



Financiado por  
la Unión Europea  
NextGenerationEU



Plan de Recuperación,  
Transformación  
y Resiliencia



VNIVERSITAT  
DE VALÈNCIA



IFIC  
INSTITUT DE FÍSICA  
CORPUSCULAR

LHCb  
TFCP

# Sustainability of real-time analysis at 5 TB/s data rate

Volodymyr Svintozelskyi

Arantza De Oyanguren Campos, Brij Kishor Jashal, Jiahui Zhuo, Valerii Kholoimov  
Volodymyr Svintozelskyi, Luca Fiorini, Álvaro Fernández and Alberto Valero

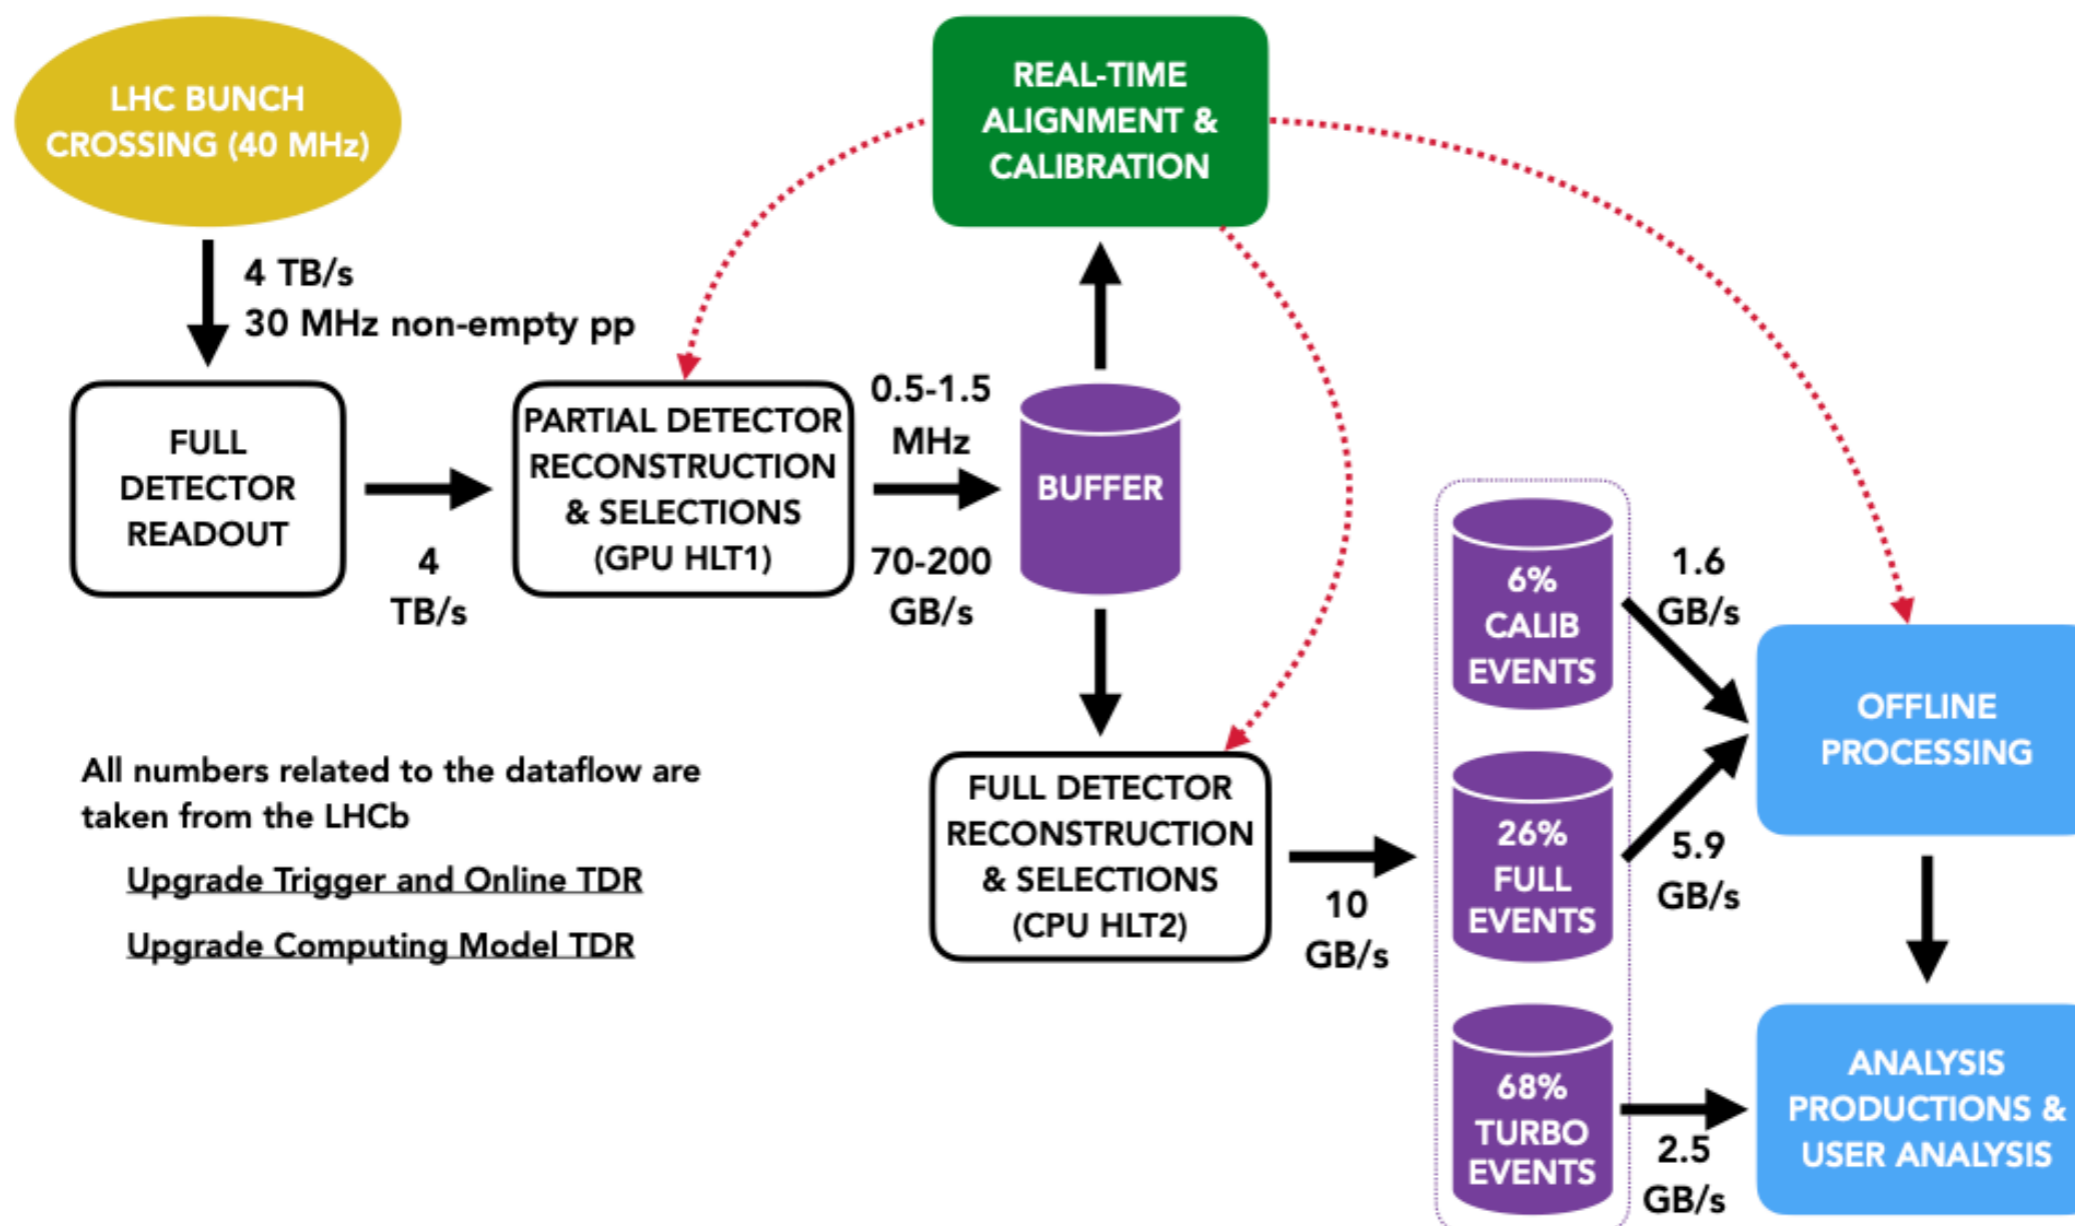
IFIC, Univ. of Valencia and CSIC (ES),  
TATA Institute of Fundamental Research (TIFR),

