Report from Analysis Ecosystems Workshop II

23–25 May 2022, Paris

Oksana Shadura
University Nebraska-Lincoln

Software & Computing round table: Analysis III: Techniques and Tools
(https://indico.jlab.org/event/505/)

Hybrid workshop in Paris, IJCLab - 70 participants in person

Workshop Agenda

Analysis Ecosystems Workshop I

- First analysis ecosystem workshop was organised almost 5 years ago in Amsterdam https://indico.cern.ch/event/613842/
- Main idea was to examine the analysis ecosystem, currently and in the future with a 5-10 year view
  - Evolution of ROOT
  - Development of analysis tools landscape
  - Connection to other sciences

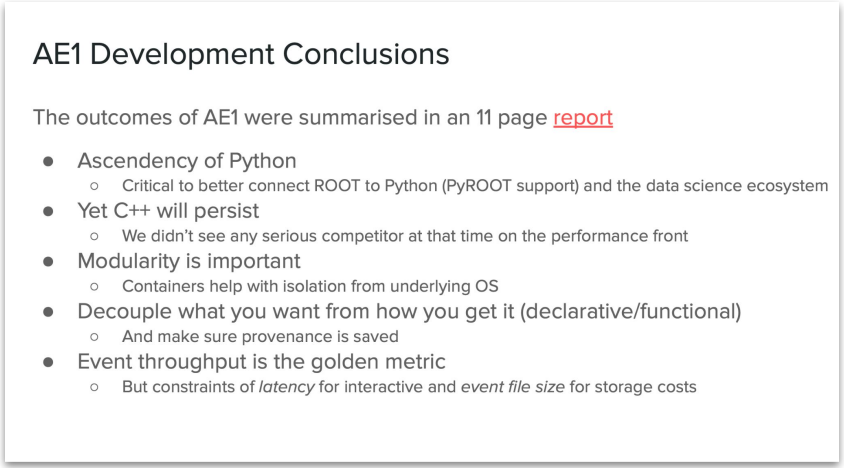

Introduction to workshop: Graeme A Stewart

AE2 Topics

Analysis Facilities

ML and Autodiff Workflows

Reduced Formats (nanoAOD, PHYSLITE)

