

Analysis Model

Recent Hot Topics

CAT Specific Issues

Analysis Model Workshop

Goals & Sampling of Discussion Topics

Amir Farbin

CERN

Structure Athena-Aware Ntuples

- Ketevi has presented ATLAS with the provocative proposal of replacing the POOL based AOD with SAN
- Prompted by:
 - Apparent slowness of current AOD
 - Strong user desire to access AOD directly in ROOT
 - Need to write AOD objects into ntuples as DPD output of analyses
- BNL working meeting to illuminate the technical issues (Sept 28-29):
 - Providing Athena read-back and full AOD functionality from SAN likely to be a lot of work.
 - The up-coming persistent AOD may provide the same functionality
 - T/P separation + optimization of our use of POOL is likely to lead to very similar performance of POOL AOD and SAN
 - T/P separated AOD is ~ 5x faster in Athena (no numbers from SAN). No optimization done yet. Much more speed improvements expected.
 - Appears that ROOT access to persistent AOD is on par with very latest SAN performance numbers
 - Adding accessor functions, better naming, and some reorganization of persistent AOD → POOL AOD identical to SAN.
 - However adding user interface requirements to persistent AOD may be in conflict with the technical motivations for T/P split
 - ➔ Until now, no requirement that persistent EDM be directly interpretable.
 - ➔ My opinion: Problem for RDO & ESD, but no issue for AOD

SAN Taskforce

- A taskforce has been setup to make a recommendation wrt the SAN proposal
- My view of goals of TF:
 - Evaluate performance/technical advantages (if any) of SAN versus optimized POOL AOD.
 - Weight this evaluation against
 - Manpower necessary to implement full AOD functionality from SAN.
 - Implications of supporting 2 parallel persistency machineries and potential inconsistencies between SAN AOD and POOL ESD.
- Analysis model implications: How much do we support non-Athena based analysis?
- SAN TF+other forums → ATLAS is now essentially re-evaluating choice of persistency technology.

AOD Analysis

- Profiling typical full analyses:
 - After lots of optimization: SUSYView/TopView run at 2-3 Hz now:
 - 55% spent in data access (StoreGate/POOL).
 - 35% spent in writing DPD (copy AOD sub-set to Ntuple).
 - 15% spent in EventView analysis. (See PAT session on Tuesday for detailed break-down).
- T/P & POOL optimization, AOD thinning, and AOD/ESD merger are likely to bring significant speed improvements to AOD analysis.
- Writing SAN from Analysis may also provide speed improvements.
- Always limited by slowest part of job
 - If read-in/write-out 10x faster: limited by complexity of analysis
 - Faster AOD → less need to copy AOD to ntuple → Do less in one job → more/faster iterations.
- Back of envelope projection: Today's SUSYView/TopView at 10-20 Hz in 1 year... much faster if these packages do less.
- We shouldn't forget
 - One point of using Athena for analysis is to re-do complex algorithms like jet finding, b-tagging, ... complexity limits speed
 - Scalability (Disk access issues/data volume): If you want to see a plot for all ATLAS data in 5 mins, you'll need to simultaneously use 1000's CPUs with their own copies of the data... even in ROOT.

Today's CAT Analysis Model

- CAT lxbatch queues + CAF castor pools + data replication to castor → rapid processing of AOD by CAT members. → We're very happy!
- No need to use GRID... yet?
- Fundamental problem: no large (TBs) network accessible disk pool for DPD
 - Jobs must read/write to castor
 - Castor not intended for interactive access
 - So multiple copies of ntuples sit on various desktop disks.
- Solutions?
 - Large NFS mountable disk (eg Trigger "space" area)
 - we require global access (ie lxbatch/lxplus)
 - Generally dislike of NFS by IT
 - BaBar members: But this is how SLAC worked!
 - AFS? → Volumes usually too small!
 - More demands on castor → their goals are not interactive user access.

