ATLAS workflow: Monte Carlo Production

Doug Gingrich
University of Alberta and TRIUMF
Junichi Tanaka
University of Tokyo

MC production workflow

- We'll explain the MC production workflow of the bulk sample requests, so-called, MC15a and MC15b campaigns.

- MC15a:

  - mc15_13TeV.
    410007.PowhegPythiaEvtGen_P2012_ttbar_hdamp172p5_allhad.merge.
    AOD.e4135_s2608_s2183_r6869_r6282

  - We have 5 different job types: e4135, s2608, s2183, r6869, r6282
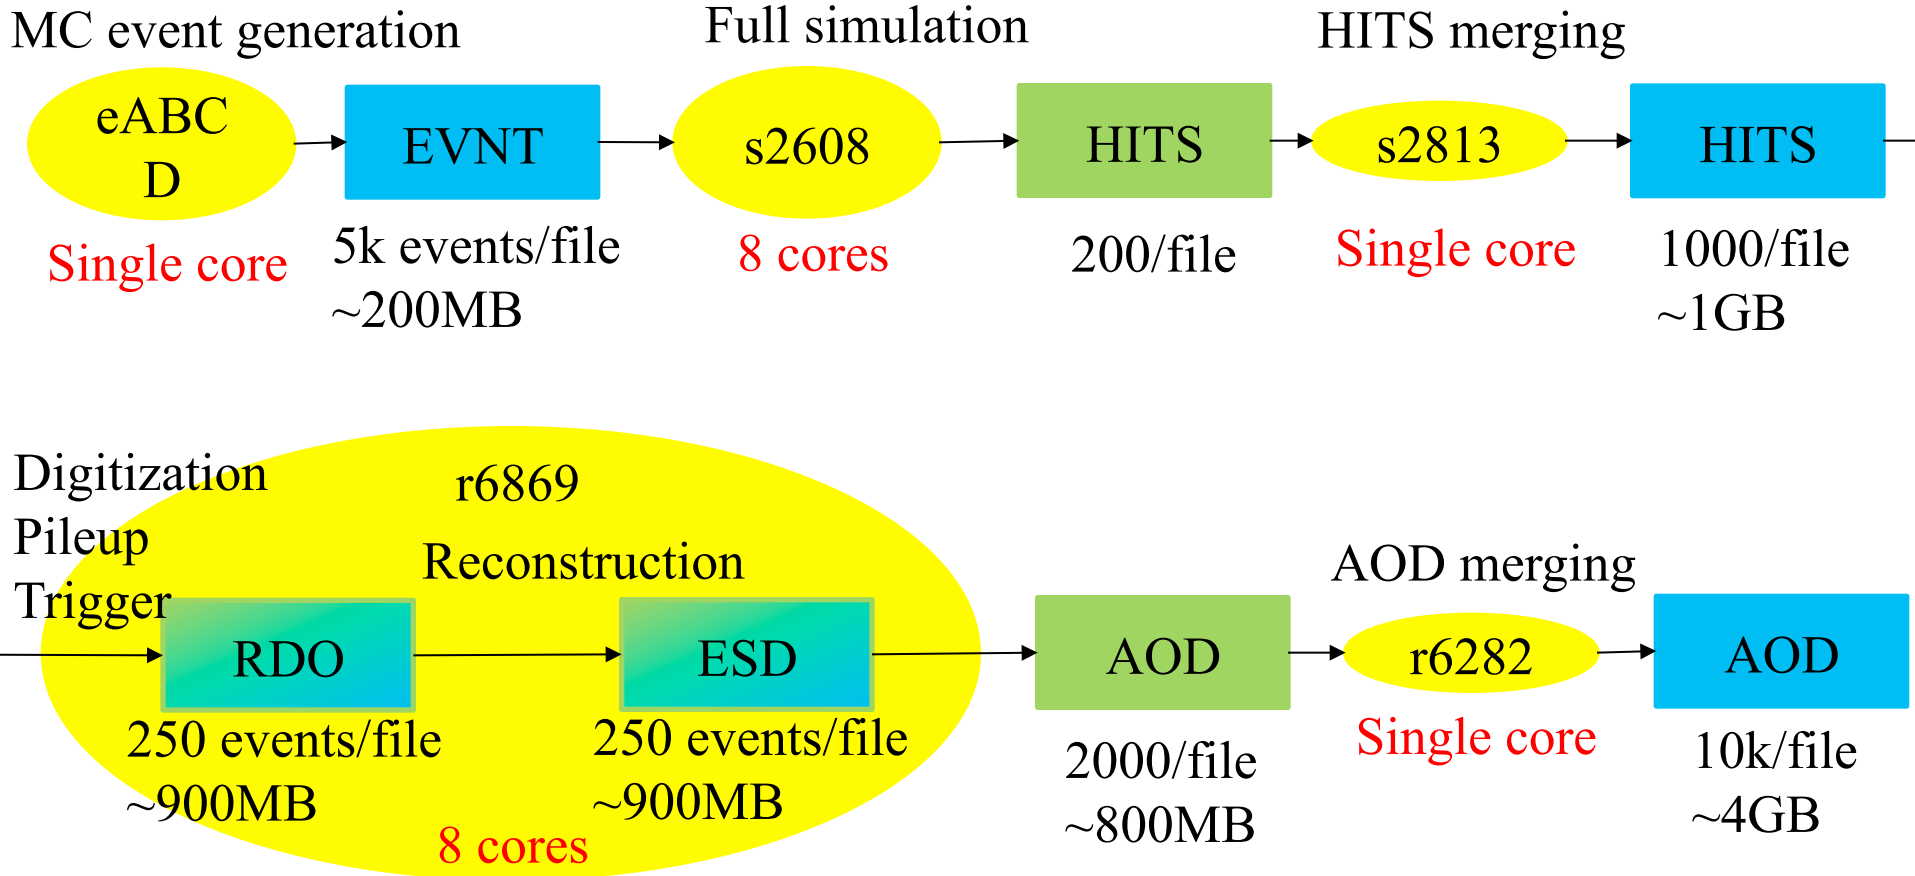
- MC15b:

