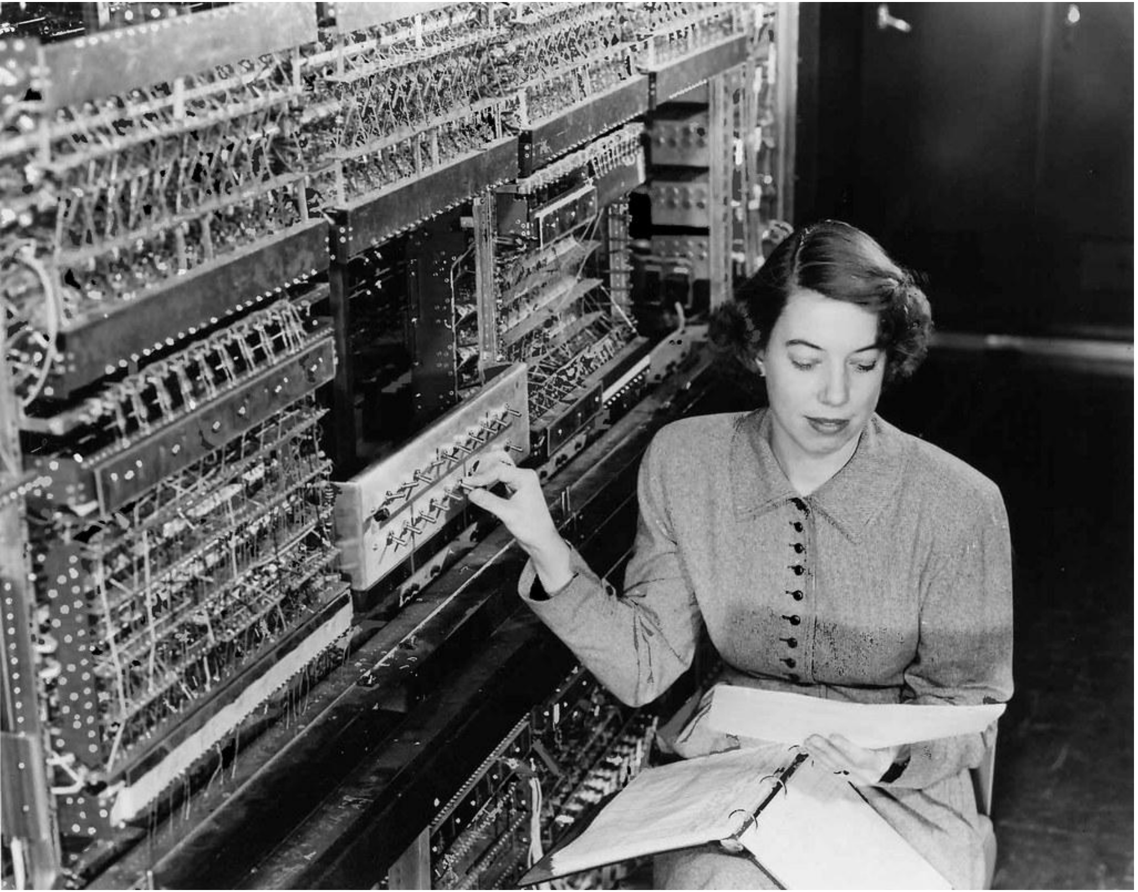

ATHENA

EIC Software Working Group Meeting 2021-07-28

The Software and Computing WG Conveners:
Andrea Bressan (University of Trieste and INFN) ,
Dmitry Romanov (Jefferson lab) ,
Sylvester Joosten (Argonne National Laboratory) ,
Whitney Armstrong (Argonne National Laboratory) ,
Wouter Deconinck (The University of Manitoba)

ATHENA Software Ecosystem

Simulation in npsim (DDG4) provides hits for Gaudi

Gaudi (originally LHCb): connect geometry and data to algorithms (clustering, tracking)

Gaudi process calls ACTS and provides it with hits and detector geometry

DD4hep (CLiC, SLD): define active and passive geometry, readout segmentation

ACTS has built-in support for DD4hep geometry

ACTS (ATLAS): provide general track reconstruction algorithms

Oversimplified software stack

DD4hep: Geant4 geometry, detector plugin library, wrappers to run Geant4

Juggler: Digitization and reconstruction software (based on Gaudi with Podio-based data model and ACTS for tracking)

Gaudi: Generic open project for building event processing frameworks. Enables modern task-based concurrent execution in a heterogeneous computing environment. Used by ATLAS and LHCb.

ACTS: Experiment-independent tracking toolkit (ACTS' geometry constructed from DD4hep via plugin)