Analysis Model Workshop

(Oct 25-28 @ CERN)

- Goals:

- Try to get global perspective of offline sw/computing from physicist (end-user) point-of-view → produce document for users outlining the current status and plans.
- Try to think through the technical steps of analysis → discuss

- Topics:

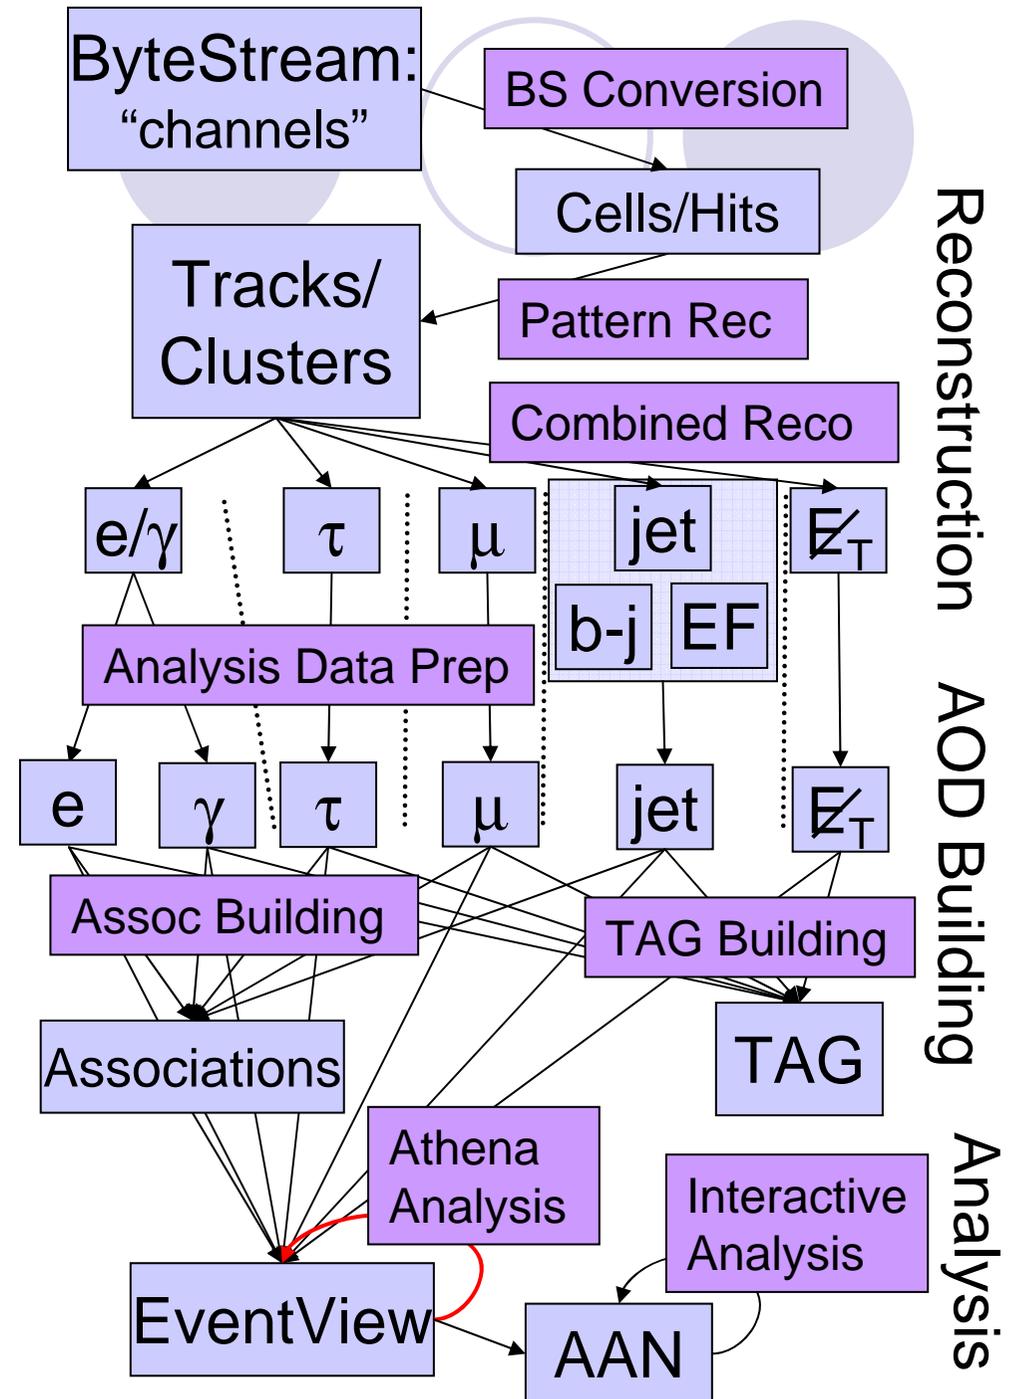
- Analysis Model talks from D0, CDF, BaBar, & CMS
- Computing Model Overview
- Data Access, MetaData, Databases
- Steams/Tags
- Trigger, Reconstruction, Data Preparation
- Event Data Model
- Physics Analysis Tools, Distributed Analysis