  - mc15_13TeV.
    343362.PowhegPythiaEvtGen_P2012_ttbar_hdamp172p5_dil_highMjj.m
    erge.AOD.e4661_s2726_r7326_r6282

  - We have 4 different job types: e4661, s2726, r7326, r6282

MC production tags

- All tags (e-tag, s-tag, r-tag, etc) now maintained in AMI.

- e-tag: EVNT (EVGEN) production.

  - We run MC event generators, for example Pythia6/8, Herwig++, Powheg, Sherpa, MadGraph, Alpgen, etc.

  - Index of e-tag is different for each request so there are many e-tags.

- s-tag: Geant4 simulation and also HITS merging.

- r-tag: Digitization and Reconstruction, and also AOD merging.

- Ideally we would have only one s/r-tag for each campaign (MC15a, MC15b, …) but in practice we have several s/r-tags depending on improvements, bug-fixes, etc.
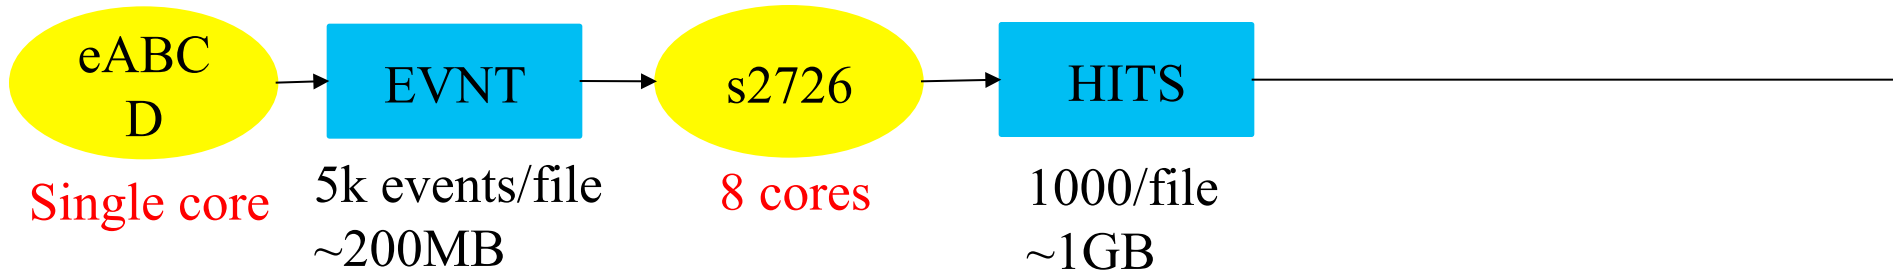
MC15a : eABCD_s2608_s2183_r6869_r6282

