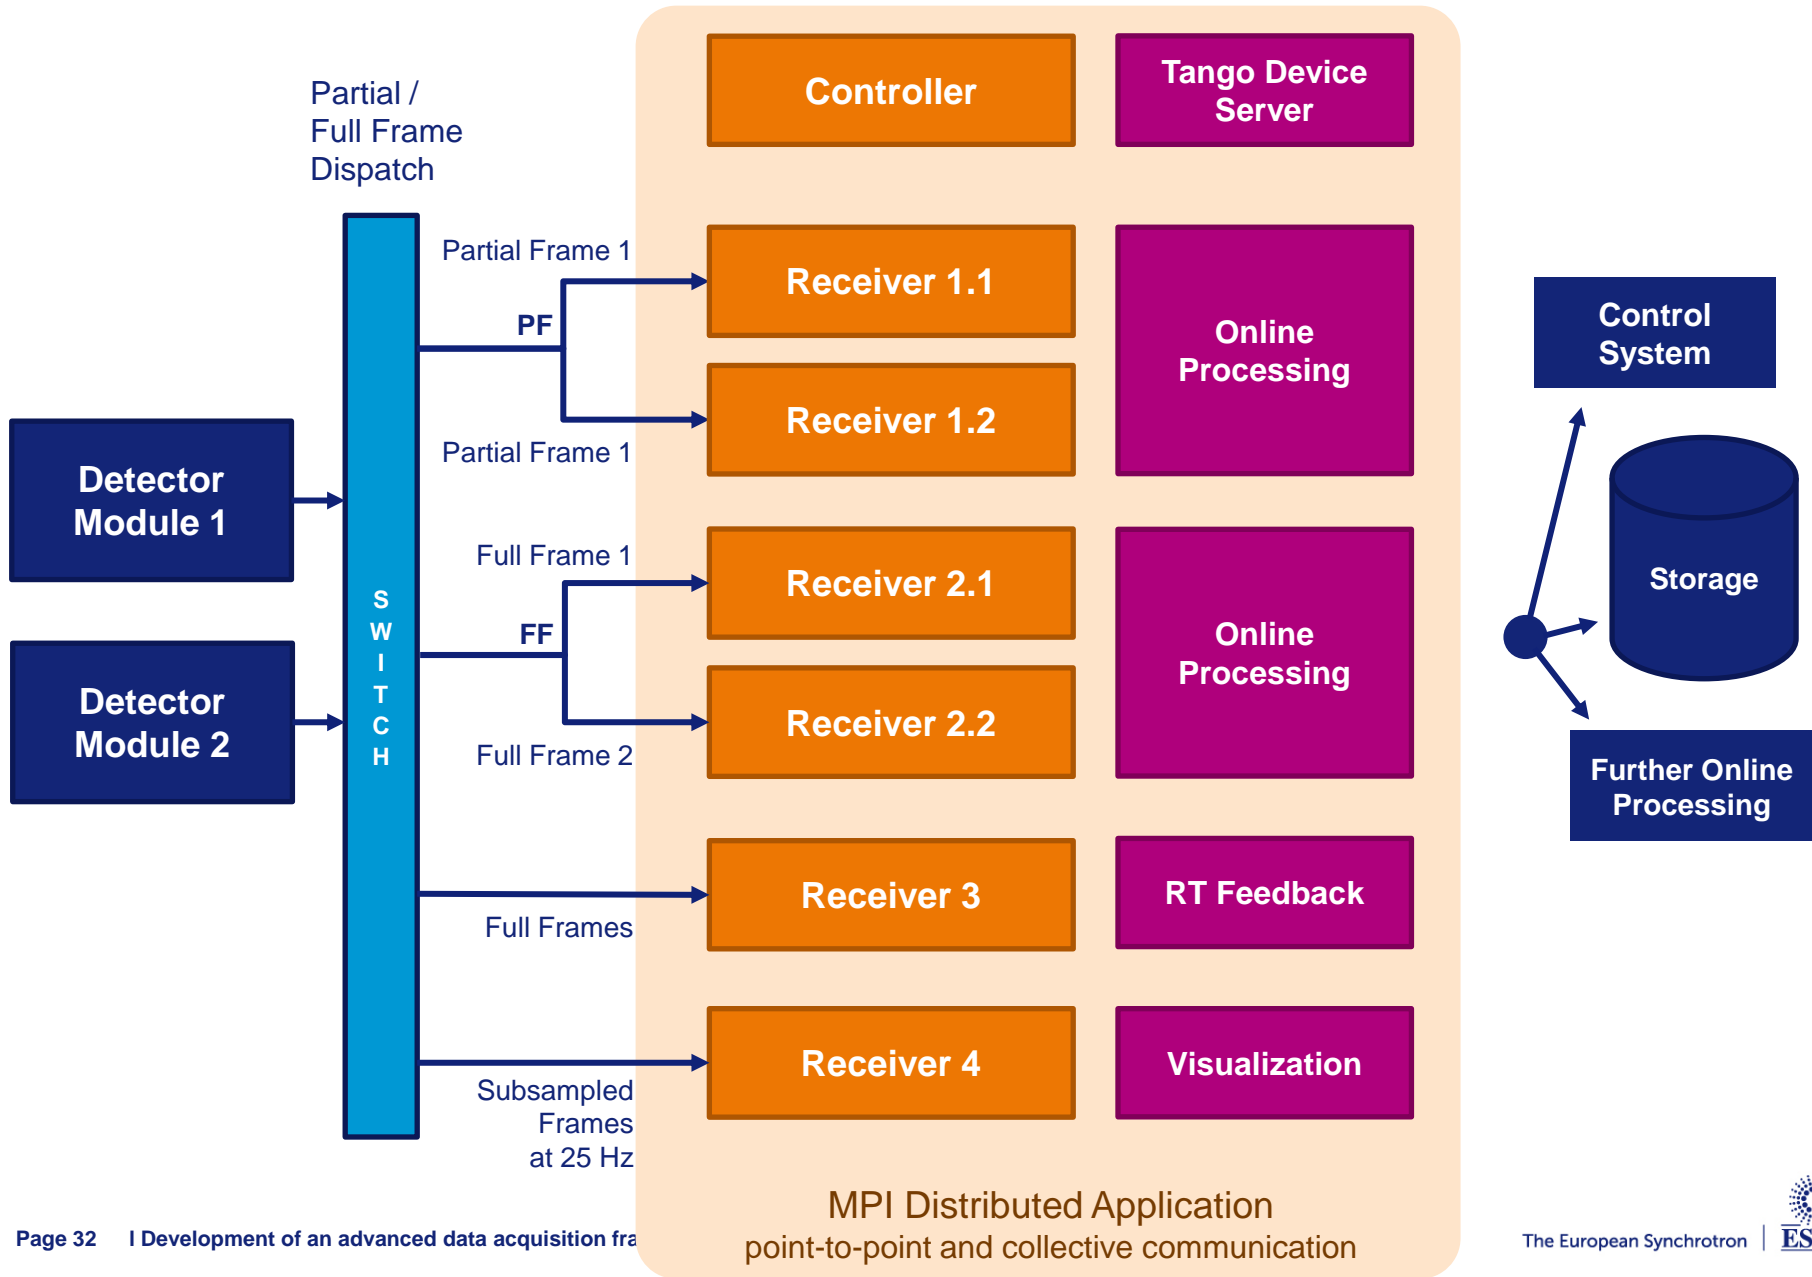
→ Decoupled DAQ, Data Transfer and Processing