UX and Declarative Languages

Metadata & Systematics

Realtime Trigger Analysis

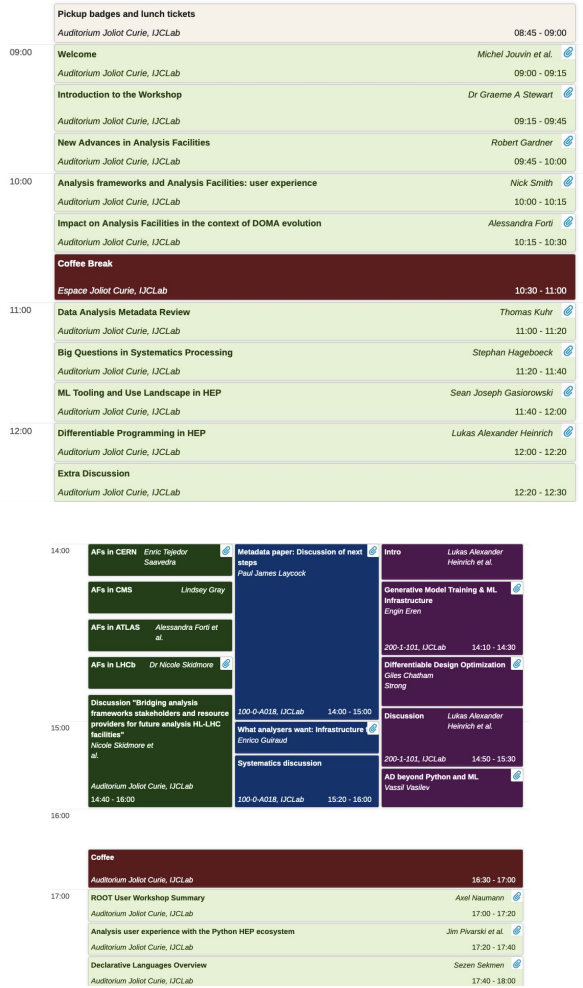

Dedicated plenary talks

Discussion sessions

Analysis facilities: from design to first prototypes

Amsterdam, 2017

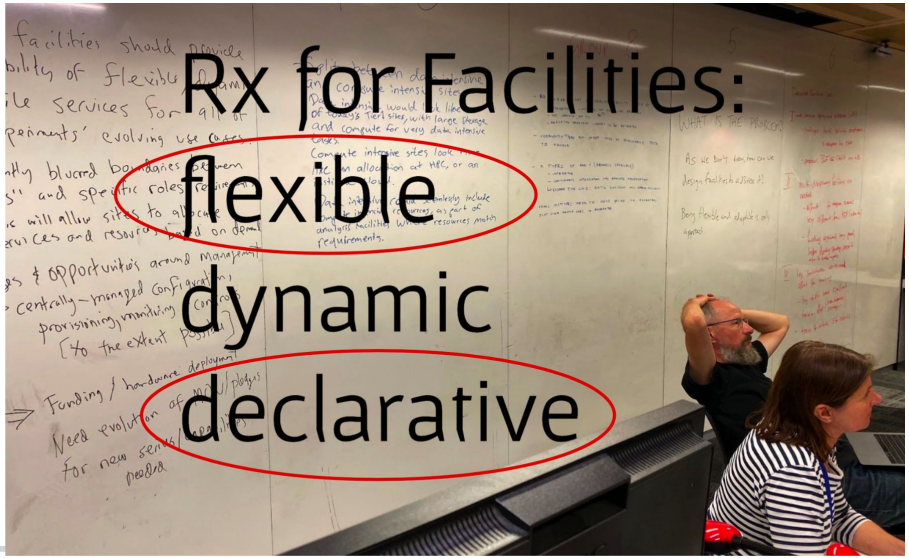

Now

Now new AFS are materializing

- Facility elements
  - service platform (substrate)
  - clustered storage
  - analysis processing
  - fast WAN for delivery
  - fast LAN for access
  - edge caches
  - interactive services
  - accelerators for ML

CAPABILITIES: can I implement new ideas quickly or even interactively? • c
I reliably get new results without lots of babysitting? • can I get my data an
software to where I need it? • can I share results / grant access with my tea
worldwide? • can I run on the hardware I need (e.g. train ML models) • can
preserve my analysis / port it to a new facility

A. Forti, HSF Analysis Facilities Forum, 21–April 22

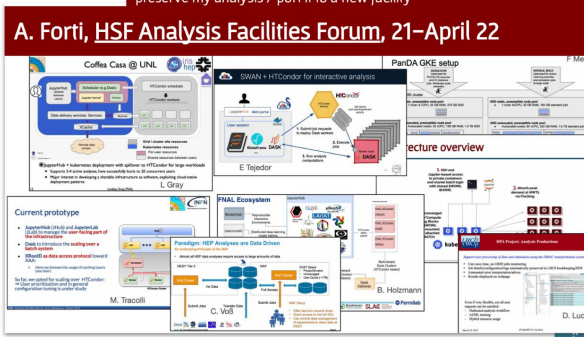

New Advances in Analysis Facilities - Robert Gardner

Analysis facilities: from the user PoV

Analysis Tools and Analysis Facilities: user experience Nick Smith

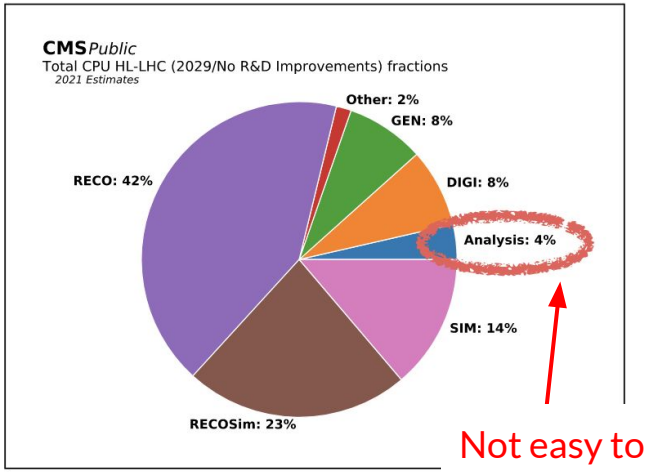

Not easy to quantify "Dark" analysis

Focused a lot on fast - it's not the bottleneck anymore. Now we have to make things easier.

