Finanziato dall'Unione europea
NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA
ICSC
Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

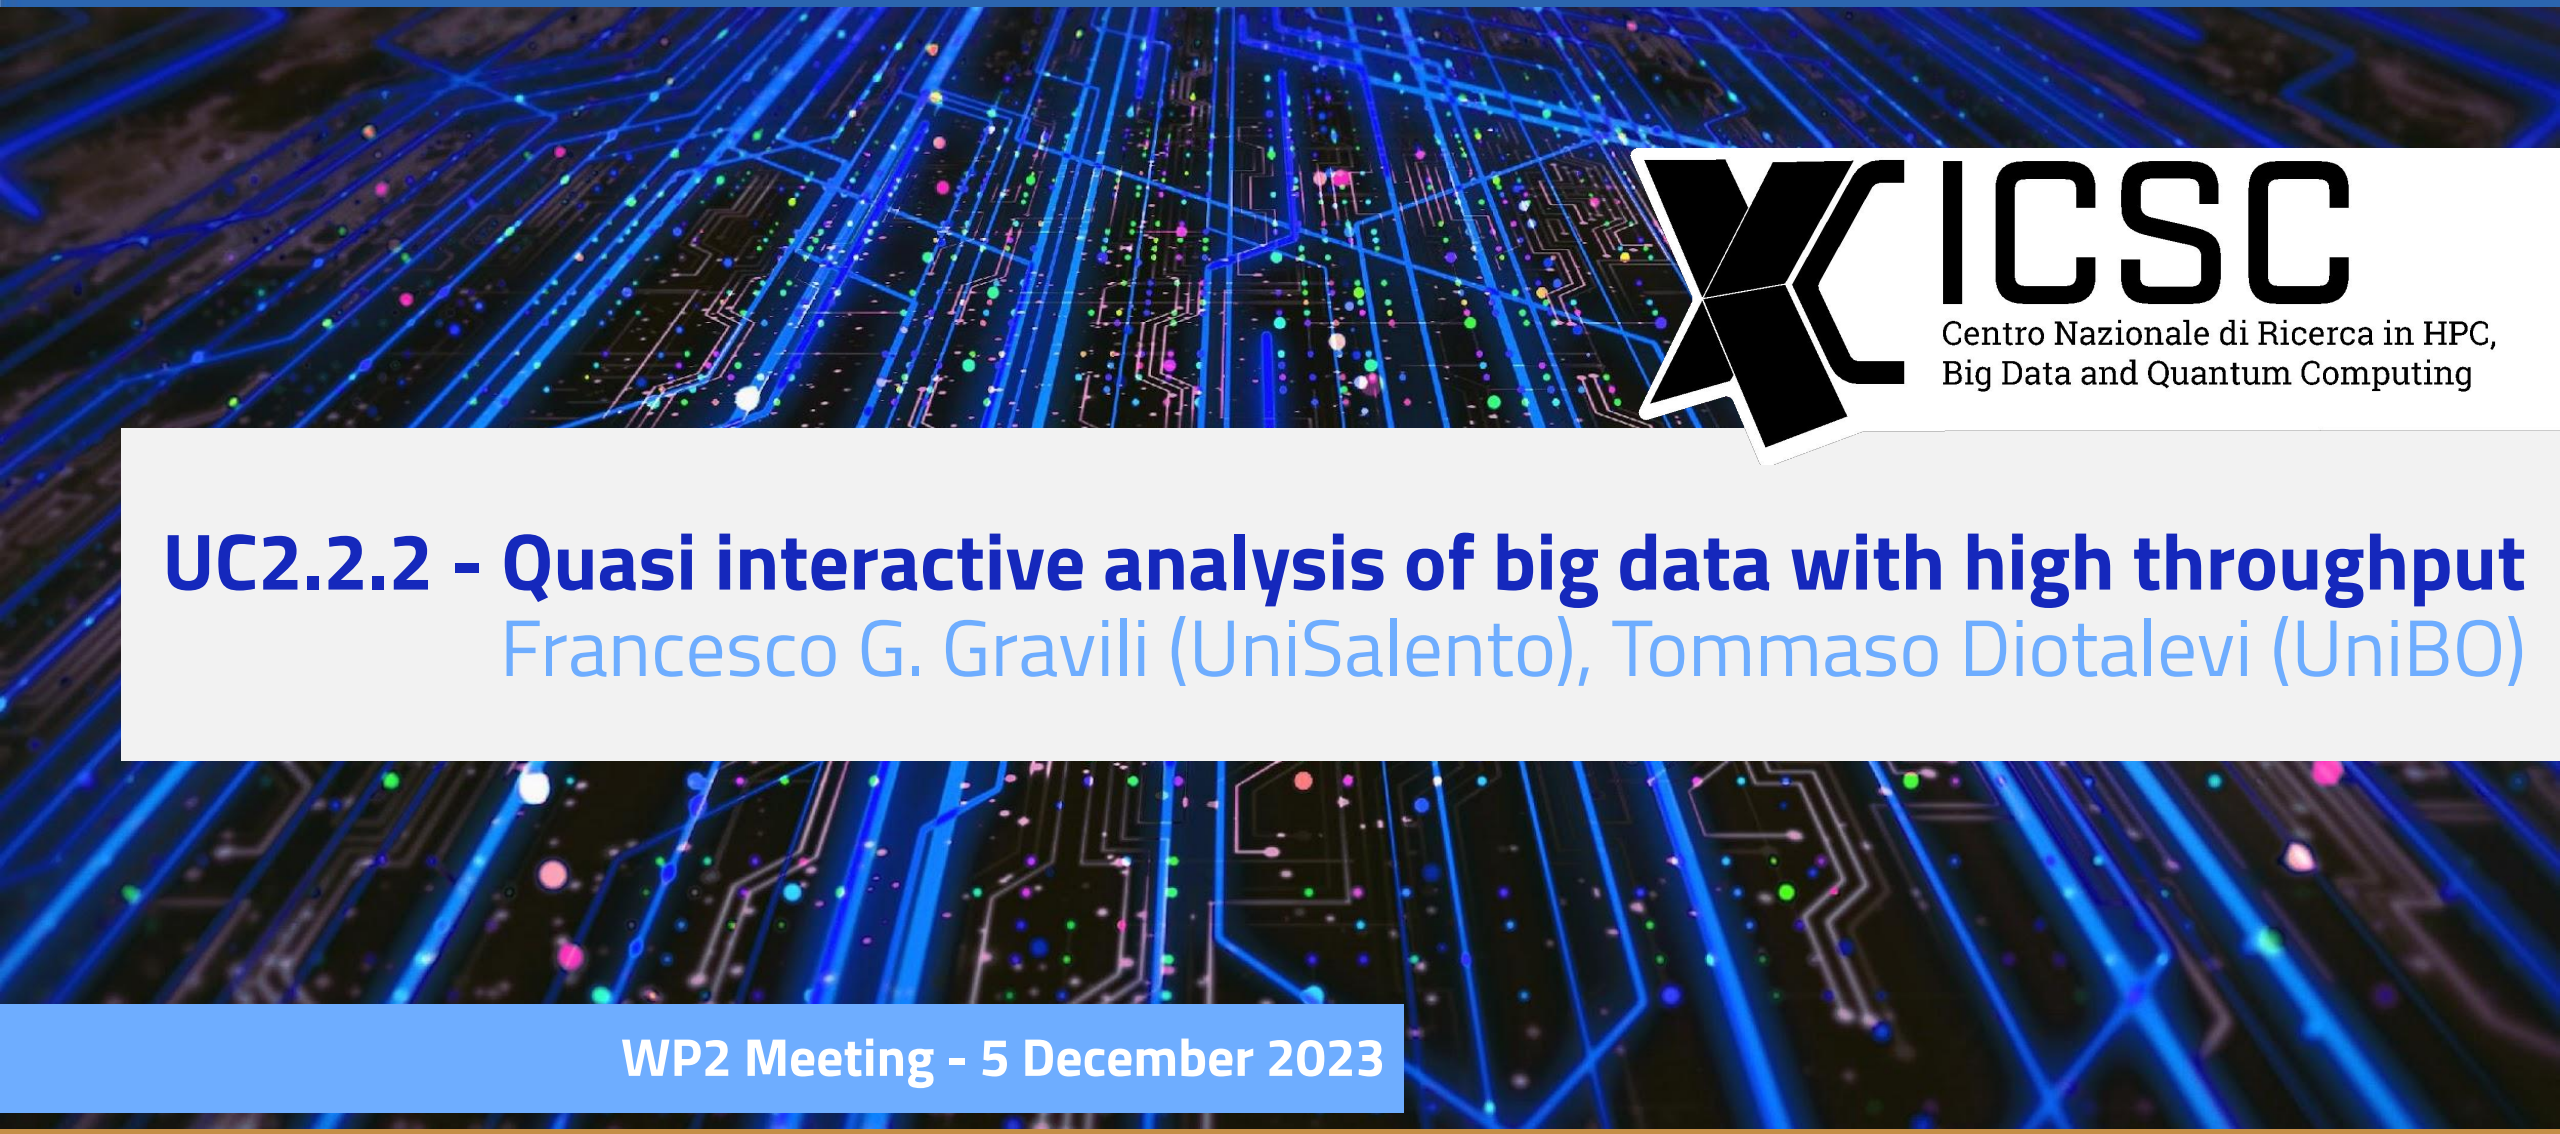

ICSC
Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

UC2.2.2 - Quasi interactive analysis of big data with high throughput

Francesco G. Gravili (UniSalento), Tommaso Diotalevi (UniBO)

WP2 Meeting - 5 December 2023

Activities ongoing

In expansion...

CMS:

- Vector Boson Scattering ssWW analysis in hadronic tau and light lepton (*INFN e Univ. Perugia*);
- Heavy Neutral Lepton search on heavy neutrinos in the D_S decays (*INFN e Univ. Bologna*); } Joint WP2 meeting (2 May)
  - Studies for performance comparison speed-up, using in-site metrics and RDF monitoring, talk WP5 (3 Nov);
- Search of rare events in $\tau \rightarrow 3\mu$ (*INFN e Polit. Bari*). Explorative phase with RDF, currently porting the old ntupliser to flat data format.
