Karlsruhe Institute of Technology

24th IEEE Real Time Conference
ICISE, Quy Nhon, Vietnam
Apr 22 – 26, 2024

# Cross-Chip Partial Reconfiguration for the Initialisation of Heterogeneous Systems

**Marvin Fuchs**, Hendrik Krause, Timo Muscheid, Lukas Scheller, Luis E. Ardila-Perez, and Oliver Sander

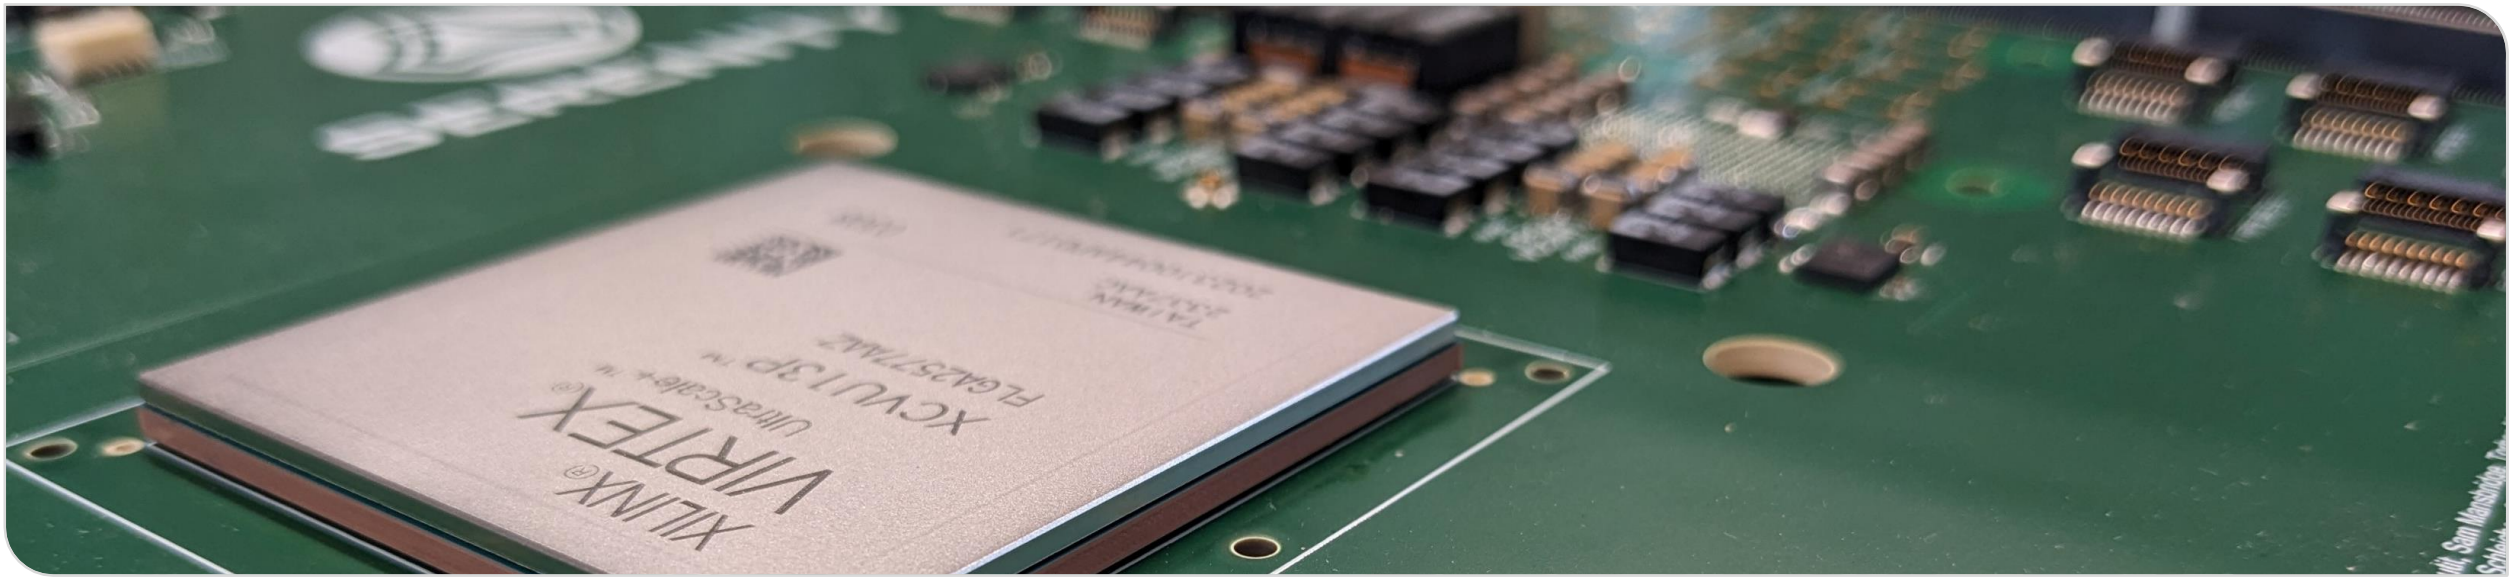

# Heterogeneous Systems

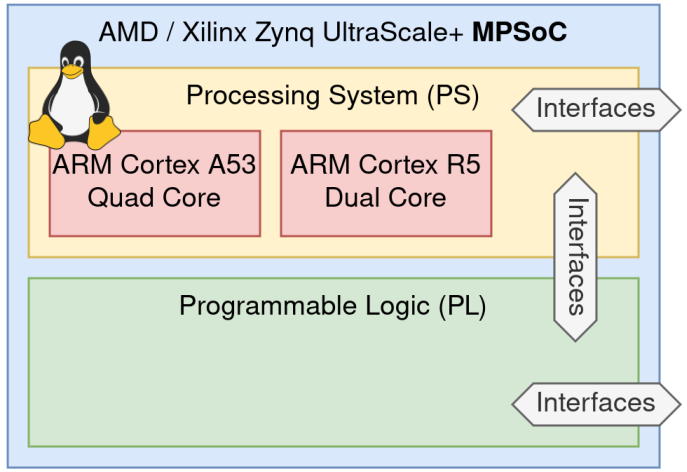


Heterogeneous device

## Definition for this talk:

- System composed of various devices from AMD / Xilinx
  - MPSoCs
  - RFSoCs
  - FPGAs

Multiple such hetero-geneous systems developed at CERN:

Global Common Modul ATCA Card

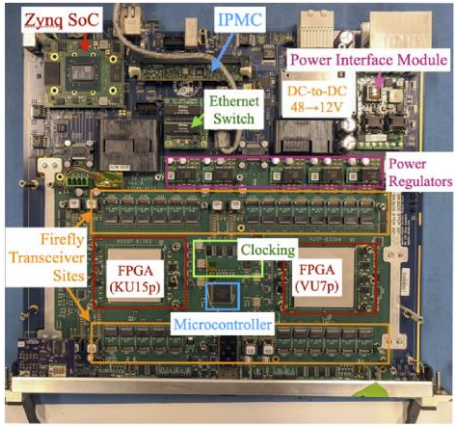


Apollo ATCA Card

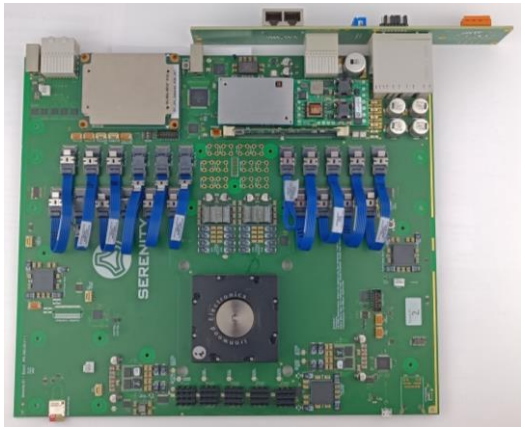

Serenity S ATCA Card

...

# Heterogeneous Systems

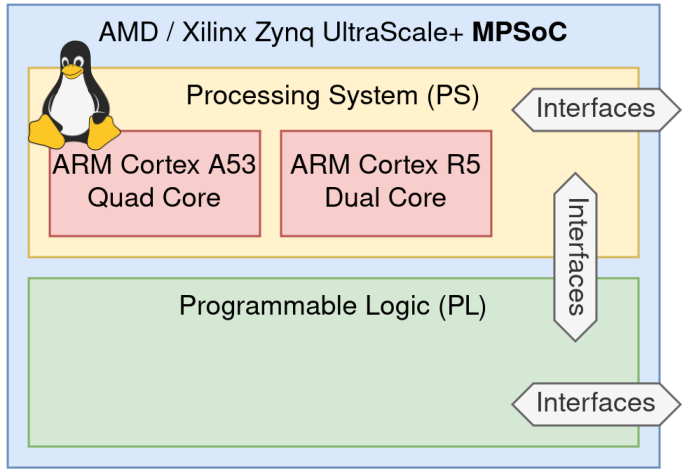


Heterogeneous device

## Definition for this talk:

- System composed of various devices from AMD / Xilinx
  - MPSoCs
  - RFSoCs
  - FPGAs

Multiple such hetero-geneous systems developed at CERN:

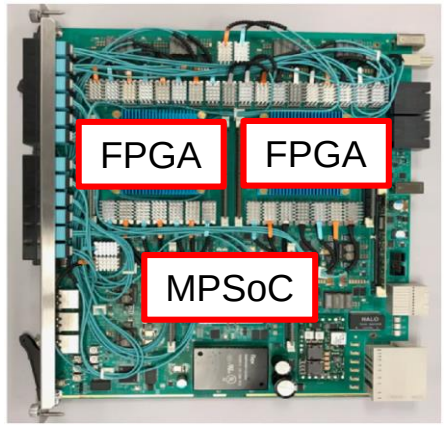


Global Common Modul ATCA Card

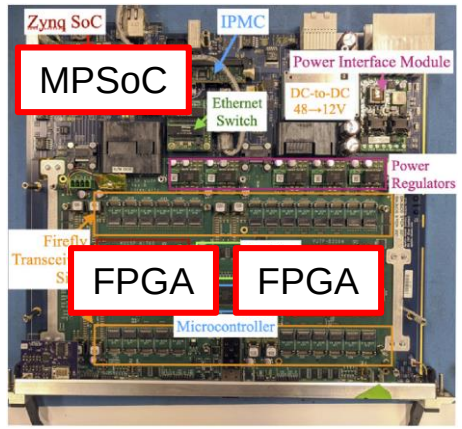


Apollo ATCA Card

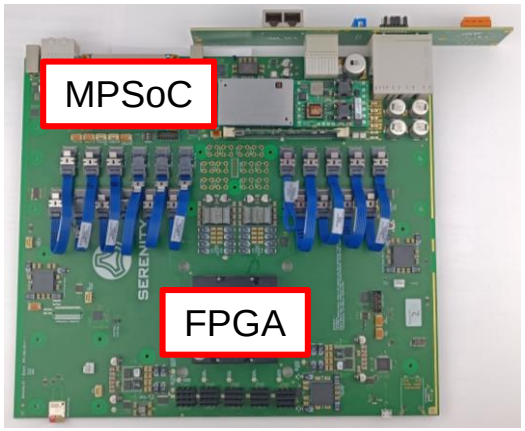


Serenity S ATCA Card

...