Analysis Model Workshop Format

- Each session composed of
 - Moderator(s): lead discussion
 - Panel members: experts on relevant topics
 - Wise-persons (Frank Paige, Paolo Calafiura, Daniel F.): report impressions and make recommendations to SPMB
 - Other experiment AM speakers: provide experience
 - General Audience: interact with experts
- Sessions start with introduction(s)/overview by session conveners.
- Then follow bulleted list of discussion topics/questions (compiled on wiki... I'll give a sample today).
- When appropriate, panel members or other experiment speakers may provide few (~3) slides to address specific discussion points.
- Unlikely that we will come to clear conclusions, but hopefully we identify the issues.

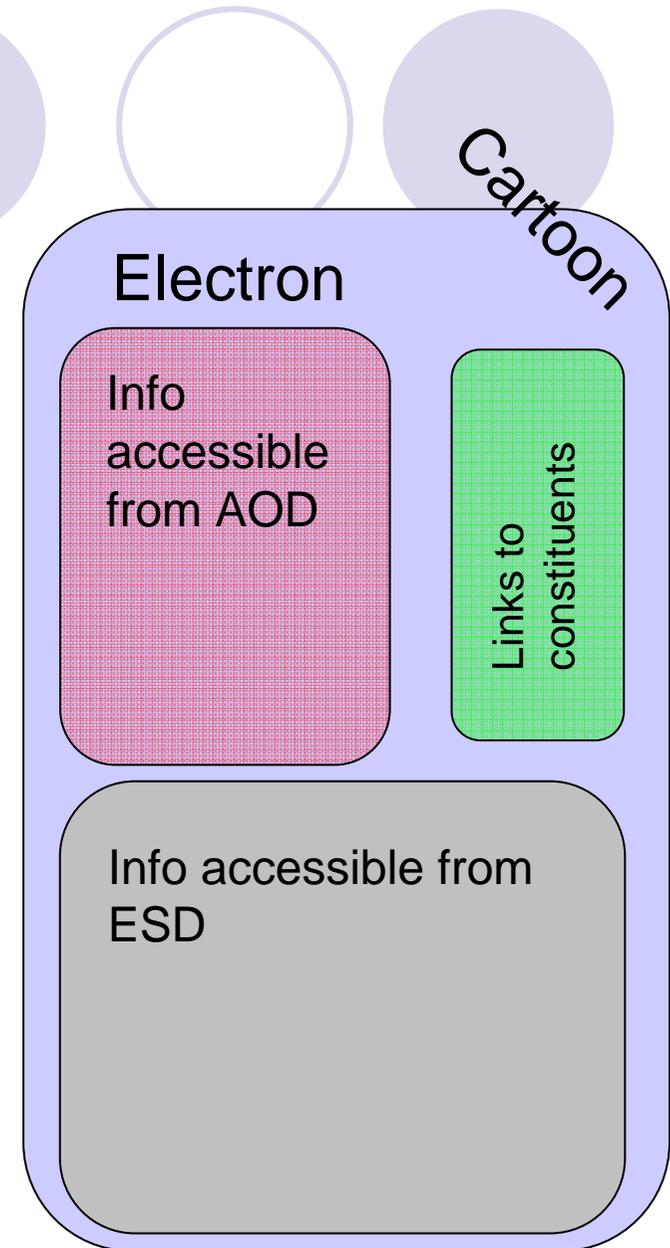
Event Data Model

- EDM design, organization, content, & size has implications on what can be done at analysis time.
- Panel members will overview the data flow (provide schematics).
- Basic principles:
 - Delay algs/corrections as much as possible
 - Example: Electron calibration based on layer weights. Can be applied on AOD when reading Electrons → minimize need to reprocess data for distribution of new calibrations
 - Example: Particle ID bit masks or likelihoods can be calculated on AOD
 - Constraints:
 - Not enough space to store all objects (eg all cells on AOD)
 - Too time consuming (eg always rerunning jet finding on AOD)
- Lots of proposed changes to make more possible at AOD