- **NEW**: Continuous Integration pipeline, triggering analysis execution on AF (*INFN, Firenze*) — REPORT TODAY
- top quark+MET analysis (*INFN e Univ. Napoli*). Still in an explorative phase, talk WP5 (24 May); — REPORT TODAY

ATLAS: — REPORT TODAY

- Search for new phenomena in events with two opposite-charge leptons, jets and missing transverse momentum (*UniSalento e INFN Lecce*). *Currently in a joint effort with INFN e Univ. Napoli.*
- *UniCal*: possible applications within several subjects: top physics, flavour tagging, trigger analysis.

Future colliders:

- Feasibility studies started on FCCee pseudo-data (*INFN e Univ. Napoli*), exploiting Naples facility. Approach already adopting *Dask, Jupyter and ROOT RDataFrame*. talk WP5 (13 Oct). — REPORT TODAY

LHCb:

- Survey still ongoing (planning to contact more institutions). If someone is interested to the UC activities, please contact us!

News from WP5

Good news from WP5!

- A first batch of computing resources (from INFN Cloud) has been acquired for implementing a first ICSC-Spoke2 playground Analysis Facility. Talk WP5 (13 Oct)
- In a joint effort between Naples and Perugia, the first configuration HELM charts have been pushed on the official ICSC-Spoke2 repository and the deployment is undergoing.
- An early-access version should be available from January 2024, to provide users with an area to start playing and, on the same time, to help developers with feedbacks and bug reports.

In the meanwhile:

- Waiting for RAC requests to be processed, to realize how much firepower we will have from the management in the next year.

Computing resources - RAC request

- During the last month, RAC (Resource Allocation Committee) asked us to fill an official request with an estimate of the resources that our use case will need in the upcoming months.