Podio: Robust data model definition to cross the boundaries between the tools

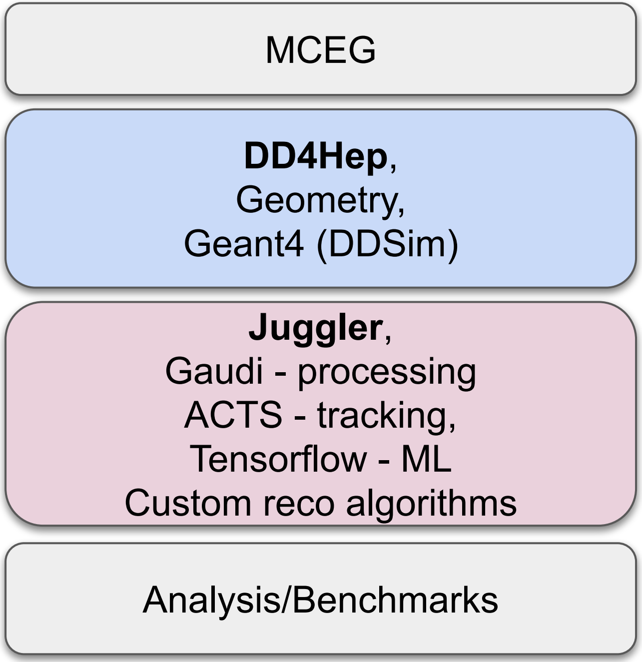

DD4Hep handling geometry problem

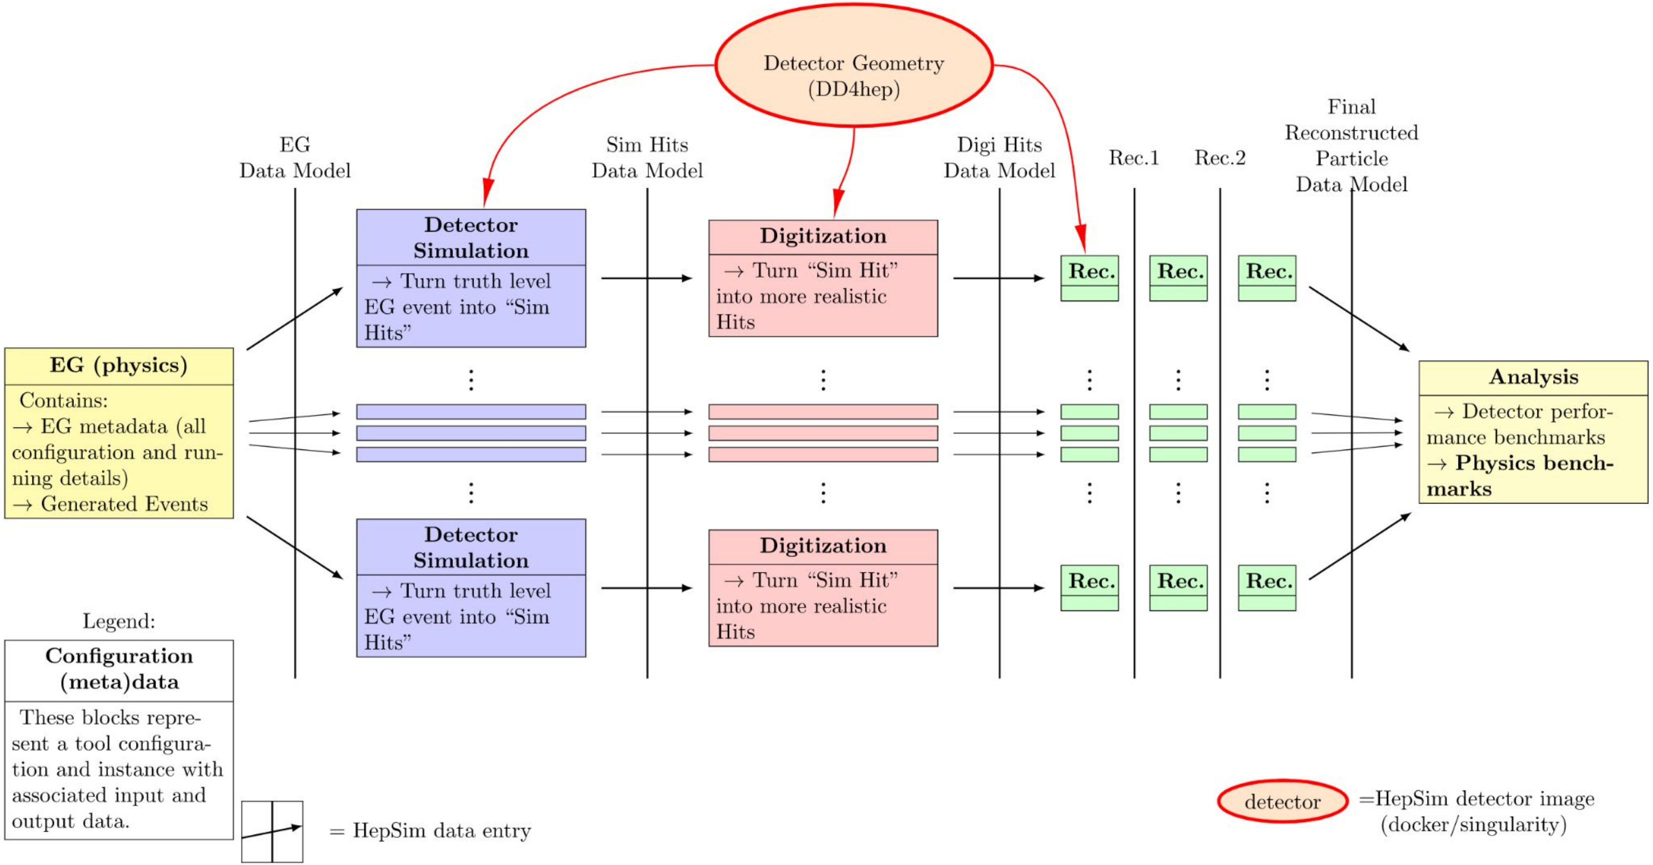

DD4Hep community

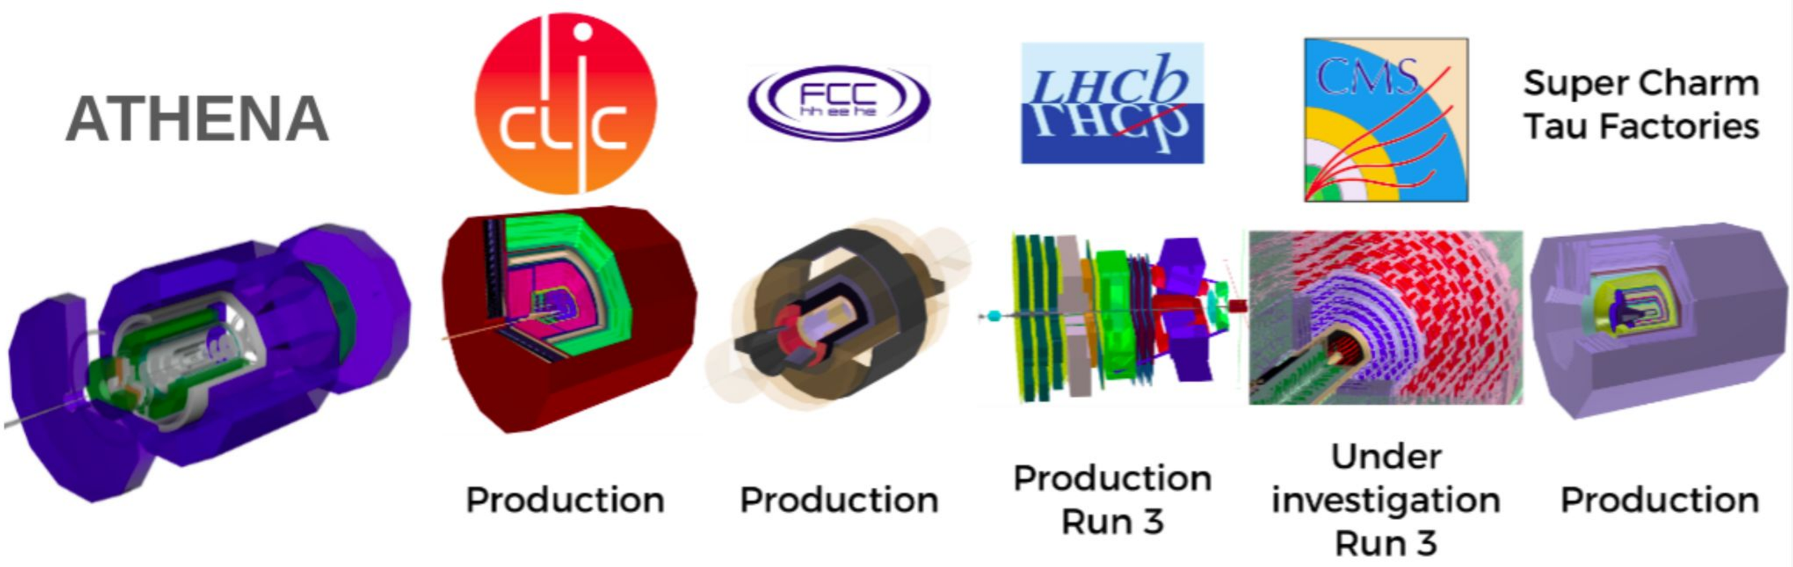

Automated workflows at eicweb

GitLab server (eicweb.phy.anl.gov)

- continuous integration
- dedicated build cluster

Runs automatically on each user commit, executing workflows running multiple tests, benchmarks and analysis

Automated containers

Both Docker and Singularity images are created nightly or on demand (commit) providing:

- reproducibility,
- production level images
- latest updates for those working locally