# Overview

- Real-time analysis and DAQ system at LHCb
- High-Low project at Valencia
- Hardware power consumption studies
- FPGAs as a way to reduce consumption?
- Consumption per HLT1 algorithm. Can we detect unoptimized code?
- Power consumption in future (WLCG example)

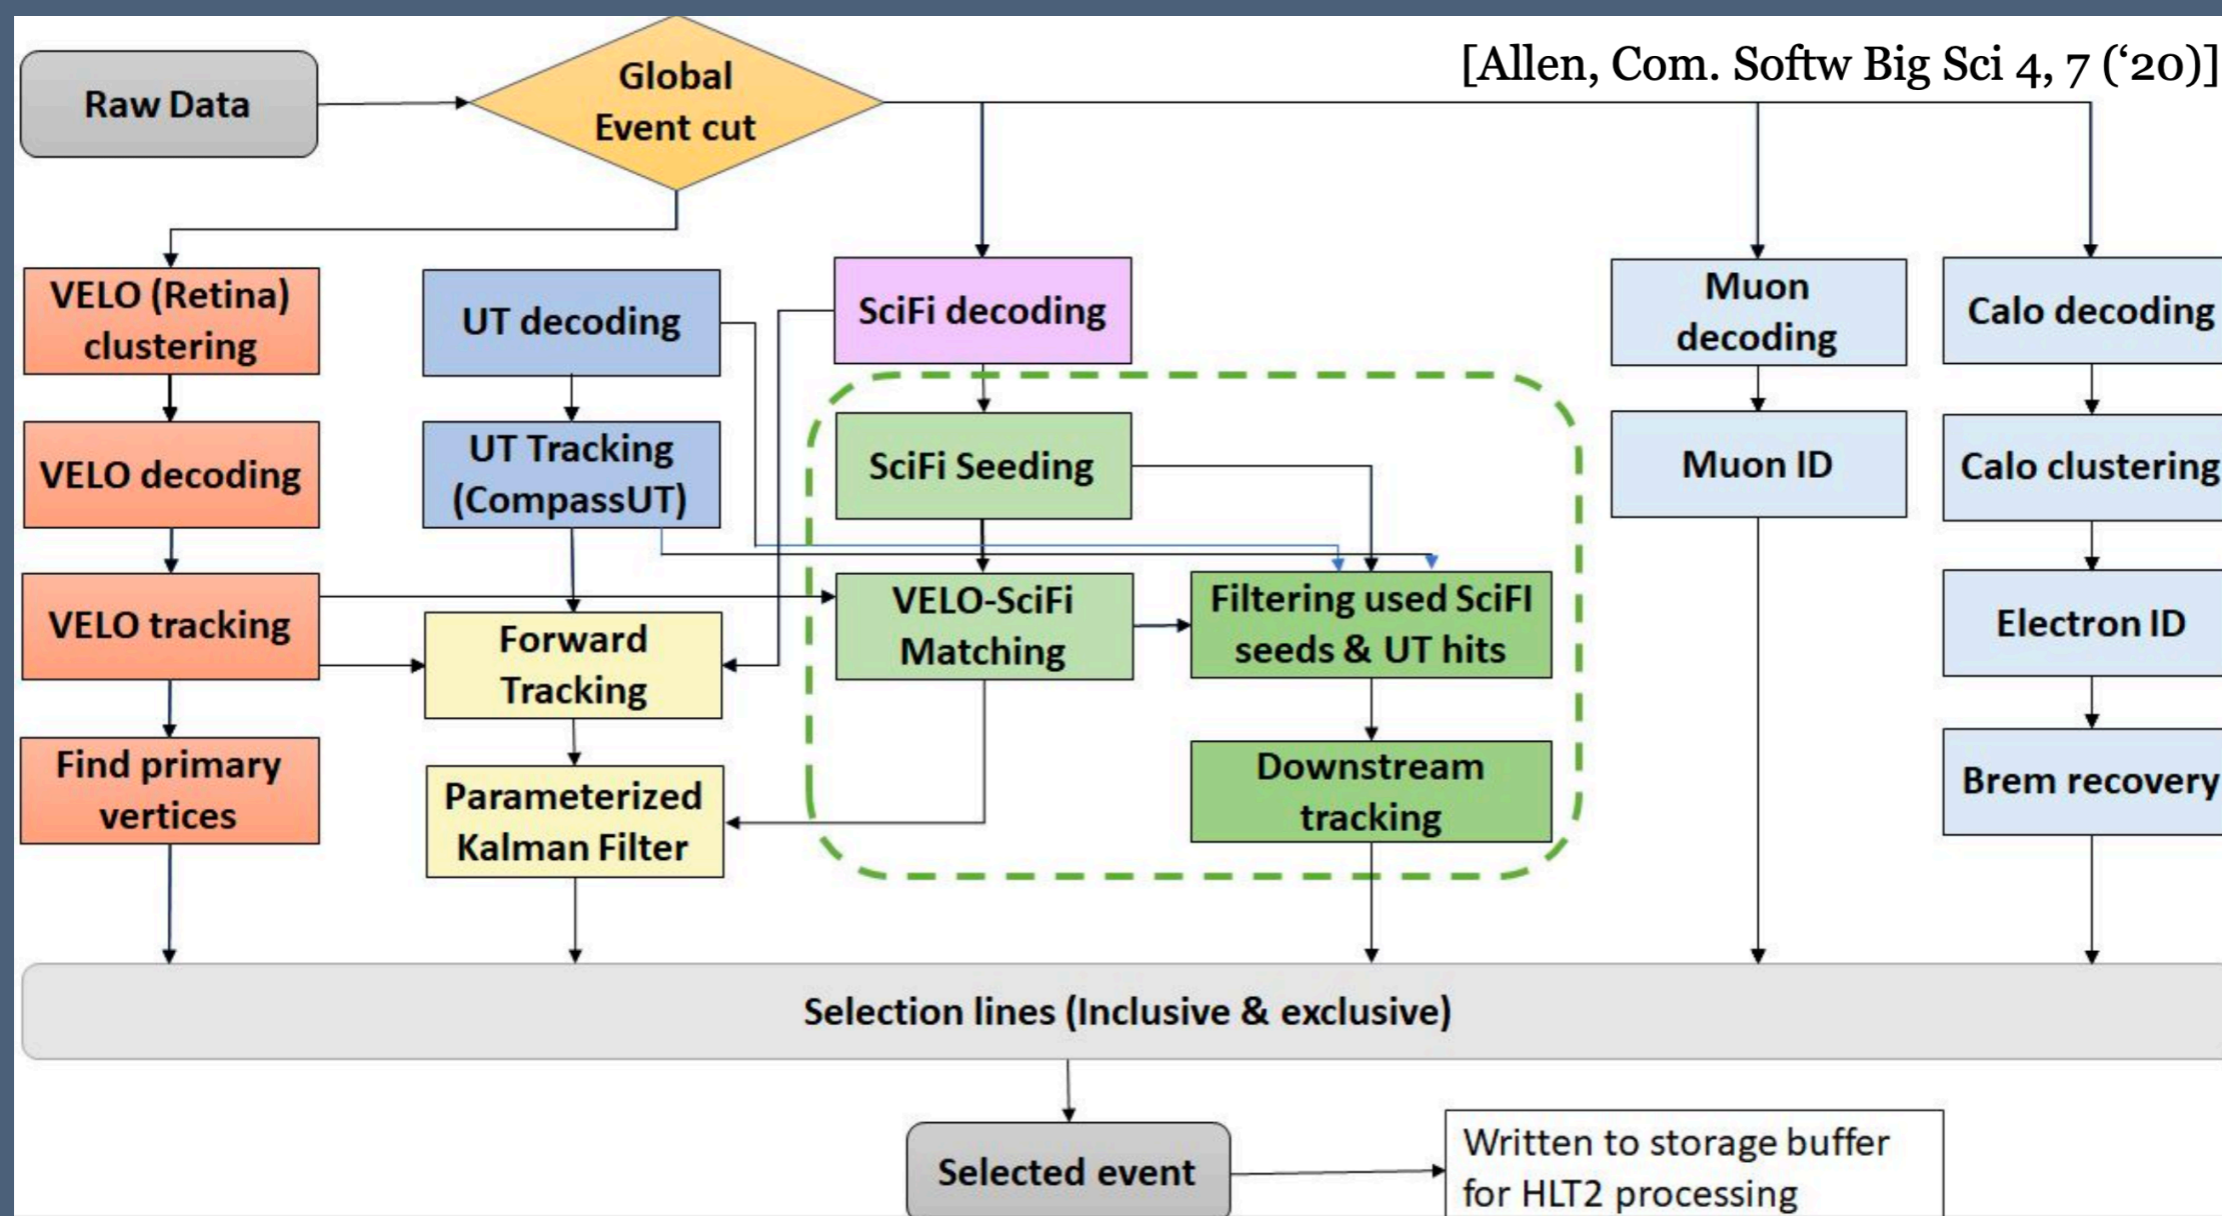
# Real-time analysis at LHCb

- No hardware-based trigger in Run 3
- Allen: the LHCb high-level trigger 1 (HLT1) application on GPUs
- Fast detector reconstruction in  $O(500)$  Nvidia RTX A5000



