

# SWARM Sim – Scalable Workflows for AI-powered Radiation-Matter Simulation

L. Anderlini<sup>1</sup>, A. Bombini<sup>1</sup>, C. Buti<sup>1,2</sup>, R. Petrini<sup>1</sup>, A. Rosa<sup>1</sup>, C. Mancini-Terracciano<sup>3</sup>, L. Arsini<sup>3</sup>



<sup>1</sup> Istituto Nazionale di Fisica Nucleare – Sezione di Firenze

<sup>2</sup> Dip. di Fisica e Astronomia, Università di Firenze

<sup>3</sup> Dip. di Fisica, Università La Sapienza di Roma



## Abstract

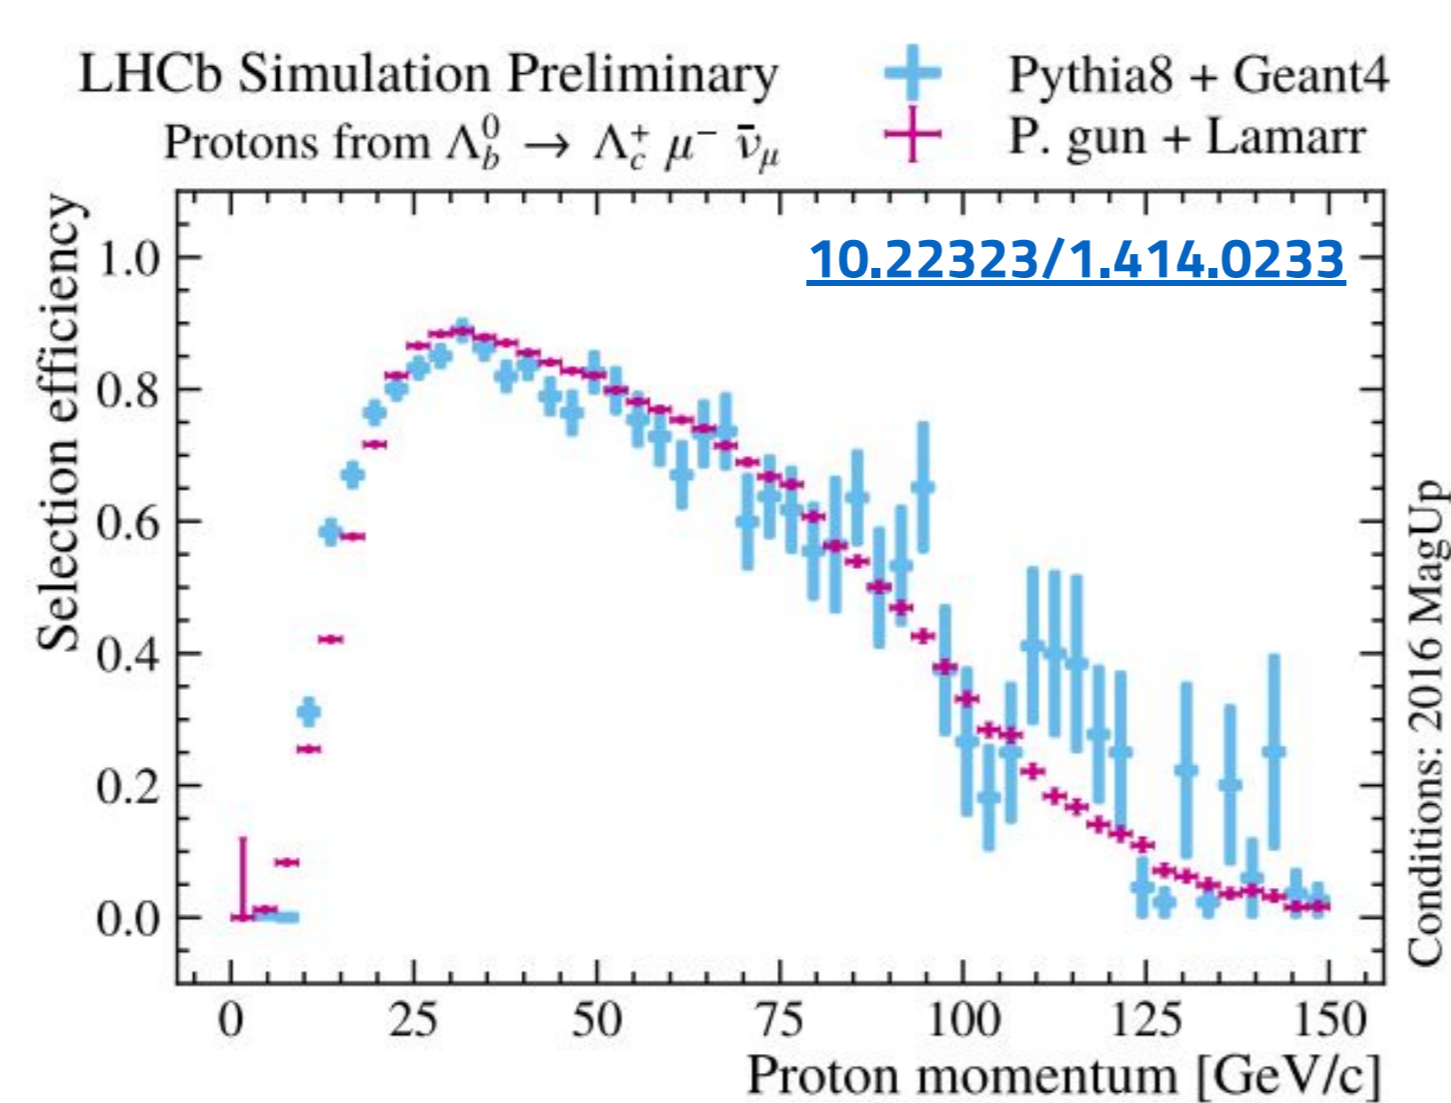
**SWARM Sim** is an **AI-powered** simulation framework for radiation-matter interactions across domains like high-energy physics and radiotherapy. It leverages *GANs*, *Graph Neural Networks*, and *Physics-Informed models* trained and/or evaluated using **nVidia GPUs**.