Reduce dead-time between acquisitions

Memory management and resource monitoring

Support multi-band, sparse and more data types

Re-built on solid foundation (with 3rd party libraries)



MPI, the Message Passing Interface, is a standardized and portable **message-passing system**.

Features

Optimal transport (shared memory, InfiniBand) of messages

- Synchronization of parallel tasks
- Sharing and moving data
- Collective operations

Process placement (affinity to NUMA node)

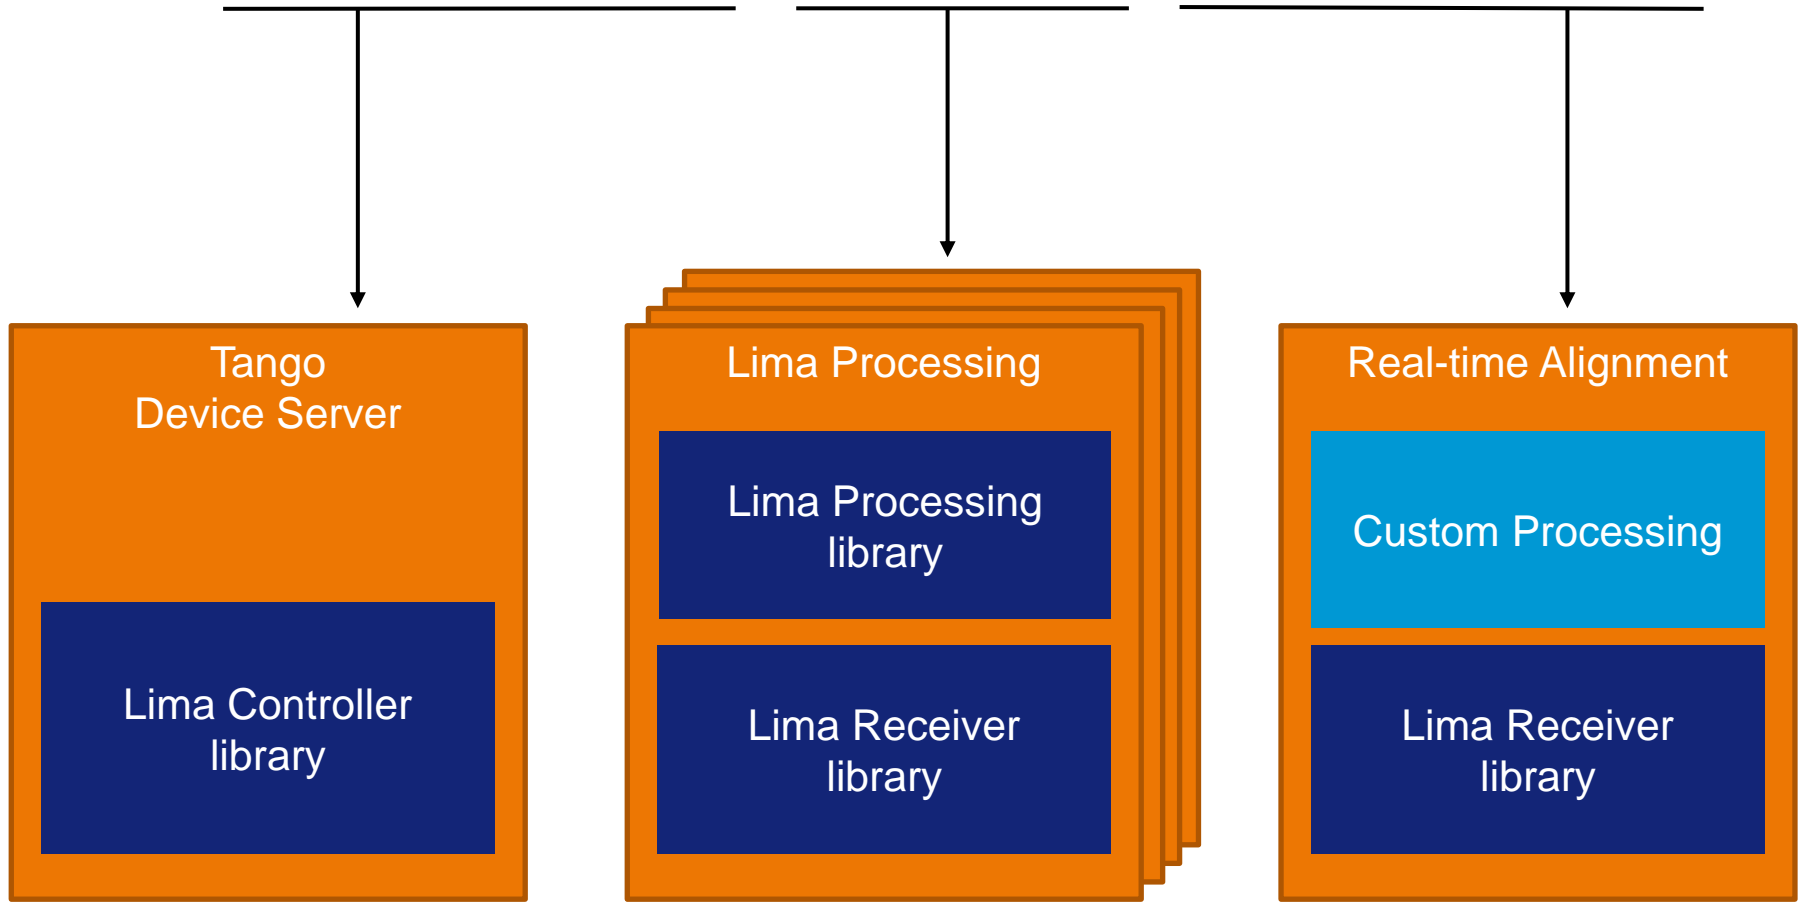
Cross platform (x86-64, POWER9), Linux, Windows...