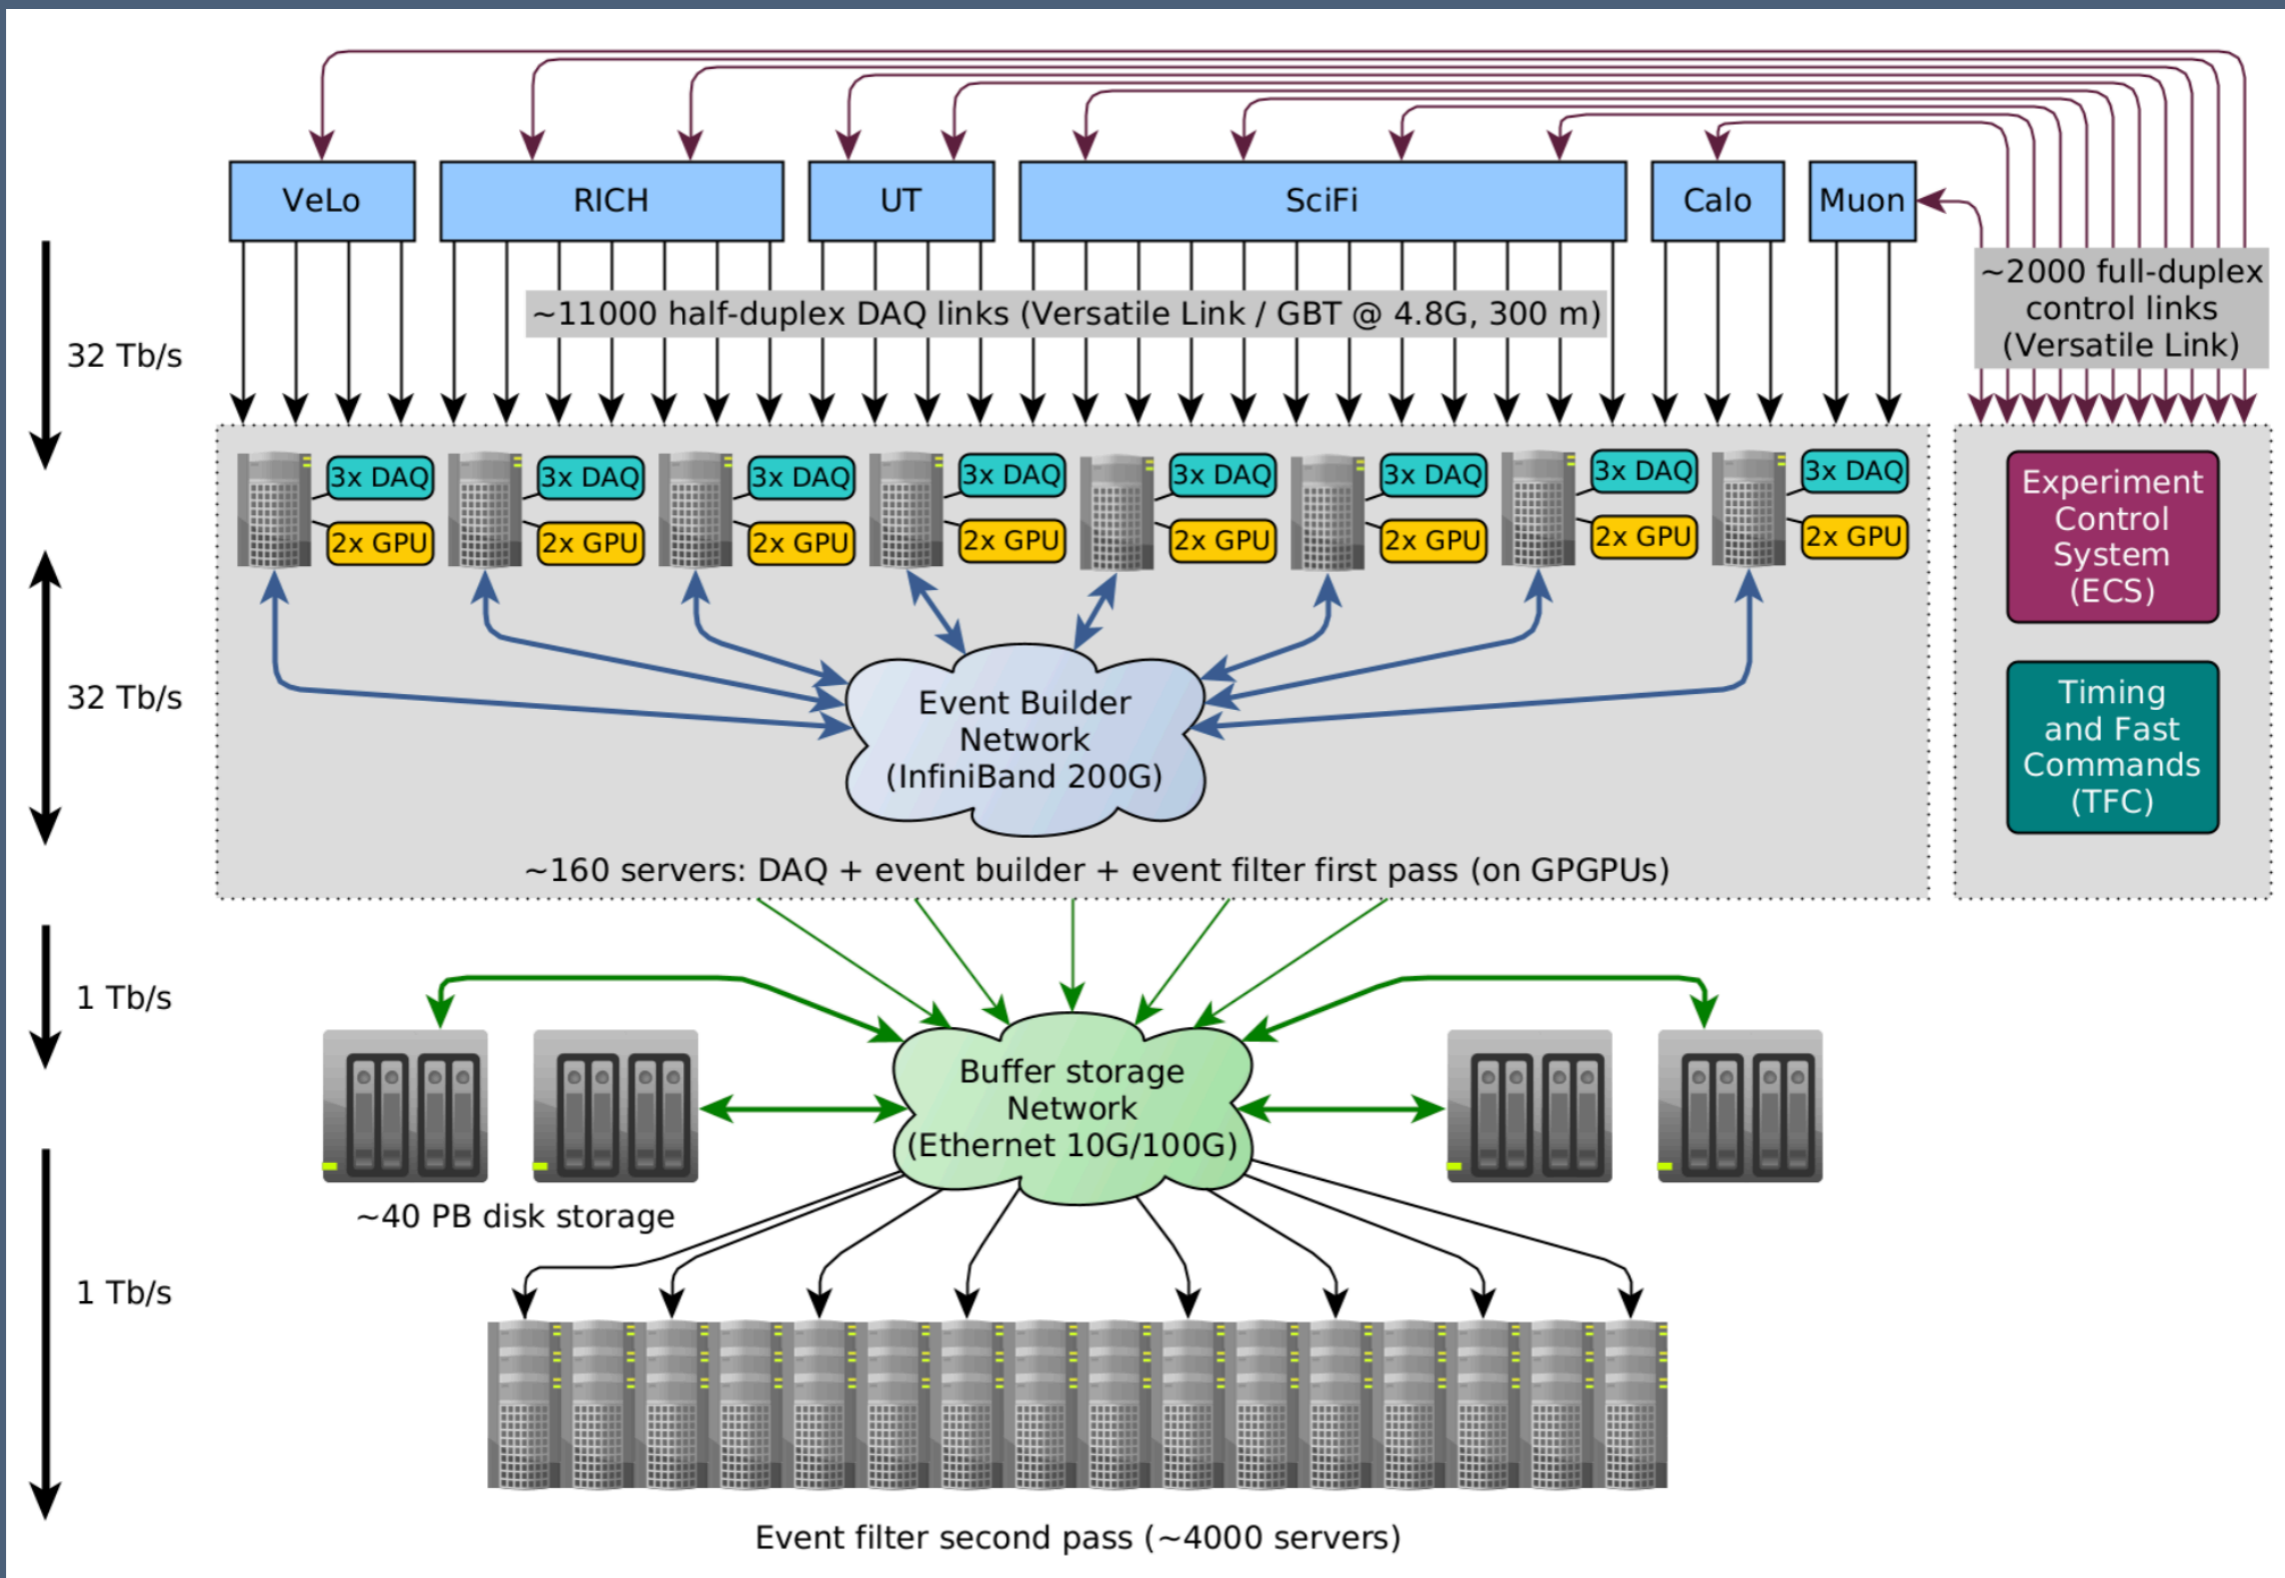
# Real-time analysis at LHCb

- More than 250 algorithms with a high level of parallelization
- Algorithms are executed in predefined sequence
- Allen is an example software, but the conclusions are general