Solutions must be:

Easy to use

Scalable

Fast

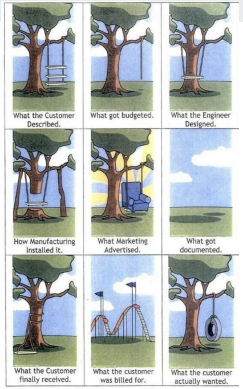

Fast

- We're all fast *enough*
  - **µs to ms per event**

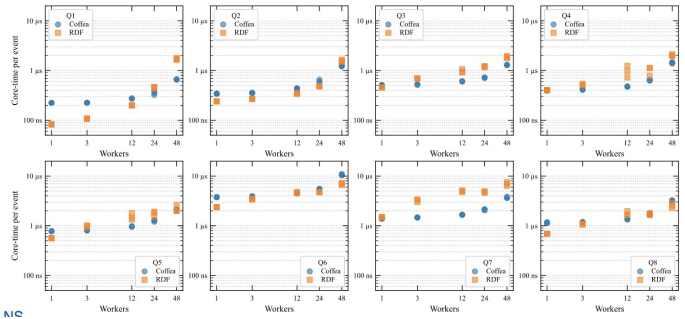

[Analysis Facilities Summary - Nicole Skidmore, Oksana Shadura]

Analysis facilities: what to take into consideration

MANCHESTER 1824

GridPP UK Computing for Particle Physics

DOMA Evolution

- HL-LHC scale of data has initiated a revision of this model introducing DOMA activities to look at
  - Network (latency/bandwidth/network orchestration),
  - Caches (hide latency, protect source storages, use diskless sites)
  - Storage QoS (reliability, vs cost vs performance)
  - Storage Deployment and Operation models (full storage, vs caches, vs remote access)
  - Access protocols (gridftp vs https, root, s3)
  - AAI (x509 vs tokens)
  - Storage type (object stores vs posix)
  - Data formats (data transformation and delivery services)
  - Introduction of DataLake and CDN concepts
- So what has all this to do with the Analysis Facilities?
  - AFs use a lot of these concepts to optimise data access for new workflows

3

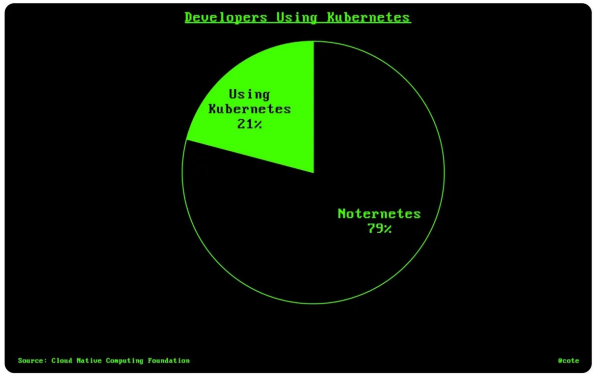

Current discussion that Kubernetes has "Crossed the Chasm" and "is entering the mainstream."

Impact on Analysis Facilities in the context of DOMA evolution - Alessandra Forti

Highlights from AF discussions

- Early testing is a good approach!
- Need to create the shared knowledge database about DOMA related technologies and its integration in AF (e.g. tokens, xcache)
- Work together with analysis frameworks developers to achieve better user experience
- How to organise possibility for users to share environment
- Importance of easy reproducible installation for Analysis Facilities (!)
- Analysis Grand Challenge could be used as a baseline test for scalability of AF and Data Lakes
- Collection of AF related metrics for better understanding what users are doing on facility
    - Working on analysis performance, resource usage and UX metrics

Lots of topics to be followed up in HSF Analysis Facilities forum!

Machine learning: “big” issues

Large models (that are also big)

This is the **number one advancement of ML in industry, and we're not using it**

- Models trained at scale with many, many parameters are showing a degree of generalisation and performance that wasn't thought possible (particularly for language)

We didn't discuss it much, mostly because we have no experience with this type of thing. -> possible industry collaboration point?