Processes talk on *communicators* (attached to a group of process)

Communications can be *point to point* (e.g. send, recv) or *collective* (e.g. broadcast, reduce)

RUNNING MPI (MPMD) APPLICATION

```
$ mpiexec -n 1 python -m lima_tds : -n 4 lima_recv : -n 1 python -m rt_align
```



lima_recv communicator

world communicator

Introduction

A. RASHPA: RDMA-based Data Acquisition framework

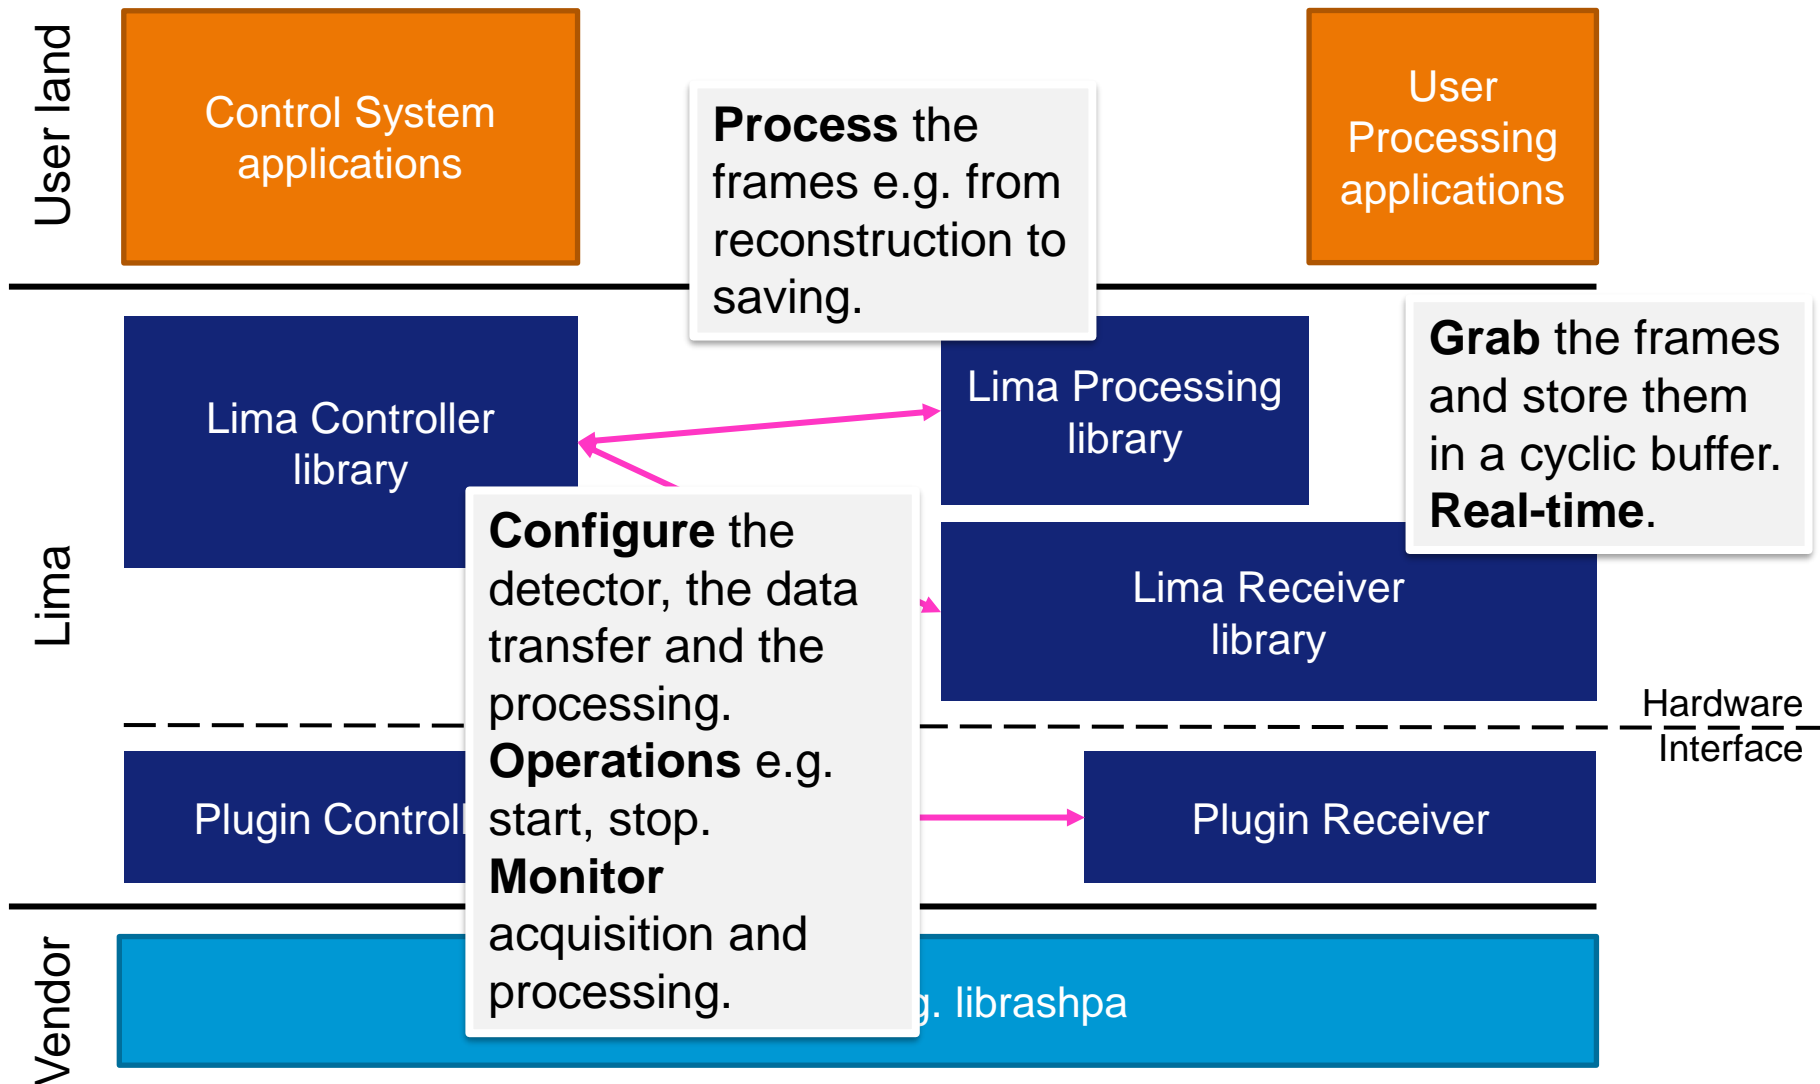
- I. Introduction
- II. Architecture
- III. Implementation
- IV. Example

