

Gábor Biró¹, Balázs Majoros^{1,2}, Gergely Gábor Barnaföldi¹
¹Wigner Research Centre for Physics, 29-33 Konkoly-Thege Miklós street, 1121, Budapest, Hungary
²Eötvös Loránd University, 1/A Pázmány Péter walkway, 1117 Budapest, Hungary

biro.gabor@wigner.hu
 majoros.balazs@wigner.hu
 barnafoldi.gergely@winger.h

Introduction

The ever growing increase of computing power necessary for the storage and data analysis of the high-energy physics experiments at CERN requires performance optimization of the existing and planned IT resources. To optimize the resource usage, the concept of Analysis Facility (AF) has been introduced. The AFs are special computing centres with a combination of CPU and fast interconnected disk storage resources, allowing for rapid turnaround of analysis tasks on a subset of data.

LHC in numbers

One of the main computing capacity consumers in the HEP software workflow is the data analysis → **WLCG**

2013 vs. now:

Data: 15 PB/y vs 200+ PB/y
Tape: 180 PB vs 740+ PB
Disk: 200 PB vs 570+ PB
HS06 hours: 2M vs 100+ B

Worldwide LHC Computing Grid and the Wigner Datacenter



- 2013-2019: Tier-0 site
- 2 independent, dedicated HV lines
- Full UPS and diesel coverage for all IT load (incl. Cooling)
- 2 x 100 Gbps to CERN
- 20 000+ CPU cores, 5.5 PB+ storage