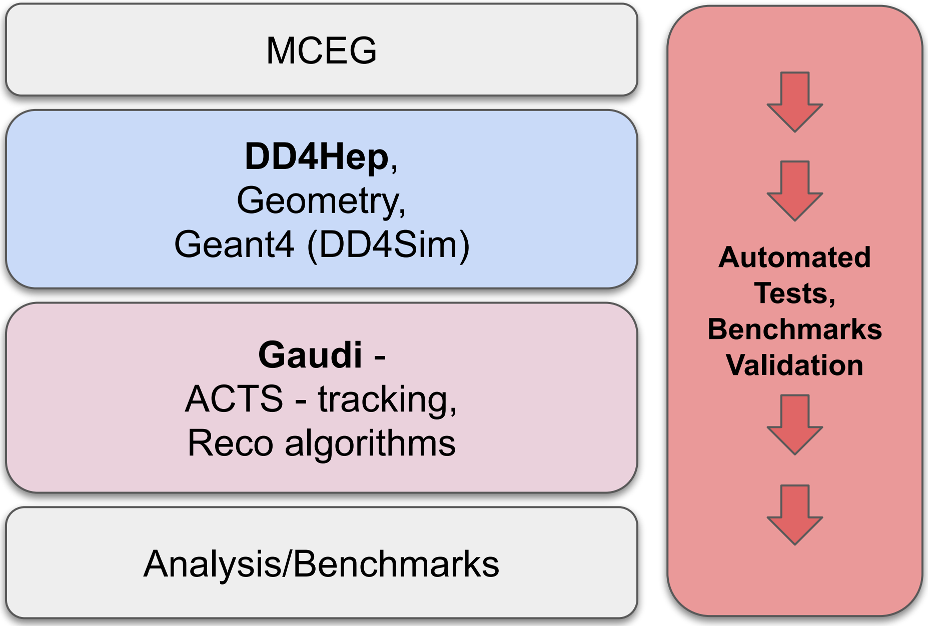

Benchmarks, documentation, conterization

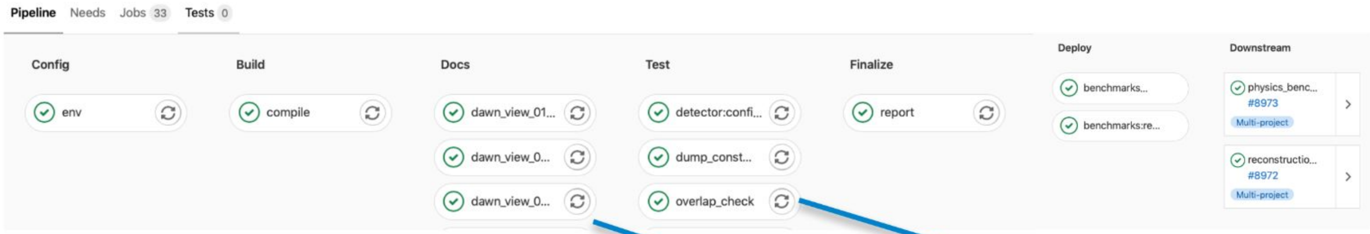

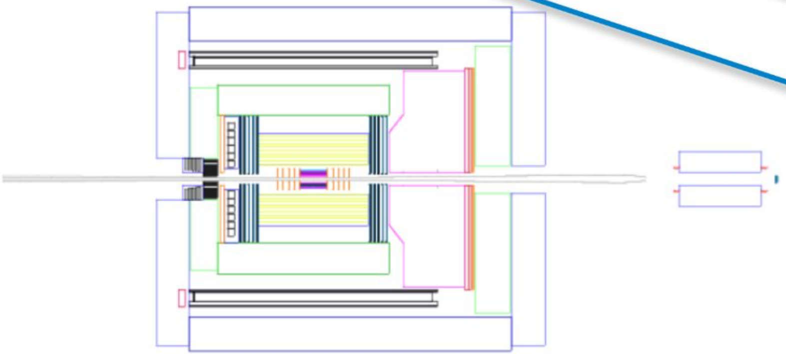

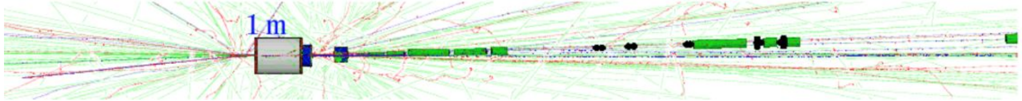

Geometry overlap checks running as part of every merge request

Automatic visualizations for detector geometries, saved as job artifacts (browsable!)