# Heterogeneous Systems

QiController:

- RFSoC based Qubit control system
- Modular PL firmware

Next step:

- Scaling to control more Qubits

➜ Heterogeneous system

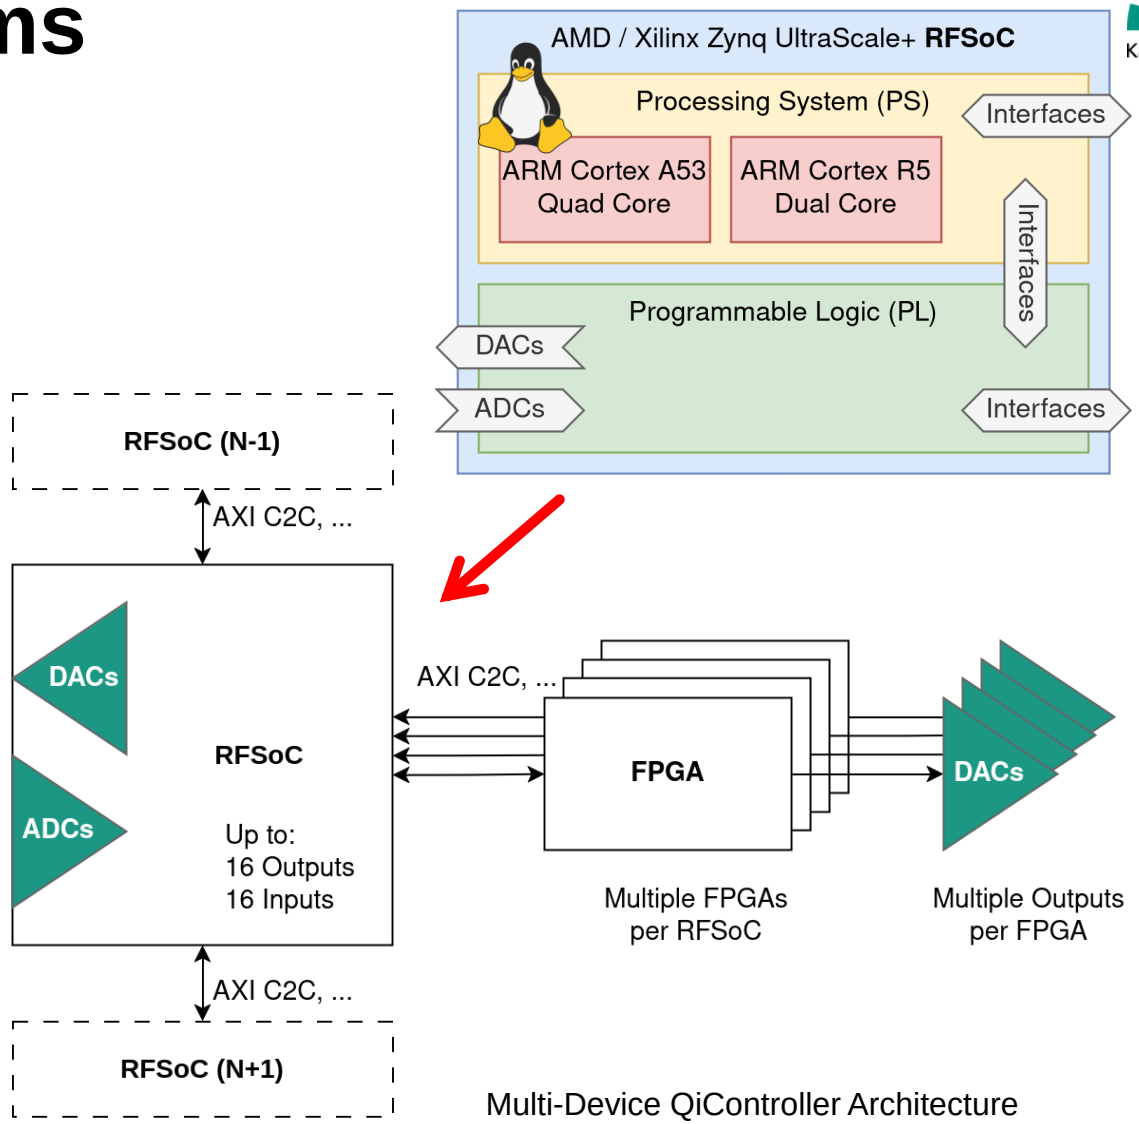


Multi-Device QiController Architecture

# Initialisation of Heterogeneous Systems

Karlsruhe Institute of Technology

Conventional Approach:

- Independent initialisation of devices
- Assembly to overall system in late boot stage

Drawbacks:

- Hard to manage
- Hard to update

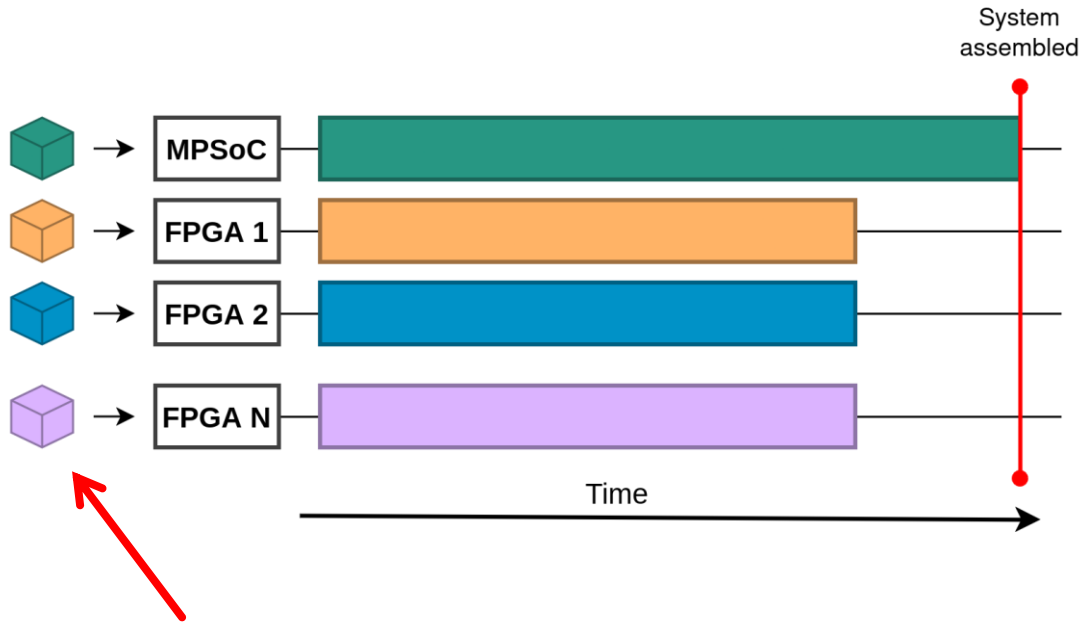


Initialisation data stored in separate places

# Initialisation of Heterogeneous Systems

Proposed Approach:

- Early formation of overall system
- Boot as one system

Advantages:

- System manageable by MPSoC / RFSoC
- Centrally controlled updates
- Scalable approach

