Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani
ICSC

Flash simulation and bleeding edge machine learning applications

December status report

Lucio Anderlini
Istituto Nazionale di Fisica Nucleare, Sezione di Firenze

External Partner

Who we are

Staff members:

- Alessandro Bombini [j], INFN
- Giuseppe Piparo[l], INFN
- Maurizio Martinelli [a], Università Milano Bicocca
- Simone Capelli [a], Università Milano Bicocca
- Federica Maria Simone [i], Politecnico di Bari
- Nicola De Filippis [i], Politecnico di Bari
- Vieri Candelise [h], Università di Trieste
- Giuseppe Della Ricca [h], Università di Trieste
- Valentina Zaccolo [k], Università di Trieste
- Mattia Faggin [k], Università di Trieste
- Lorenzo Rinaldi [e], Università di Bologna
- Piergiulio Lenzi [g], Università di Firenze
- Vitaliano Ciulli [g], Università di Firenze
- Sharam Rahatlou[h], Università Roma 1
- Daniele del Re [h], Università Roma 1
- Lorenzo Capriotti [f], Università di Ferrara
- Francesco Conventi [e], Università di Napoli
- Francesco Cirotto [e], Università di Napoli

PhD students:

- Francesco Vaselli [c], Scuola Normale Superiore di Pisa
- Matteo Barbetti [b], Università di Firenze
- Muhammad Numan Anwar [j], Politecnico di Bari
- Benedetta Camaiani [g], Università di Firenze
- Alkis Papanastassiou [g], Università di Firenze
- Antonio D'Avanzo [e], Università di Napoli

External collaborators:

- Andrea Rizzi [c], Università di Pisa

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani
ICSC

Products since last update

- New presentations
  - Alkis Papanastassiou, *"Anomaly detection with Autoencoders for Data Quality Monitoring in HEP"*, From Physics to Medicine: XAI workshop, 2023-11-21, Milano
    https://indico.cern.ch/event/1312529/timetable/#10-anomaly-detection-with-auto
    Flagship's student
  - Benedetta Camaiani, "Example of Adaptation domain in High Energy Physics", From Physics to Medicine: XAI workshop, 2023-11-20, Milano
    https://indico.cern.ch/event/1312529/timetable/#3-example-of-adaptation-domain
    Flagship's student
  - Lucio Anderlini, "Generative models at the LHC", ALPACA Workshop, 2023-11-23, Trento
    https://indico.ectstar.eu/event/184/contributions/4356/
    With ICSC credits
  - Stefano Giagu, "Introduction to transformers", Fifth hackathon of ML_INFN, 2023-11-14, Pisa
    https://agenda.infn.it/event/37650/contributions/212823
    With ICSC credits
  - Andrea Rizzi, "An overview of Machine Learning in High Energy Physics", Fifth hackathon of ML_INFN, 2023-11-16, Pisa
    https://agenda.infn.it/event/37650/contributions/211047/
    Flagship's ext. member

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani
ICSC

KPIs

KPI ID	Description	Acceptance threshold	2023-09-26
KPI2.2.1.1	N_{MC} billion events obtained from ML-based simulation, as demonstrated by official links in experiments' simulation databases	N_{MC} >= 1	0 events (completed: 0%)
KPI2.2.1.2	N_{EXP} experiments have tested a machine-learning based simulation	N_{EXP} >= 2	0 experiment (completed: 0%)
KPI2.2.1.3	Machine-learning use-cases tested in the context of the CN were presented at N_{CONF} international and national events	N_{CONF} >= 3	3 use-cases (since Sept. '23) (completed: 100%)
KPI2.2.1.4	N_{UC} different machine-learning use-cases were tested in the context of the CN and made available in git repositories	N_{UC} >= 5	3 use-cases (completed: 60%)