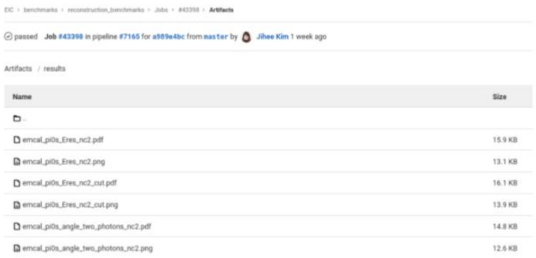

Web Integration: GeoViewer, XRootD

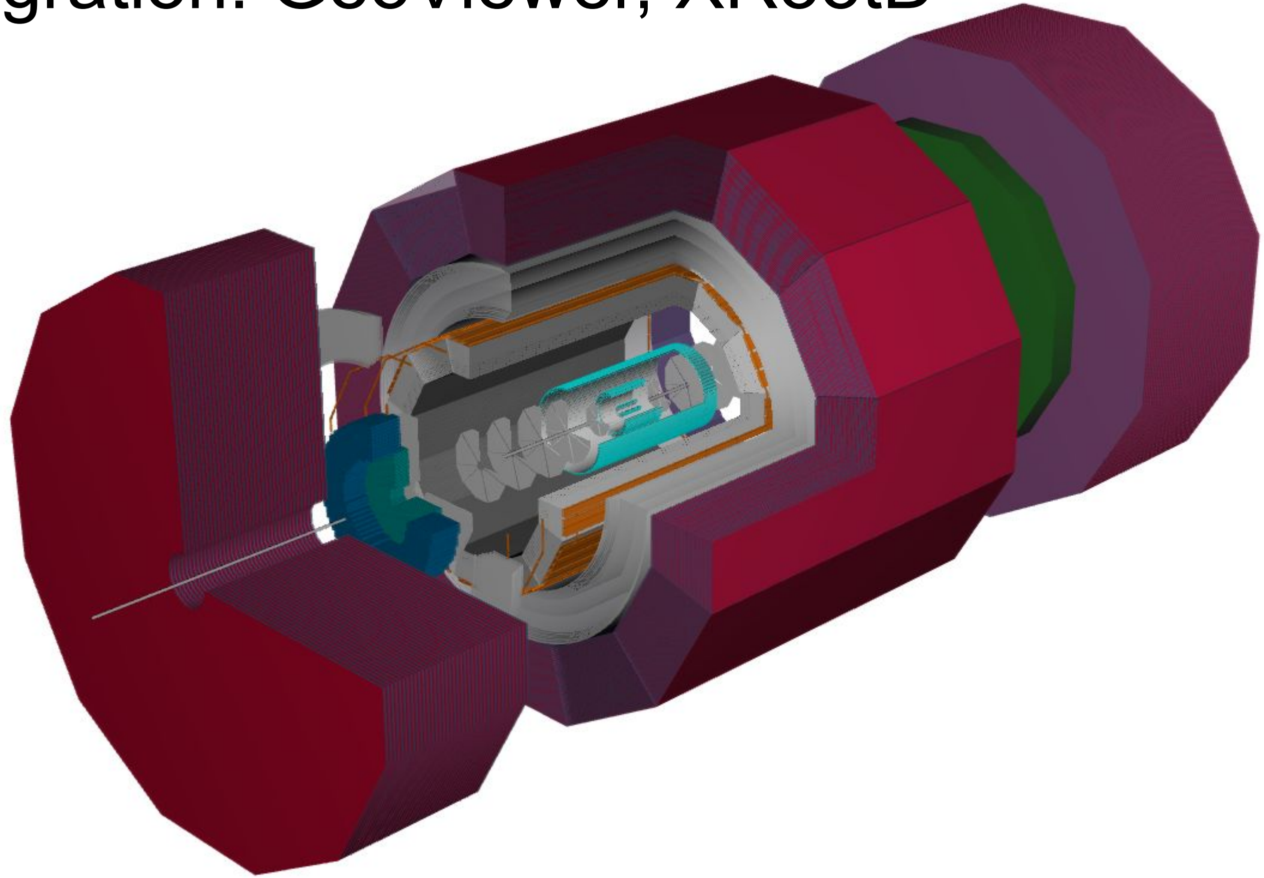

Live link

Calorimetry WG

- Ready to study impact of magnet on HCAL:
  - ✅ Realistic HCAL geometry
  - ✅ Solenoid & Helmholtz
  - ✅ HCAL clustering and energy calibration
- ECAL system well-developed:
  - Barrel ECAL:
    - ✅ Barrel SiW imaging calorimeter
    - ✅ Barrel hybrid SiW + WSciFi calorimeter
  - Electron-endcap ECAL:
    - ✅ Crystal calorimeter
    - ✅ Glass calorimeter
    - ✅ Hybrid electron endcap for baseline
      - 🚧 Optimize geometry implementation
    - 🚧 WSciFi calorimeter
- Geometry:
  - 🚧 Finalize/validate geometry for "baseline" setup

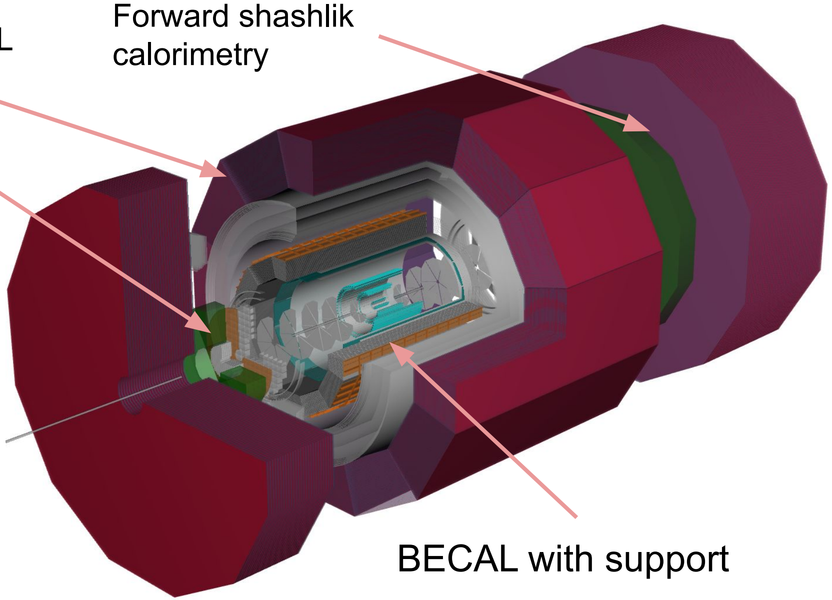

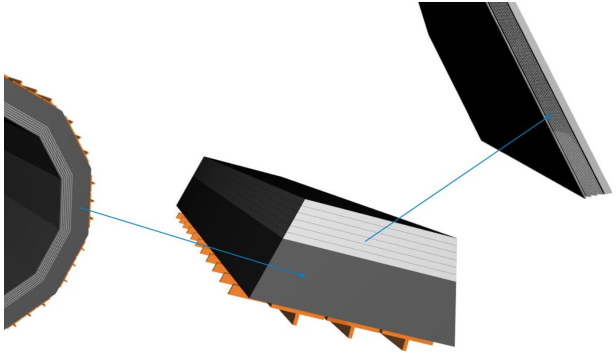

Calorimetry delegate: Vladimir Berdnikov
S&C WG contacts: Wouter Deconinck & Sylvester Joosten