# LHCb DAQ system for Run3



# The HIGH-LOW project at Valencia

## Design of High-Performance Algorithms for Low-Power Sustainable Hardware for LHC Experiments and Their Upgrades

- Transversal project: ATLAS and LHCb experiments
- PIs: Luca Fiorini (ATLAS), & Arantza Oyanguren (LHCb)
- Funded by the Spanish Ministry of Science and Innovation (*TED2021-130852B-I00*)
- About 10 people (physicists + engineers + students)

*Aim:* Benchmarking new hardware architectures and developing fast and high efficient algorithms with reduced power consumption



### HL hardware:

- Rack K RETEX LOGIC-2 A600 42U F1000 PH
- APC Metered Rack PDU ZeroU 2G AP8
- SWITCH D-LINK DXS-1210-28T 24x 10GB
- T10G Dual Xeon Scalable HPC 10xGPU PCIe
  - 2 x Intel® Xeon® Gold 5318Y 2,10GHz 24 Cores 3.40 GHz
  - 8 x 32GB DDR4 3200MHz ECC REG
  - 1 x SSD Samsung 990 PRO 2TB M.2 NVMe 2280 PCIe 4.
  - 1 x Controladora BROADCOM MegaRAID 9560-16i PCIe
  - 8 x HD 10TB SAS 12Gb/s 7.200 rpm
  - 1 x NVIDIA® RTX™ A5000 24GB GB GDDR6 ECC
  - 1 x NVIDIA® RTX™ A6000 Ada Generation 48GB GB GDDR6 ECC
  - 1 x HBA Broadcom N210GBT Dual 10GbE RJ45



# The HIGH-LOW project at Valencia



**IFIC**  
INSTITUT DE FÍSICA  
CORPUSCULAR





# Hardware power efficiency



What is the effect of faster hardware on overall power consumption?



The LHCb decision to run HLT1 on GPUs may have saved up to a factor 10 in energy consumption!

# Hardware utilisation

Higher number of GPU streams leads to faster CPU heating → higher consumption




\* Part of Allen algorithms needs to be executed on CPU, which leads to CPU heating and cooler fan speed ramp up



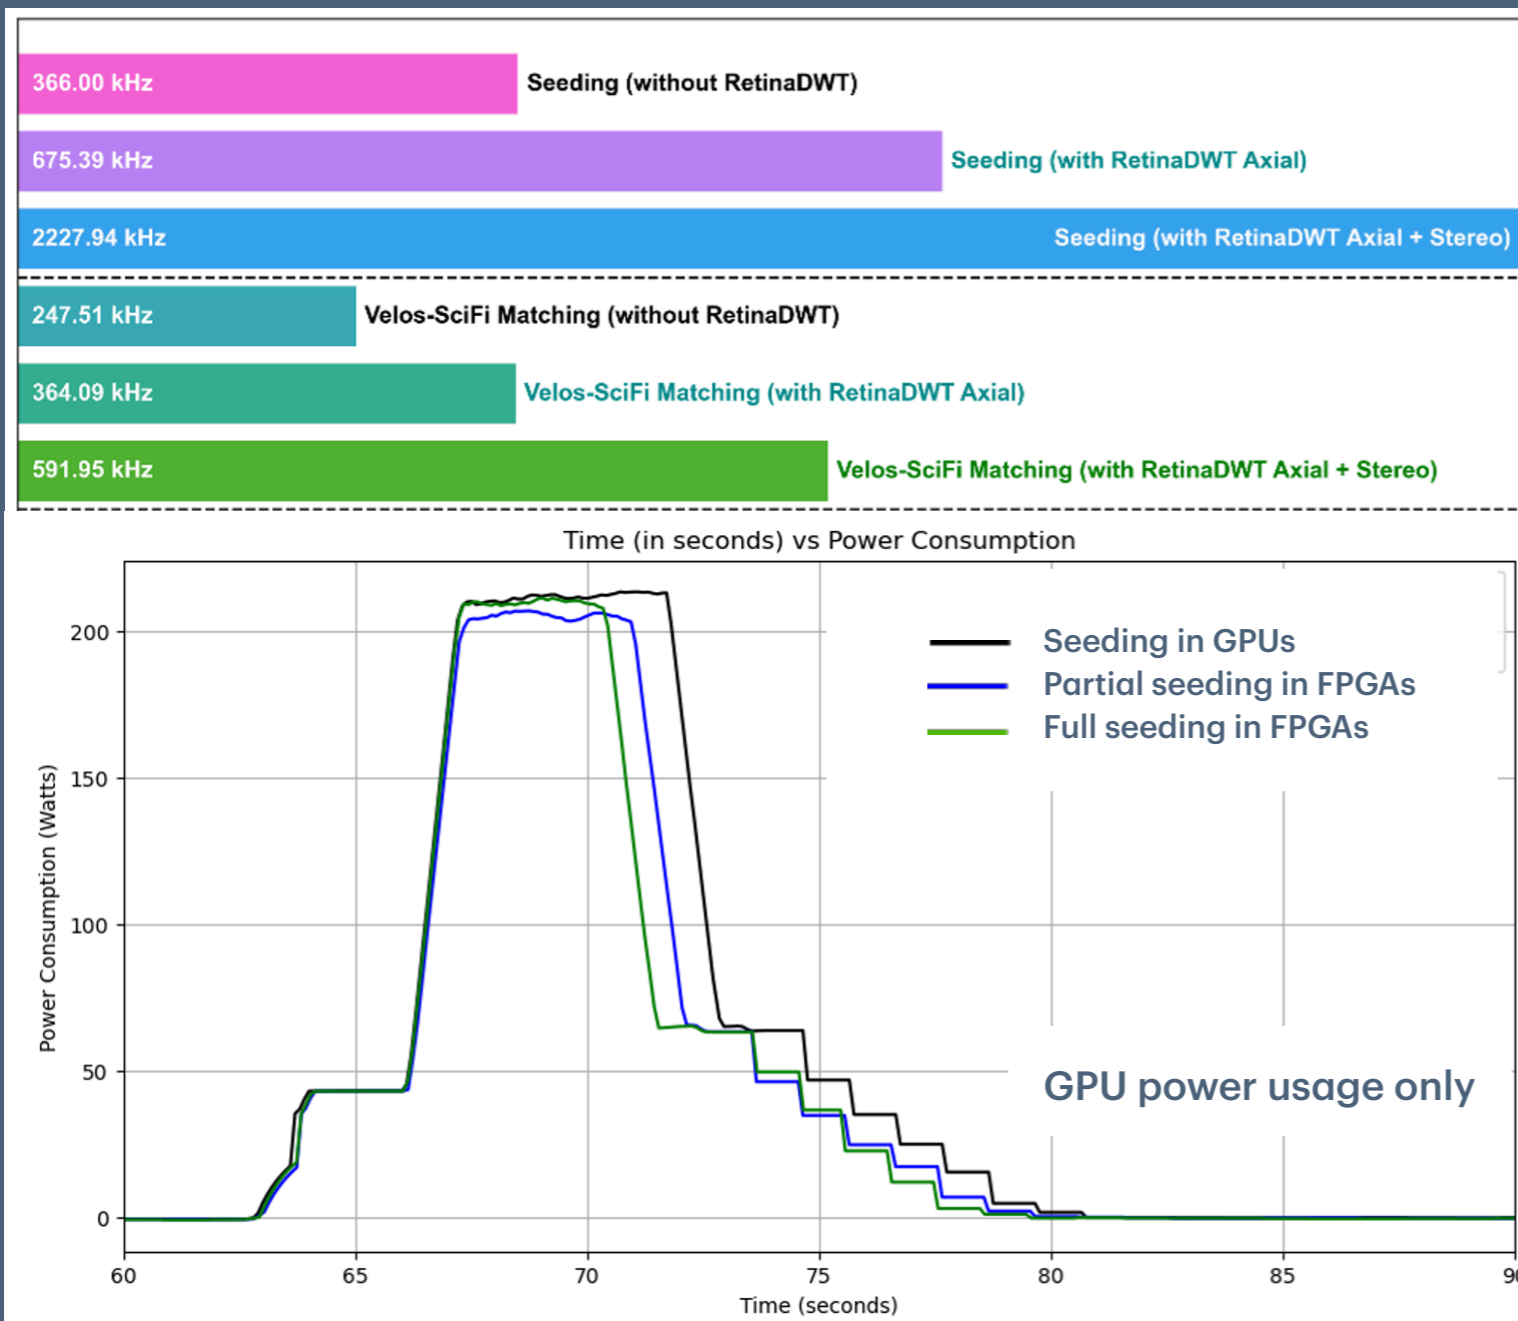
# FPGA's for less consumption?