- Based on our predictions from the flagship document, we can summarise the request with:

> - CPU systems, without the requirement of fast inter-node infiniband communication (HTC-like). This type of resources will be used during the peak phases, where different users (with applications / data analysis) will exploit the maximum of the computational power at their disposal, concurrently. Furthermore, these nodes must have storage access, with the latter being accessible through WebDAV and XRootD protocols.
> - Cloud VMs, for the services instantiation (Jupyter Hub/Lab, Dask, HTCondor). Such resources will require storage as well, implemented in the form of volumes for the cloud services, although considered as a "temporary type" of storage.
>
> The numerical details of the requested resources (CPU and storage), as well as their split and the time development, will be listed on the next pages. In absolute values, this Usecase will expect around 14000000 CPU total core hours and an overall storage of 750TB.

ICSC Spoke 2 Annual Meeting

During the **ICSC Spoke 2 Annual Meeting** (Cineca, Casalecchio di Reno (BO), 18-20 December 2023, agenda) there will be lightning talks (10 minutes) to present specific works, in particular for newly hired personnel.

- As a flagship WP2 usecase, our idea was to submit an abstract providing an overview and status of the work, on the scientific side (mainly analysis applications). **Deadline: December, 10th**

- In case our abstract will be accepted for a lightning talk, we need to wrap-up all the applications and list all the new people hired under Spoke 2, with an effort on this usecase.
  The next two slides contain the most up-to-date information in our possession (will be presented on the event)
  Please provide us a feedback, in case of incorrect/incomplete data.

- A draft version of those documents has already circulated among the flagship members last week, thanks for the changes and help!

Flagship Usecase Map

feedback welcome please

The **blue** ones, have <u>at least one confirmed</u> analysis application:

- Univ. of Bologna (BO): detector performance and b-physics analysis;
- Polit. of Bari (BA): b-physics and Higgs analysis;
- Univ. of Salento (LE): SUSY analysis;
- Univ. of Naples (NA): top analysis and future colliders;
- INFN Perugia (PG): Standard Model physics.
- INFN Firenze (FI): CI pipeline submitting on AF

- The **red** ones, have expressed interest but currently the effort is not clear: **please help us fill those** (**UniBari, UniCalabria, UniFerrara**)

- The **green** line represents the infrastructural support WP5 backbone:
  - INFN Perugia (PG);
  - Univ. of Naples / INFN (NA).

