Flagship WP2.2 and shared with WP2.4 WP2.5

Piergiulio Lenzi, Alberto Annovi

with extensive input from the flagship PIs and WP2.4 WP2.5 coordinators

25 luglio 2023

Flagship usecase documents

- Tutti disponibili sotto “DocsWP2” qui:
  - https://drive.google.com/drive/folders/13_sGaGi7JHNmT3WsPB_S6TdMVUTBDh-x?usp=sharing

- UC2.2.1 Advanced Machine Learning: Flash simulation and other bleeding edge applications:
  - https://docs.google.com/document/d/1puKKyCKsLJbIhJ3dRnRjfiJWki1si_A2J2NJtavEgLg/edit?usp=sharing
- UC2.2.2 Quasi interactive analysis of big data with high throughput:
  - https://docs.google.com/document/d/1n5w8fGmuFTSMSOhne3W1m8MWI6ysEeM5qWu-EY2lK1k/edit?usp=sharing
- UC2.2.3 Development of ultra-fast algorithms running on FPGAs:
  - https://docs.google.com/document/d/1i0qke_GiJ9VOyZIDAtc-HnnLANsSBORTJt9qw_nEnqA/edit?usp=sharing
- UC2.2.4 Porting of algorithms to GPUs:
  - https://docs.google.com/document/d/14PKxndBzLfXSUM76aWfNoZISIYL1OW6UtFMfIDDPfkw/edit?usp=sharing
- UC2.2.5 Validation of HEP reconstruction code on ARM:
  - https://docs.google.com/document/d/1967M1tSrI3LrzdJ5led3z2b9mliRr7I6l4HBQfm9eEE/edit?usp=sharing

Landscape recognition document

- Link to the landscape recognition document (main editors are Piergiulio Lenzi, Alberto Annovi and Tommaso Boccali) with the status and directions of scientific computing in experimental HEP
- It is still a bit of a patchwork, and needs some harmonization
  - Figures are missing
  - Some text on the challenges in non-LHC experiments is missing
- Guesstimate of flagship-person association here

Flagship: “Advanced Machine Learning: Flash simulation and other bleeding edge applications”

Flagship: “Advanced Machine Learning: Flash simulation and other bleeding edge applications”

- Use case ID: UC2.2.1
- Flagship PI: Lucio Anderlini (FI)
- Required personnel
  - Expected total manpower needs: 7 FTE
    - Secured (already existing & committed): 3
    - To be recruited: 4 FTE

- Required resources
  - see the document

- Synergies:
  - Data management and cloud-based resource provisioning will require strong synergy with WP5.
  - Activities in WP6 on the porting of Geant4 on GPU are also synergic as long as DNN evaluated on GPUs are used as universal approximators of complicated functions.
- Link: https://docs.google.com/document/d/1puKKyCKsLJbIhJ3dRnRjfiJWki1si_A2J2NJtavEgLg/edit?usp=sharing

Flagship: “Advanced Machine Learning: Flash simulation and other bleeding edge applications”

Brief summary form the description in the document:

- One of the most computing demanding tasks in Particle Physics consists of the simulation of the detector simulation
- Preliminary estimates indicate that without a faster option to produce simulated samples, the community would compromise on the physics-reach
- Flash simulation approaches dramatically reduce the required computing need, however they requires scaling up the computing resources used for the development and training campaigns
- We collect in this use-case also other machine-learning applications relevant to WP2 which would strongly benefit from accessing the setup prepared to enable the research on machine-learning-based detector simulation.

Flagship: “Advanced Machine Learning: Flash simulation and other bleeding edge applications”

Milestones & Deliverables

Month 22. Initial findings and report on the development status.

- The activities are completely outlined and integrated with the hardware made available to the project. A public software repository is identified and populated with the developed solutions.
- Deliverables:
  - Detailed plan
  - Status report
  - Public software repository

Month 36. Simulation validation and performance studies

- The impact of the solutions developed with ICSC resources on the scientific collaborations and communities is documented in a final report. The software is released in a public repository with documentation.
- Deliverables:
  - Performance report with comparisons with previous solutions
  - Software released on public repository

Flagship: “Advanced Machine Learning: Flash simulation and other bleeding edge applications”

KPIs

KPI ID	Description	Acceptance threshold
KPI2.2.1.1	N_{MC} billion events obtained from ML-based simulation	N_{MC} >= 1
KPI2.2.1.2	N_{EXP} experiments have developed a machine-learning based simulation	N_{EXP} >= 2
KPI2.2.1.3	Machine-learning use-cases developed with CN resources were presented at N_{CONF} international and national events	N_{CONF} >= 5
KPI2.2.1.4	N_{UC} different machine-learning use-cases were developed using CN resources and made available in git repositories	N_{UC} >= 5