**But how to implement?**

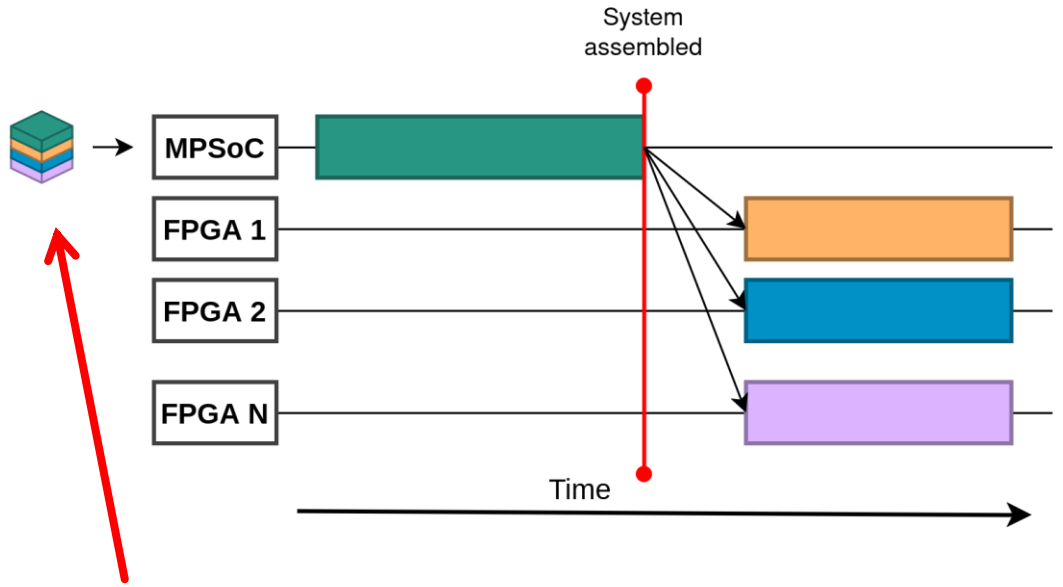


All initialisation data stored in the same place

# Initialising an FPGA from an MPSoC

- Multiple approaches
  - JTAG
  - PCIe
  - AXI Chip2Chip

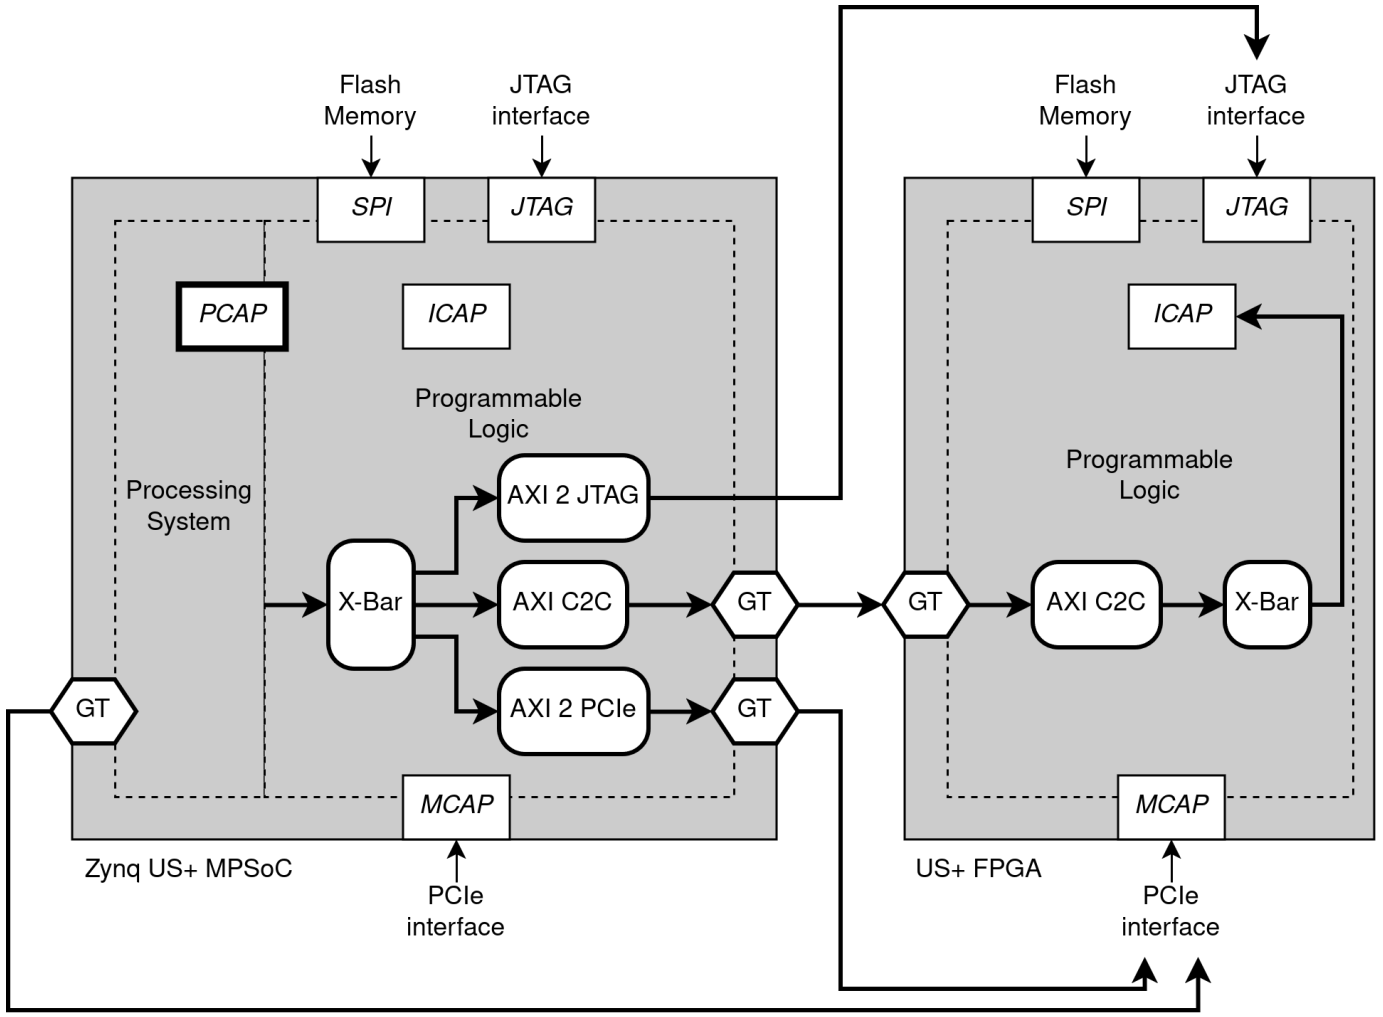

# Initialising an FPGA from an MPSoC

- Multiple approaches
  - JTAG
  - PCIe
  - AXI Chip2Chip

**Our Choice:** AXI Chip2Chip

- Reuse of existing connection
- ICAP is the fastest interface
- Requires generic pre-initialisation from flash