- Offloading of some computation tasks to FPGA
- Real-time reconstruction on FPGAs (“artificial retina”)
- VELO clustering is already implemented for Run3 in FPGAs !
- Tracking in development for Run5 (~2030)



# FPGA's for less consumption?

The same idea can be extended to FPGA:



Seeding algorithm for making tracklets in the last LHCb tracker (SciFi) in FPGAs:

Throughput increases by 30%

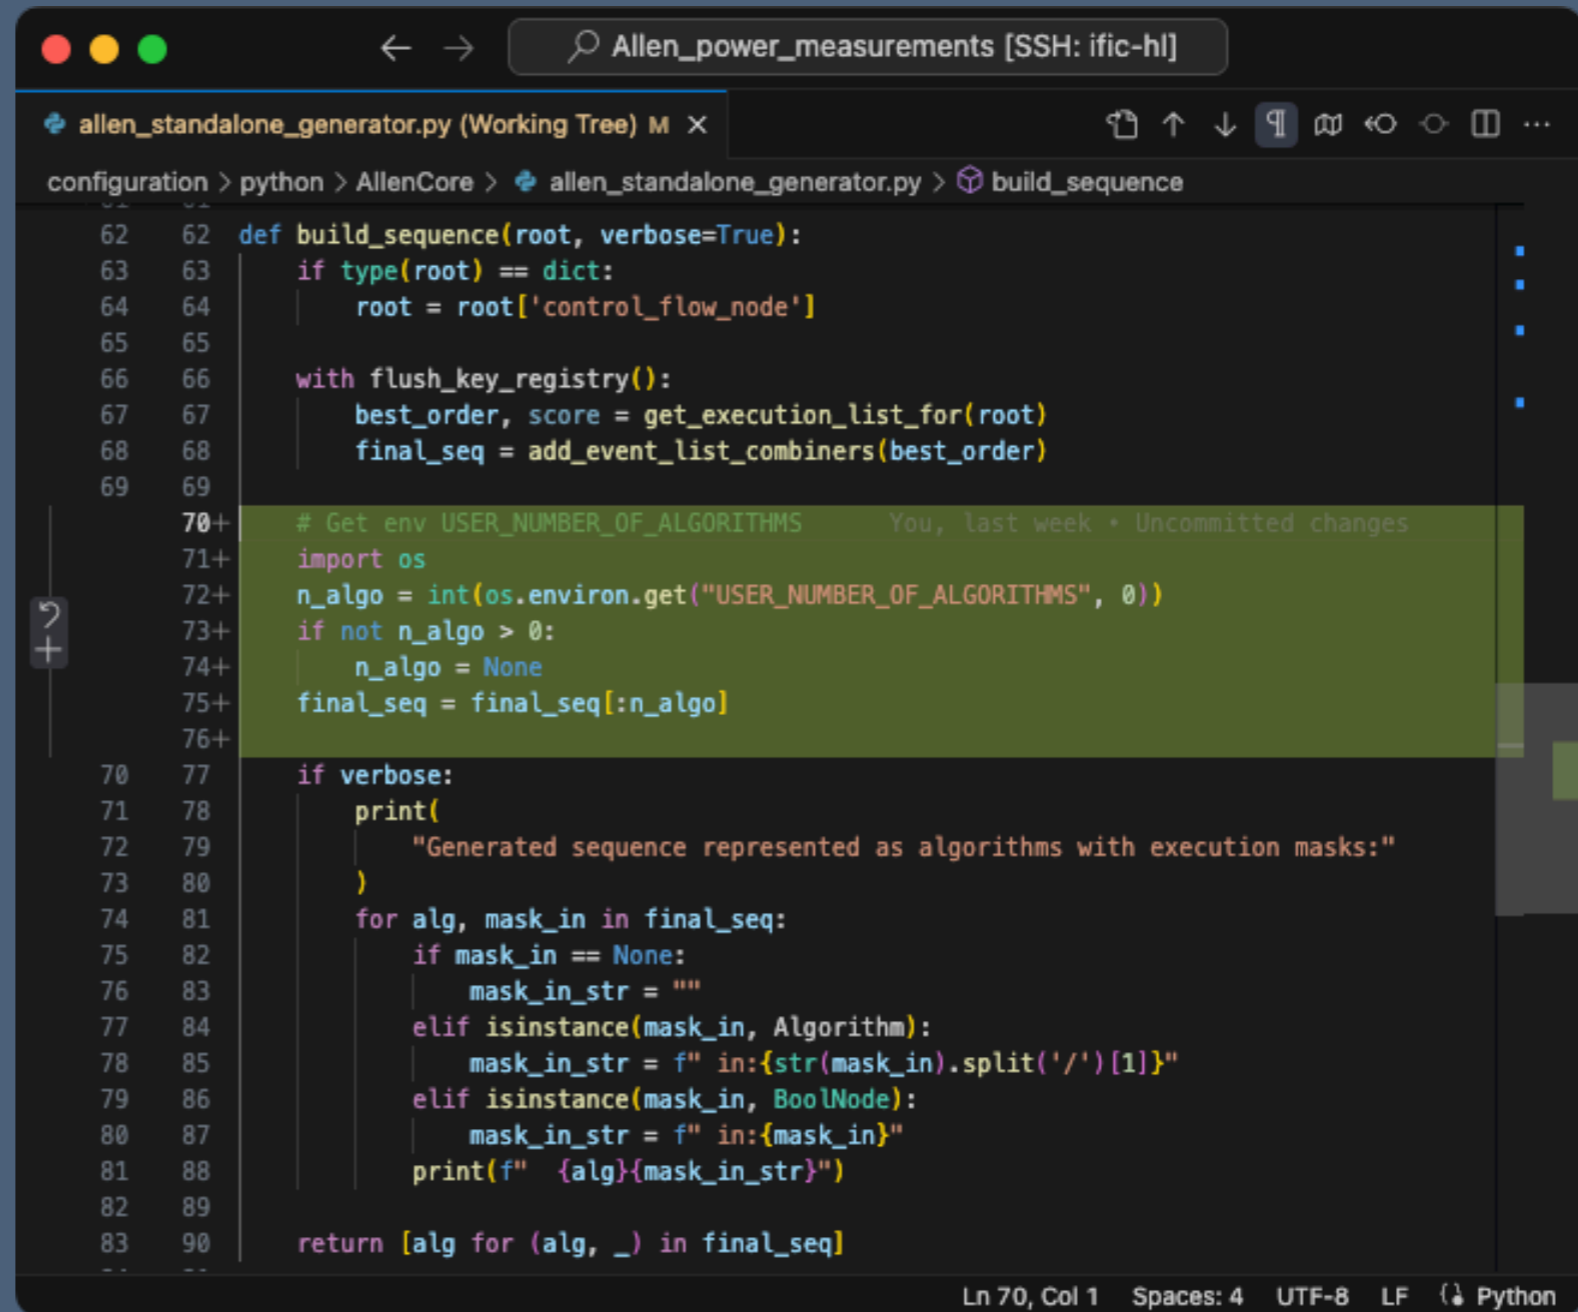
→ Saving 6.2 mW · s/event

(for 30 MHz rate: 186kW/s)