B. LIMA2: Distributed Acquisition for High Performance Detectors

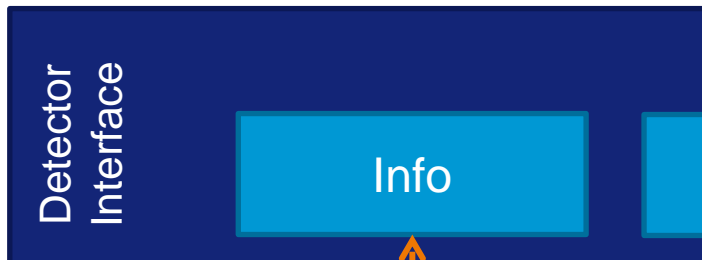
- I. Motivations
- II. Software Stack**
- III. Data transfers
- IV. Processing
- V. Plugins

Conclusion & perspectives

SOFTWARE STACK



Abstract Interface for Detectors



Configuration

`prepare(acq_params, xfer_params)`

Operations

`start()`
`stop()`
`reset()`

Monitoring

`status()`

Acquisition

`get_frames()`

Configuration

`prepare`

Operations

`start()`
`stop()`
`reset()`

Monitoring

`status()`
`nb_acq_frames()`

Detector static infos

- Name
- Pixel size
- Versions

Capabilities

- Data info (dimensions, bit depths ...)
- Binning, ROI, Flip
- Accumulation
- Shutter
- Sync / Triggers (soft, hard, gate...)
- Multiband (datasets)
- Data transfers (datachannels, reconstruction)
- Buffer (CPU/GPU, allocation, ownership..)