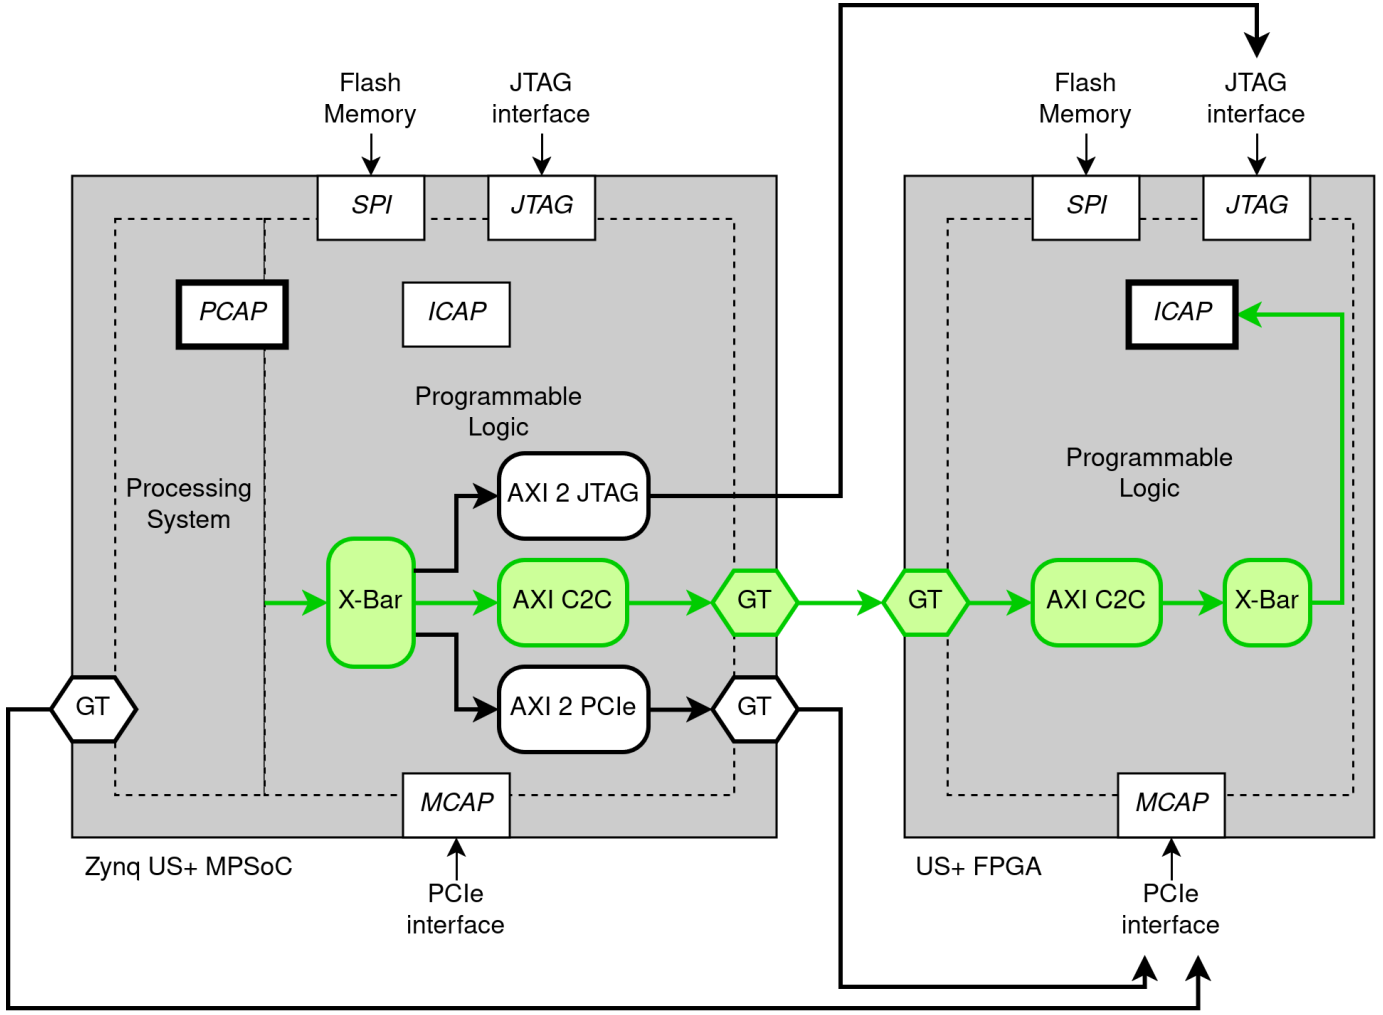

# System Start-Up

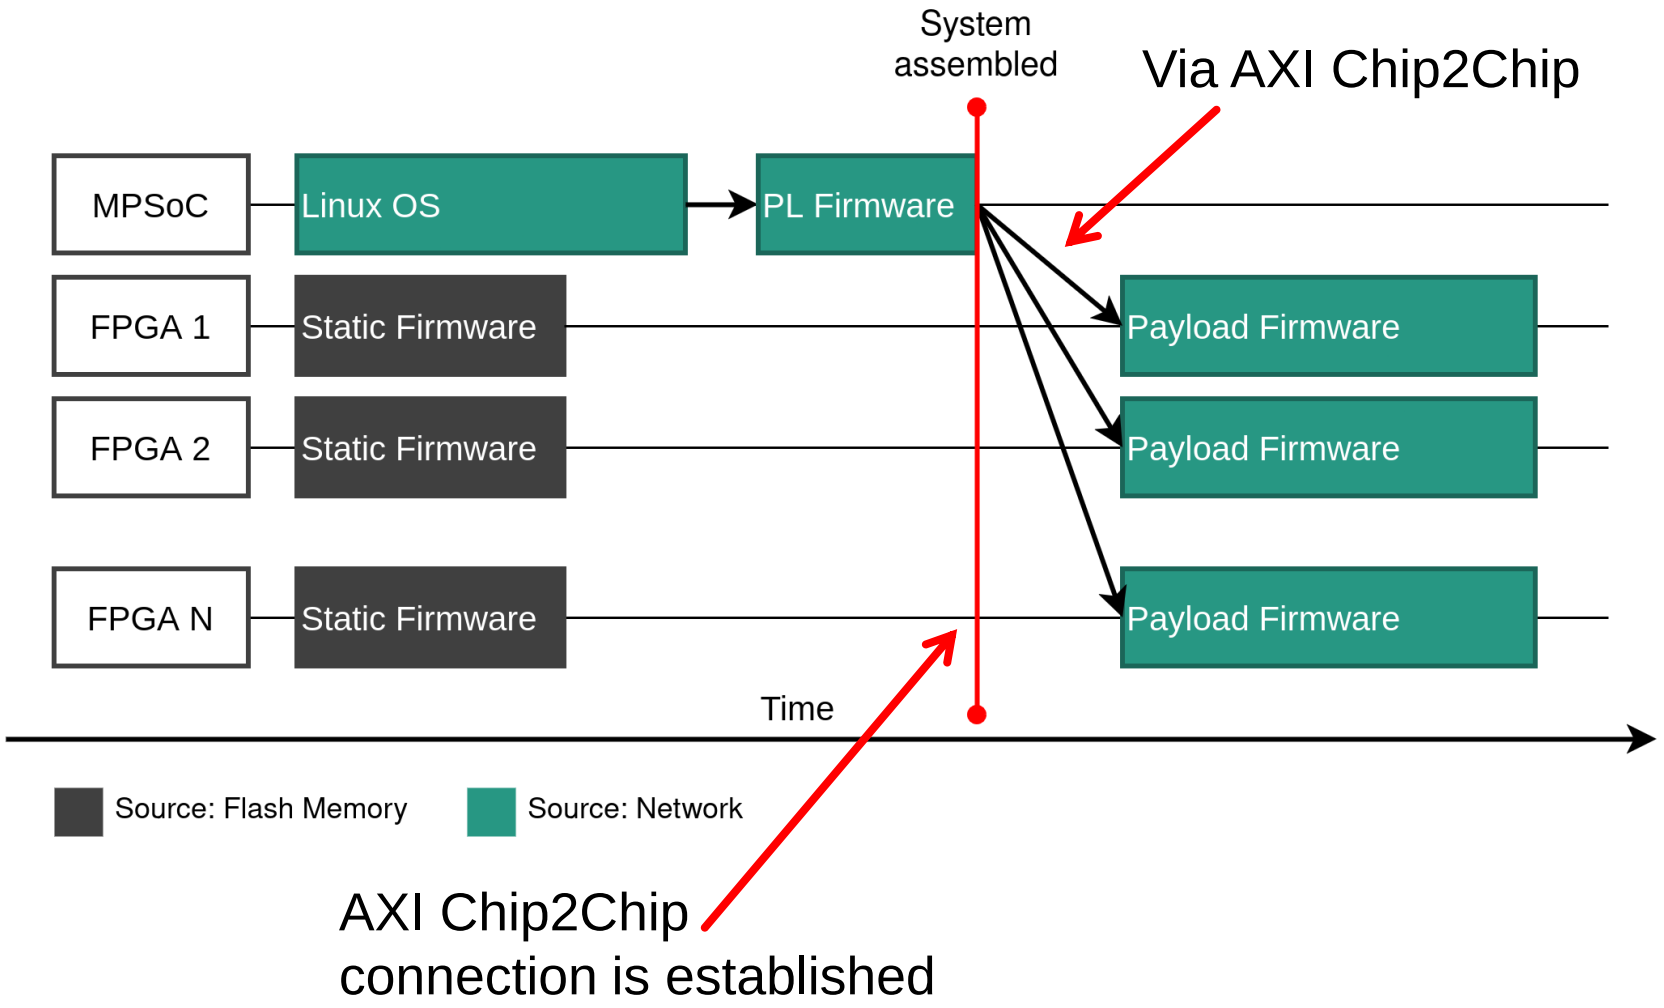


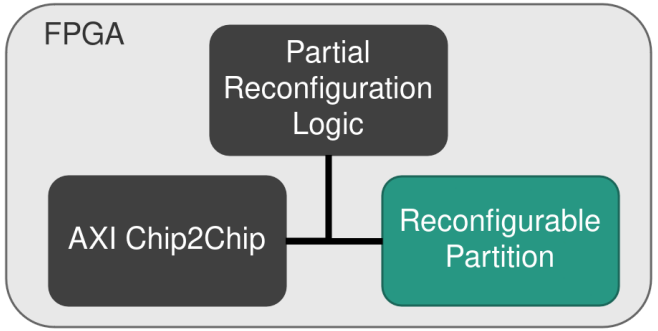

# Three Test Setups

HWICAP IP Core
Data transfer: **AXI4-Lite**
Max. burst length: **1**

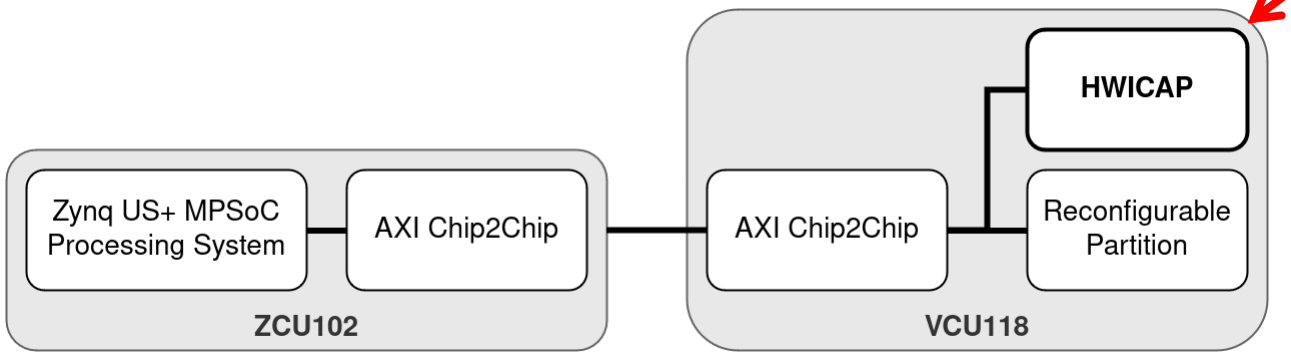


**HWICAP Chip2Chip Setup**

HBICAP IP Core
Data transfer: **AXI4**
Max. burst length: **256**

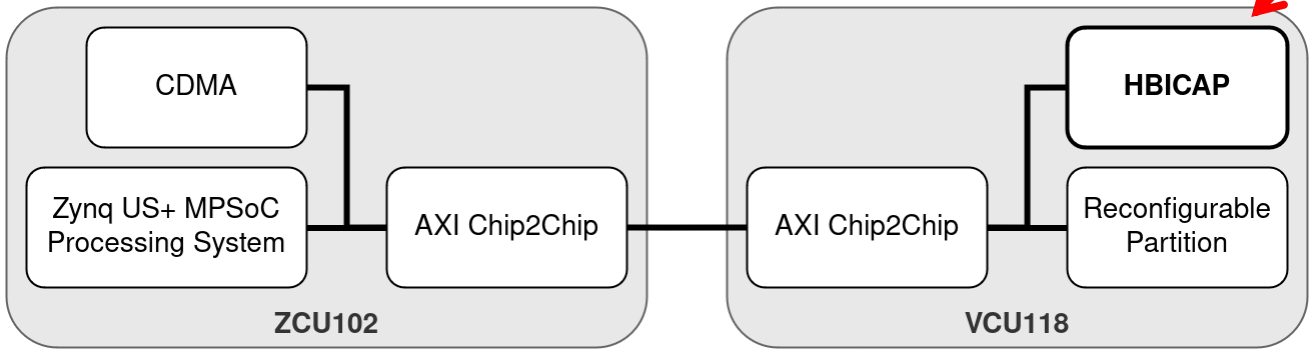


**HBICAP Chip2Chip Setup**

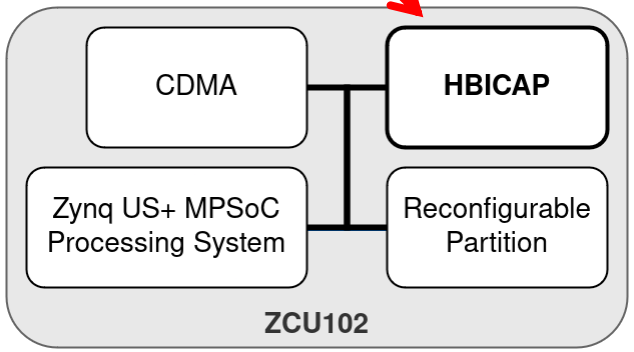


**HBICAP Standalone Setup**

# How can Linux control this?

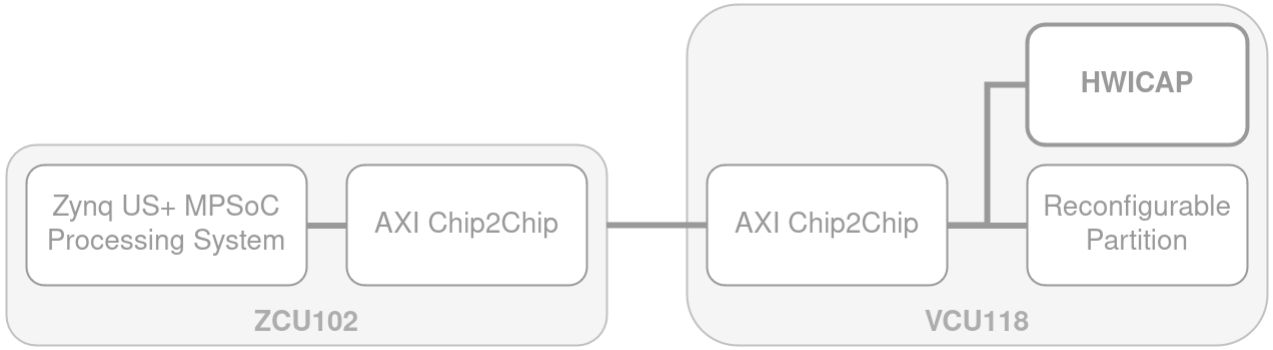


HWICAP Chip2Chip Setup

**Let‘s use this as an example**

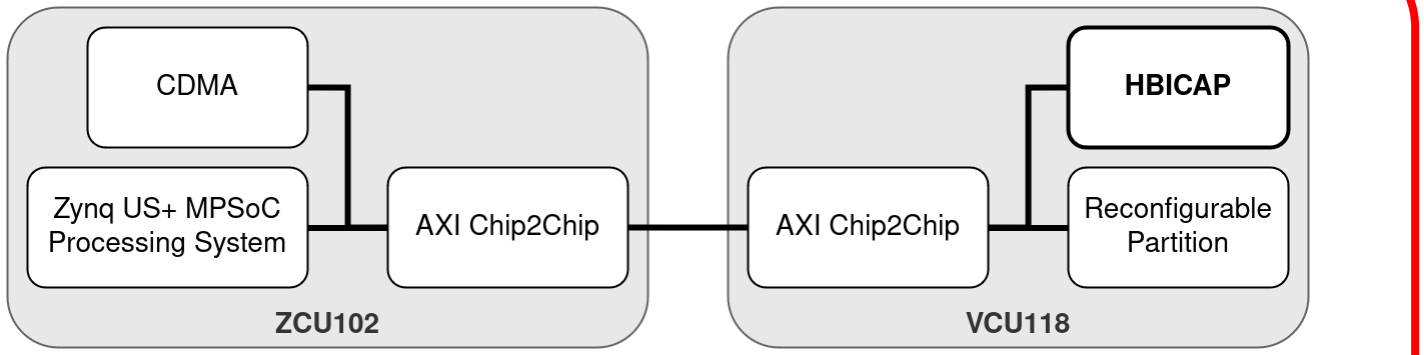


**HBICAP Chip2Chip Setup**

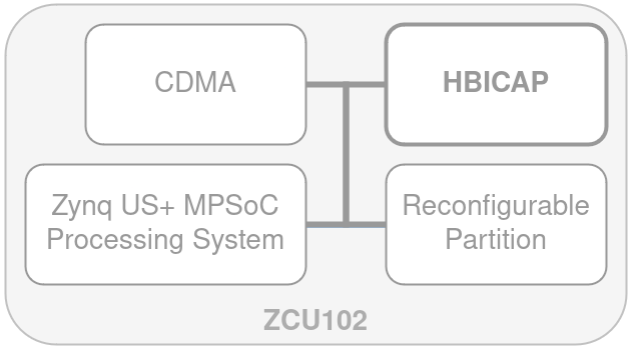


HBICAP Standalone Setup

# FPGA Subsystem in Linux

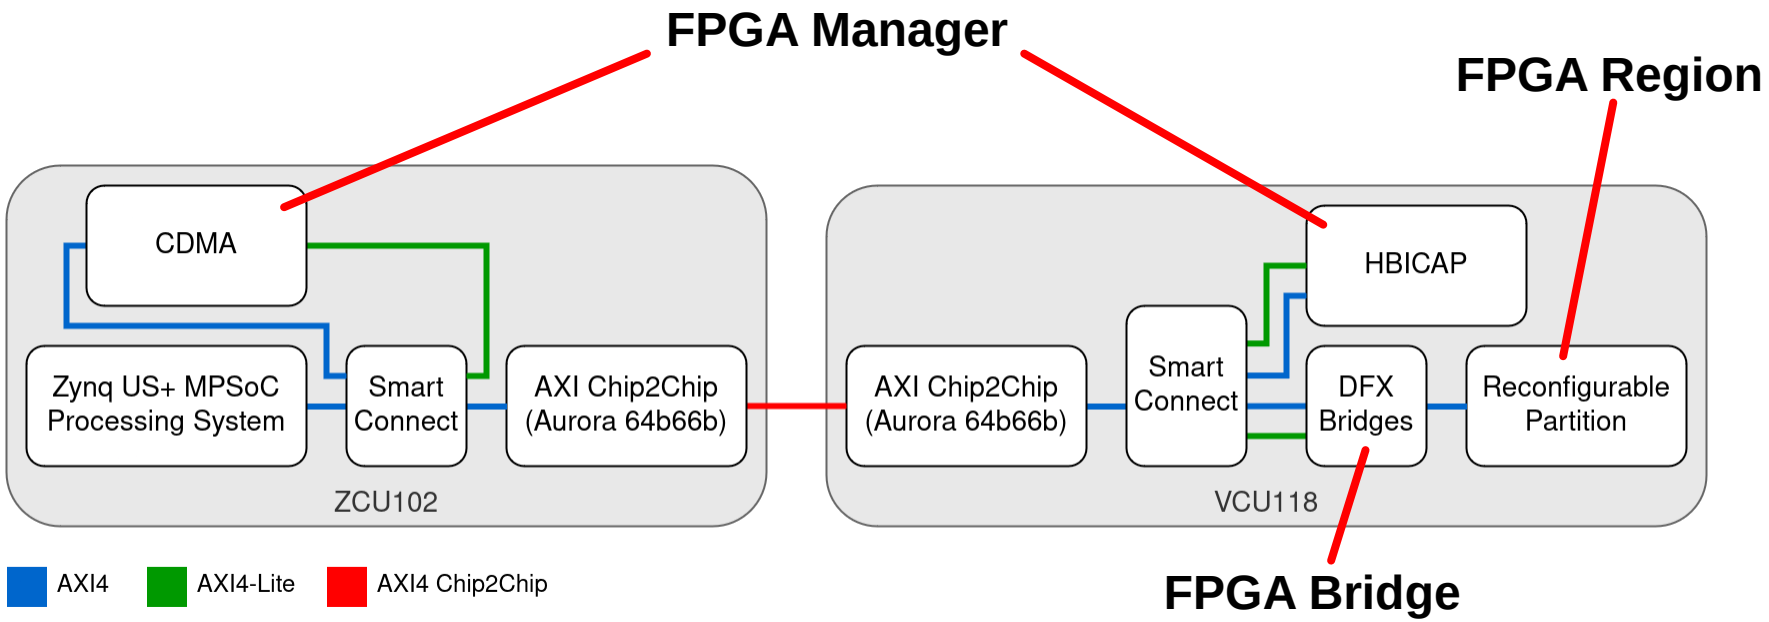

# FPGA Subsystem in Linux

```
