Reduced data formats & real-time analysis

mostly ATLAS/CMS/LHCb-centric due to time constraints

CATERINA DOGLIONI - LUND UNIVERSITY

Inputs, discussion and inspiration:
ATLAS TLA team (especially W. Kalderon, E. Tolley, A. Boveia),
C. Fitzpatrick, V. Gligorov,
real-time analysis speakers at JLab HSF workshop
(W. Kalderon, A. Pearce, D. Sperka, D. Rohr),
participants of Institut Pascal Real-Time Analysis Workshop

@CatDogLund
http://www.hep.lu.se/staff/doglioni/

Too much interesting data at the LHC?

- If **everything was recorded** at the LHC:
  - up to 30 million collisions/second (MHz)
  - 1-1.5 MB/data per collision
  - 30 MHz * 1 MB = 30 TB/s
  - 30 TB/s * 10e+6 s/year (day & night) ~ 0.05 ZB/year
- facebook
  - 600 TB/day ~ 200 PB/year [Facebook 2014]
- **“There’s always a bigger fish”**
  [C. Tully’s talk @ siRTDM18]
- But bigger fish also have bigger money…
  cost-effectiveness even more important for scientific instruments!

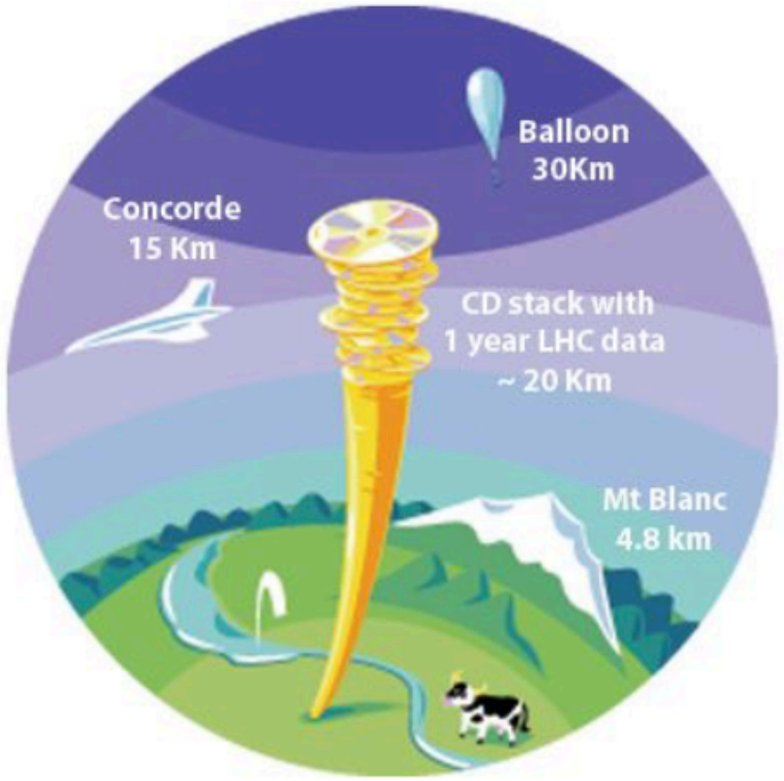

More timely analogies in
https://what-if.xkcd.com/31/

LHC experiments need to select & record ”interesting” events **in real-time** (milli/microseconds) but the **definition of interesting varies**

Online | Offline

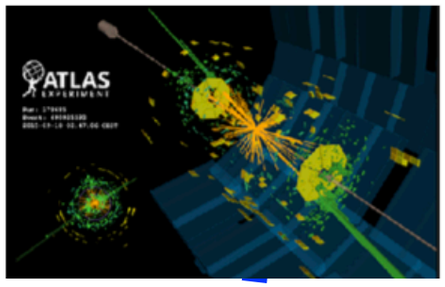

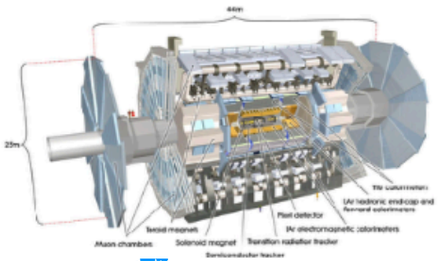

Event selection (trigger) → Object reconstruction and calibration → Data analysis

Data produced by the (HL-)LHC
(multiplied by large number)