**Reproducibility** and **scalability** are achieved through containerization and continuous integration, in a cloud-native environment backed by the **Kubernetes** platform operated by the **AI\_INF**N initiative.

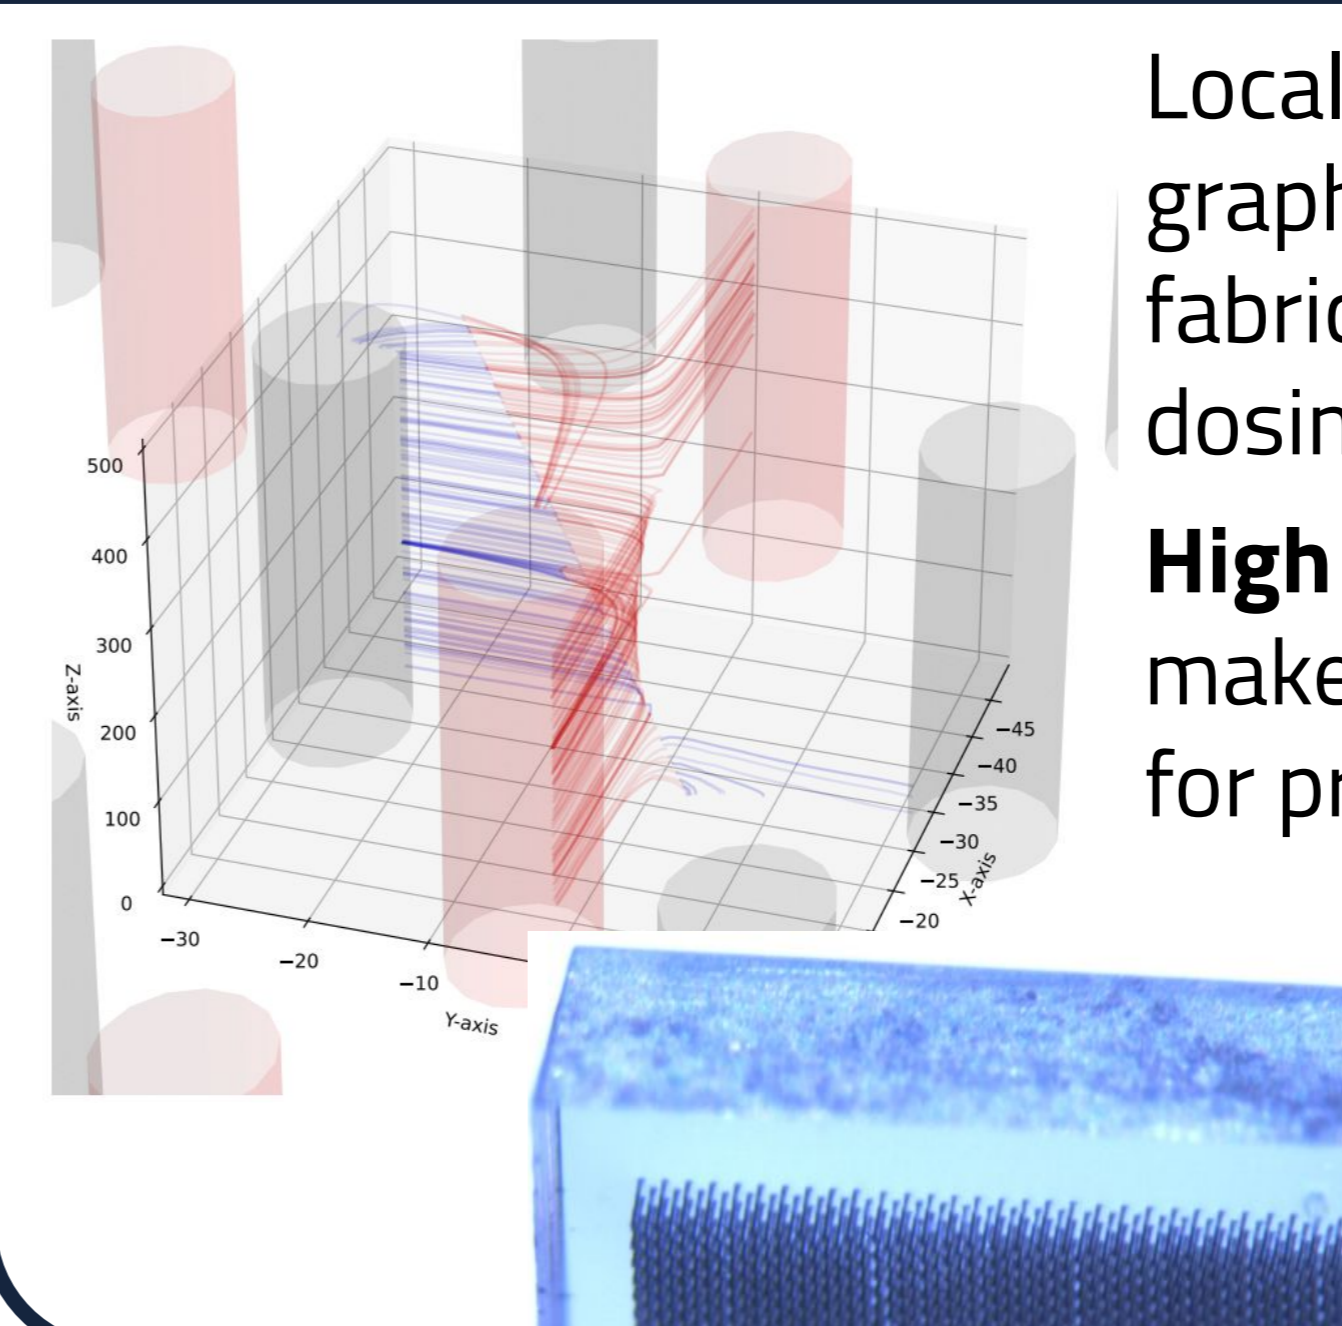
The most computing-intensive workloads are **offloaded** to traditional **High Performance Computing (HPC)** and **High Throughput Computing (HTC)** resources. The bridge is provided by **InterLink** and **CVMFS Unpacked**, technologies developed by **INFN** in the context of the **TeRABIT project**.

## Physics Case 1: LHCb Flash Simulation

**GANs** are trained on official **LHCb simulations** to model detector response, enabling fast estimation of **detection and reconstruction biases** for new heavy-hadron decay channels. Models are **trained on GPU** and **transpiled in in-process C** functions to be deployed in the LHCb Simulation software stack.



## Physics Case 2: 3D Diamond detectors



Local phase transitions from diamond to graphite induced by laser pulses enable fabrication of **full-carbon sensors** for dosimetry and HEP applications.

**High resistivity of graphitized electrodes** makes traditional simulations inadequate for predicting signal shapes.

We explore **Physics-Informed Neural Networks** and innovative **spectral methods** to overcome these limitations.