Flagship: “Porting of algorithms to GPUs”

Open points:

- First draft available to be commented

Flagship: “Quasi interactive analysis of big data with high throughput”

Flagship: “Quasi interactive analysis of big data with high throughput”

- Use case ID: UC2.2.2
- Flagship PI: Tommaso Diotalevi (BO), Francesco G. Gravili (LE)
- Required personnel
  - Expected total manpower needs: 11 FTE
    - Secured (already existing & committed): 8.5 FTE
    - To be recruited: 2.5 FTE

- Required resources
  - see the document

- Synergies:
  - WP5: testing the framework, together with the AF developers, and improving the technical documentation
- Link: https://docs.google.com/document/d/1n5w8fGmuFTSMSOhne3W1m8MWl6ysEeM5qWu-EY2lK1k/edit?usp=sharing

Flagship: “Quasi interactive analysis of big data with high throughput”

Brief summary form the description in the document:

- Change of paradigm for analysis: from batch jobs followed to manual harvesting of the results to interactive analysis facilities where the interaction with the compute is completely decoupled
- Leveraging industry standards: HTCondor, DASK, Apache Spark, Jupyter Notebooks
- Modern data reduction tools such as RDataFrame and Coffea are developing in parallel
- Existing example on INFN resources for CMS
- Replica for ATLAS being developed in Naples
- The aim is to put in production an experiment-independent facility based on ICSC resources

Flagship: “Quasi interactive analysis of big data with high throughput”

Milestones & Deliverables

- Identifying target analyses (from different experimental collaborations) and porting a first complete analysis on a first operational AF testbed: 0-6 months;
- Adding more experimental applications to the existing infrastructure testbed as well as evaluating performance indicators for speed-up compared to traditional approaches: 6-12 months;
- Harmonization of the different efforts into a single infrastructure, also consolidating the relative documentation, aiming for a multitenant service. Concurrently, organizing a first hands-on workshop on AF usage for users: 12-XX Months;
- Scaling the local testbed facilities to a national-wide data lake infrastructure. Scientific reports at conferences and/or workshops: XX-YY Months

Flagship: “Quasi interactive analysis of big data with high throughput”

KPIs

KPI ID	Description	Acceptance threshold
KPI2.2.1.1	Implementation of N data analyses in the AF	$N \geq 2$
KPI2.2.1.2	Reference documentation of the AF	≥ 1 dedicated web pages
KPI2.2.1.3	Hands-on workshops for AF users	≥ 1 workshops
KPI2.2.1.4	Scaling up the testbed AF infrastructure, serving k tenants, for a total of N data analyses	$\geq (200 \cdot N)$ cores
KPI2.2.1.5	Talks at conferences/workshops about AF activities	≥ 1 talk

Flagship: “Quasi interactive analysis of big data with high throughput”

Open points:

- First draft available to be commented
- Involvement of some institutions to be checked

Flagship: “Development of ultra-fast algorithms running on FPGAs”

Flagship: “Development of ultra-fast algorithms running on FPGAs”

- Use case ID: UC2.2.3
- Flagship PI: Bernardino Spisso (NA), Simone Gennai (MiB)
- Required personnel
  - See document

- Required resources
  - See document

- Synergies:
  - WP5: Exploitation of GPU resources (single node or multi node) ie. for ML trainings
  - WP4: FPGA and GPU training
- Link: https://docs.google.com/document/d/1i0qke_GiJ9VOyZIDAtc-HnnLANsSBORTJt9qw_nEnqA/edit?usp=sharing

Flagship: “Development of ultra-fast algorithms running on FPGAs”

Brief summary form the description in the document:

- This use case explores the potential applications of FPGA technology in HEP, with a specific focus on the development of reconstruction algorithms, trigger algorithms, and data acquisition.
- Additionally, there is a need to interface FPGAs with fast transfer protocols to efficiently transmit processed data to computing servers.
- It is crucial to emphasize the critical need for personnel training in FPGA technology to fully leverage its potential in HEP research. By equipping researchers and engineers with FPGA expertise, the field can maximize the benefits offered by FPGA technology in terms of speed, efficiency, and real-time processing capabilities.