Introduction

A. RASHPA: RDMA-based Data Acquisition framework


- I. Introduction
- II. Architecture
- III. Implementation
- IV. Example

B. LIMA2: Distributed Acquisition for High Performance Detectors


- I. Motivations
- II. Software Stack
- III. Data transfers
- IV. Processing
- V. Plugins

Conclusion & perspectives

DATA TRANSFER - FULL FRAMES




Full frame dispatch
(round robin)



Full frame chunk dispatch
(round robin)

For feeding a GPU
For frame accumulation



Typical data transfer for multi-module detector without Rashpa

Requires a reconstruction, depending on the use case, either

- At the beginning of the processing using collective « gather » operation - where each receivers gather the full frame in turn
- When saving to file assembling with an HDF5 Virtual Dataset

Introduction

A. RASHPA: RDMA-based Data Acquisition framework

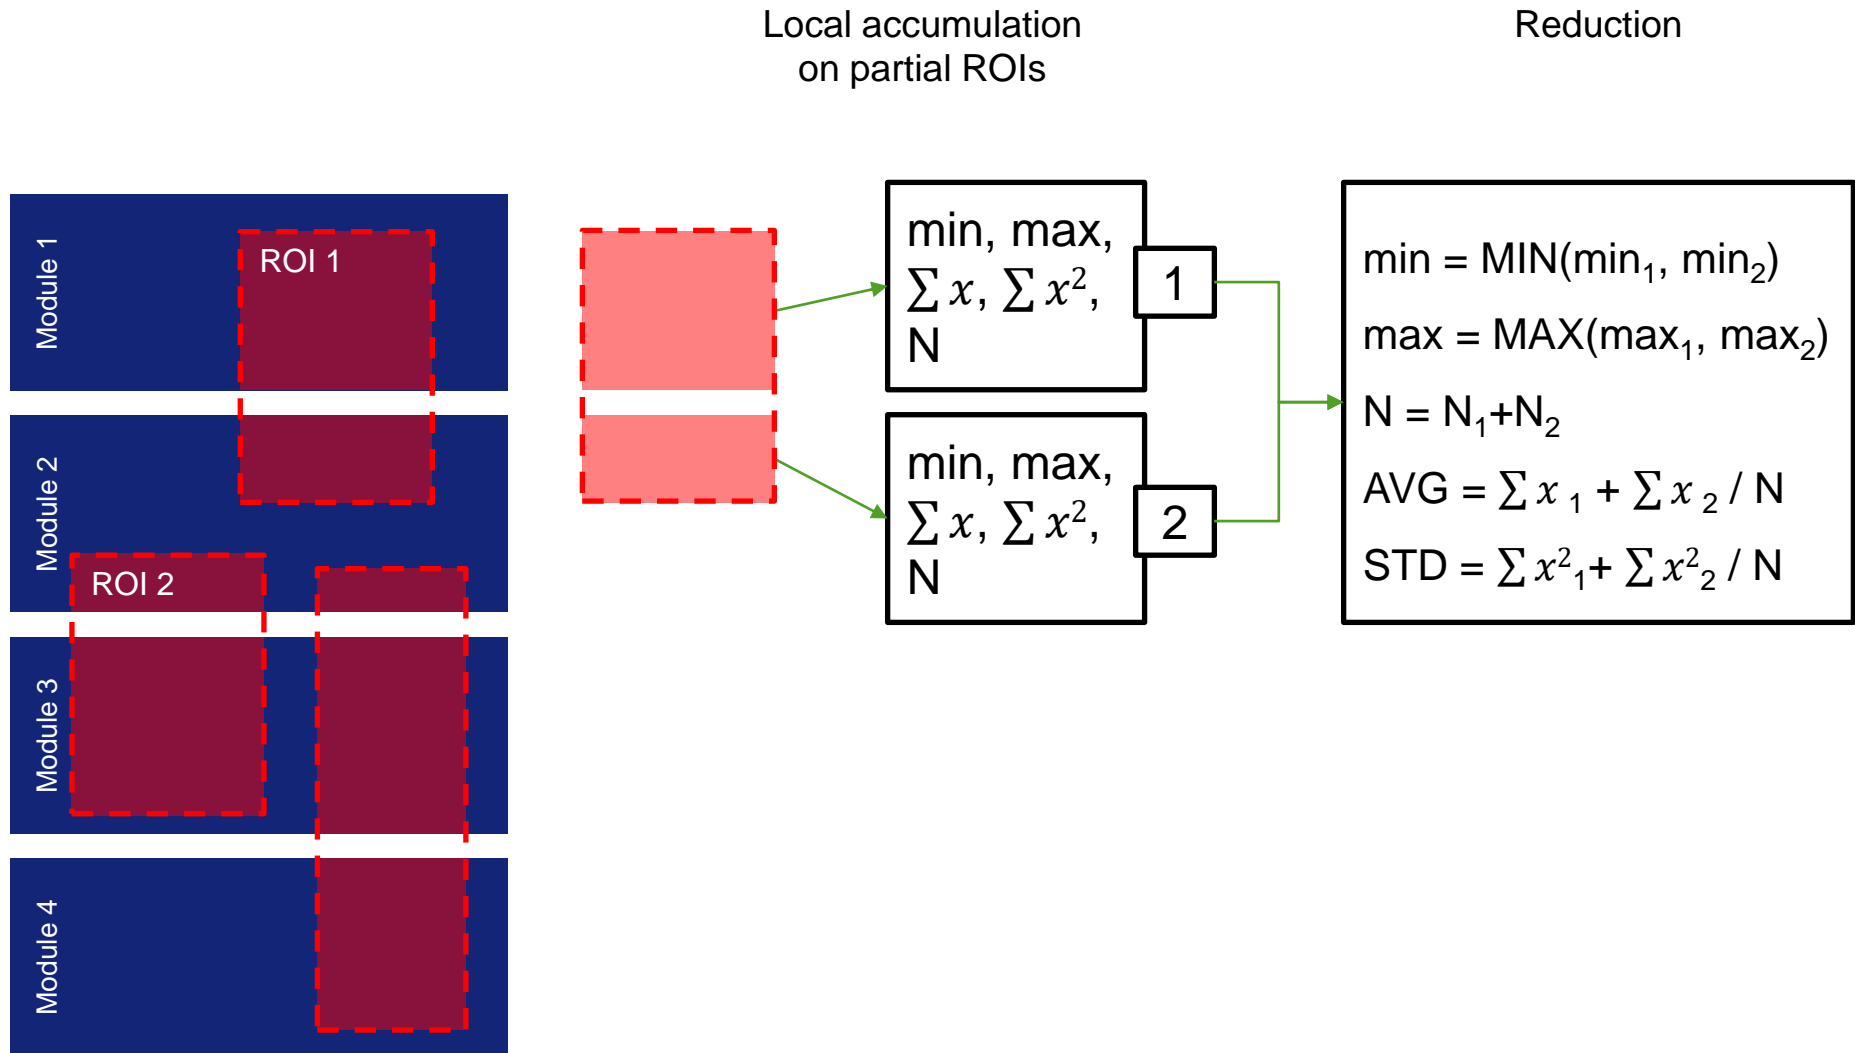
- I. Introduction
- II. Architecture
- III. Implementation
- IV. Example