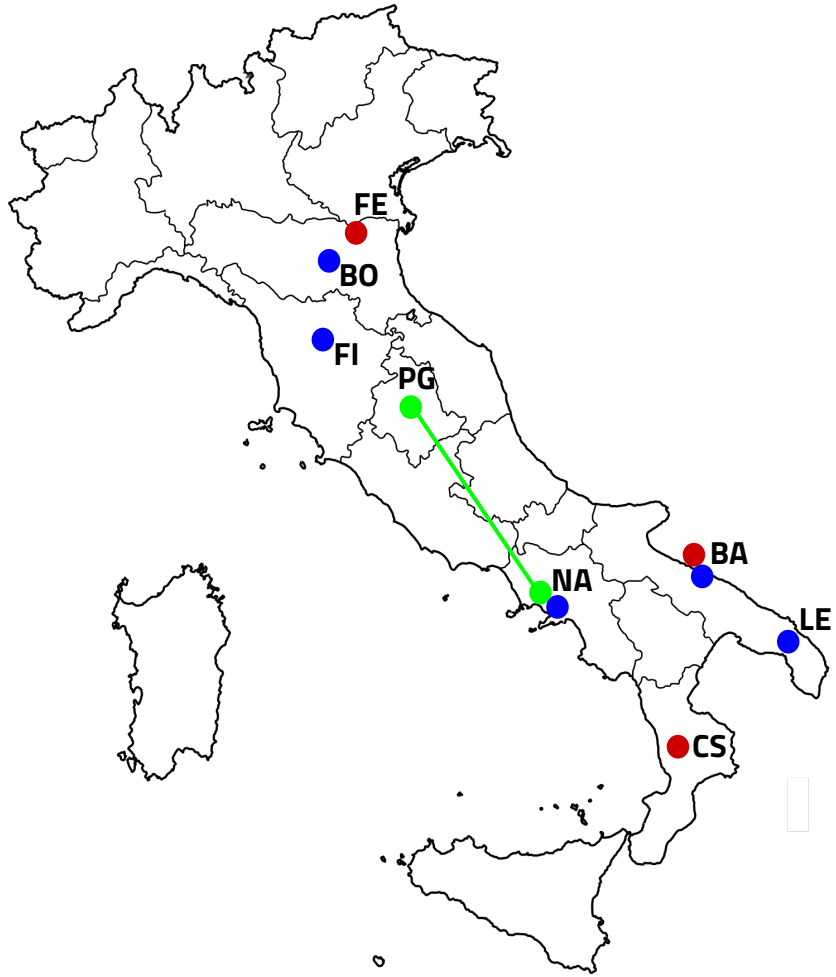

New hired personnel under UC2.2.2

feedback welcome please

- Univ. of Bologna (BO):
  - Tommaso Diotalevi, RTDa (100% Spoke 2)
- Polit. of Bari (BA):
  - Federica Maria Simone, RTDa (100% Spoke2)
  - Muhammad Anwar, Dottorando (100% Spoke 2)
- Naples (NA):
  - Adelina D'Onofrio, tecnologa INFN (100% Spoke2) - WP2/WP5
- Perugia (PG):
  - Tommaso Tedeschi, borsista INFN (100% Spoke2) - WP2/WP5
- Univ. of Firenze (FI):
  - Matteo Bartolini, assegnista

...others??

- Univ. Calabria (CS) (?) - WP5 (?)
- Univ. of Bari (BA)

Conference opportunities for our flagship

- ACAT abstract submission deadline is the 5th of January:
  - Our idea was to submit an abstract related to the physics benefits regarding the usage of the AF, leveraging on the different applications gathered in our use case, maybe with a cumulative benchmarking study.
  - Just an idea… Any other ideas are more than welcome!
  - This can help the progress of our KPIs, as well as helping new hired with publication quotas etc…
  - Details and problems in terms of collaboration, signers, and PNRR technicalities must be discussed.

Thank you

cn1-spoke2-wp2-analysisfacility@lists.infn.it (click here to subscribe or get in touch in case of issues)

RAC resource request

7.1 Resources granted by INFN-CLOUD (PaaS)

vCPU (number of vCores and requested allocation time)

Number	Time	vCORE	Notes
*	*	*	time in hours
4	17520	32	

Number of requested GPU and allocation time

Number	Time	GPU	Notes
*	*	*	NO GPU REQUIRED. The numbers were entered due to a problem with the form
1	1	1	

vCPU (number of vCores and requested allocation time)

RAM per VCORE	Notes
*	Memory in GB
4	

☐ Software used or required, including preferred Cloud services

I.e.: Kubernetes-as-a-service, Jupyter Notebook as a Service, Private Container Image Registry, Spark and Grafana as a service, Dropbox-like sync-and-share service

Software used or required	Notes
	Total core hours: 4*17520*32 = 2242560. This kind of resource does not follow time evolution, to be considered constant for the entire project duration.

7.2 Resources granted by INFN-GRID (Batch processing)

CPU (number of Cores and requested allocation time)

Number	Time	vCORE	Notes
*	*	*	time in hours
21	**17520**	**32**	

RAM Requirements

RAM per VCORE	Notes
*	Memory in GB
4	

Software used or required	Notes
	Total core hours: 21*17520*32 = 11773440. This kind of resource follows a time evolution, where: **- first 6 months: 876000 core hours, to be considered as 1000 cores * 20% of time (1000 cores distributed in *5* data analyses);** **- following 4 months: 730000 core hours, to be considered as 1000 cores * 25% of time (1000 cores distributed in *5* data analyses)** **- following 4 months: 730000 core hours, to be considered as 1000 cores * 25% of time (1000 cores distributed in *5* data analyses)** **- last 10 months: 9437440 core hours, to be considered as 5000 cores * ~26% of time (5000 cores distributed in *5* data analyses)** **These nodes, on N CPU systems without the requirement of fast inter-node infiniband communication (HTC-like), must have access to permanent storage.**

KPI - Key Performance Indicator

These parameters, which are measurable objects, will reflect the final success of the entire use case.