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani
ICSC

KPIs

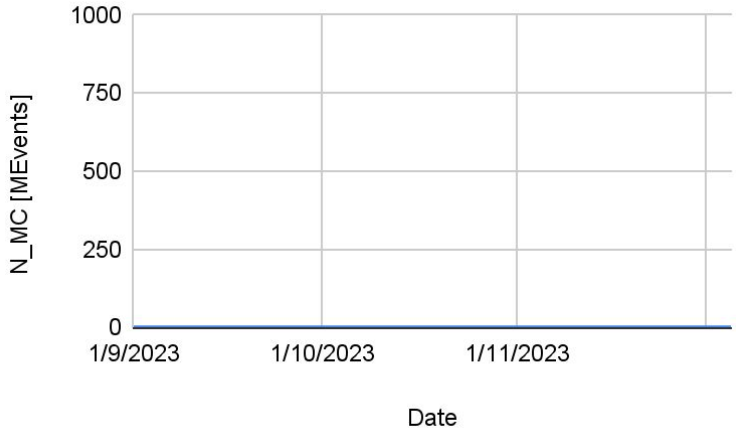

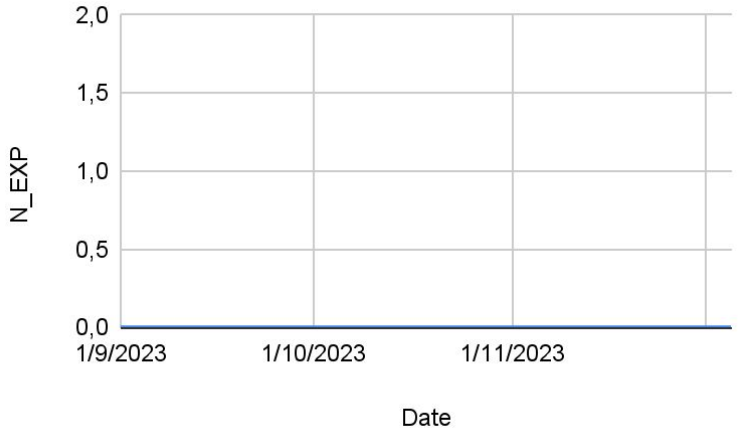

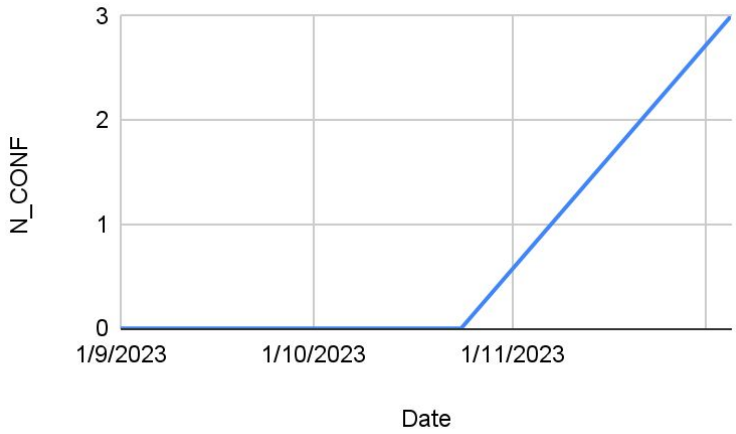

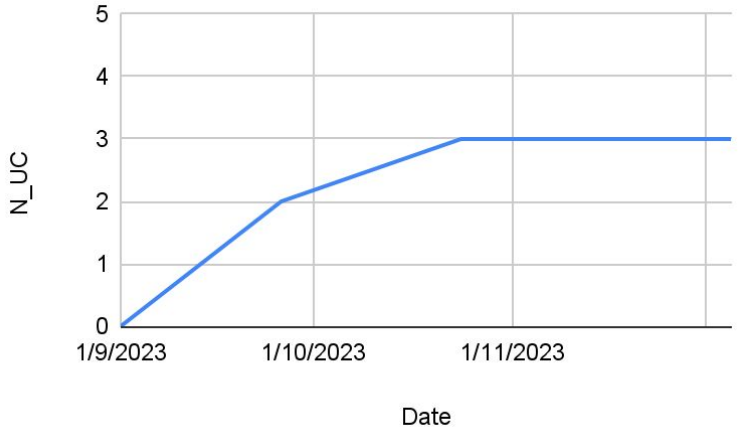

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani
ICSC

Update on the infrastructure

In the programme of the flagship, while waiting for HPC bubbles to become available we have to:

- develop the software to **train and validate the models** on ICSC resources
  - Work is ongoing, but still internal to the collaborations and won't be updated today
- commission the **infrastructure to run on HPC bubbles and Leonardo**
  - Develop and test a provisioning model for interactive development of GPU-powered ML software
  - Develop and test an batch system to use the GPUs dedicated to interactive applications when not used (opportunistic batch usage)
  - Develop and test an on-demand batch system with offloading to HPC centers (Leonardo)

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani
ICSC

Interactive usage

- We are running a cluster for interactive development machine since June, provisioning a single **A100 GPU partitioned in 7 instances**
- It is used for development and training on a regular basis. **We start suffering from the limitation to 7 GPUs** (we'll upgrade it in January).
- A second cluster is provisioning 2 non-split A100 for studies requiring more than 32 GB of GPU memory (Virgo and Innovation Grant ENI-PIML).
- TODO: *Improve documentation, stability and transparent scalability.*