Would require very carefully selected task definition, lots of training data, and a metric buttload of compute

Sprouts in the shape of text 'Imagen' coming out of a fairytale book. | A photo of a Shiba Inu dog with a backpack riding a bike. It is wearing sunglasses and a beach hat. | A high contrast portrait of a very happy fuzzy panda dressed as a chef in a high end kitchen making dough. There is a painting of flowers on the wall behind him.

Teddy bears swimming at the Olympics 400m Butterfly event. | A cute corgi lives in a house made out of sushi. | A cute sloth holding a small treasure chest. A bright golden glow is coming from the chest.

Foundation models in HEP

large

Big question mark: is it worth it?

How does it compare to existing methods?

- Summary statistic-based methods, e.g. INFERNO, neos, have shown some improvement on toy problems and open data
- Also come with an additional factor of compute, and batching concerns (whole analysis in one batch?)
- Not many other uses with comparisons already done! [that came up today]

Can it scale? *We don't know!*

So one thing that is clear: we need toy problems that have a degree of realism for comparison, along with expert-tuned benchmarks from standard methods

4

[ML tools and differentiable computing workflows](#) - Lukas Alexander Heinrich, Nathan Daniel Simpson

Differentiable programing: what we should expect?

Machine learning:
more tracking of communalities and benchmarking!

Common, realistic benchmarks!

- Standard ML community has very established problems that are used as metrics for all new methods
  - Handwritten digits (MNIST), CIFAR-10, etc.
- We're taking steps towards making our workflows available to people for this
  - TrackML competition in 2018
  - Calorimeter challenge
  - Your idea here!
- Also could help consolidate training data for large models

Experiment tracking

Clear that we don't do enough of this (and it's not centralised)
- tools exist, like tensorboard, weights and biases, comet.ml

We could also use this in general for analysis optimisation!
- e.g. seeing how your cutflow opt is doing over time, tracking a bunch of principled metrics along the way

ML tools and differentiable computing workflows - Lukas Alexander Heinrich, Nathan Daniel Simpson

Machine learning: the need of dedicated facilities!

Platforms

Again, this is not a centralised thing (some people have their own GPU clusters, some don't)

Some effort at CERN, e.g. ml.cern.ch, which provides a jupyter entrypoint to using GPUs + Kubeflow for sophisticated tracking of training workflows

- Also work on a VSCode frontend

Nice if more effort put into this kind of thing + visibility, especially for people still training on lxplus CPU [they exist, and we should end their suffering]

ML tools and differentiable computing workflows - Lukas Alexander Heinrich, Nathan Daniel Simpson

User experience and declarative languages: key items to improve

Pain points in analysis user experience, ordered

1. **Systematics**
   - Recurring topic throughout this workshop: this is not solved
2. **Metadata**
   - Finding & handling information
3. **Scale-out**
   - Prototyping vs scale-out, different implementations / details on different sites
   - Need for consistent environments across all resources

Analysis User Experience and Declarative Languages Alexander Held, Jonas Rembser

User experience and declarative languages: interoperability!

- Identified **histograms** as key area for improving interoperability
- We should ensure that **Python bindings** are interoperable
- Demand for general **statistical model format** based on JSON
- Discussion about **data interoperability** at the level of individual **columns** (also in memory)

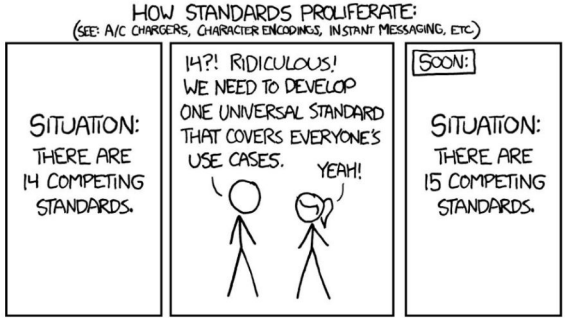

[Analysis User Experience and Declarative Languages](#) Alexander Held, Jonas Rembser

User experience and declarative languages: some ideas what to improve in future

Towards a better future - what would we like?