B. LIMA2: Distributed Acquisition for High Performance Detectors

- I. Motivations
- II. Software Stack
- III. Data transfers
- IV. Processing**
- V. Plugins

Conclusion & perspectives

COOPERATIVE PARALLEL ALGORITHMS - ROI COUNTERS

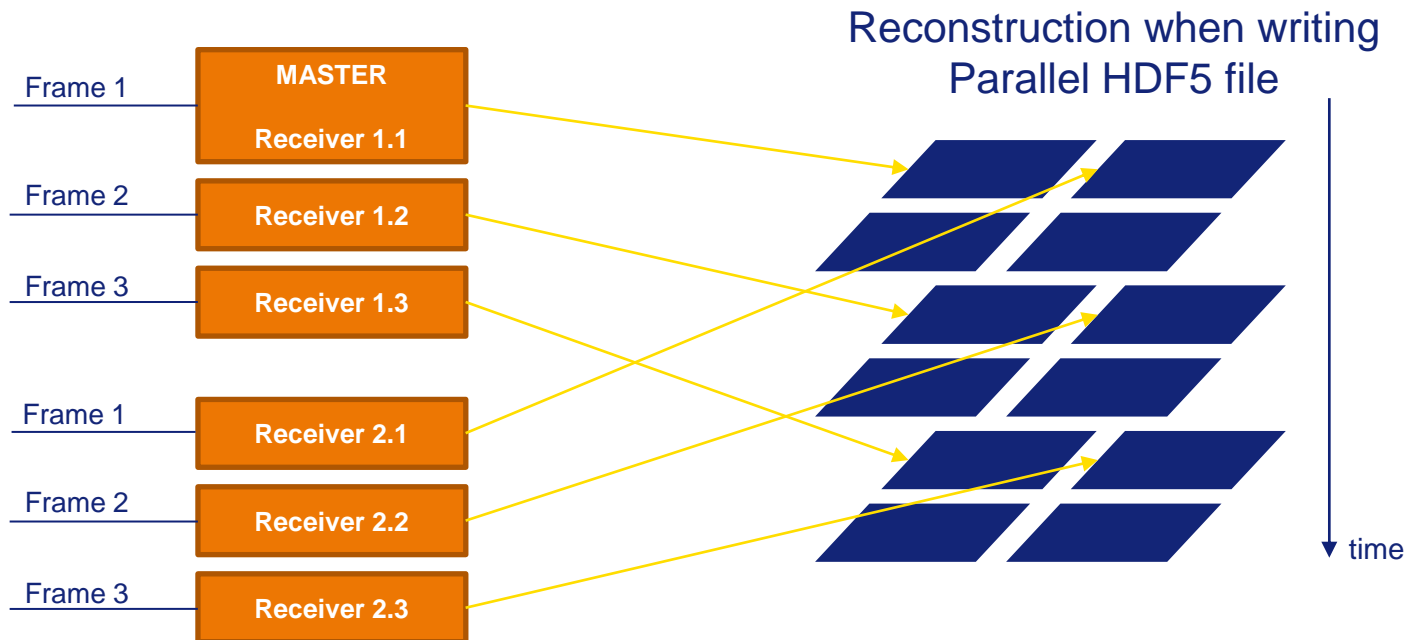


At the program level

- Concurrent reads or writes from multiple processes to a common file
- Build large blocks, so that reads/writes in I/O system will be large
 - Requests from different processes may be merged together
 - Particularly effective when the accesses of different processes are non-contiguous and interleaved


At the system level

- A parallel file system and hardware that support such concurrent access (e.g. GPFS)



Data flow programming paradigm with TBB Flow Graph library

Expressive, Flexible, Fast and scalable



TBB nodes	Description
checkpoint_node	Count frames and log idx / time
io_hdf5_node	Save frames to HDF5 format
roi_flip_rot_node	Apply geometric transformations
reconstruction_node	Reconstruct frames according to layout
roi_counters_node	Get statistics on ROIs
background_sub_node	Background subtraction
flatfield_node	Flatfield correction
mask_node	Mask defective pixels
python_node	Delegate processing to a Python interpreter
mpi_node	Delegate processing to another MPI process

Introduction

A. RASHPA: RDMA-based Data Acquisition framework

- I. Introduction
- II. Architecture
- III. Implementation
- IV. Example

B. LIMA2: Distributed Acquisition for High Performance Detectors

- I. Motivations
- II. Software Stack
- III. Data transfers
- IV. Processing
- V. **Plugins**

Conclusion & perspectives

On the Detector side

Implementing new **Detector Plugins**.

- SmartPix
- PSI Detectors (starting with JungFrau)
- REMU / Simulator

On the Processing side

Implementing new **Processing Nodes** (including in Python)

Implementing new **Processing Graphs**

Implementing new **Processing Systems** on top of the receiver API

Introduction


A. RASHPA: RDMA-based Data AcQuisition framework

- I. Introduction
- II. Architecture
- III. Implementation
- IV. Example

B. LIMA2: Distributed Acquisition for High Performance Detectors

- I. Motivations
- II. Software Stack
- III. Data transfers
- IV. Processing
- V. Plugins

Conclusion & perspectives



Lima2 is **Work in Progress**

Open-source, available on ESRF's Gitlab instance

Open to **Cooperation** and Contribution

Integration of SmartPix and PSI JungFrau for Q1-Q2 2021

More data transfer and processing options to be implemented

Rashpa **v1.0 release** foreseen for Q1 2021

Main DAQ framework for the SmartPix detector

Support within Lima2 under development

Demonstration on a beam-line foreseen in 2021

Thank you !

Development of an advanced data acquisition framework for 2D X-Ray Detectors

RASHPA team:

Aurélien BIDEAUD, Pablo FAJARDO, Nicolas JANVIER,
Wassim MANSOUR, Raphaël PONSARD

LIMA2 team:

Laurent CLAUSTRE, Samuel DEBIONNE, Alejandro HOMS