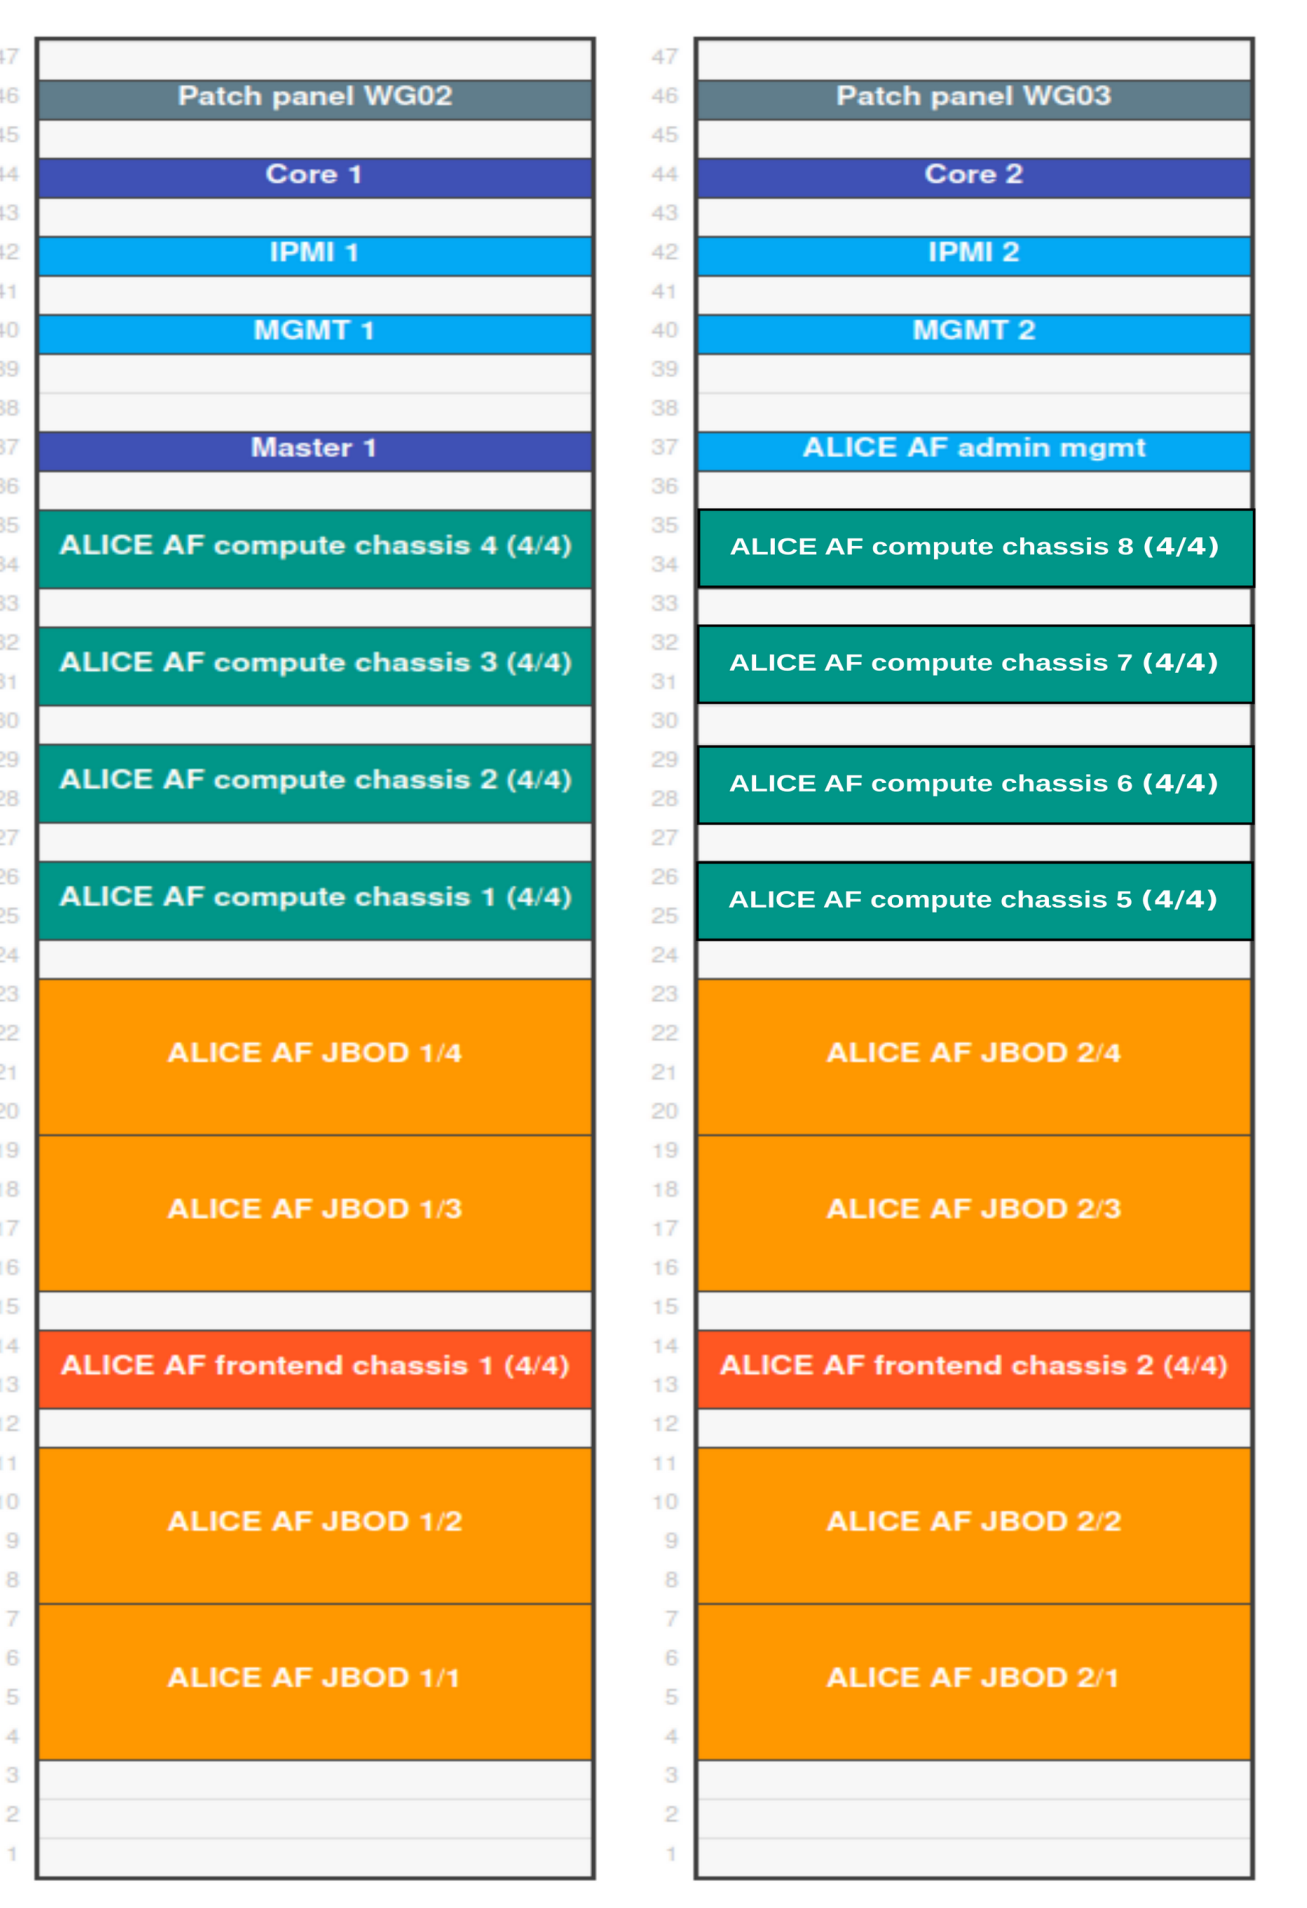