- Automatic (graph-based) **optimization** (à la RDF Vary()) to enable users to focus on physics instead of optimization
- **Grouping of columns into objects** to allow physics reasoning
- **Factoring out** the **analysis chores** to mini-frameworks / libraries
- More **documentation and learning material**

Analysis User Experience and Declarative Languages Alexander Held, Jonas Rembser

Reduced data formats:

Reduced Formats in ATLAS

AOD (300-600 kB/event) → PHYS (30-50 kB/event) → PHYSLITE (10-15 kB/event)

Common formats aiming to be used by 80% of the analysis (PHYS in run-3, PHYSLITE in run-4)

PHYSLITE will be frequently produced using latest recommnedations for calibrations etc.

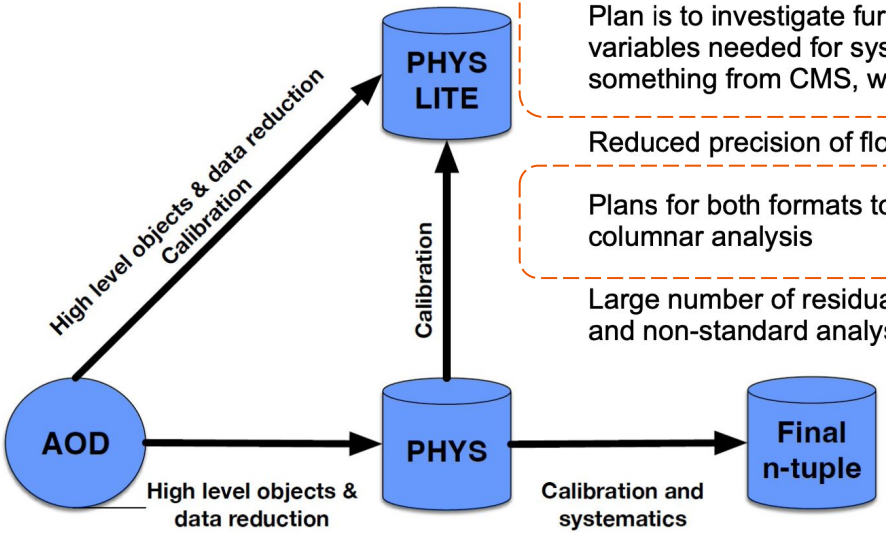

Plan is to investigate further reduction by dropping variables needed for systematics → ATLAS could learn something from CMS, where systematics are parametrizec

Reduced precision of floats is available but not yet applied

Plans for both formats to move to RNTuple to facilitate columnar analysis

Large number of residual formats needed for CP activities and non-standard analyses

Details in James talk

Analysis on reduced formats or specialist inputs, Allison Reinsvold Hall, Jana Schaarschmidt, Loukas Gouskos

Reduced data formats:

Reduced Formats in CMS

AOD (~500 kb/event) → **MiniAOD** (~50 kB/event) → **NanoAOD** (~2 kB/event)

MiniAOD format contains slimmed object collections, PFlow and tracks

NanoAOD is a flat ntuple, strictly controlled to keep size small, containing high-level objects

MiniAOD and NanoAOD serve 85% of all current analysis!

Floats stored with limited precision (based on detector resolution)

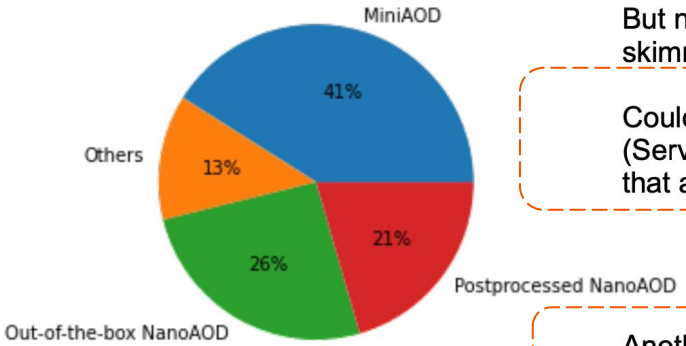

But nearly half of NanoAODs are customized (either skimmed or extended with extra info)

Could avoid the „full-copy“ overlaps by central service (ServiceX, Crab, Dask, regular Batch, …) that allows people to write extra columns („LegoAOD“)

Details in Lindsays talk

Another possibility: Object stores, eg. to avoid copying columns across processing tiers (→ see talk by Nick)