- The activities are then divided into three main topics:
  - Trigger, DAQ and on-line processing
  - High level FPGA tools
  - Low level FPGA tools

Flagship: “Development of ultra-fast algorithms running on FPGAs”

KPIs

The KPI, for those activities which adopt them, are the same for all the tree sub-group

KPI ID	Description	Acceptance threshold
KPI2.2.1.1	Development of triggering algorithms, on-line analyses, data acquisition on FPGA	10%
KPI2.2.1.2	Development of new high level tools for FPGA	reduced by XX% with respect to [33]
KPI2.2.1.3	Development of new low level tools for FPGA	1
KPI2.2.1.4	Training: create a pool of FPGA programming experts	5

Flagship: “Development of ultra-fast algorithms running on FPGAs”

Open points:

- First draft available to be commented
- Collection of information from involved people and groups in progress
- Personnel, computing resources, milestones and deliverables TBD once more information will be collected

Flagship: “Porting of algorithms to GPUs”

Flagship: “Porting of algorithms to GPUs”

- Use case ID: UC2.2.4
- Flagship PI: Adriano Di Florio (BA)
- Required personnel
  - Expected total manpower needs: 8 FTE (o per mese)
    - Secured (already existing & committed): 5.5 FTE.
    - To be recruited: 2.5 FTE
- Required resources
  - to be defined
- Synergies:
  - porting code to GPUs WP4
- Link: https://docs.google.com/document/d/14PKxndBzLfXSUM76aWfNoZlSIYL1OW6UtFMfIDDPfkw/edit?usp=sharing

Flagship: “Porting of algorithms to GPUs”

Brief summary form the description in the document:

- Intro on HL-LHC computing challenge and the benefits of Heterogeneous computing
- The CMS software framework (CMSSW) has been extended in Run-3 to offload part of the physics reconstruction to NVIDIA GPUs.
- Alpaka has been adopted for portability. It supports serial and parallel execution on CPUs, and extremely parallel execution on GPUs.
- This flagship will foster a more extensive use of GPUs. It will train more personnel to write dedicated GPU code. It will facilitate inter-experiment exchange of experience. Finally effort is allocated to further develop the offloading of code to GPUs.

Flagship: “Porting of algorithms to GPUs”

- KPI →
- Milestones & Deliverables
  - TBD

KPI ID	Description	Acceptance threshold
KPI2.2.1.1	Porting di almeno un algoritmo di ricostruzione offline di XX esperimento di LHC a GPU (sarebbe CMS)	XX>=1
KPI2.2.1.?	Infrastructure di (cross)validation per XX platforms (AMD, nVIDIA, CPU) per un esperimento di LHC [presumibilmente CMS]	XX>=3
KPI2.2.1.2	Creare un pool di almeno XX persone in grado di sviluppare codice con ALPAKA (WP4)	XX>=10
KPI2.2.1.3	Porting di almeno un algoritmo di trigger di XX esperimento di LHC a GPU (sarebbe ATLAS)	XX>=1

Flagship: “Porting of algorithms to GPUs”

Open points:

- First draft available to be commented
- Personnel: complete with a list of institute and available people ?
- Required ICSC resources TBD
- Milestones TBD

Flagship: “Validation of HEP reconstruction code on ARM”

Flagship: “Validation of HEP reconstruction code on ARM”

- Use case ID: UC2.2.5
- Flagship PI: Francesco Noferini (BO)
- Required personnel
  - Expected total manpower needs: 1 FTE
    - Secured (already existing & committed): 1
    - To be recruited: 0
- Required resources are already secured
  - 200 hours of 4 arm machines installed at CNAF Tier-1
  - Access to experiments storage at CNAF Tier-1 guaranteed
- Synergies:
  - synergy within an open-call with E4
  - to be completed
- Link: https://docs.google.com/document/d/1967M1tSrI3LrzdJ5led3z2b9mliRr7I6l4HBQfm9eEE/edit?usp=sharing

Flagship: “Validation of HEP reconstruction code on ARM”

Brief summary form the description in the document:

- The use of ARM architecture in HEP represents a big opportunity to move towards more sustainable computing models in terms of energy saving.
- Power consumption reduction at the level of 30-60% was recently shown
- All LHC experiments are already providing software builds able to run with this architecture.
- Validation in full GRID submission chain is required and still missing.
- The goal of this use case is to provide such a validation at least for two major LHC experiments in the CNAF Tier-1 facility.

Flagship: “Validation of HEP reconstruction code on ARM”

- Milestones & Deliverables
  - The activities on the use case will include:
  - First period (tentatively month 1-12): configuring ARM machines at CNAF to provide access to the experiment software and storage (qui dobbiamo capire interplay tra xrootd e gpfs)
  - Second Period (tentatively month 12-24): testing and validating the use case (data reconstruction and simulation) , setting and validating submission via HTCondor CE and
  - moving to full production
- KPIs

KPI ID	Description	Acceptance threshold
KPI2.2.1.1	Software validation on ARM in the full GRID chain	50% (2/4 LHC experiments)
KPI2.2.1.2	Presentation at conferences	>=2
KPI2.2.1.3	Technical notes	>=2

Flagship: “Validation of HEP reconstruction code on ARM”

Open points:

- First draft available to be commented

Backup