Far-forward & Far-backward WGs

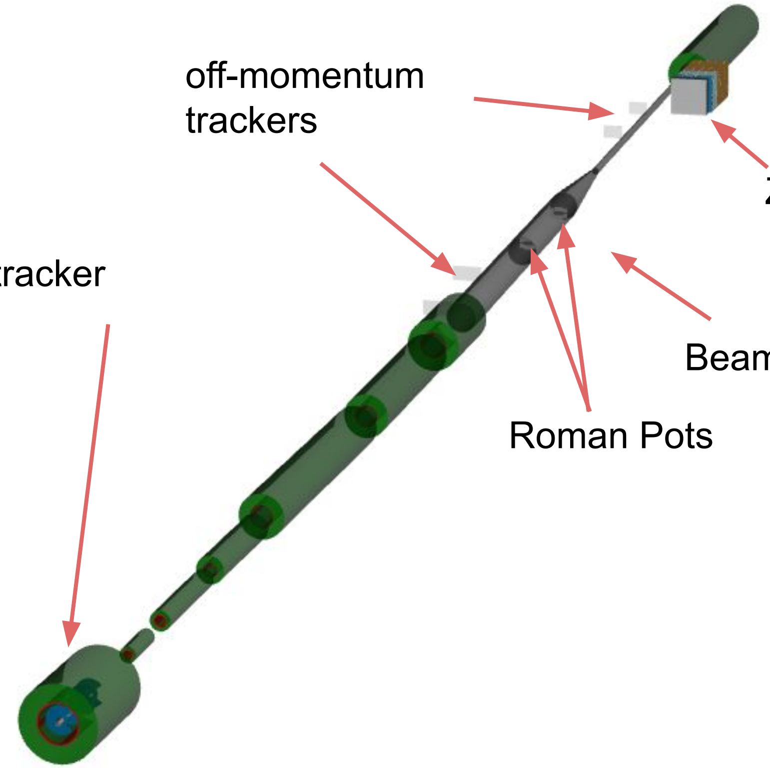

- FF being worked on by Alex Jentsch
  - ✅ Magnets and detectors updated to latest design
  - 🚧 Flip IR orientation to right-handed coordinate system - **"big flip"** being prepared, right now (tentative merge before Monday)
  - ✅ 🆕 FF beamline elements
    - 🚧 Validate/tune materials
- FB
  - 🚧 IR implementation

FF delegate: Alex Jentsch
FB delegate: Jaroslav Adam
S&C WG contact: Whitney Armstrong

PID WG

- dRICH (Christopher Dilks, Chao Peng)
  - ✅ 🆕 Base geometry plugin ready
  - 🚧 mirror alignment issue being solved by Christopher
- mRICH (Murad Sarsour, Whitney Armstrong)
  - ✅ received baseline realistic geometry (implemented in fun4All) from Murad
  - ✅ implement realistic detector in DD4hep
  - 🚧 Fix issue with optical photons getting trapped in Fresnel lens
- DIRC (Grzegorz Kalicy, Dmitry Romanov)
  - ✅ Converted geometry (some polishing needed)
  - 🚧 Refactor the code for DD4hep
  - 🚧 Make initial validation benchmark
- TOF (Zhenyu Ye)
  - ✅ LGAD implementation with realistic services
- GridPix (Sanghwa Park)
  - 🚧 Initial implementation