axi_hbicap_0_region_0: axi_hbicap@1080010000 {
    compatible = "xlnx,hbicap-fpga";
    reg = <0x10 0x80010000 0x00 0x00001000>,
          <0x10 0x80011000 0x00 0x00001000>,
          <0x00 0xA0000000 0x00 0x00001000>;
};
```

FPGA Manager

```
region_0: region_0 {
    compatible = "fpga-region";
    fpga-mgr = <&axi_hbicap_0_region_0>;
    fpga-bridges = <&dfx_sdm_0_region_0>;
};
```

FPGA Region

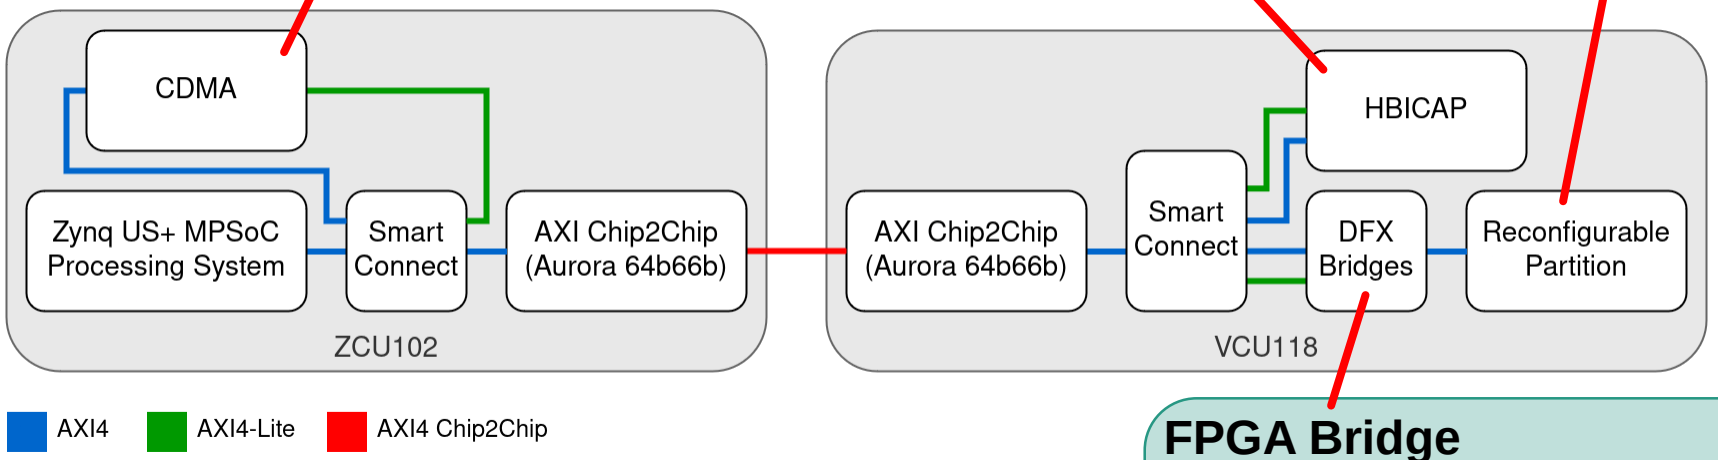


FPGA Bridge

```
dfx_sdm_0_region_0: dfx_axi_shutdown_manager@1080000000 {
    compatible = "xlnx,dfx-axi-shutdown-manager-1.00";
    reg = <0x10 0x80000000 0x00 0x00000080>;
    bridge-enable = <1>;
};
```

# FPGA Subsystem in Linux

```
axi_hbicap_0_region_0: axi_hbicap@1080010000 {
    compatible = "xlnx,hbicap-fpga";
    reg = <0x10 0x80010000 0x00 0x00001000>,
          <0x10 0x80011000 0x00 0x00001000>,
          <0x00 0xA0000000 0x00 0x00001000>;
};
```