Analysis on reduced formats or specialist inputs, Allison Reinsvold Hall, Jana Schaarschmidt, Loukas Gouskos

Metadata & Systematics

DAWG - Metadata paper review

John De Stefano (BNL, IT)
Elizabeth Gallas (Oxford, ATLAS, databases)
Giacomo Govi (INFN Padova, CMS)
Thomas Kuhr (LMU Munich, Belle II)
Igor Mandrichenko (FNAL, IT)
Tibor Simko (CERN, IT, reusable analyses)

The Paper

Constraints on future analysis metadata systems in High Energy Physics

T. J. Khoo[5], A. Reinsvold Hall[10], N. Skidmore[16], S. Alderweireldt[15], J. Anders[13], C. Burr[3], W. Buttinger[9], P. David[11], L. Gouskos[3], L. Gray[4], S. Hageböck[3], A. Krasznahorkay[3], P. Laycock[1], A. Lister[14], Z. Marshall[6], A. B. Meyer[2], T. Novak[2], S. Rappoccio[12], M. Ritter[7], E. Rodrigues[8], J. Rumsevicius[3], L. Sexton-Kennedy[4], N. Smith[4], G. A. Stewart[3], and S. Wertz[11]

https://arxiv.org/abs/2203.00463

What next?

Request a detailed summary of the panel's opinions (short arXiv doc?)

Suggestions for Next Step

- Provide a detailed discussion of use cases
- Describe what problems should be addressed, not what the solution is
- Avoid mixing the discussion of the design of a system and how it is used
- Sharpen metadata scopes definitions
- Derive requirements from use cases, assign them to metadata scopes
- Discuss arguments for or against common solutions (across metadata scopes and experiments)

Analysis Ecosystems Workshop 23.05.2022 page 6

What will they do?