Use hybrid systems to take benefits of each one

# Per-algorithm power consumption

- Control number of executed algorithms in sequence
- Run power consumption test for first  $N$  and  $N + 1$  algorithms
- The difference is the consumption gain from algorithm  $N$
- Expected to be proportional to total execution time



```
Allen_power_measurements [SSH: ific-hl]
allen_standalone_generator.py (Working Tree) M x
configuration > python > AllenCore > allen_standalone_generator.py > build_sequence

62 62 def build_sequence(root, verbose=True):
63 63     if type(root) == dict:
64 64         root = root['control_flow_node']
65 65
66 66     with flush_key_registry():
67 67         best_order, score = get_execution_list_for(root)
68 68         final_seq = add_event_list_combiners(best_order)
69 69
70+ # Get env USER_NUMBER_OF_ALGORITHMS
71+ import os
72+ n_algo = int(os.environ.get("USER_NUMBER_OF_ALGORITHMS", 0))
73+ if not n_algo > 0:
74+     n_algo = None
75+ final_seq = final_seq[:n_algo]
76+
70 77     if verbose:
71 78         print(
72 79             "Generated sequence represented as algorithms with execution masks:"
73 80         )
74 81     for alg, mask_in in final_seq:
75 82         if mask_in == None:
76 83             mask_in_str = ""
77 84         elif isinstance(mask_in, Algorithm):
78 85             mask_in_str = f" in:{str(mask_in).split('/')[1]}"
79 86         elif isinstance(mask_in, BoolNode):
80 87             mask_in_str = f" in:{mask_in}"
81 88         print(f" {alg}{mask_in_str}")
82 89
83 90     return [alg for (alg, _) in final_seq]
```

Ln 70, Col 1 Spaces: 4 UTF-8 LF Python

# Per-algorithm power consumption

- Control number of executed algorithms in



```
Allen_power_measurements [SSH: ific-hl]
allen_standalone_generator.py (Working Tree) M
configuration > python > AllenCore > allen_standalone_generator.py > build_sequence

62 62 def build_sequence(root, verbose=True):
63 63     if type(root) == dict:
64 64         root = root['control_flow_node']
65 65
66 66     with flush_key_registry():
67 67         best_order, score = get_execution_list_for(root)
68 68         final_seq = add_event_list_combiners(best_order)
69 69
70 70     return [alg for (alg, _) in final_seq]

NUMBER_OF_ALGORITHMS You, last week • Uncommitted changes
environ.get("USER_NUMBER_OF_ALGORITHMS", 0)
> 0:
None
final_seq[:n_algo]

Generated sequence represented as algorithms with execution masks:"
alg, mask_in in final_seq:
if mask_in == None:
    mask_in_str = ""
elif isinstance(mask_in, Algorithm):
    mask_in_str = f" in:{str(mask_in).split('/')[1]}"
elif isinstance(mask_in, BoolNode):
    mask_in_str = f" in:{mask_in}"
print(f" {alg}{mask_in_str}")

return [alg for (alg, _) in final_seq]


Ln 70, Col 1 Spaces: 4 UTF-8 LF Python
```

- Expected to be proportional to total execution

Inspired by Allen algorithm breakdown!

# Per-algorithm power consumption

- Due to time limitations, only specific  $N$  were chosen for studies
- The HighLow machine is a shared resource - required to run overnight to minimize effect from other users (but still there)
- Multiple runs are necessary for stat. error estimation
- No anomaly is detected → no significant performance drops?