FPGA Manager

```
region_0: region_0 {
    compatible = "fpga-region";
    fpga-mgr = <&axi_hbicap_0_region_0>;
    fpga-bridges = <&dfx_sdm_0_region_0>;
};
```

FPGA Region

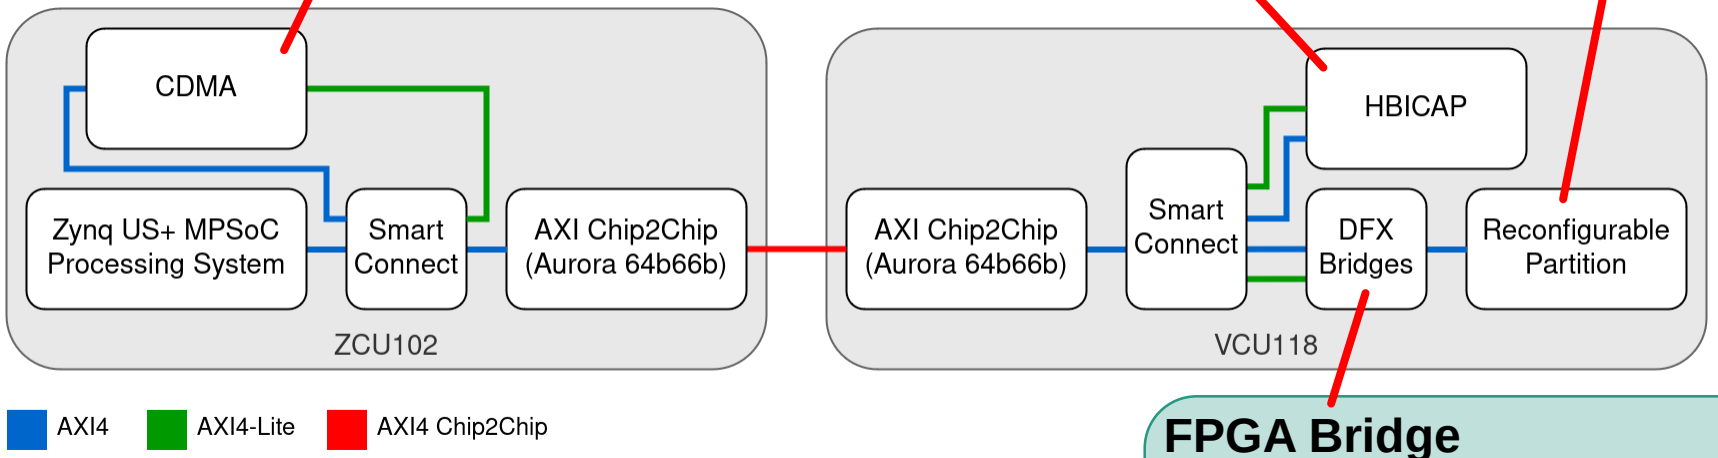


Custom FPGA-Manager-drivers for HBICAP and HWICAP

FPGA Bridge

```
dfx_sdm_0_region_0: dfx_axi_shutdown_manager@1080000000 {
    compatible = "xlnx,dfx-axi-shutdown-manager-1.00";
    reg = <0x10 0x80000000 0x00 0x00000080>;
    bridge-enable = <1>;
};
```

# Performance: Throughput vs. RP Size

## HBICAP

- Throughput depends on reconfigurable partition (RP) size
- Max. ~**180 MB/s**
- Only minor fluctuations
- Impact of Chip2Chip neglectable

## HWICAP

- Significantly slower
- Only minor fluctuations

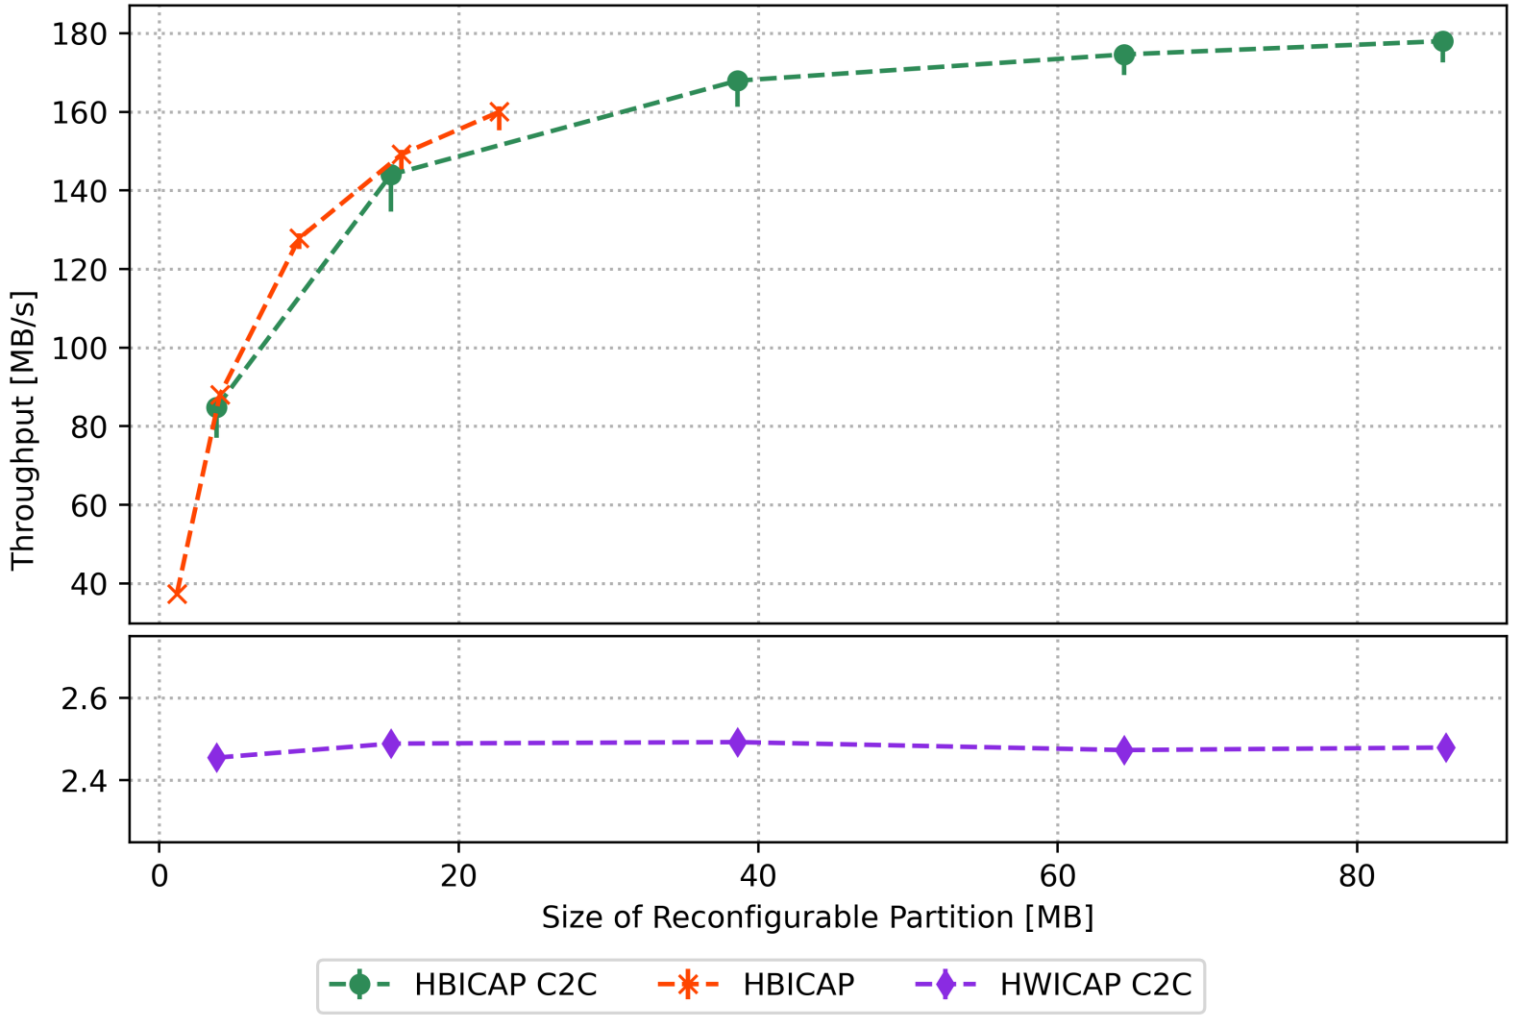

# Performance: Consecutive Reconfigurations

- Throughput almost constant from second reconfiguration on
- Only minor fluctuations

- **First reconfiguration always slow**
  - Cause so far unknown
  - Still under investigation