MC event generation   Full simulation   HITS merging

eABCD → EVNT → s2608 → HITS → s2813 → HITS →

Single core   5k events/file ~200MB   8 cores   200/file   Single core   1000/file ~1GB

Digitization
Pileup
Trigger

r6869

Reconstruction

→ RDO → ESD → AOD → r6282 → AOD

250 events/file ~900MB   250 events/file ~900MB   2000/file ~800MB   Single core   10k/file ~4GB

8 cores

In some requests, we need to keep RDO and/or ESD as persistent data.

job   Transient data   Persistent data

MC15b : eABCD_s2726_r7326_r6282

MC event generation   Full simulation

eABCD → EVNT → s2726 → HITS →

Single core

5k events/file
~200MB

8 cores

1000/file
~1GB

Digitization
Pileup
Trigger

r7326

Reconstruction

RDO → ESD → AOD → r6282 → AOD

250 events/file
~900MB

250 events/file
~900MB

8 cores

AOD merging

Single core

2000/file
~800MB

10k/file
~4GB

In some requests, we need to keep
RDO and/or ESD as persistent data.
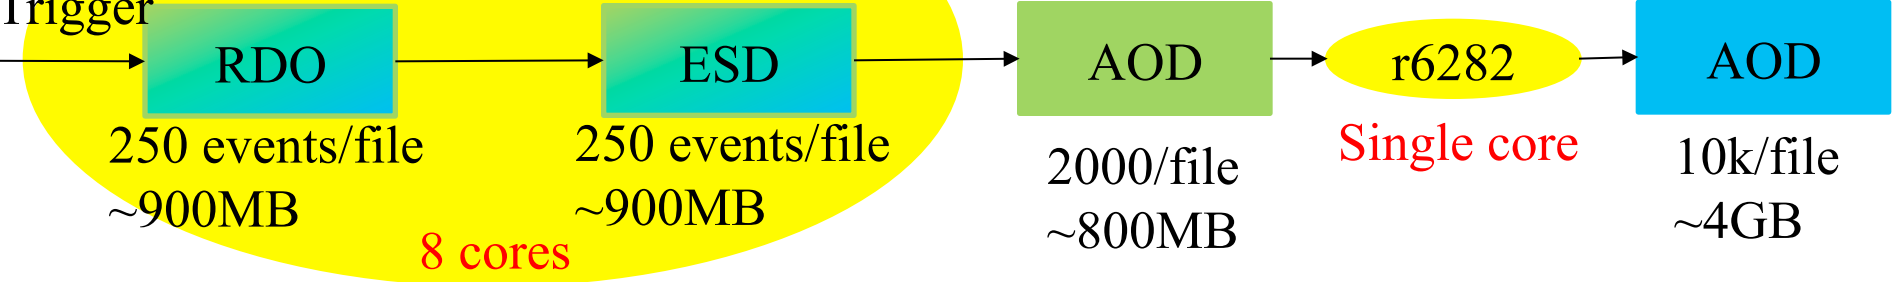
job   Transient data   Persistent data

Future Plans

- We don't have any plan to change the workflow for the bulk production.

  - We'll use MC15b-type workflow for MC15c and MC16.

- We'll have different (at least two) types of "new" tasks to produce samples.

  - Embedding tasks: Replace some data information with MC information.

  - Overlay tasks: Add data events (pile-up) to MC events.

  - Both tasks are still begin discussed and tested.

  - We'll be able to provide workflows for these tasks after tests.

- Comment on "dark logs" from EVNT production (MadGraph events).

  - We realized that the size of log files is important for healthy grid operations.

  - We'll check it carefully for future production.