User	Admin	Last Activity	Running (5)
Add Users			Stop All
anderlinil	admin	a day ago	start server
dalpra	admin	7 months ago	start server
petrinir	admin	Never	start server
postuma	admin	20 days ago	start server
sgaravat	admin	Never	start server
spiga	admin	5 months ago	start server
verlato	admin	5 months ago	start server
aespis		12 days ago	start server
alberto.sciuto		Never	start server
apaccagn		an hour ago	stop server
aviolle		2 months ago	start server
barbetti_matteo		6 months ago	start server
bardelli_giulio		5 days ago	start server
bombini		a few seconds ago	stop server
buti		a few seconds ago	stop server
ciangottini		5 months ago	start server
flizzi		24 days ago	start server
fnapolit		Never	start server
gianelle		Never	start server
grdeluca		2 months ago	start server
kellotomas		18 hours ago	stop server
lizzo_mattia		2 months ago	start server
lomarini		a month ago	start server
malai		21 hours ago	stop server
passaleva		3 months ago	start server
scapelli		6 days ago	start server
smanti		20 hours ago	start server
viliani		5 months ago	start server
zafara		a month ago	start server

Displaying users 1 - 29 of 29

Opportunistic usage of resources

In the longer term, **jobs longer than a few hours should be submitted as batch jobs**, queued and executed on resources left free from interactive usage.

We have a prototype of queueing system based on kueue (native Kubernetes batch).

Eviction of batch jobs upon connection of an interactive user: **in place**.
Fair-share configuration: **drafted** (it works, needs tuning).

Custom WebUI integrated in JupyterLab (very rough).

A lot to improve on user experience…

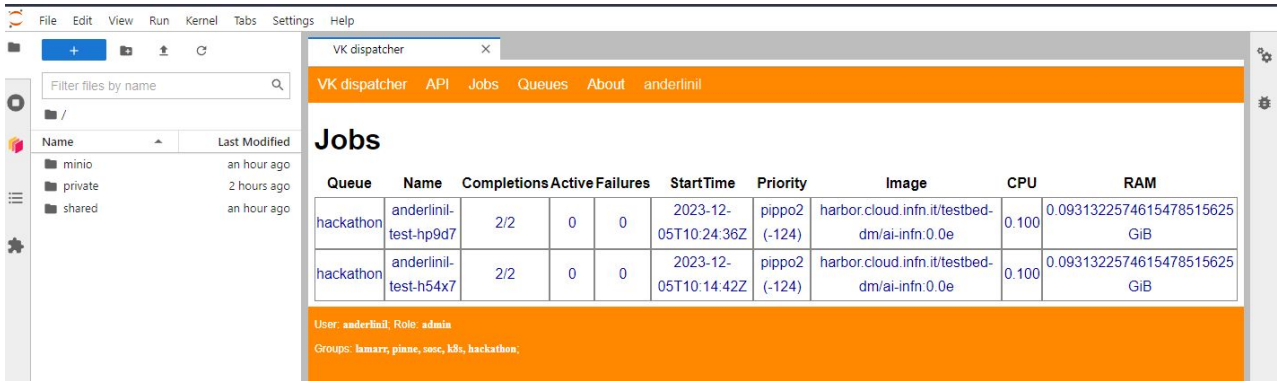

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani
ICSC

Opportunistic usage of resources: *filesystem snapshots*

The kubernetes-native submission system enables **cloning the user environment and file-system and running the job** as if the user was in front of the notebook.

While very simple and effective, this may lead to problems if the user filesystem is changed (in interactive mode) between the submission and the execution of the job.

We have a first prototype of a mechanism to take **"snapshots" of the filesystem at submission time**, and mount it at during the execution (job output retrieval is still to be implemented).

More testing is needed.

Offloading to Leonardo… work in progress…

Jobs not relying on shared filesystem (but using snapshots) should be ready to be submitted through *virtual kubelets* (InterTwin, Perugia, Spoke2-WP5).

The **caching of the snapshots**, especially for the software environments, should happen with a dedicated mechanism to avoid sending all the environments for each single job.

Caching mechanism based on docker and cvmfs is in place and needs to be tested.
At submission time, a custom docker image is generated copying the public information from the filesystem, and is uploaded to harbor.cloud.infn.it.
The upload operation triggers the deployment of the image on cvmfs *(M. Verlato, PD)*.
The remote kubelet can use the image from cvmfs with per-center, layered caching.

Early testing is in progress.

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani
ICSC

Take part

Developments are ongoing in ML_INFN, that will become AI_INFN since January.

If you wish to contribute, please join the kick-off meeting, on December 15.

https://agenda.infn.it/event/38651/