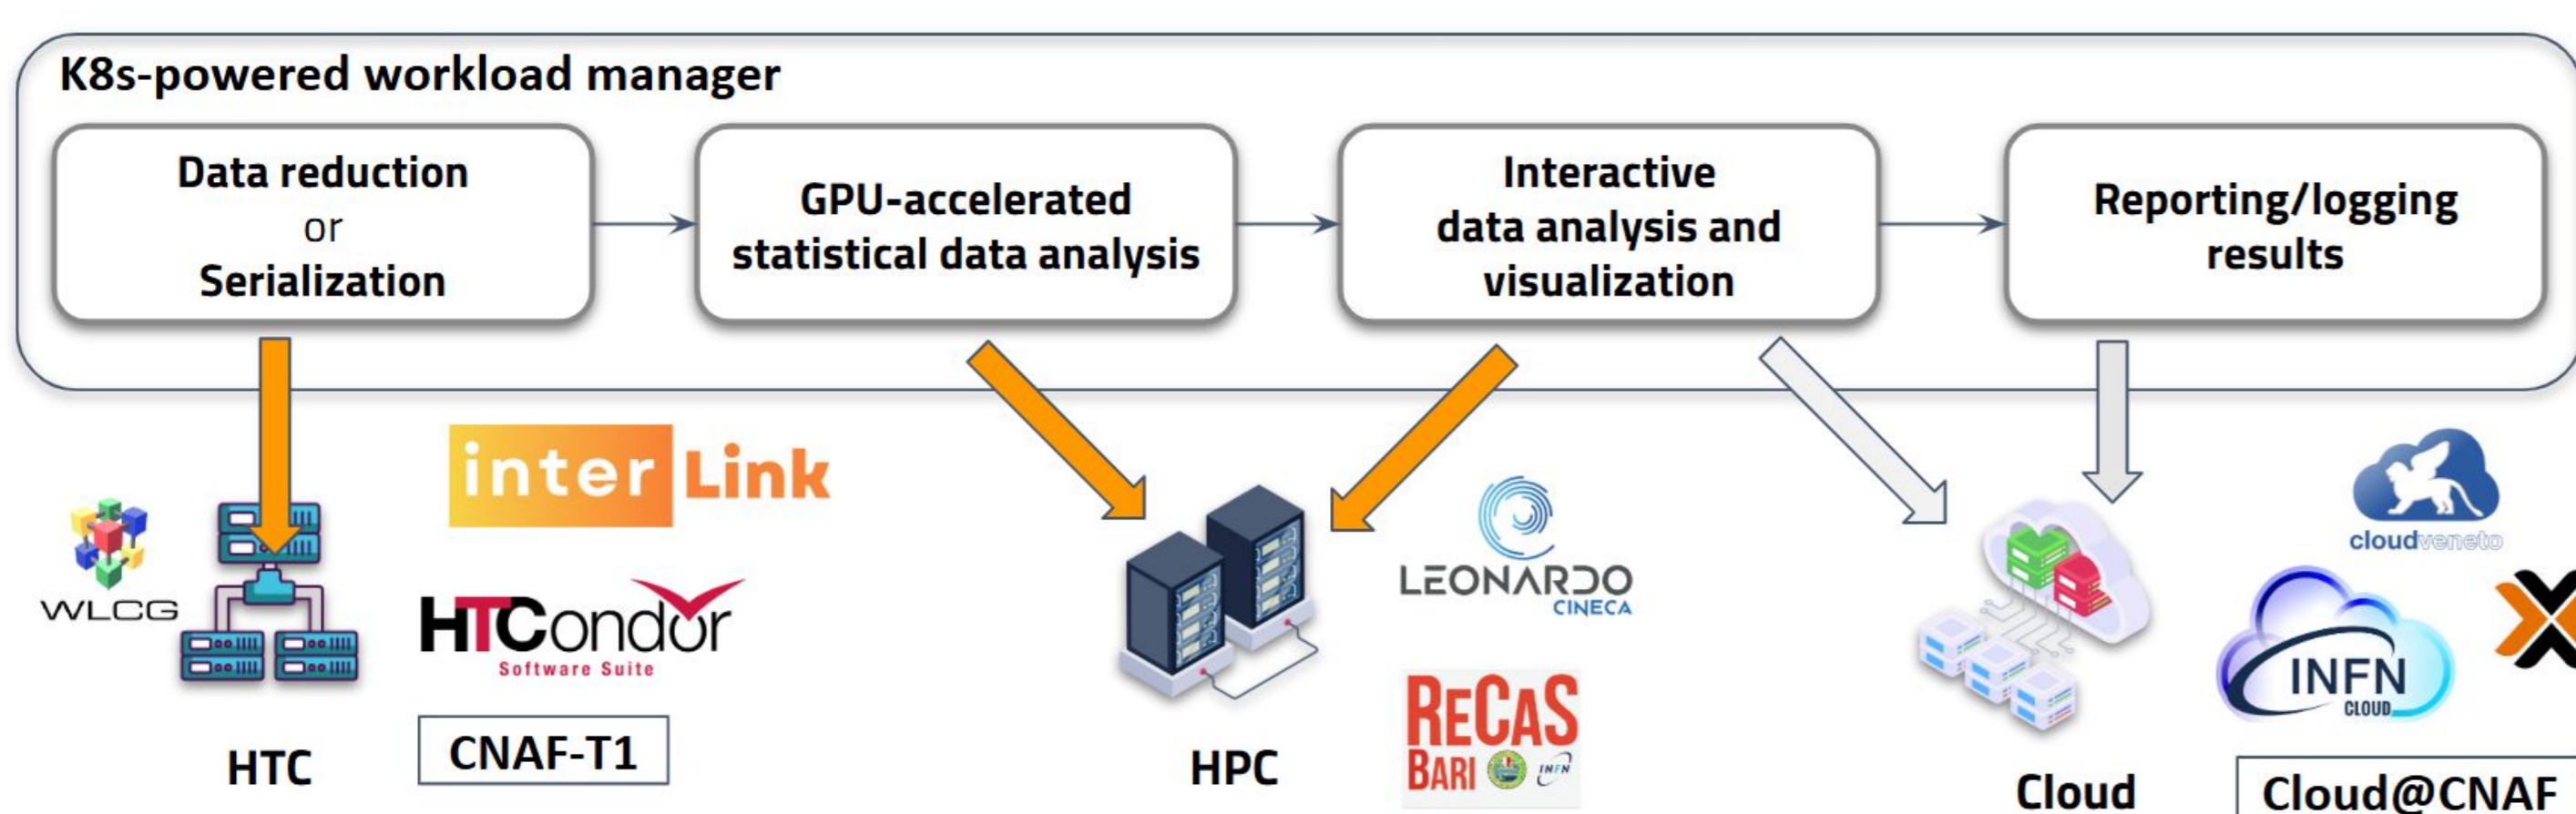
## Accessing the TeERABIT HPC Bubbles: leveraging the Padova setup with InterLink