Start with use cases

across *all* areas

Covering end-to-end, i.e. conditions, analysis metadata, analysis preservation…

Define metadata scopes mapped to use cases

Look for commonality across use cases

Bookkeeping and systematics handling, Paul James Laycock, Teng Jian Khoo

Metadata & Systematics

Vision: The deliverables

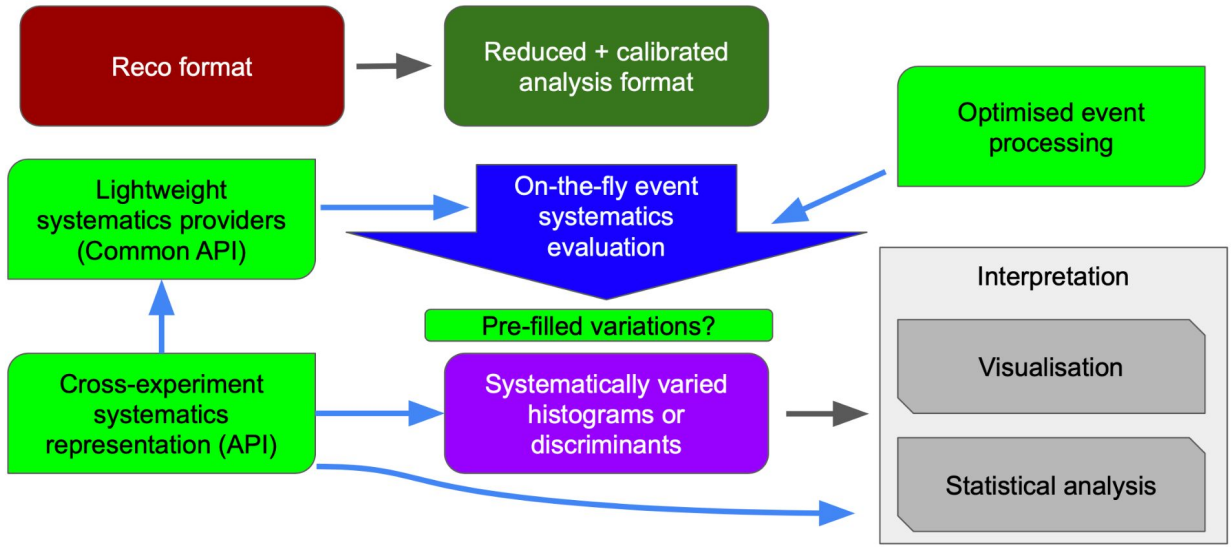

Bookkeeping and systematics handling, Paul James Laycock, Teng Jian Khoo

Realtime Trigger Analysis

LHC challenge: too much data for too few resources

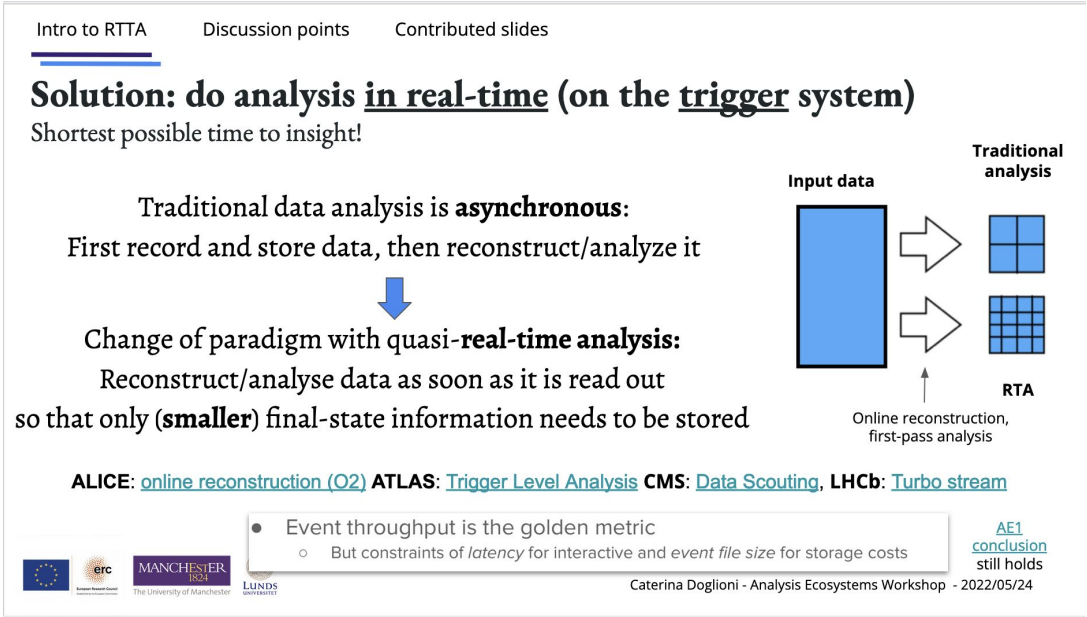

[Real-Time and Trigger Analysis Overview]

Realtime Trigger Analysis

RTTA → Building analysis ecosystem, on constrained resources

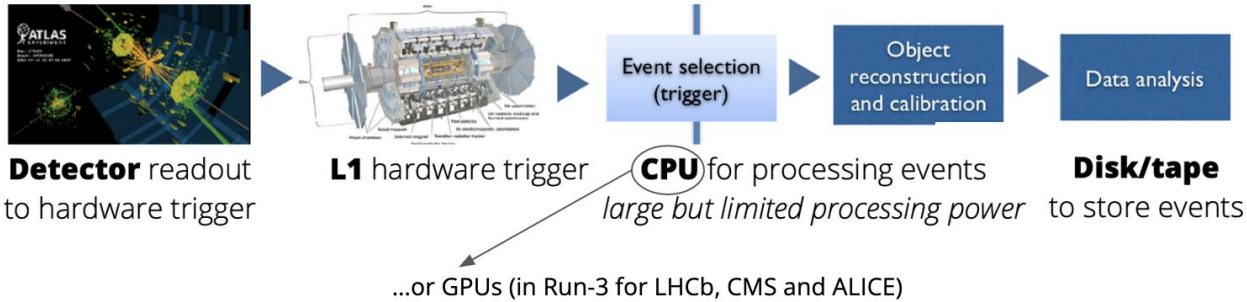

What kind of analysis ecosystem / analysis facilities do we want to build for HL-LHC and beyond?