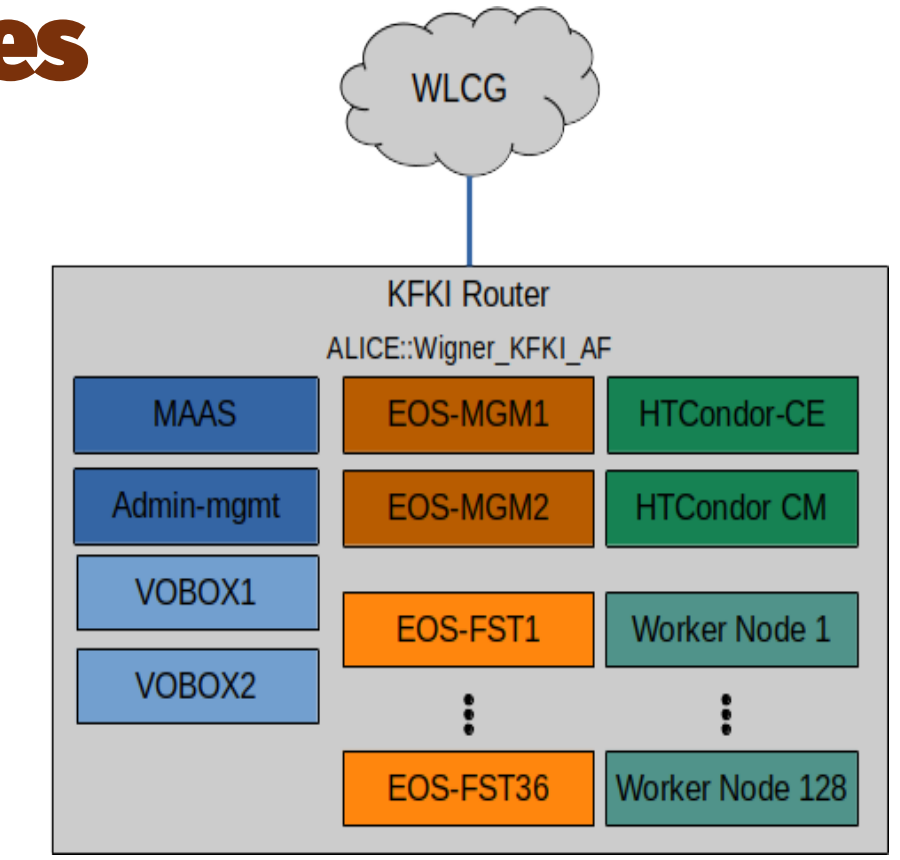
A new, specialized facility