ai.cloud.infn.it

SWARM Sim builds on the **AI\_INF**N Platform, hosted by **INFN Cloud** and powered by **Kubernetes** and **JupyterHub**, bringing together **physicists, ML practitioners, and IT researchers** to advance data-intensive applications.

AI\_INFN federates **INFN CNAF Tier-1**, **ReCaS Bari Tier-2**, **CINECA Leonardo**, and the **TeRABIT HPC Bubble in Padova** via **interLink**



**Software distribution via CVMFS Unpacked** `/cvmfs/docker-graphdriver`  
Software stacks combining **Python** and **C/C++** user code are packaged directly on the platform and published to the **INFN Cloud Harbor** where they get distributed to sites through **CVMFS Unpacked**.

**Snakemake Workflows** `/snakemake/snakemake`  
Complex workflows described as **Directed Acyclic Graphs (DAGs)** with Snakemake can be **distributed** through the **federated** computing centers relying on a central **S3 endpoint** to exchange temporary data.

**Integration with GitHub Workflows** `/landerlini/apptainer-gha-runner`  
Specialized **GitHub Actions** were developed to execute CI/CD workloads using Apptainer (rather than OCI runtimes). **InterLink** ships steps requiring **HPC to TeRABIT resources**, Azure is used for the rest.

## Conclusion and Future Work

**SWARM Sim** demonstrated coordinated access to Italian heterogeneous computing ecosystem, a major challenge for modern scientific workflows.

Integration of **cloud-native** and **traditional batch resources** was achieved via **InterLink**, enabling distributed data-intensive workflows managed with Snakemake or GitHub Actions, with **CVMFS Unpacked** emerging as a key component for reliable software distribution.

Future work will focus on **productionizing the infrastructure** to deliver customizable, feature-rich *Virtual Research Environments* for small-experiment communities, easing adoption of a future computing landscape spanning **CPUs** (in multiple flavors), **GPUs**, **FPGAs**, and emerging **Quantum Processors**.

**Acknowledgements.** We gratefully acknowledge **M. Sgaravatto**, **M. Verlato**, **D. Spiga**, **D. Ciangottini** and **G. Bianchini** for developing and supporting the infrastructure.

