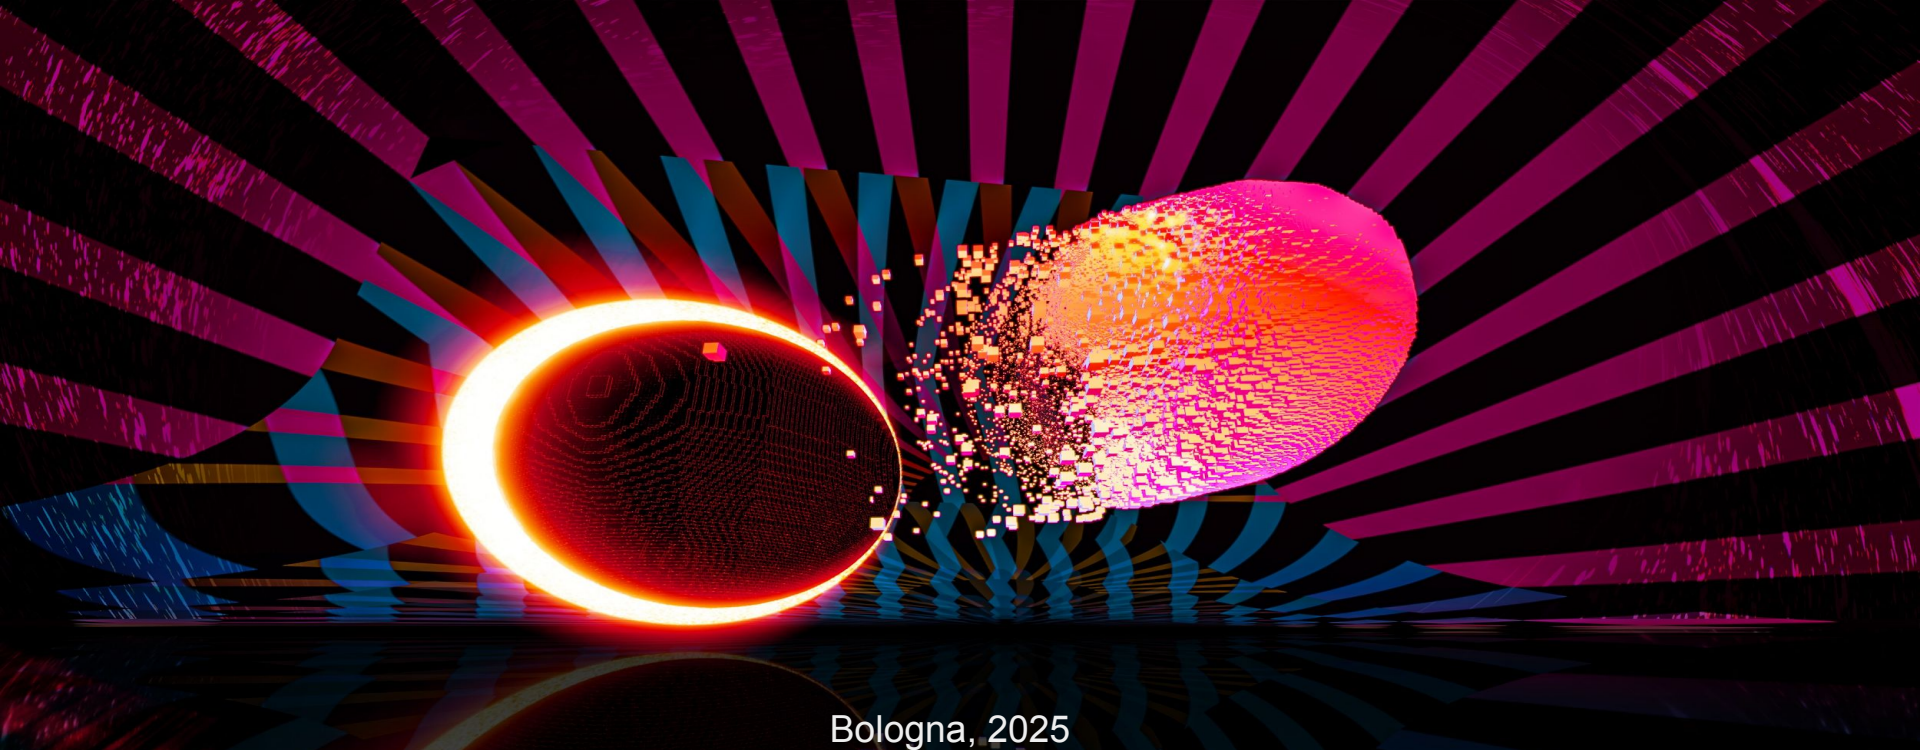

Bologna, 2025

EINSTEIN TELESCOPE

Benchmarking RIFT with HEPScore

Georgy Skorobogatov, EIB Div 3

Goals

- Obtain a GPU-utilizing workload representative of those running on current detectors
- Modify the inputs of the selected workload to better resemble those expected for ET
- Benchmark the workload on various hardware configurations; focus on GPU

Benchmarking

We are in contact with HEPiX WG.

Their tasks:

- Development of a benchmarking framework
  - HEP Workloads – Individual reference HEP applications; multiple domains
  - HEPScore – benchmark suite; runs HEP Workloads, collects results, and computes HEPscore; support for both CPU and GPU jobs
  - HEP Benchmark suite – orchestrator of multiple benchmarks
- Organize exchanges on CPU/GPU benchmarking topics

~20 sites contributing
~110 distinct configurations (CPU models, etc.)

“*GPU workloads exist, but we are still far from having an HEPScore for CPU+GPU*”

Workload

RIFT – an algorithm to perform Rapid parameter inference on gravitational wave sources via Iterative FiTting

- Based on CuPy
- The largest GPU consumer in the IGWN pool
- "Integrate likelihood extrinsic" (ILE) jobs are the main GPU executable used by the RIFT pipeline

Workload Integration into HEPScore

1. Get approval from the HEPiX WG
2. Use any of the existing workloads as a template (e.g. Python-based, relying on GPU, not using CVMFS)
3. Create a *Docker.append* for setting up the environment (not possible to specify an existing image)
4. Update the script for running a single job
5. Update the script for parsing the output results and score calculation
6. Adapt the workload to support running multiple concurrent copies
7. Submit a merge request to the HEP Workloads GitLab

Note: direct contributions require a CERN account

Current Status and Future Work

- Working with the HEPiX development team on integration of RIFT to the HEPScore
- Developed a draft of the workload:
  - Created a Dockerfile.append that sets up dependencies, clones the RIFT, and generates synthetic data – to be validated
  - Updated the script for running a single job – works only with one copy
  - Score calculation is implemented but needs to be reviewed
- Next steps:
  - To verify what other metrics, logs, or results need to be collected
  - Whole workload needs to be rewritten using a different simpler template
  - Input data, parameters, and constants have to be updated so that they make sense for ET

References

1. Workload – https://github.com/GeorgySk/rift/
2. HEP-Workloads – https://gitlab.cern.ch/hep-benchmarks/hep-workloads
3. HEPScore – https://gitlab.cern.ch/hep-benchmarks/hep-score
4. RIFT – https://git.ligo.org/rapidpe-rift/rift
5. HEPiX Benchmarking WG – https://w3.hepix.org/benchmarking.html
6. HEPiX report – https://indico.cern.ch/.../HEPiX...giordano.pdf