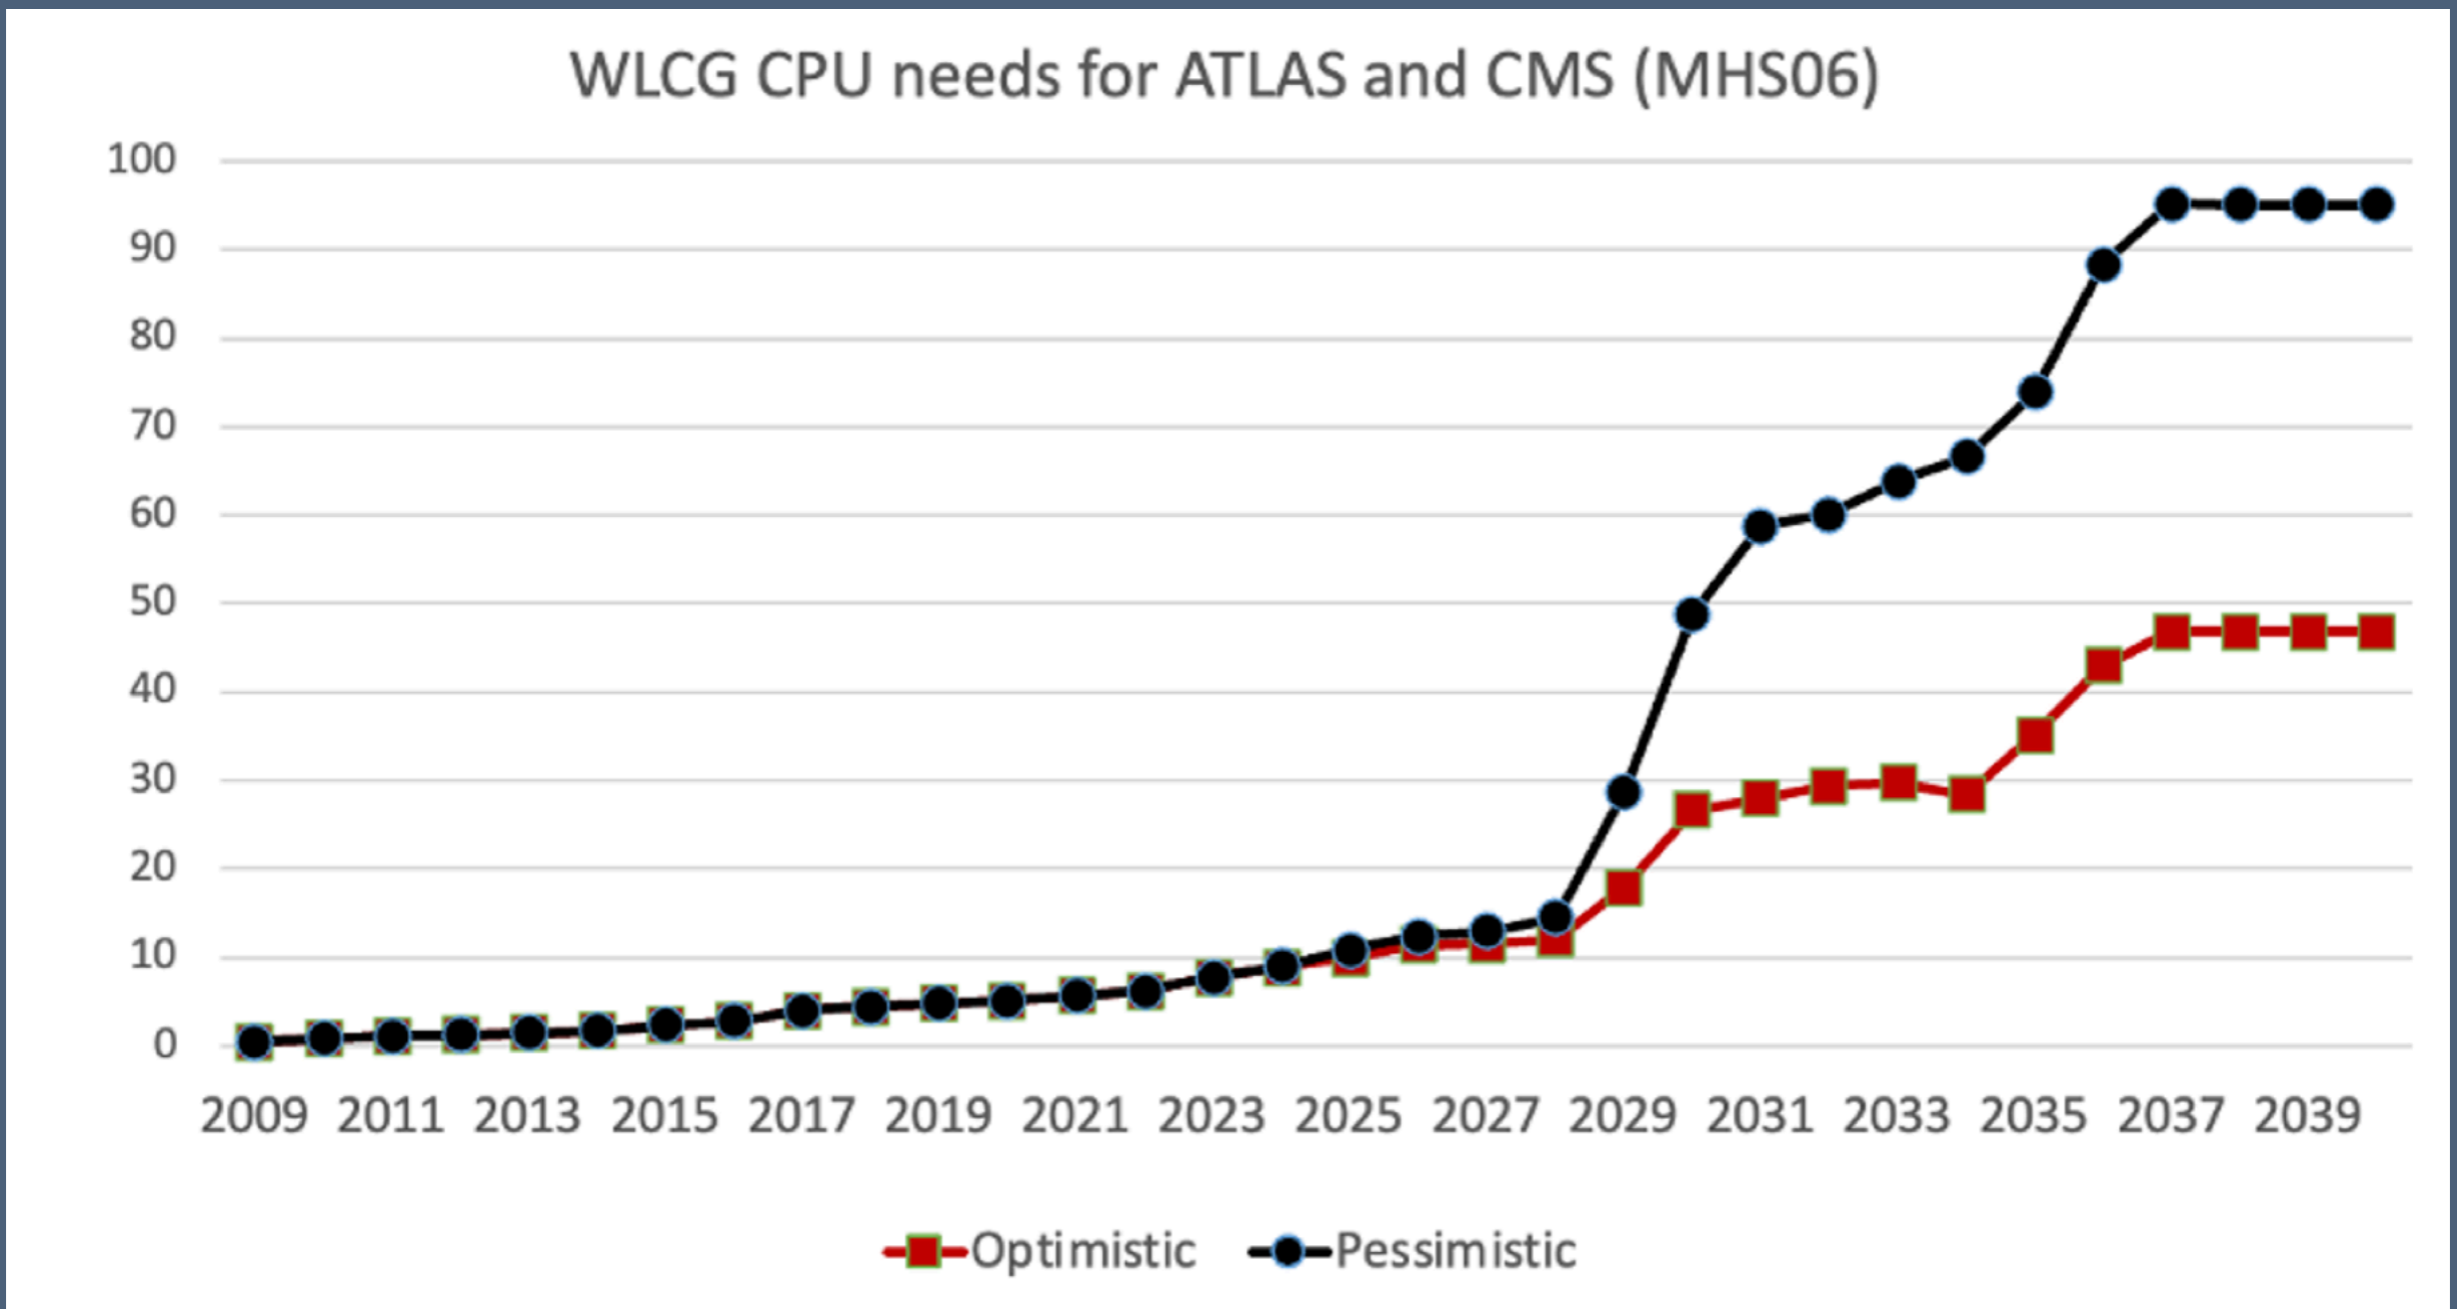




# Estimation of WLCG CPU requirements

In Millions of HS06 of ATLAS and CMS

A similar prediction needs to be made for the trigger systems



# Summary

- Move towards heterogeneous computing systems
- Use the best and more efficient available hardware (vs €)
- Optimize the utilization of the hardware
- **Keep an eye on your power consumption!**
- Planned to make Allen support more hardware architectures (ARM; Intel GPUs)
- Need to make predictions for the trigger system consumption for HL-LHC
- Work is still ongoing 😊

Thanks for your attention!