Re-utilizing the Tier-0@ Budapest hardwares
First design: 2019Q4
Today: 9 racks optimized for maximal data throughput

Aim: maximize the throughput via optimizing the layout. With a flexible, easy to 'copy-paste' design

- Storage node: 24 x 3 TB**
Computing node:
- 32 vCPU (AMD Opteron 6276)
 - 64 GB RAM
 - 160 GB swap
 - 2 TB local storage

High-speed internal communication between the nodes:
HP ProCurve 6600-24XG SFP+ 10 GbE ports
High bandwidth of 10 Gbps within a cell and also between different cells


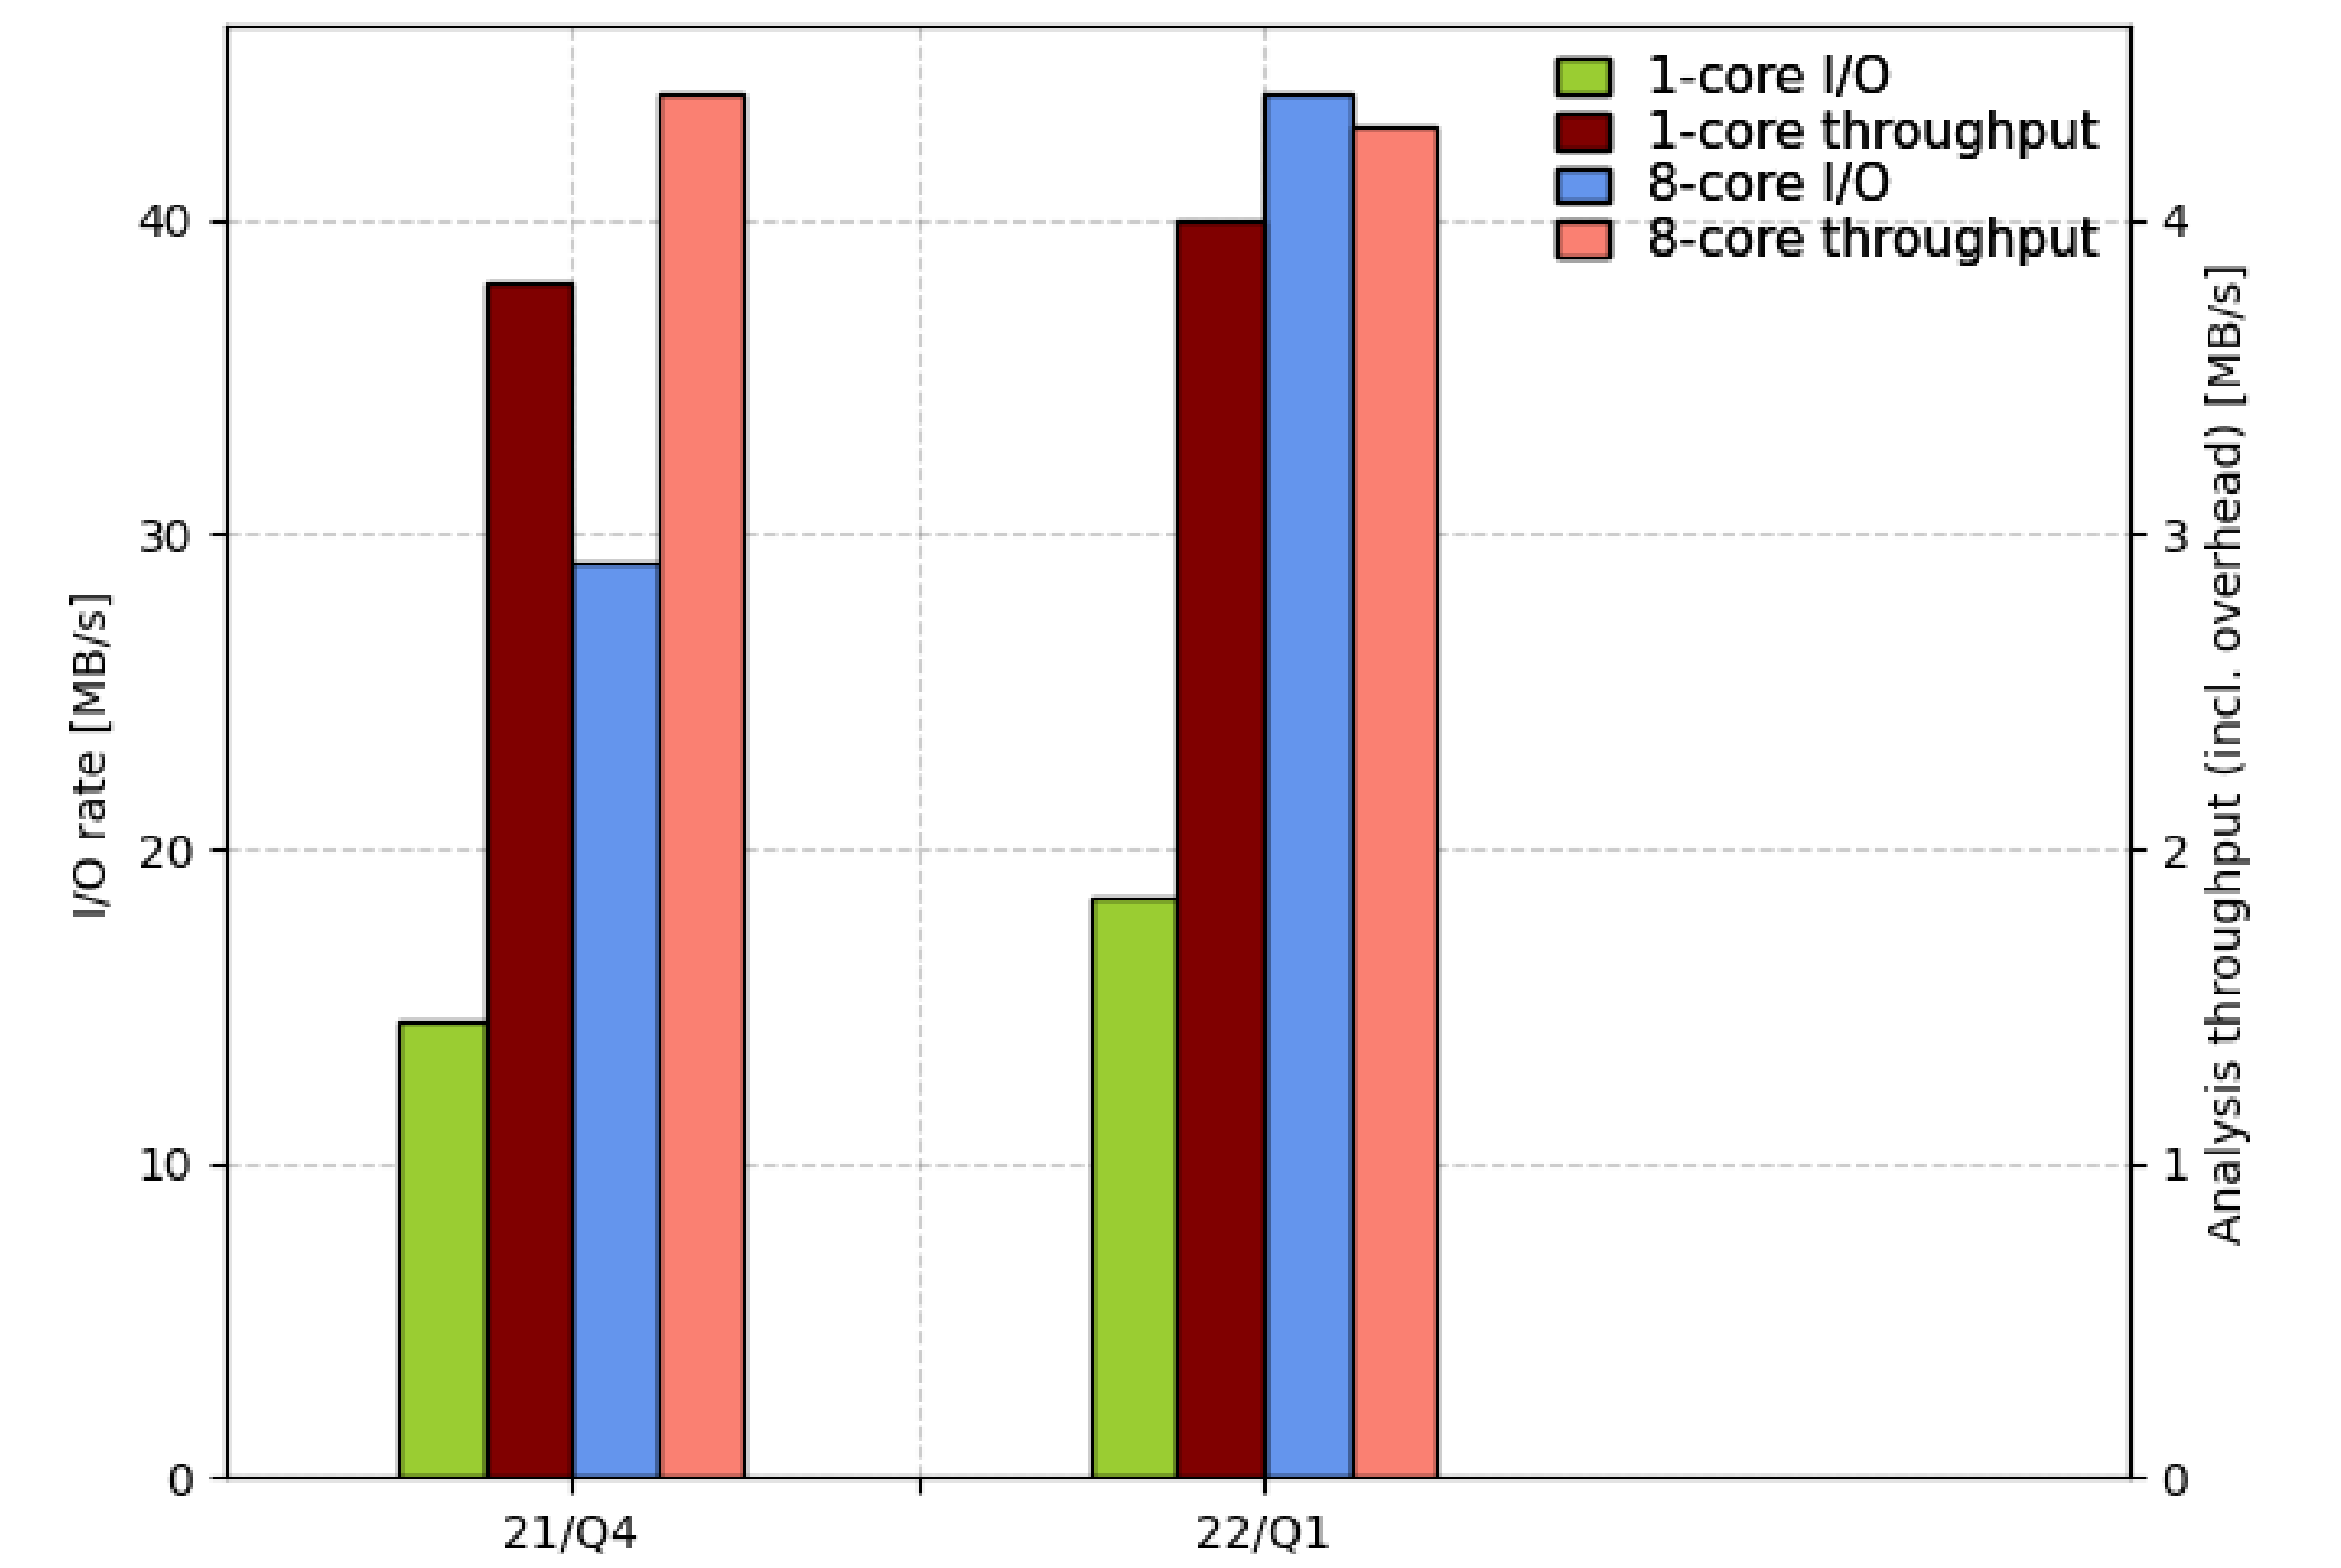


Utilization of the resources

Softwares:

- CentOS7, CVMFS, JaliEn
- Storage:
 - EOS (v4.8.62) + Xrootd (v4.12.8) config & monitoring
 - 2 redundant MGM nodes w/ QuarkDB backend
 - 36 FST nodes, with 24 x 3 TB for each node
- Raw capacity: ~2.6 PB
- Usable capacity: ~1.3 PB
- Computing:
 - Configured with HTCondor(-CE), 1 single-core queue and 1 multi-core queue
 - 128 worker nodes – 4096 vCPU

Since the start of the operation:
19M hours of delivered CPU time



Estimated daily throughput, scaled to the 8-rack setup:
0.16-1.11 PB/day

Summary

Dedicated facility for high analysis throughput
Allows for optimization of the analysis process and the codes before the analysis is performed on the large data samples on the WLCG Grid
Upcoming developments: 4 more racks and network upgrades



Acknowledgement

This work was supported by the Hungarian National Research Fund OTKA grant K135515, NIKFIH 2019-21.11-TÉT-2019-00078, 2019-21.11-TÉT-2019-00050; 2020-21.1-ED-2021-00179 and the Wigner Scientific Computing Laboratory (WSCLAB) (the former Wigner GPU Laboratory).

References

- [1] ALICE-PUBLIC-2021-007
- [2] J. Phys. Conf. Ser. 513 (2014) 012037.
- [3] J. Phys. Conf. Ser. 331 (2011) 052015.
- [4] arXiv:2109.09594