DEPLOYMENT OF ATLAS CALORIMETER FAST SIMULATION TRAINING ON HPCs THROUGH CONTAINER TECHNOLOGY

22nd International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT)
Stony Brook University, Stony Brook, NY, United States, 11-15 March 2024

THE CONTEXT

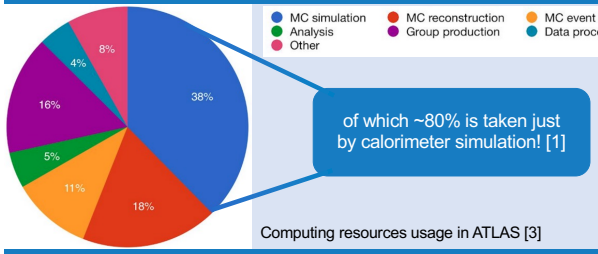

Computing resources usage in ATLAS [3]

- **Simulation of detector response = major computing challenge** at LHC experiments (blue on chart).
  - **Calorimeter simulation = most demanding part**, to increase with HL-LHC!
- Solutions already developed → **fast simulation tools** with Machine Learning, faster than Geant4 at simulating calorimeter response but keeping high accuracy.
  - ***AtlFast3***: the ATLAS fast simulation tool [1] (in production for Run 3), combining classic parametric approach + **GANs** (Generative Adversarial Networks) [2]. GANs require a lot of resources to train!
  - **Further idea: deploy training on other resources** than the CERN batch system (LXBATCH) or the Worldwide LHC Computing Grid, for more resource saving + extra performance **boost** → **BoloGANtainer!**

TECHNICAL DETAILS on BOLOGANTAINER

- **BoloGANtainer containerizes AtlFast3 GAN training** for deployment on other machines than the usual ones (recipe on right).
  - **Extends the official ATLAS CentOS 7 image**, which replicates the (current until recently) OS and software of LXBATCH.
  - CVMFS-independent, requires CUDA-11 and CuDNN for GPU usage.
  - **Additional software directly installed and not linked via CVMFS** → the program runs on nodes w/out CVMFS nor Internet connection.

PERFORMANCE

- Deployed on the systems in table.
  - Training **on Leonardo runs in ~half the time** needed on LXBATCH. **Remarkable speedup given by A100**. **Using supercomputers brings great advantage!**
- **Plots below** [5]: χ^2/NDF per iteration on Leonardo (best iteration in red).
  - **Pion training** (left) is stable and the **high energy photons** one (center) also rather is, it is rather unstable for **low energy photons** (right). Analogous results found for LXBATCH.
- **To do**: deployment onto further resources (also cloud), architectures (ARM) and for more particle types, code optimization (both general and to take advantage of multi-GPU nodes).

The Leonardo supercomputer at CINECA in Bologna, Italy [4]

REFERENCES

[1] The ATLAS Coll., *AtlFast3: The Next Generation of Fast Simulation in ATLAS*, *Comput Softw Big Sci* **6**, 7 (2022)
[2] The ATLAS Coll., *Fast simulation of the ATLAS calorimeter system with Generative Adversarial Networks*, ATL-SOFT-PUB-2020-006
[3] The ATLAS Coll., *ATLAS HL-LHC Computing Conceptual Design Report*, CERN-LHCC-2020-015 (2020)
[4] leonardo-supercomputer.cineca.eu
[5] Public plots: atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/SIM-2024-004
[6] TOP500 Ranking - November 2023, top500.org/lists/top500/2023/11/

Run on Leonardo

```
BootStrap: docker
From: atlas/centos7-atlasos

%setup
 mkdir ${APPTAINER_ROOTFS}/data

%files
 < copies GAN training code into the container >

%post
  rpm -Uvh
https://packages.microsoft.com/config/centos/7/packages-
microsoft-prod.rpm
  yum install -y dotnet-sdk-7.0 man man-pages libXpm gcc-
c++ git root python3-root
  pip3 install --upgrade pip
  pip3 install --upgrade setuptools
  pip3 install tensorflow
  pip3 install pandas

%environment

%runscript
  if [ "$1" = "init" ] ; then
    < prepares directory for output and logfiles >
  elif  [ "$1" = "train" ] ; then
    if [ "$2" != "pions" ] && [ "$2" != "photons" ] ; then
      < warns about invalid option and exits >
    else
      < launches training script >
    fi
  elif [ "$1" = "bestiter" ] ; then
    if [ "$2" != "pions" ] && [ "$2" != "photons" ] ; then
      < warns about invalid option and exits >
    else
      < launches iteration evaluation script >
    fi
  else
    < warns about invalid option and exits >
  fi
```

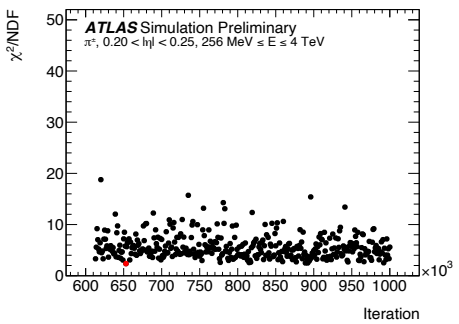

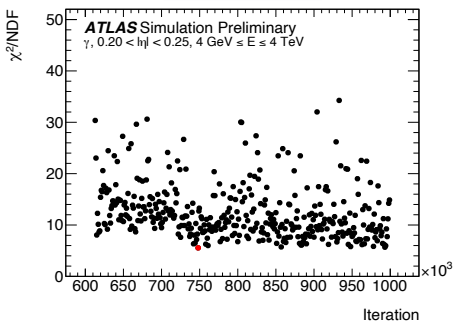

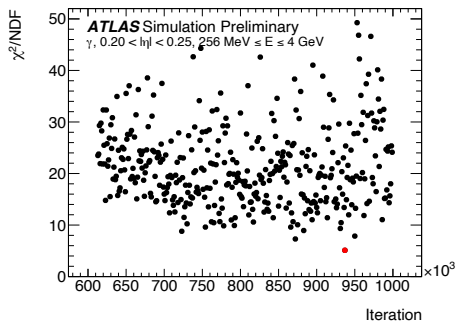

Resource	Type and Owner	Hardware and Software	Pion Results One GAN trained for all energies	Photon Results Two GANs trained, one for energies ≤ 4 GeV, one for above
LXBATCH (reference cluster)	CERN batch system	CentOS 7 (for the used nodes), CVMFS, HTCondor, V100 GPUs	Run time: 12 h χ^2/NDF ~ 2	Run time: 1 d 6-7 h χ^2/NDF ~ 5
Leonardo	The 6th most powerful cluster on the TOP500 ranking [6], at CINECA	RHEL 8.7, no CVMFS, SLURM, A100 GPUs, isolated nodes	Run time: 6-7 h χ^2/NDF ~ 2	Run time: 10-11 h χ^2/NDF ~ 5

Federico Andrea Guillaume Corchia (University and INFN, Bologna, IT)
Michele Faucci Giannelli (University and INFN, Rome Tor Vergata, IT)
Joshua Falco Beirer (CERN)
Lorenzo Rinaldi (University and INFN, Bologna, IT)
Rui Zhang (University of Wisconsin-Madison, US)
on behalf of ATLAS Computing Activity