- So far we are using multi-cores for simulation and digitization + reconstruction but we aim to use it in EVNT production for some generator/physics processes.
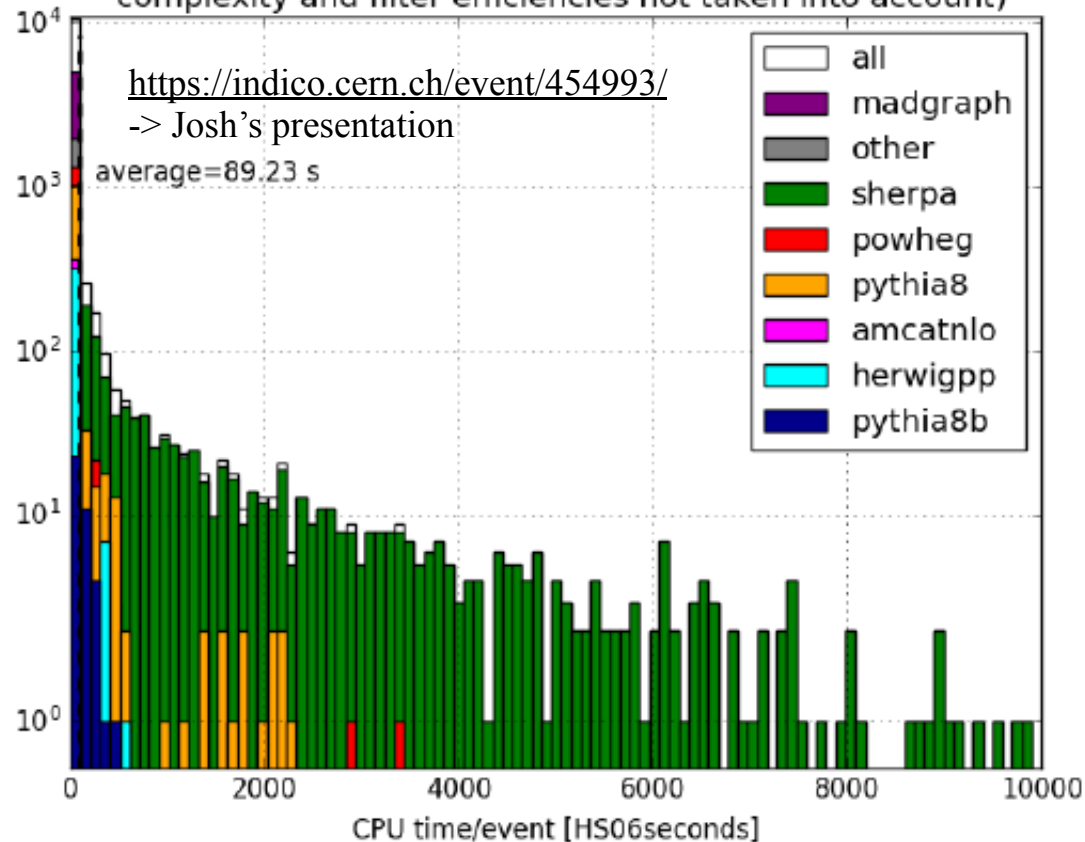
Backups

Resources

- Full simulation.
    - We are asked to have 2400 HS06sec/event by ADC.
    - We can finish one job (1000 events/job) within 8 hours@8core node.
- Other jobs (Fast Simulation, Digitation + Reconstruction, Merging).
    - No request from ADC except EVNT production.
    - Our standard setting (the number of events per job) looks OK in term of CPU, memory, disk space, etc.

CPU for EVNT Production

We ask requesters to finish one job within one-day CPU time.
5k events/file means that the average CPU time per event is 170 HS06sec.

CPU time/event for 2015 MC event generation at $\sqrt{s} = 13$ TeV
(All physics processes included, correlations with process complexity and filter efficiencies not taken into account)

https://indico.cern.ch/event/454993/
-> Josh's presentation

average=89.23 s

- all
- madgraph
- other
- sherpa
- powheg
- pythia8
- amcatnlo
- herwigpp
- pythia8b

CPU time/event [HS06seconds]

- Most of tasks satisfy this requirement.
- However because of heavy CPU time requirement for some physics processes, we need to optimize the number of events per a job.
- Longer jobs might be killed due to CPU (walltime) limit in WNs but we may want to have a longer queue to avoid unnecessary iteration (many attempts) if possible.
- We'd like to have 8 (or more) cores for EVNT production.