Real-Time and Trigger Analysis Overview

BOF session: Let's Define the Steps of an Analysis Workflow

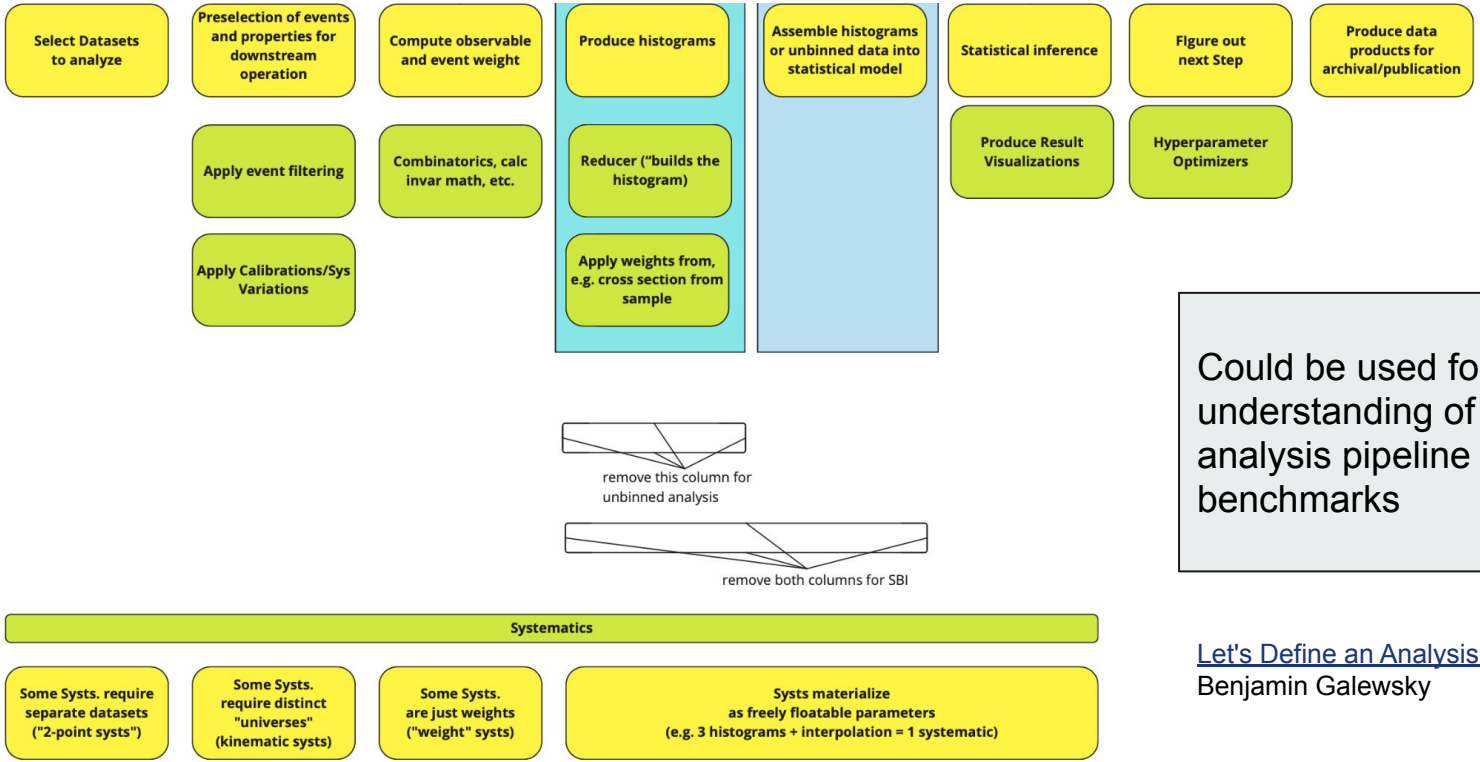

Could be used for better understanding of "typical" analysis pipeline or for synthetic benchmarks

[Let's Define an Analysis Workflow](#)
Benjamin Galewsky

Next steps: community whitepaper

- A lot of time for discussions and a community brainstorming about missing features or future plans!
- As a result, there is an ongoing work about the community report/whitepaper: AE2 Report
    - Main idea is **to document R&D** needed in next 5 years to improve the HEP ecosystem
    - Previous report is available here: https://zenodo.org/record/6599290