

Software developments

Frédéric Derue, LPNHE Paris
Arnaud Duperrin, CPPM Marseille

Calcul ATLAS France (CAF) meeting
16th September 2021

See slides « ATLAS document from James at S&C plenary 10/09/2021 [link] for LHCC review in November + also see slides for Simulation & Generation

Core Software & HCAF (i)

- Complete athenaMT migration
- Lossy compression
- New (to us) storage technologies:
 - HDF5
 - ROOT 7 columnar data (RNTuples)
- Re-evaluate ATLAS data formats
 - Organization, compression, contents

Core Software		
1	Pileup-digitization in AthenaMT production ready	Dec 2022
1.1	Ensure reproducibility of MT production of presampled MB RDO files	
2	Complete investigation of lossy compression techniques	Dec 2021
2.1	Lossy compression of the ID track covariance matrix in the primary AODs	Dec 2021
2.2	Lossy compression of DAOD	Dec 2021
3	Lossy compression of primary AODs	Dec 2023
4	Implement I/O roadmap metadata recommendations	
4.1	Multi-threaded in-file metadata handling	Jun 2022
4.2	Redesign of the metadata handling infrastructure (better support for fine-grained workflows)	Dec 2022
5	Evaluation of data formats well-suited for massively parallel I/O (HPCs)	
5.1	Storing intermediate EventService Simulation data in HDF5	Mar 2022
6	Re-evaluation of simulation data formats (e.g. EVNT, HITS, RDO)	
6.1	Technical review of file storage format, compression, etc	
6.2	Content review of data format	
7	Migration to ROOT 7	Dec 2026
7.1	ROOT and LCG release contributions, testing, feedback	
7.2	StorageSvc capable to write a subset of DAOD_PHYS(lite) data to RNTuple	Dec 2023
7.3	Migration of Athena from Root 6 to 7	Mar 2025

Core Software & HCAF (ii)

- **Technology decision**
 - CUDA vs kokkos vs std::par...
 - Single source required?
- GPU resource management
 - Memory, I/O, ...
- Develop realistic heterogeneous simu/reco applications to test
 - Multi-event kernels
 - Accelerator-friendly EDM
 - *Requires investment from simu/reco communities*
 - **Can't impact baseline work!**
- Scheduling heterogeneous applications across nodes
 - Raythena, hpx, ...

Heterogeneous Computing and Accelerators		
9	GPU management techniques and infrastructure in Athena	Dec 2021
10	HL-LHC Technology decision: CUDA or one of its less-proprietary competitors.	Dec 2023
10.1	Full parallelization pattern recommendation to collaboration	Mar 2024
10.2	Design patterns/tutorial on GPU migration	Mar 2024
11	GPU Memory management	
11.1	First (Vecmem) prototype [Attila, Stephen, Q1 2022]	Mar 2022
12	GPU Kernel scheduling	
12.1	Basic support for Kernel scheduling in athena	Jun 2021
12.2	Integration with Gaudi Scheduler	
13	Intra-node scheduling, targeting HPCs and grid	
13.1	Raythena/HPX-based scheduler prototype	Sep 2022
14	Infrastructure for processing data across multiple events on an accelerator	
14.1	Proof-of-concept prototype	Dec 2021
15	Develop Multi-algorithm heterogeneous applications	
15.1	ACTS-based multi-algorithm workflow	Dec 2021
15.2	Calorimeter clustering	
15.3	FastCaloSim GPU merged into master	Jun 2021
15.4	ML inference	
16	Make our EDM classes accelerator-friendly.	Dec 2022
16.1	Support for reduced/mixed precision in ATLAS EDM	
16.2	

Paolo and James - 10 Sept 2021

Simulation, Digitization, and Detector Description

- (Very) detailed planning on [jira](#)
- FastChain key to satisfy CPU constraints for Run 4
- R&D priorities (beyond baseline)
 - DNN surrogate models for fast sim (e.g. FastCaloGAN)
 - Port select simulation algorithms to GPUs
 - Tune (fast) simulation to analysis data
- **Review of simulation strategy**
 - Decision point on "G4 on GPU" "FastCaloSim on GPU", ISF retirement