AOD/ESD Merger

- Currently, an analysis performed on particle object in AOD (eg Electron) will need to be rewritten to work on comb reco object in ESD (eg e/gamma)
- Japan PAT WS: proposed common AOD/ESD interface.
- Borrow ideas from CaloCluster split store: the same object give access to more info in ESD than AOD.
- Other advantages of breaking obj into pieces:
 - Read only the necessary pieces
→ faster access
 - Easily promote/demote info between AOD/ESD.
- AM WS: discuss implications of doing this to all objects
- E/gamma WS in Nov to merge e/gamma+Electron+Photon



Some other EDM Discussion Topics

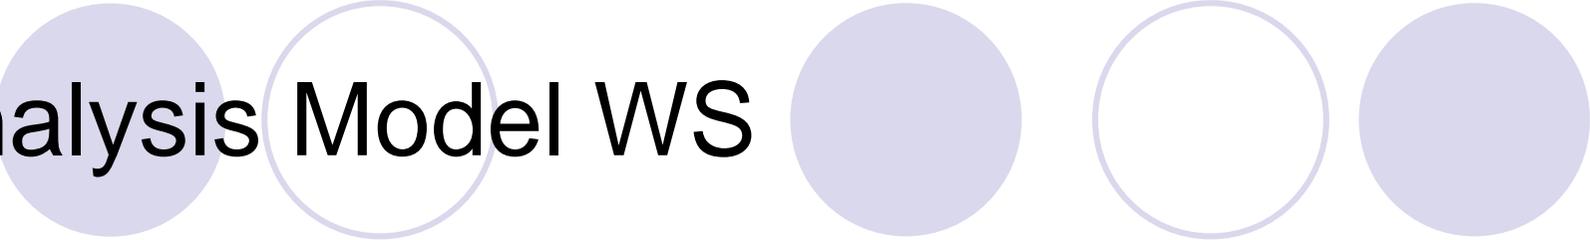
- Uses of AOD Thinning: Use polymorphic containers to flag subset of objects used in analysis and written out new AOD → next AOD job 10x faster.
- Cells and Hits for select objects (eg electrons) in AOD
- What are appropriate selections for:
 - ESD containers: output of combined reco?
 - AOD containers: subset good for analysis?
 - TAG: Good for selection?
- Get a sense of start up versus steady state from other experiments.
- Identify other (event-level?) quantities necessary for analysis. Example: Localized Activity:
 - N of tracks/hits/segments in NxN bins of eta/phi → find leaking jets
 - Sum cell Energy (per sampling?) in NxN bins of eta/phi → underlying event

Trigger Discussion

- Different types of trigger aware analysis:
 - Optimize trigger menu or alg for specific physics channels.
 - Analyze/extract trigger efficiency from control samples.
 - Determine trigger eff/back rates for specific physics channel.
- 2 general types of trigger objects available for analysis: trigger decision & objects reco'd/ID'd by trigger.
- Use of these objects depends on type of trigger-aware analysis and trigger eff information provided by trigger experts:
 - One possibility: Electron trigger eff wrt offline for e25i as function of p_T /eta/isolation. → no need for general use of trigger objs.
 - Another: Electron trigger eff wrt to HLT Electron trigger obj from simulation → explicit dependency on trigger simulation and obj
 - Another: Trigger simulation good enough (or tuned) to give the correct answer → just ask the MC for the eff
- Discuss how to combine “eff tables” + changing trigger menus/conditions + simulation → estimate eff, back-ground rates.

Cond DB Discussion

- Overview of what tasks require cond DB access and what is stored in cond DB.
- What additional high-level info is appropriate for storing in cond db
 - Ex: H1 or electron layer weights?
 - Ex: isEM cut values, likelihood parameterizations?
- Mechanisms to keep track of applied corrections.
- “Local” access to DB... Laptop?
- Luminosity Block:
 - We need following mappings: data → LB, LB → lumi, lumi → LB → data.
 - LB info access in athena/ROOT analysis
 - Integration w/ data access (I want all data w/ $X < \text{lumi} < Y$)
 - Lumi calculation versioning (the calculation will be refined over time)



Analysis Model WS

- You may look at:

<https://twiki.cern.ch/twiki/bin/view/Atlas/AnalysisModelWorkshop> for more details of discussion topics.

- Please attend the WS, actively participate, and perhaps serve on panels.