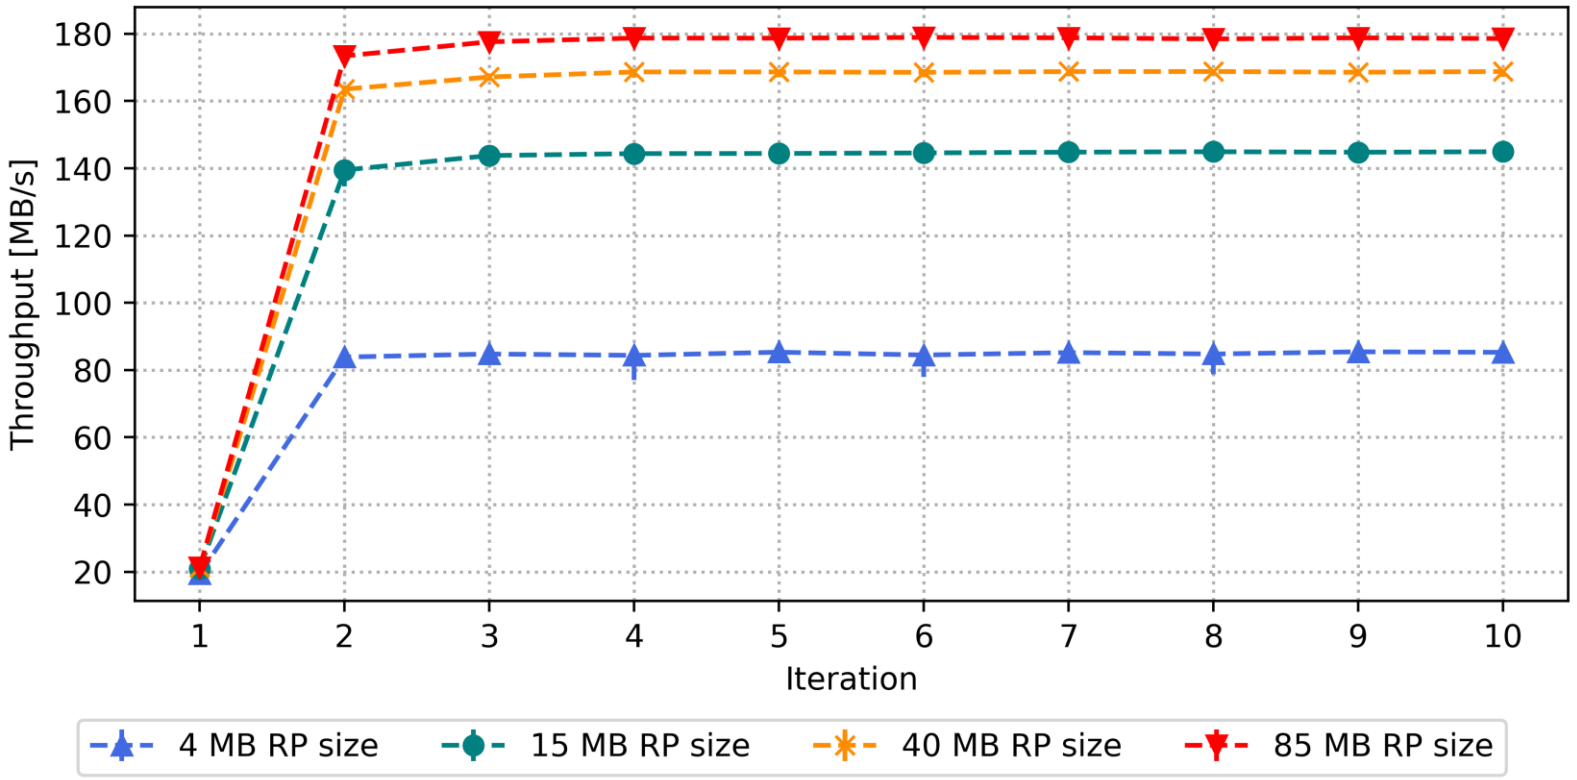

# Outlook

- Further investigation of the slow first reconfiguration
  - Focus on the Processing System of the MPSoC and especially on Linux
- Exploration and optimization in larger systems with more FPGAs connected
  - So far only 2 FPGAs connected to 1 MPSoC
  - The goal is 10 FPGAs and 1 RFSoC on custom HW
- Exploration of open source ICAP controllers
  - Adaptations needed

# Conclusion

- We have developed a method based on **partial reconfiguration** for initializing heterogeneous systems **via AXI Chip2Chip**.
- The proposal can **reduce maintenance effort** and facilitate the **distribution of updates**.
- In test setups, the concept has proven to be **fast, reliable, and stable** and will soon be scaled up to the full QiController system with up to 10 FPGAs per RFSoC.
- Many **ATCA cards at CERN** are also compatible with this method, and this field of application will also be explored.

# Backup

# Test Setup

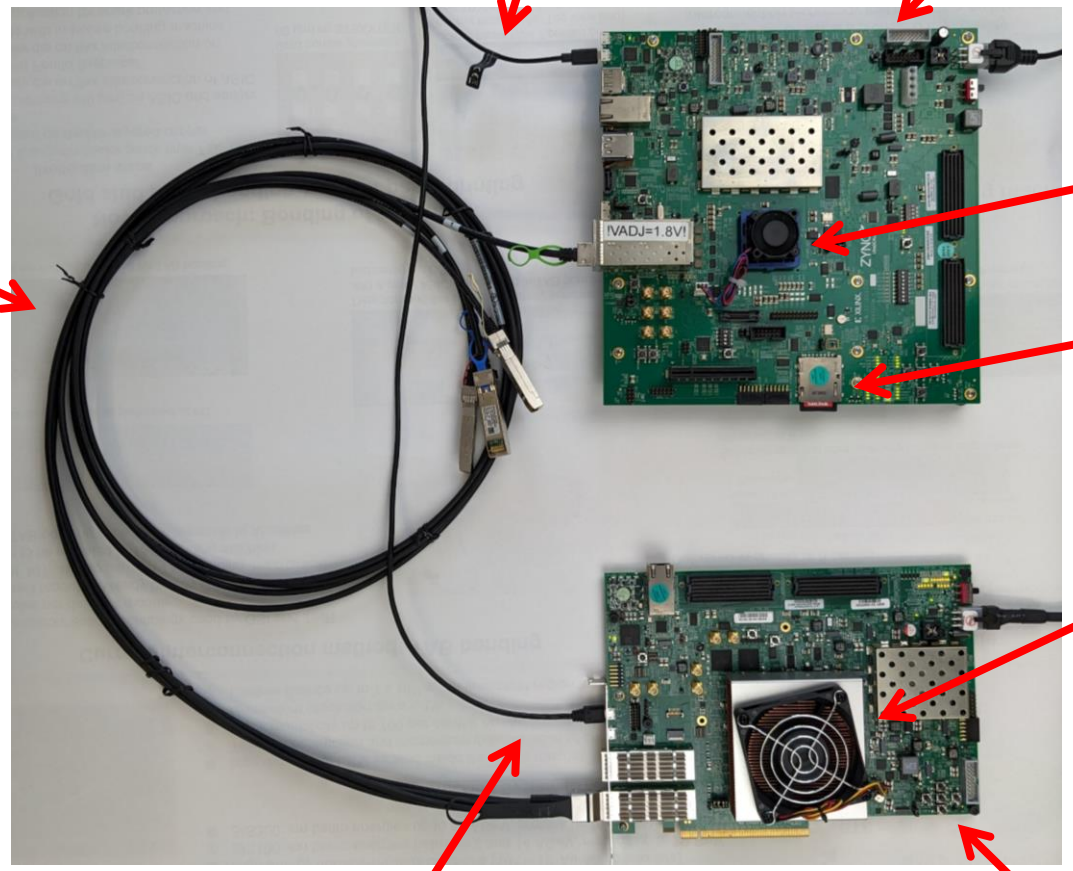
UART
AMD/Xilinx ZCU102
QSFP to SFP+ Cable
(AXI Chip2Chip)
UltraScale+ MPSoC
SD Card
UltraScale+ FPGA
JTAG
AMD/Xilinx VCU118

