






**SPRACE**

# **A new firmware Command Processor Block for the test of the Phase-2 CMS modules**


**LUIGI CALLIGARIS (SPRACE/UNESP)**

# Module testing infrastructure

- Outer Tracker Modules
  - Pixel-Strip (PS) modules → inner layers
  - Strip-Strip (2S) modules → outer layers
  - Different type of sensors & ASICs
- Inner Tracker modules
  - One type of readout chip (RD53B CROC)
  - Different sensors, different arrangements
- Common testing infrastructure
  - FC7 AMC FPGA card + adapter cards
  - Computer running testing software
- Prototype tests & final module QC



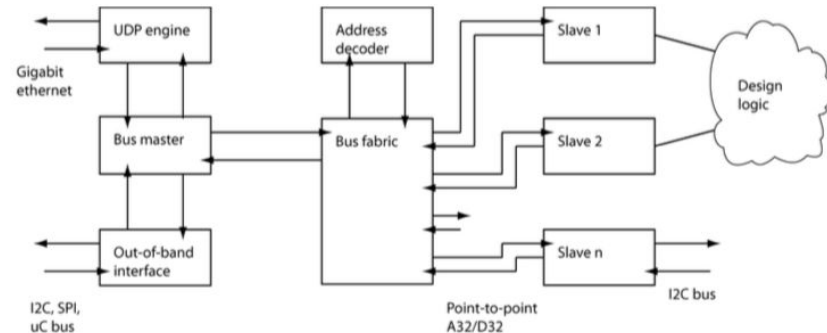

1x2



2x2

# Phase-2 Acquisition Control Framework (Ph2-ACF)

- Software implementing high-level test functions
  - Channel testing, calibration, histogramming
  - Modular, to accommodate different sensor modules
- Used for IT and OT development and testing
  - Common tools, shared procedures, standardization
- Interface with FPGA via IPBus/uHAL over Ethernet
  - Simple, 32-bit addresses mapping to 32-bit registers
  - Uses UDP protocol, low resource utilization
  - Reliable, guaranteed writes/reads
    - Some latency/bandwidth limitations

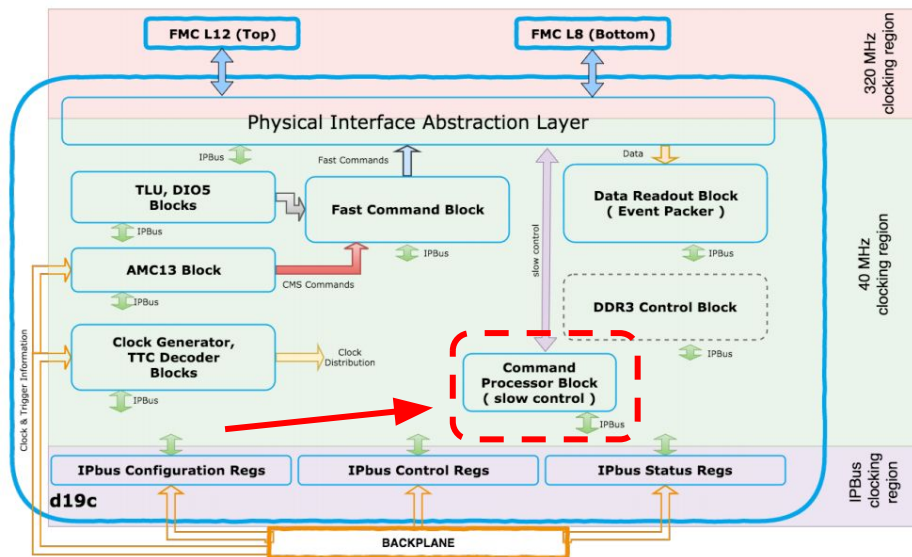


# The d19c/uDTC gateway



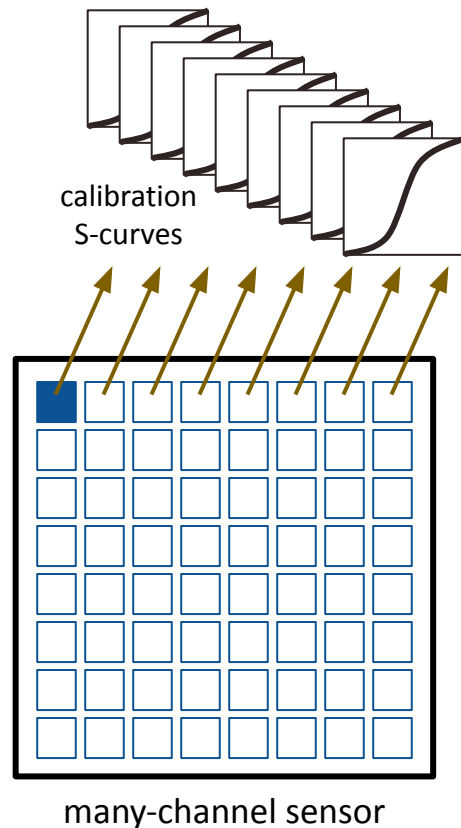
FC7 (Kintex-7 420T)