Data I like

Data I would have used to like if I had been a data hipster

Data you like

What we can record after the trigger

Data ~nobody likes

This works for a number of LHC measurements (& searches...)

https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2019-024

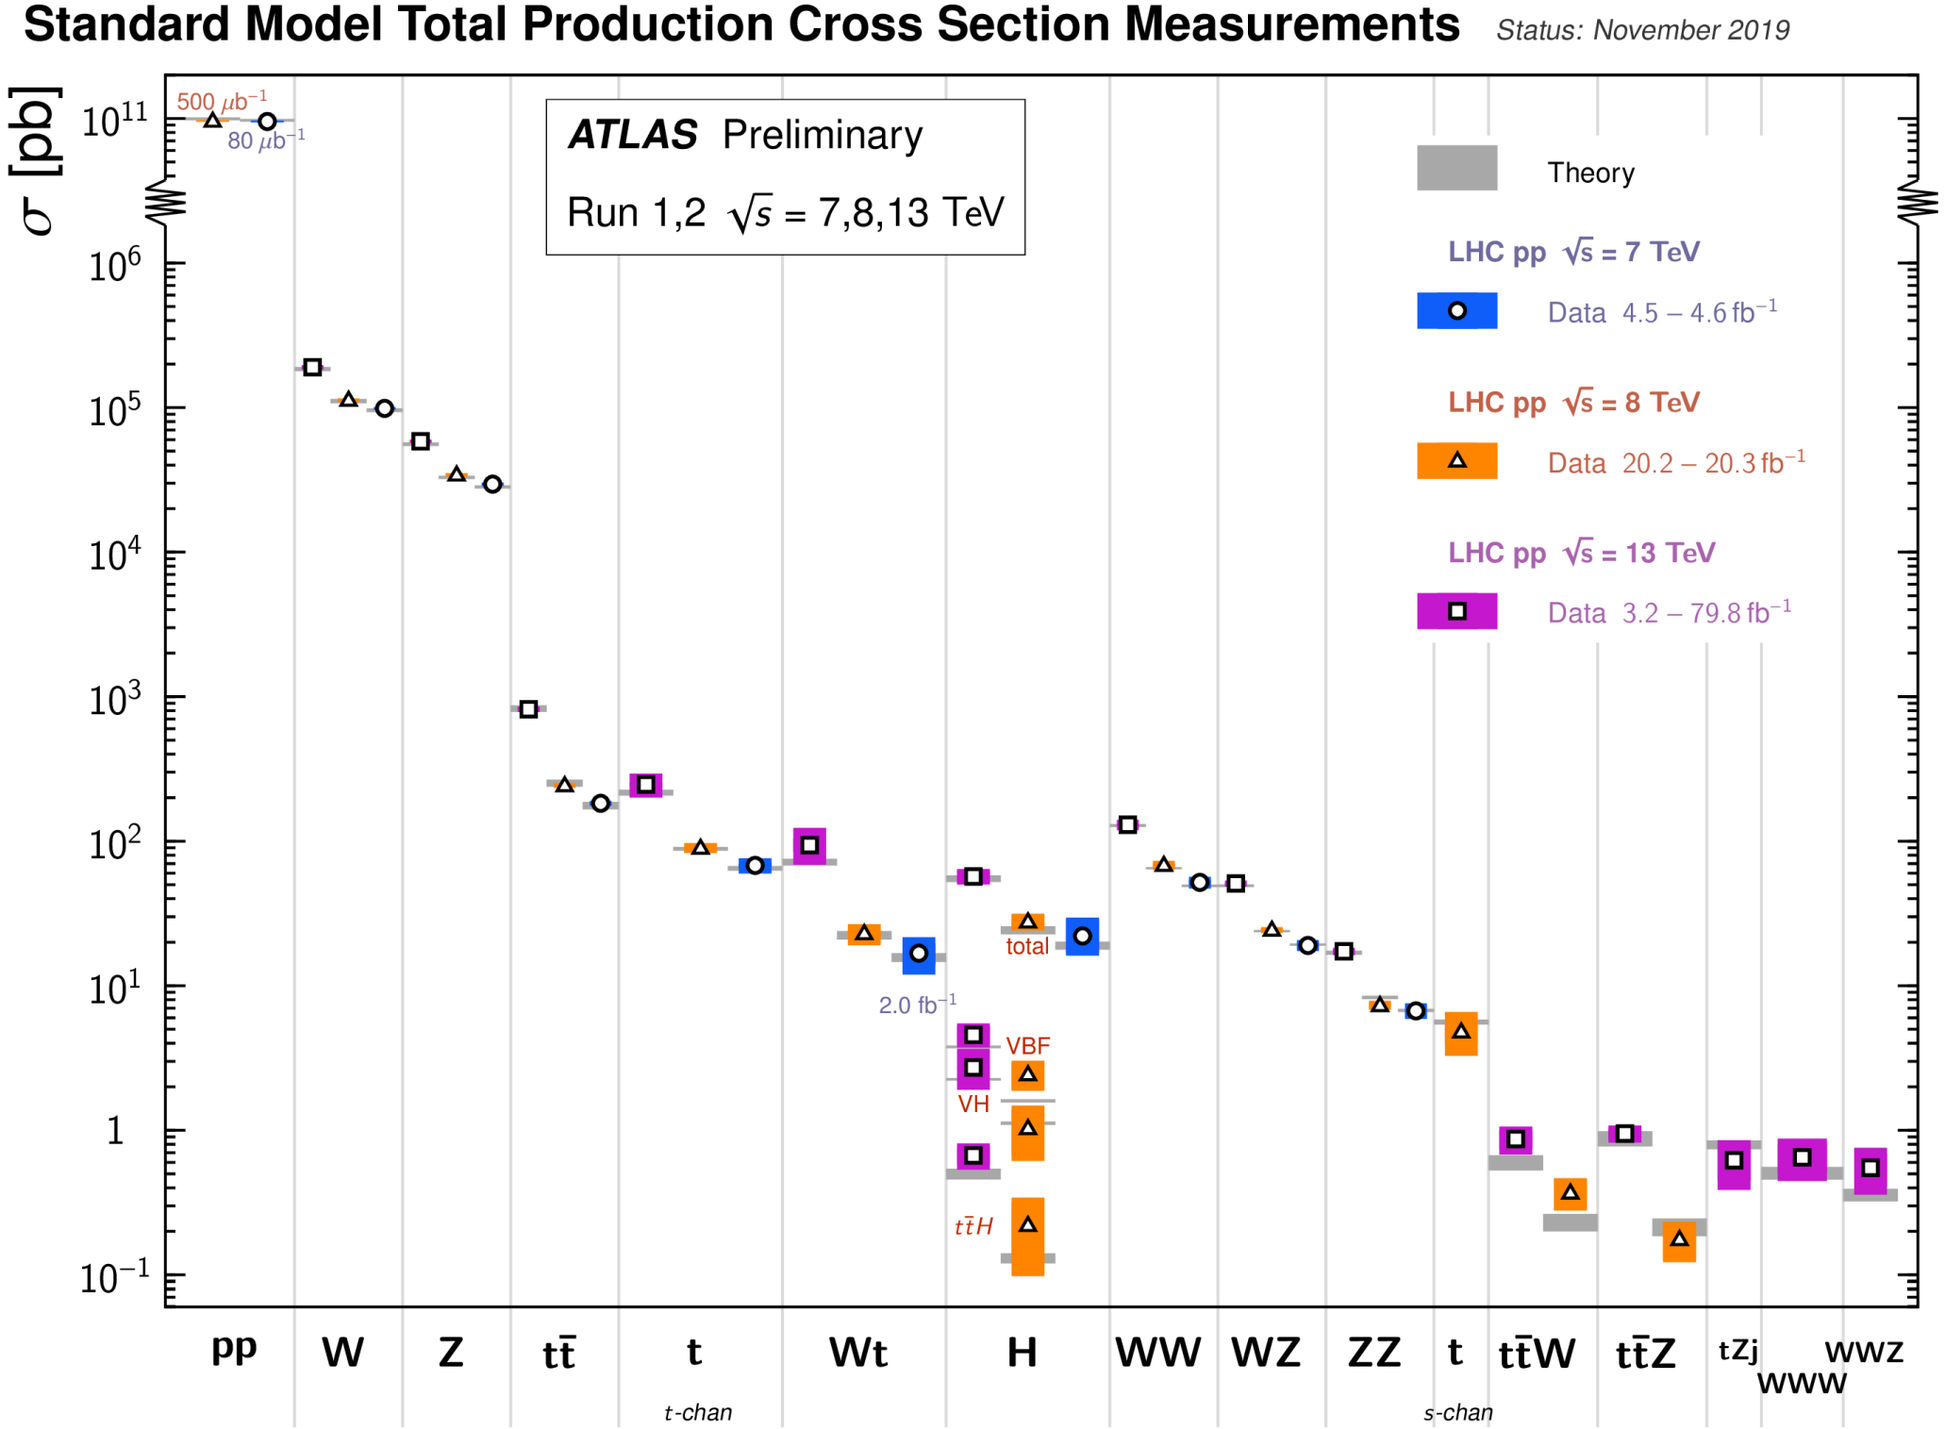

...but are we missing something?