KPI ID	Description	Acceptance threshold	Current Status
KPI2.2.2.1	Implementation of *N* data analyses in the AF	$N \geq 2$	> 50%
KPI2.2.2.2	Reference documentation of the AF	≥ 1 dedicated web site	0%
KPI2.2.2.3	Hands-on workshops for AF users	≥ 1 workshops	0%
KPI2.2.2.4	Scaling up the testbed AF infrastructure, serving *k* tenants, for a total of *N* data analyses	≥ (200·*N*) cores	0%
KPI2.2.2.5	Talks at conferences/workshops about AF activities	≥ 1 talk	100%*

* Talk at CCR in Loano (22 May).
For other conference opportunities, here the WP2 Tracking document. Please contact us for abstract submissions, to coordinate a "flagship collaboration"!

Participating Institutions

- **PIs:** Tommaso Diotalevi (Università di Bologna), Francesco G. Gravili (Università del Salento)
- **INFN** (including all sections)
- **Università del Salento**:
  - Search for new phenomena in events with two opposite-charge leptons, jets and missing transverse momentum, using LHC Run2 data. Porting of the analysis in the AF framework.
- **Università di Bologna:**
  - Search of a heavy neutral lepton (bump hunt) in the D_S meson decays with Run2 data. Porting of the analysis in the AF framework;
  - CMS Muon detector performance analysis: target quasi-interactive performance studies of phase space corners using large datasets; driven by the need to accurately assess analysis systematics (e.g. high-energy muons).
- **Politecnico di Bari:**
  - Testbed implementation of the AF in the ReCaS computing center
- **Università degli Studi di Bari:**
  - Charm and multi-Charm baryon measurements.
- **Università degli Studi di Napoli:**
  - Top quark + MET search. Full analysis implementation with Run3 data.
  - Analysis Facility deployment on local cloud infrastructure.
- **Università degli Studi di Firenze:**
  - Developing common analysis tools based on analysis facilities for the CMS experiment.
- **Università degli Studi Ferrara:**
  - Adopting tools for quasi interactive analysis for the IDEA experiment at the FCC-ee.
- **Università della Calabria.**

Please, check the correctness of the contact names in the document!

Analysis Facility testbeds

State of the art:

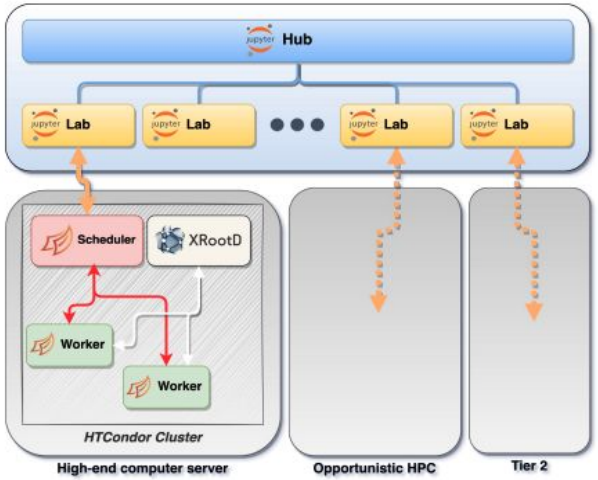

INFN-CMS *Analysis Facility*:

- Access to a single HUB and authentication via token (INDIGO-IAM)
- Based on standard industry technologies
- Customisable python kernel
- Workarea fully containarisable
- Overlay based on HTCondor (also available standalone)
- DASK library (python) for distributed computing
  - Scale the execution from 1 to N cores
- Possible implementation on heterogeneous resources (HTC/HPC/Cloud)
- Data access configurable with WLCG (xrootd, WebDAV, ...)

Similar approach for the INFN@Naples Analysis Facility:

- Jupyter Hub/Lab for interactive environment, supporting multiple accesses
- DASK library for scalability
- Kubernetes for container management and connection
  - Ongoing tests on this container technology to achieve a higher scalability