- FPGA gateway for FC7 AMC cards
  - Supports IT and OT sensor modules (optical & electrical connection)
  - Modular source HDL → can generate configuration-specific bitfiles
    - e.g. one 2S module connected electrically, four PS module connected optically (IpGBT), ...
- CMS-wide effort involving many institutes in the tracker effort
- Command Processor Block
  - Module (slow) control
    - Configuration registers of the ASICs
    - Discriminator thresholds
    - Some R/W ops are time-critical



# Calibration is a latency-sensitive operation

- OT sensors → binary hit info
- IT sensors → binary hit info + ToT (log charge)
- Calibration S-curve of the sensor
  - Dependency of hit occupancy vs. DAC threshold
- The CMS sensor modules host a large # channels
  - 2S (1'016 ch), PS (32'128 ch), IT (up to 307'200 ch)
- O(100) of DAC values scanned to build S-curve
  - A single module may require sweep of millions of DAC values
  - Many modules (up to 16) can be served by one FC7 card
  - The threshold scan NEEDS TO to be optimized



# Latency issue in IPBus transactions

- Thousands of modules will need quality control testing
  - Initial attempts at PS module calibration required O(1 hour)
  - This was unsustainable → it would take too long
- Issue tracked down to IPBus latency in slow control operations:
  - Manipulation of the registers in the IpGBT...
  - ... to drive the I2C masters of the IpGBT...
  - ... to manipulate the registers in the ASICs → tens of IPB transactions / ASIC reg
  - Small latencies add up quickly when performing millions of operations

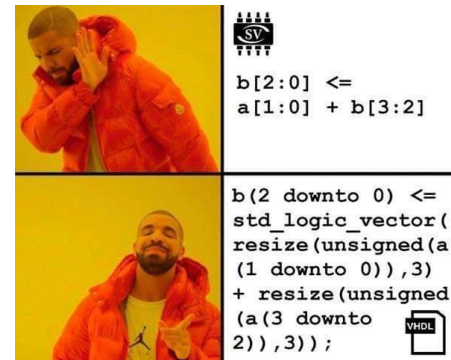
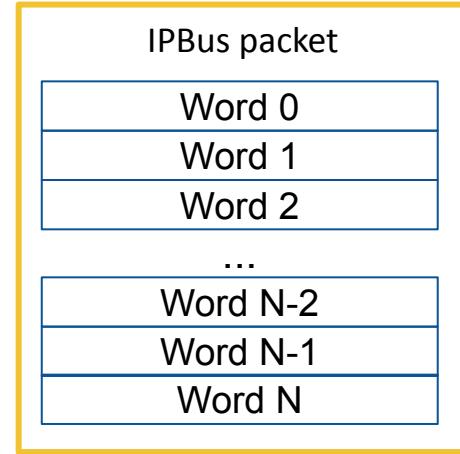
From IPBus  
manual



- For one software client controlling one device, the single-word read/write latency is approximately 250 micro-seconds
  - Although this single-word latency is larger than VME/PCIe-based control, for multiple transactions or large block transfers this is compensated by concatenating multiple (up to order of 100) transactions into each packet, and by having multiple packets in flight around the system at any given time
  - Hence, optimal performance will be obtained if network dispatches are only performed when necessary.
- The 1-client-to-1-device block read/write throughput for payloads larger than 1 Mbyte is above 0.5 Gbit/s
- The total block read/write throughput is above 0.75 Gbit/s for three or more boards in a single MicroTCA crate.