# Test Firmware

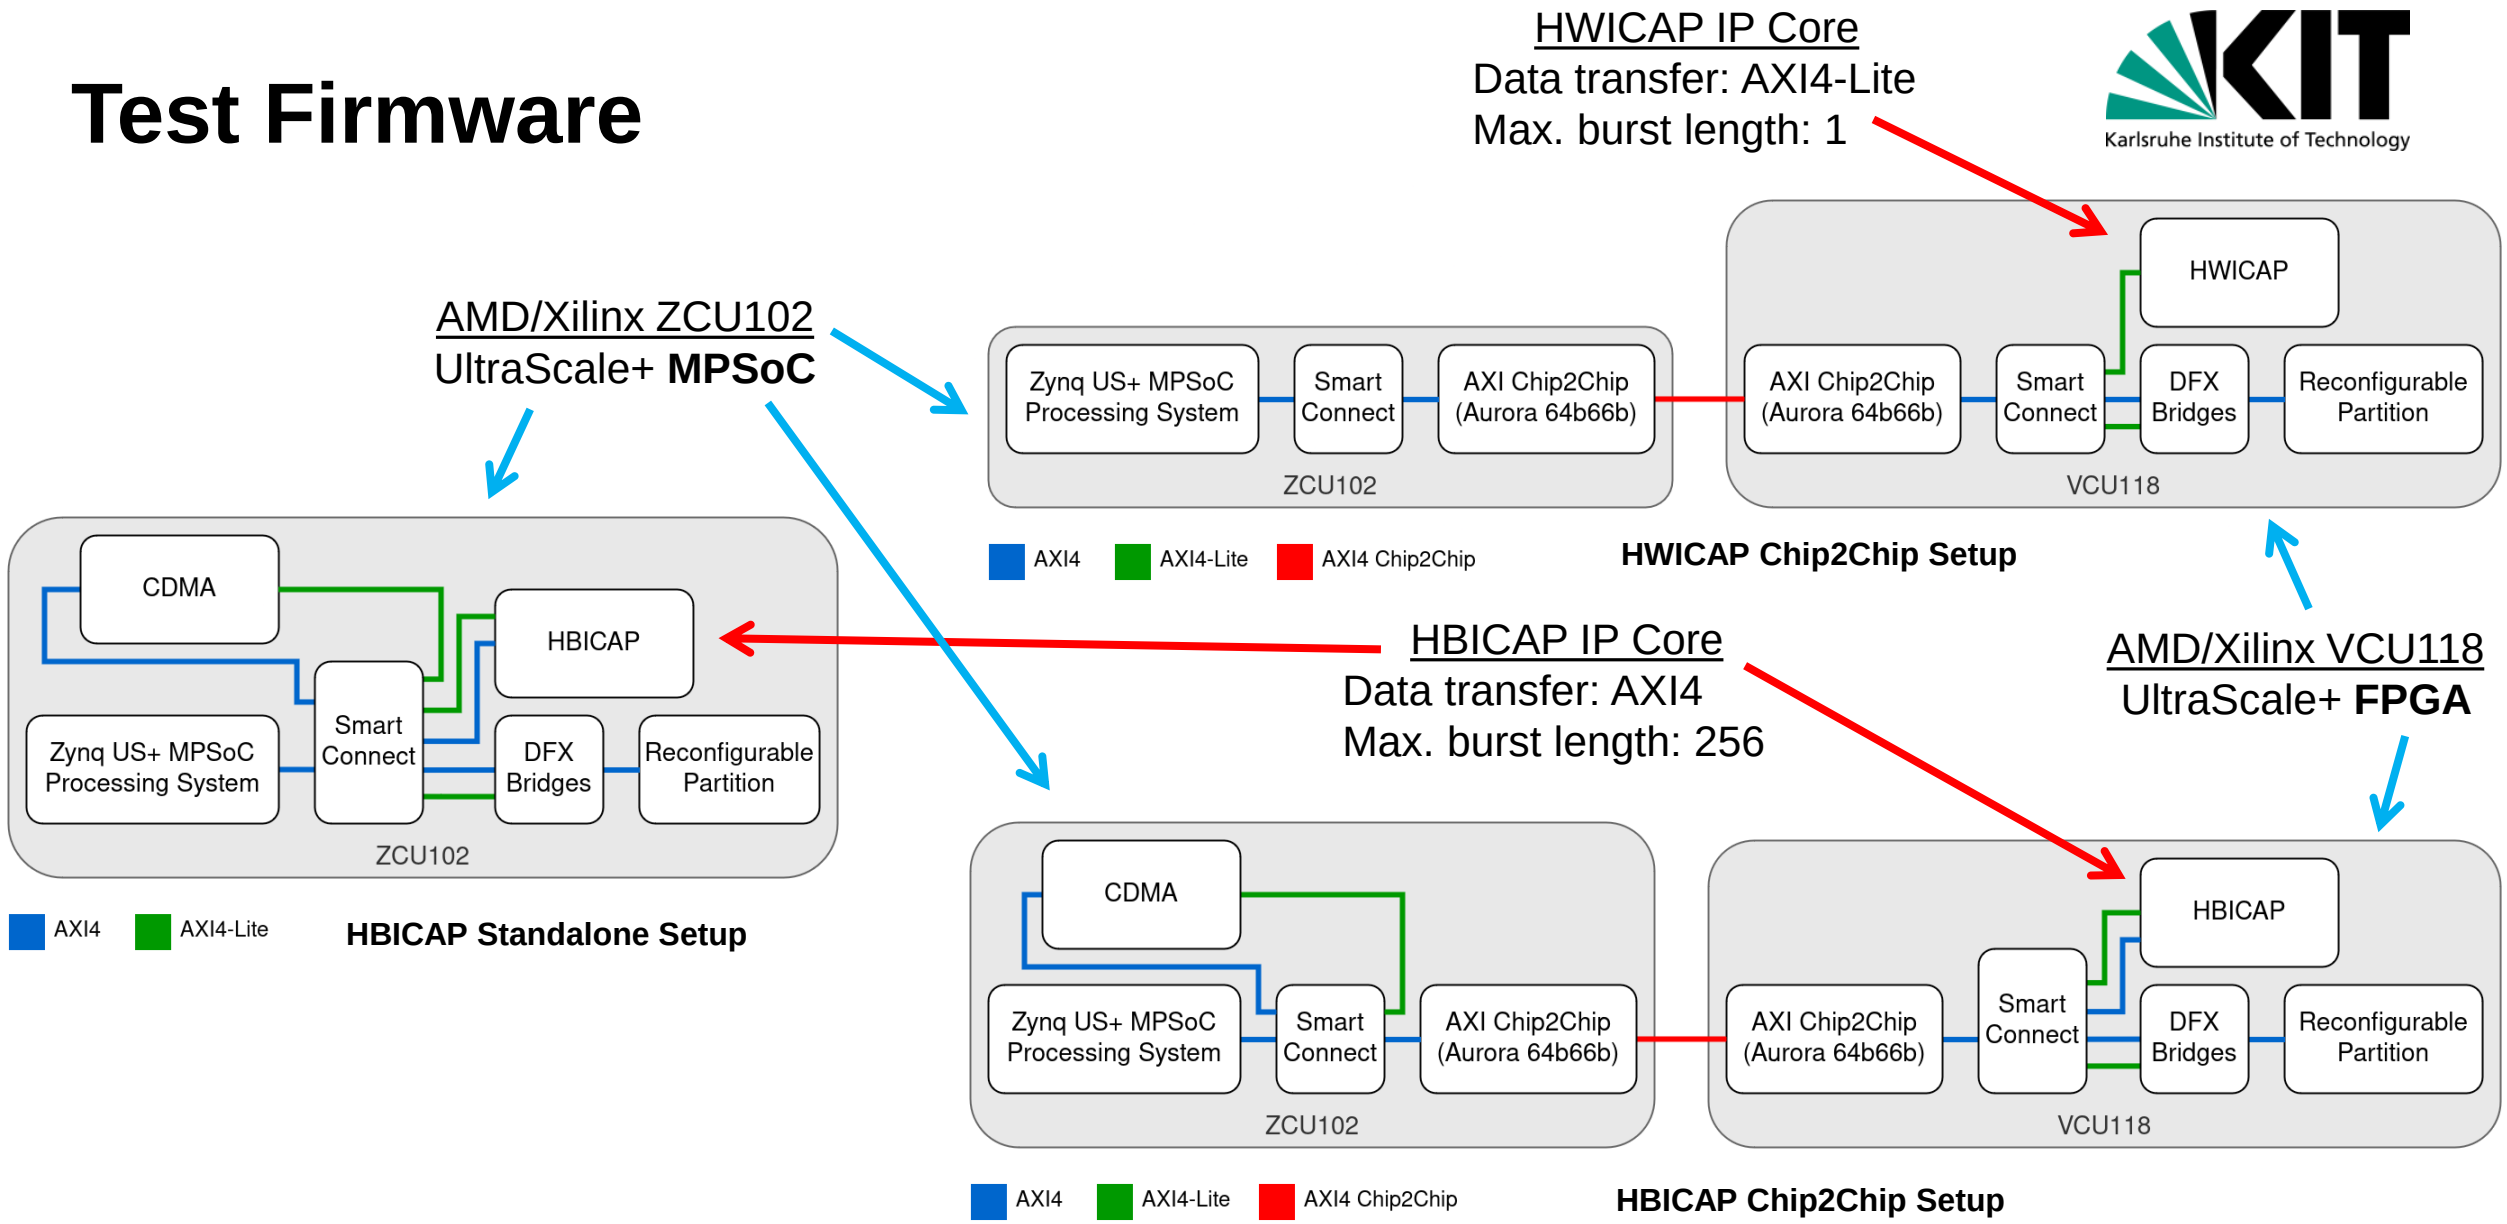

# Heterogeneous Systems

- Composed of various devices from AMD / Xilinx
  - MPSoCs
  - RFSoCs
  - FPGAs
- Architecture for Qbit control system
- Multiple systems developed at CERN

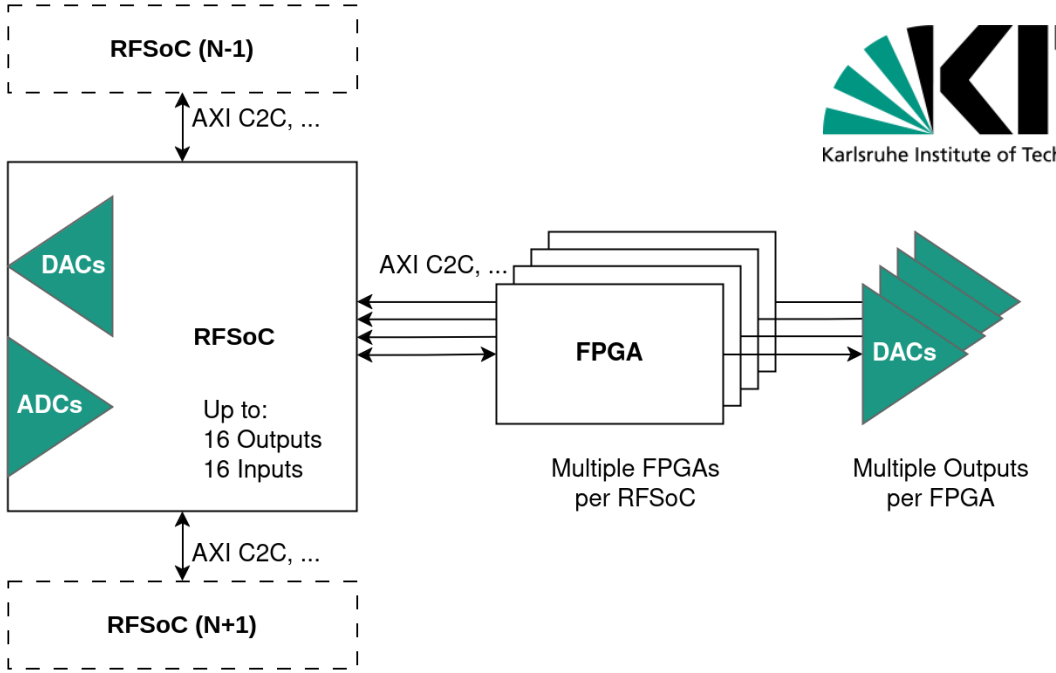


Multi-Device QiController Architecture

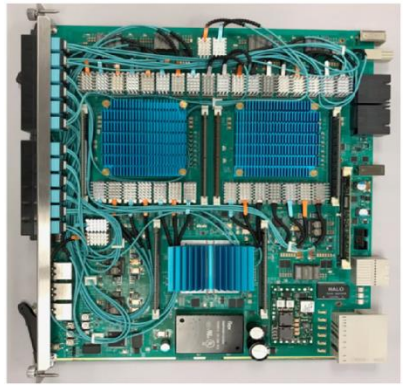

Global Common Modul ATCA Card

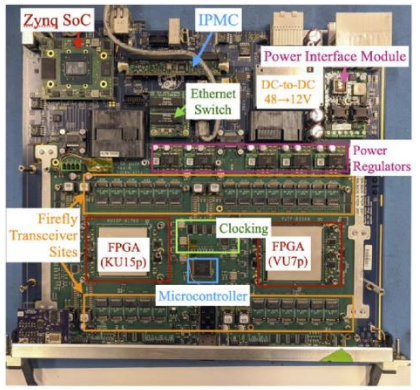


Apollo ATCA Card

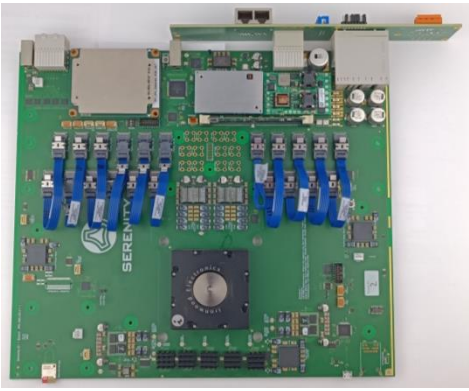

Serenity S ATCA Card