dRICH geometry implementation in DD4Hep

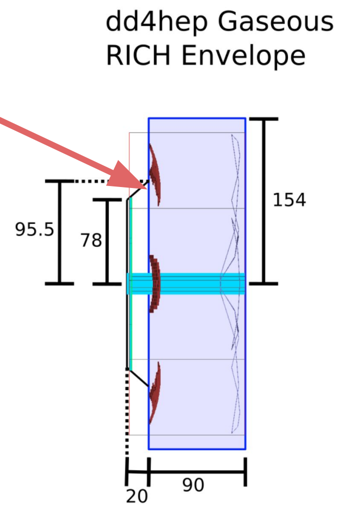

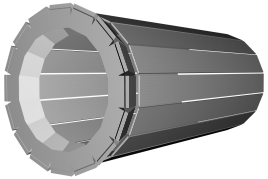

DIRC geometry ported

mRICH geometry (w/ optical surfaces) complete

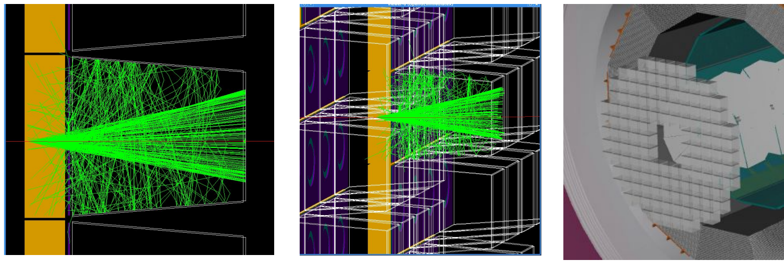

PID delegate: Zhenyu Ye
S&C WG contact: Dmitry Romanov

Tracking WG

Outer LGAD layer not part of the “0-0-0” setup

Barrel staves as in ITS2 TDR

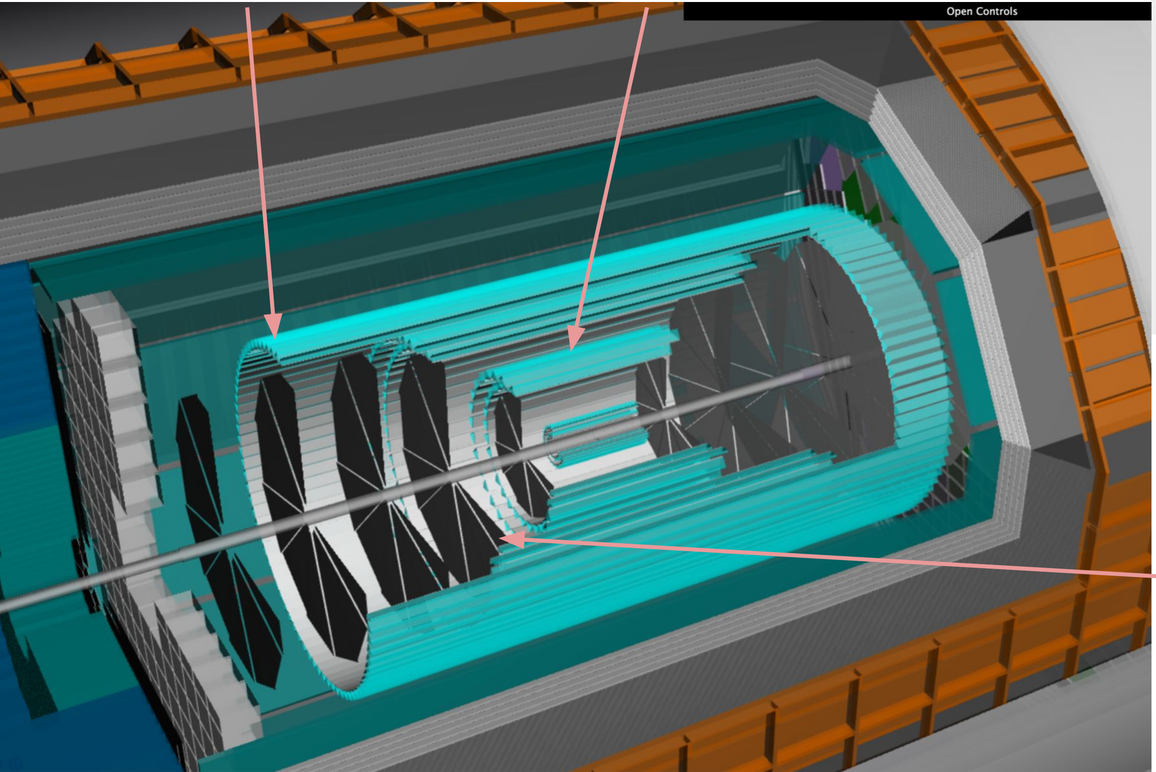

- 🚧 silicon tracker
    - 🚧 Material validation (Shujie Li)
    - 🚧 🆕Vertex layers to be changed to cylindrical geometry
    - 🚧 Validate barrel geometry
    - ❌ Support cones
- ✅🆕 barrel MMGAS (Francesco Bossu)
    - ✅ also have barrel μRWEL
- ✅ GEM

Disks are wedges with sensitive layer and average material backing. Needs better constraints from WG

Tracking delegate: Matt Posik
S&C WG contact: Sylvester Joosten

Reconstruction Status

✅ Calorimetry
- ✅ Algorithms
  - ✅ Simple Clustering
  - ✅ Island Clustering (2D)
  - ✅ 2+1D Clustering
  - ✅ Topological Clustering (3D)
- ✅ Clustering benchmarks

🚧 PID
- 🚧 Algorithms
  - ✅ Fuzzy-K ring clustering
  - 🚧 MRICH clustering (Murad)
  - 🚧 DIRC clustering
- ❌ PID benchmarks

🚧 Far Forward & Far Backward
- ❌ Integrate B0 with tracker
- 🚧 Matrix transform for Roman Pot & OMD reconstruction
- ❌ low Q2 tagger

🚧 Tracking
- 🚧 Algorithms
  - ✅ Proof-of-concept working again!
  - 🚧 Finalize tracking for baseline
  - 🚧 Tracking benchmarks
  - ❌ Incorporate B0
  - ❌ Setup realistic vertex reconstruction
- 🚧 Tracking Benchmarks
  - ✅ Basic benchmarks working
  - ❌ Tracking with realistic background

🚧 Global
- Event builder (produces ReconstructedParticle)
  - ✅ Dummy event builder to test reco chain
  - 🚧 Implementation of full realistic event builder
- 🚧 Finalize data model
- 🚧 Tutorial on full reconstruction

Ways to get involved NOW (from 🚧 to 👷 to ✅)
- Develop test cases for reconstruction (e.g. here or here)
- Get familiar with running code locally using containers
- Implement material geometry (often without coding)
- Clearly document detailed geometry in issue tracker

Resource Utilization

From Bottom's Up Resource Estimates

- **Identification of synergies**
- **Streamlining of run plans**
- **Prototyping and small runs**