# Moving the intelligence into the gateway

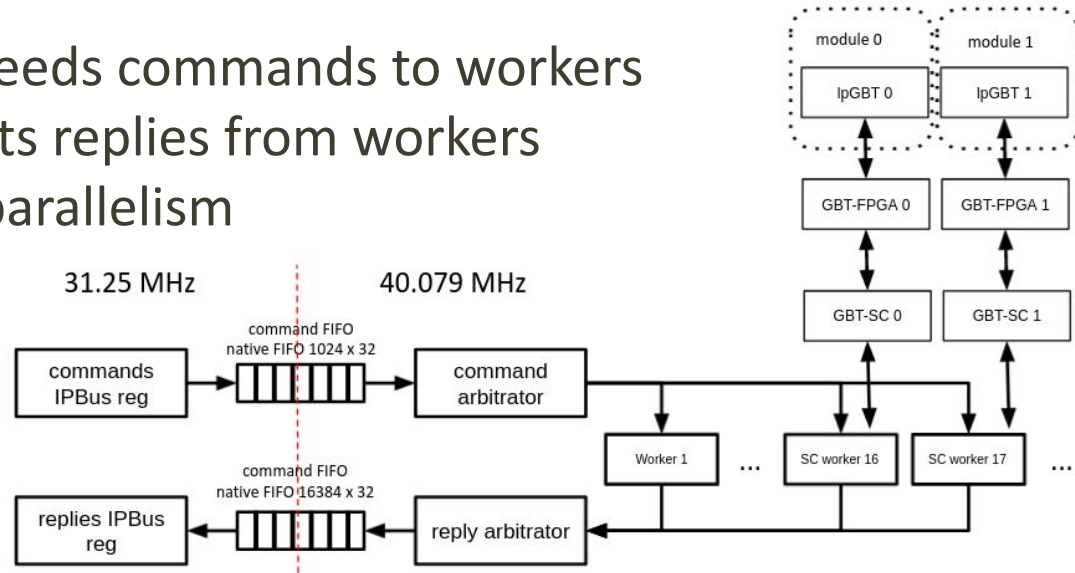
- Speed up the slow control operations
  - Move intelligence of the code into the VHDL gateway
- New Command Processor Block
  - Receive commands via a single IPBus FIFO
    - This avoids register synchronization issues
    - Allows to pack multiple words (up to ~300) in one transaction
  - Offer replies via a single IPBus FIFO
    - Uses just one BRAM for the CDC and reply accumulation
    - The controlling PC can perform block reads and sort locally
  - Modular structure → optimize out unneeded logic
    - e.g. do not use resources for optical link logic when



sometimes VHDL can be... "fun"

# Design of the CPB

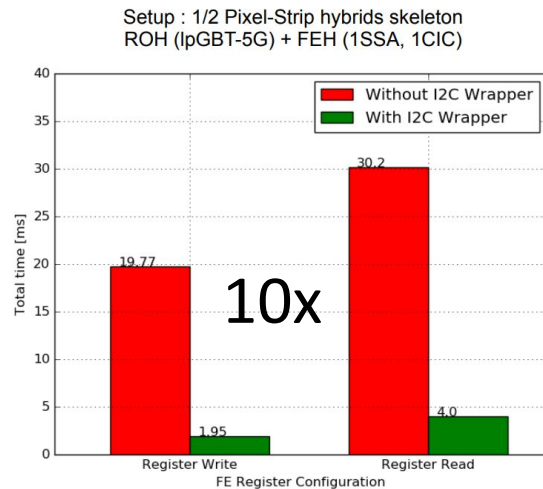
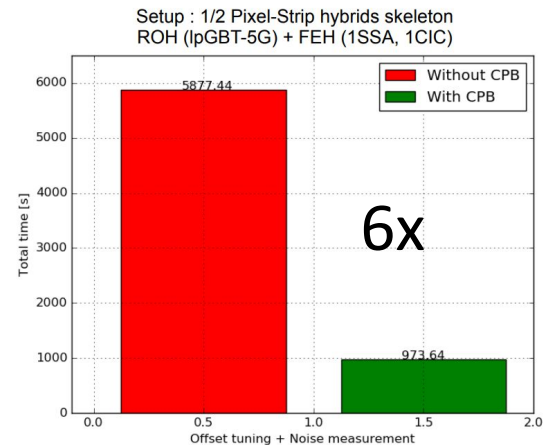
- Clock-domain crossing buffers (IPBus <--> LHC clk domains)
- Local CPB bus with up to 254 “workers”
  - Each worker is an independent core with a specific function
  - e.g. control of a single IpGBT in a module
- Command arbitrator → feeds commands to workers
- Reply arbitrator → collects replies from workers
- This architecture allows parallelism
  - Once given a command, workers can process and reply back asynchronously
  - Workers can be intelligent





# Performance improvement

- Implementation of the IpGBT worker for OT
  - Repacking of multi-reg operations on IpGBT in one cmd
    - 6x speedup
  - Adding more intelligence into the worker
    - I2C transactions driven by the gateway instead of PC
    - Further 10x speedup
- What was taking more than 1 hr now takes minutes
- The structure of the CPB makes it scalable
  - Multi-module tests expected soon
  - The asynchronous worker arrangement should scale well



# Summary

- Continuing SPRACE contribution to the uDTC/d19c gateway effort
- Design of the New Command Processor Block
  - Overall latency improvement of the order of 60x
  - Critical for testing and QC of CMS tracker modules
- Lessons learned
  - ETH-IPbus is very flexible, but has high costs in terms of latency
  - This aspect can be mitigated by moving logic into the gateway
- Outlook
  - Future gateway may benefit from use of quicker PC-FPGA interfaces
  - Examples
    - PCI-e for x86 PC architectures
    - AXI Chip2Chip for ZYNQ architectures



Thanks from the SPRACE team!

A pesquisa aqui apresentada recebe o suporte da FAPESP (processos 2018/18955-0 e 2019/18166-9)