	Simulation/Digitization/Detector Description	
1	Updates required for MC+MC Overlay for Run4	12/2022
2	Track Overlay in FastChain	12/2022
3	Review or Run 4 Simulation Strategy	6/2023
4	Pile-up Digitization in AthenaMT	12/2023
5	Make data overlay useable for p-p collision simulation	12/2024
6	Fast simulation development for Run4	9/2026
7	Geometry updates for the Run4 MC Campaign	9/2026
8	Run4 optimisation of the performance of ATLAS full Geant4 simulation	9/2026
9	Testing Geant4 Versions and configurations for Run4	10/2026
10	Validation of full and fast simulation for Run4	12/2026

Analysis Model Group

- DAOD_PHYS and PHYSLITE needed to satisfy disk storage constraints
 - Early Run 3 demonstration will be key to ensure adoption
 - Prove PHYSLITE usability for wide range of analyses (esp. systematics)
- PHYSLITE will enable prototyping of **columnar analysis**
 - data formats (e.g. RNTuple)
 - analysis infrastructure (e.g. coffea)
 - Prototyping & review of columnar data operations for end-to-end analysis

	Analysis	
A	Baseline DAOD_PHYSLITE with run 3 s/w (event loop based) & ROOT6	Q3 2023
A.1	First bulk production of prototype DAOD_PHYSLITE	Q4 2021
A.2	Mechanism for evaluation of systematic uncertainties with PHYSLITE	Q2 2022
A.3	Demonstrator for full analysis on PHYSLITE for target analyses	Q4 2022
A.4	Recommendations for application of lossy compression	Q4 2022
A.5	Finalised list of PHYSLITE contents	Q2 2023
A.6	Development/roll-out of docs/training for run 3 analysis w/ PHYSLITE	Q3 2023
B	PHYSLITE working with ROOT 7	Q4 2023
B.1	Implementation of RNTuple and revised xAOD in PHYSLITE	Q4 2023
C	Prototyping & review of columnar data operations for end-to-end analysis	Q2 2024
C.1	Tests of basic reading performance using TTree version of PHYSLITE	Q4 2022
C.2	Prototyping of tools for columnar CP operations and other systematics	Q4 2023
C.3	Adoption of ROOT7 data structures	Q1 2024
C.4	Performance and ease-of-use assessment leading to decision on adoption	Q2 2024
D	Development of columnar analysis infrastructure	Q1 2026
D.1	Prototyping of framework for orchestrating columnar CP operations	Q3 2024
D.2	...	Q1 2025
E	Accomodate special analyses in the run 4 analysis model	Q2 2026
E.1	Assessment of which run 4 analyses are incompatible with PHYSLITE	Q4 2022
E.2	...	Q2 2023

Reconstruction

Still gathering information, will have more complete, detailed milestones soon

- Waiting to hear back from
 - Muon, calo, tracking

Select decision points

- Scope of ACTS in athena
 - Multiple decisions (ID, mus, combined)

- GPU-friendly tracking


- Depends on acts-GPU, GNN tracking evaluations
 - Connected with HLT Software Tracking

- DL algorithms for Jets/Etmiss/Tau (ParticleFlow)

- Main focus of e-gamma is improving forward e/γ and photon conversion reconstruction
 - Will explore ML possibilities (building of supercluster from topological clusters, initial selection of topoclusters)
 - Very moderate CPU consumers

Reconstruction		
0	Merge 21.9 and master	3/2022
1	Feature freeze (3 years !)	3/2025
2	Performance freeze:	3/2026
3	Ready for data taking	3/2027

ATLAS Preliminary
2020 Computing Model -CPU: 2030: Aggressive R&D



Paolo and James - 10 Sept 2021

- **Last release 22.0.41.1**
 - final version 22.0.41.2 build last week (?)

- **Performance WG**
 - have finished the tuning and validation of the release at the reconstruction level (AOD) for the reprocessing of Run 2.
 - continuation of work at the derivation level

- **Physics analyses**
 - validation of software/framework for PAV1 is done for rel 22
 - PAV2 samples will soon be launched

- **Containers**
 - Jets+MET+b-tagging containers are not saved anymore at AOD level, but done at derivation level (DAOD_PHYSVAL)
 - derivations produced after 6th July cannot read older AODs

- **Recent talks**
 - A. Duperrin, This month in Ftag software, status and plans, Aug 2021 [link]