To Multi-TB Production Runs since June

- **Primarily consumers: singles, (SI)DIS**

Large scale ATHENA data productions

- Input: HepMC files preferred (`mcconv` developed for other formats)
- Full simulation with current detector model, all bells and whistles:
  - Typical: 0.25 to 3.0 s/event, <500MB RAM RSS, 30 kB to 750 kB output size/event
  - ScFi Barrel Ecal: example of how the impact of implementation on simulation timing was mitigated with judicious choice of intermediate non-physical volumes
  - Full ROOT files on S3 under `ATHENA/FULL/` (but likely only need reco files)
- Full reconstruction (reconstruction_benchmarks/benchmarks/full):
  - Calorimetry clustering (Ecal, ScFi, Hcal), tracking (up to inner tracker), RICH hits/digi
  - Reco ROOT files on S3 under `ATHENA/RECO/` and `sci-xrootd.jlab.org`
    - Working on jsroot and file browser support on `sci-xrootd`
- Full simulation: ~weekly repetition; reconstruction: every few days
- Written to work on any slurm batch system; performed at Compute Canada
  - Trial runs on OSG at the ~2k job scale for single particle events

compute canada | calcul canada

Computing Resource Utilization Analysis

- **Single particle simulations** with General Particle Source
  CPU time per core, event size: both scale empirically as $E^{0.85}$
  - pi+: 20 ms/ev · $E^{0.85}$, 3.4 kB/ev · $E^{0.85}$
  - neutron: 20 ms/ev · $E^{0.85}$, 2.2 kB/ev · $E^{0.85}$
  - gamma: 17 ms/ev · $E^{0.85}$, 5.7 kB/ev · $E^{0.85}$
  - pi0: 17 ms/ev · $E^{0.85}$, 6.1 kB/ev · $E^{0.85}$
- **Multi-particle simulations**: Pythia8, DIS NC/CC $Q^2 > 10\ GeV^2$
  - 25 mRad, 5x41 GeV: 28.4 tracks/ev, 666 ms/ev, 186.1 kB/ev
  - 25 mRad, 18x275 GeV: 38.8 tracks/ev, 3010 ms/ev, 566.1 kB/ev
  - Empirical scaling with s · min(Q)
- Benchmarks on HS06 ~12 nodes
- Since these studies: slight increase in time per event in past weeks (scintillating fibers), to ~35 ms/ev · $E^{0.85}$

EIC AI/ML in ATHENA

- Current use of AI/ML
  - ✅ e/π PID with 3D shower profiles from imaging calorimeter in center barrel region.
- Near-term anticipated use:
  - 🚧 ACTS: Track finding
  - 🚧 PID: Pattern recognition in RICH, DIRC0
  - ✅ Calorimetry clustering (2D, 2+1D and 3D clustering)
  - 🚧 DNN-based fast simulation
  - 🚧 DNN-based detector optimization (Bayesian optimization)
  - 🚧 DNN-based reconstruction
- Implications on computing infrastructure:
  - 🚧 Many exascale GPU accelerators, but lack of support in current software tools limited by IO/memory bandwidth

✅ = working, 🚧 = in progress, ❌ = planned

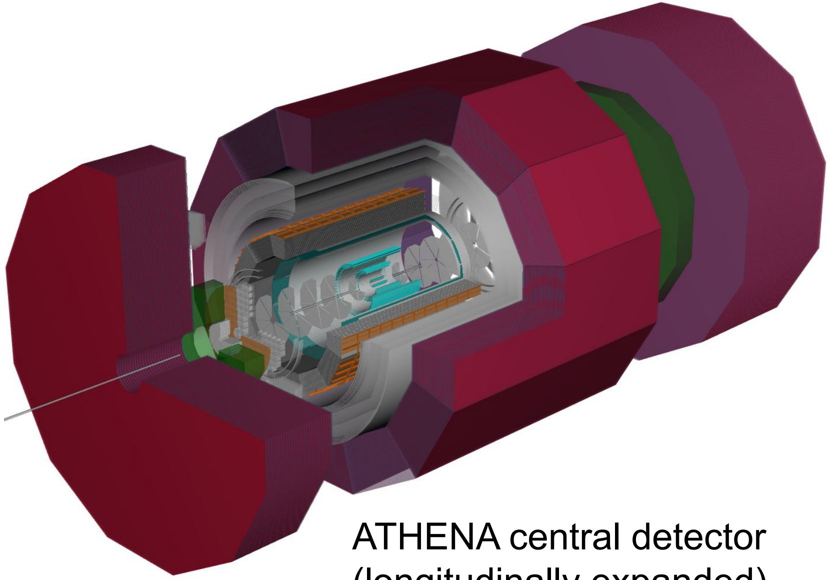

ATHENA central detector (longitudinally expanded)

Full simulation/reconstruction team

Whitney Armstrong, Miguel Arratia, Wouter Deconinck, Sylvester Joosten, Jihee Kim, Chao Peng, Tomas Polakovic, Dmitry Romanov, Marshall Scott, Zhenyu Ye, Ziyue Zhang, Maria Żurek

...and a rapidly growing amount ATHENA collaborators!

Software & Computing Conveners:

Whitney Armstrong, Andrea Bressan(*), Wouter Deconinck, Sylvester Joosten, Dmitry Romanov

(*)- liaison to EICUG software group