What we can trigger on

ATLAS/CMS example: dijet resonances ca 2013

Dijet Resonances: Constraints on Coupling Values vs. Mass, 2013

A. Boveia, LBL workshop

Dobrescu, Yu Phys Rev D 88 035021 (**2013**)

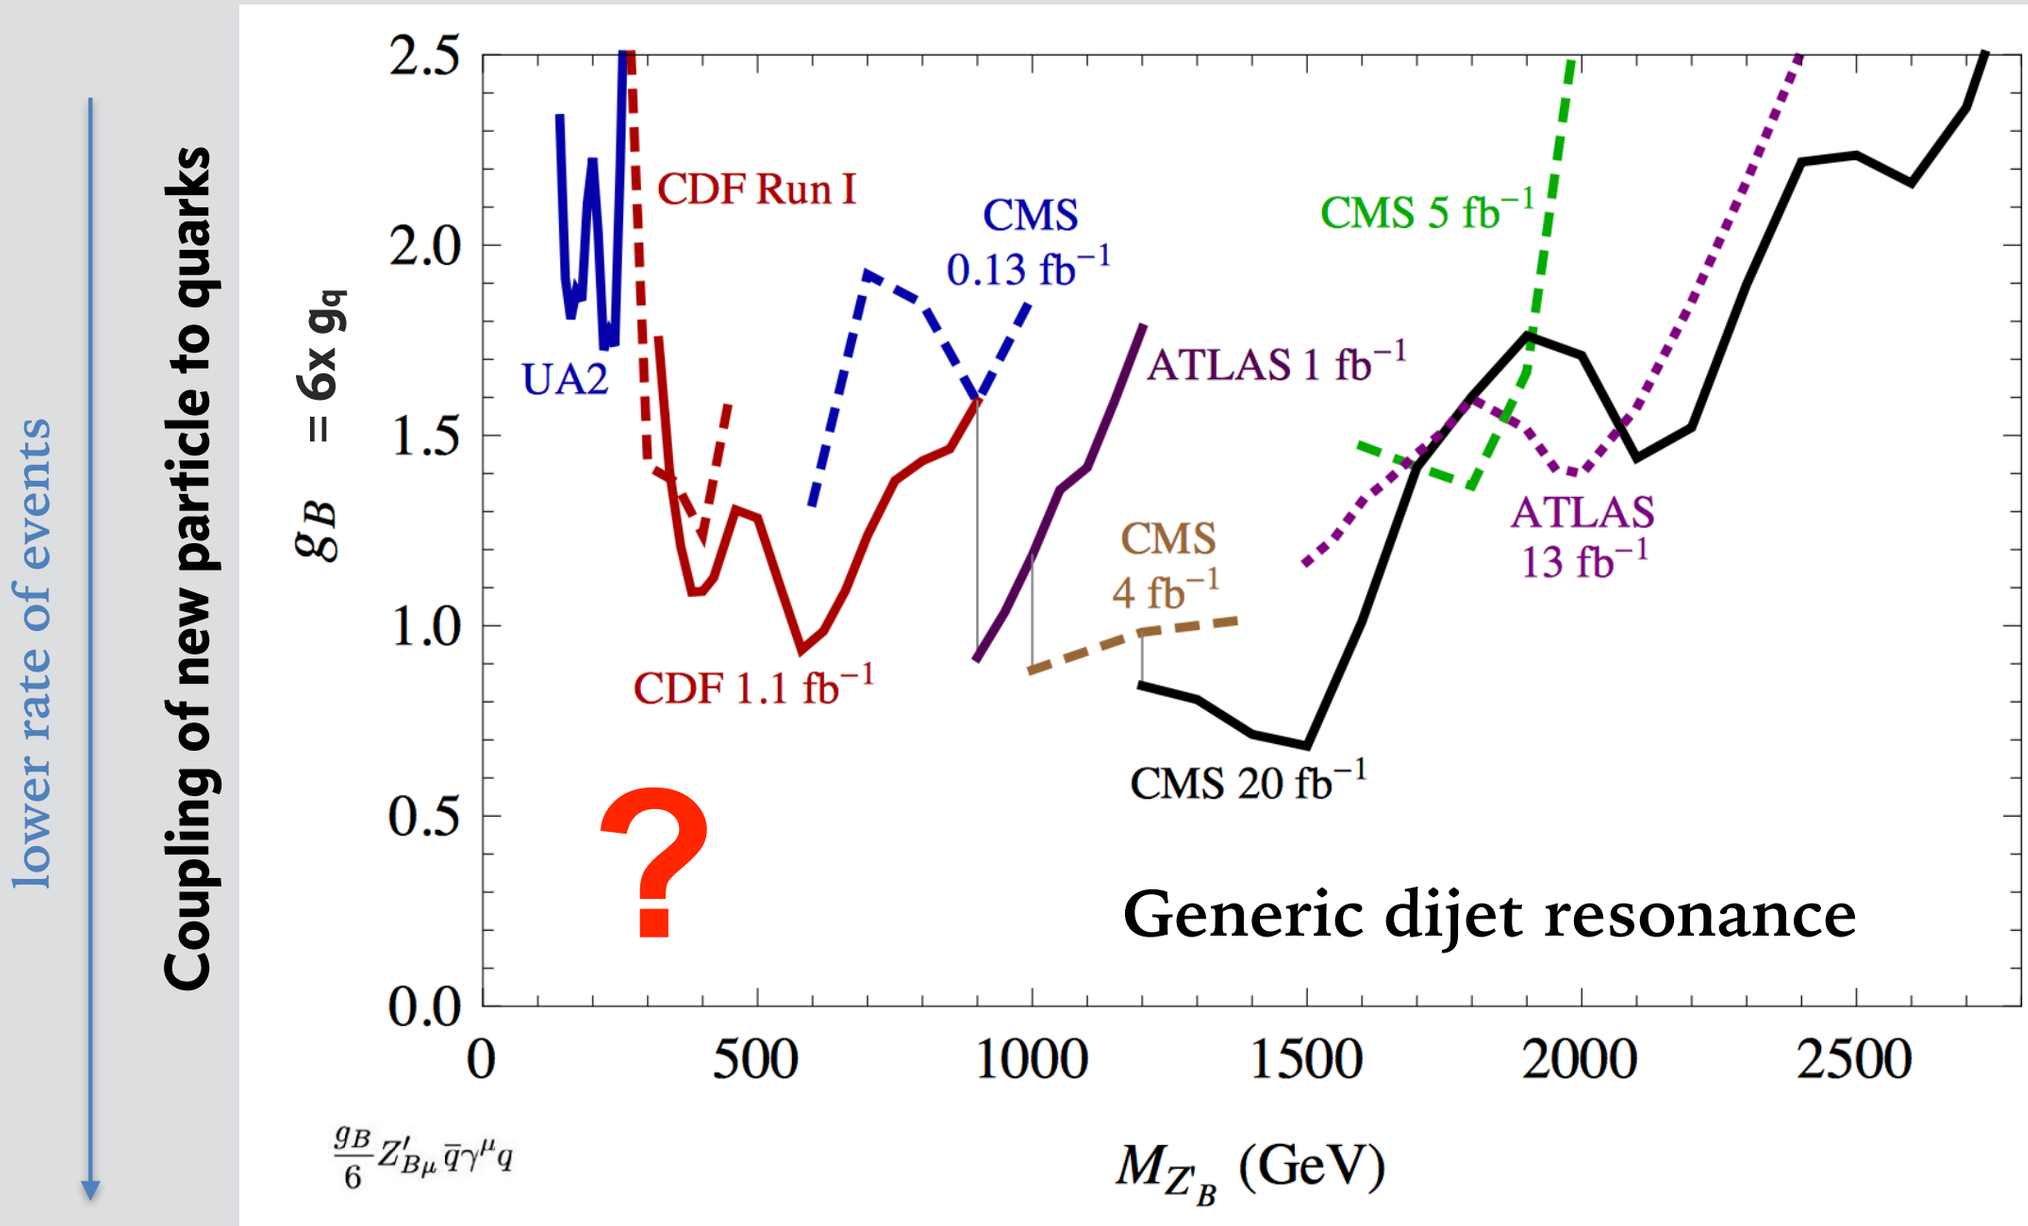

the LHC was not probing for di-jet resonances at the EW scale

43

Schematic

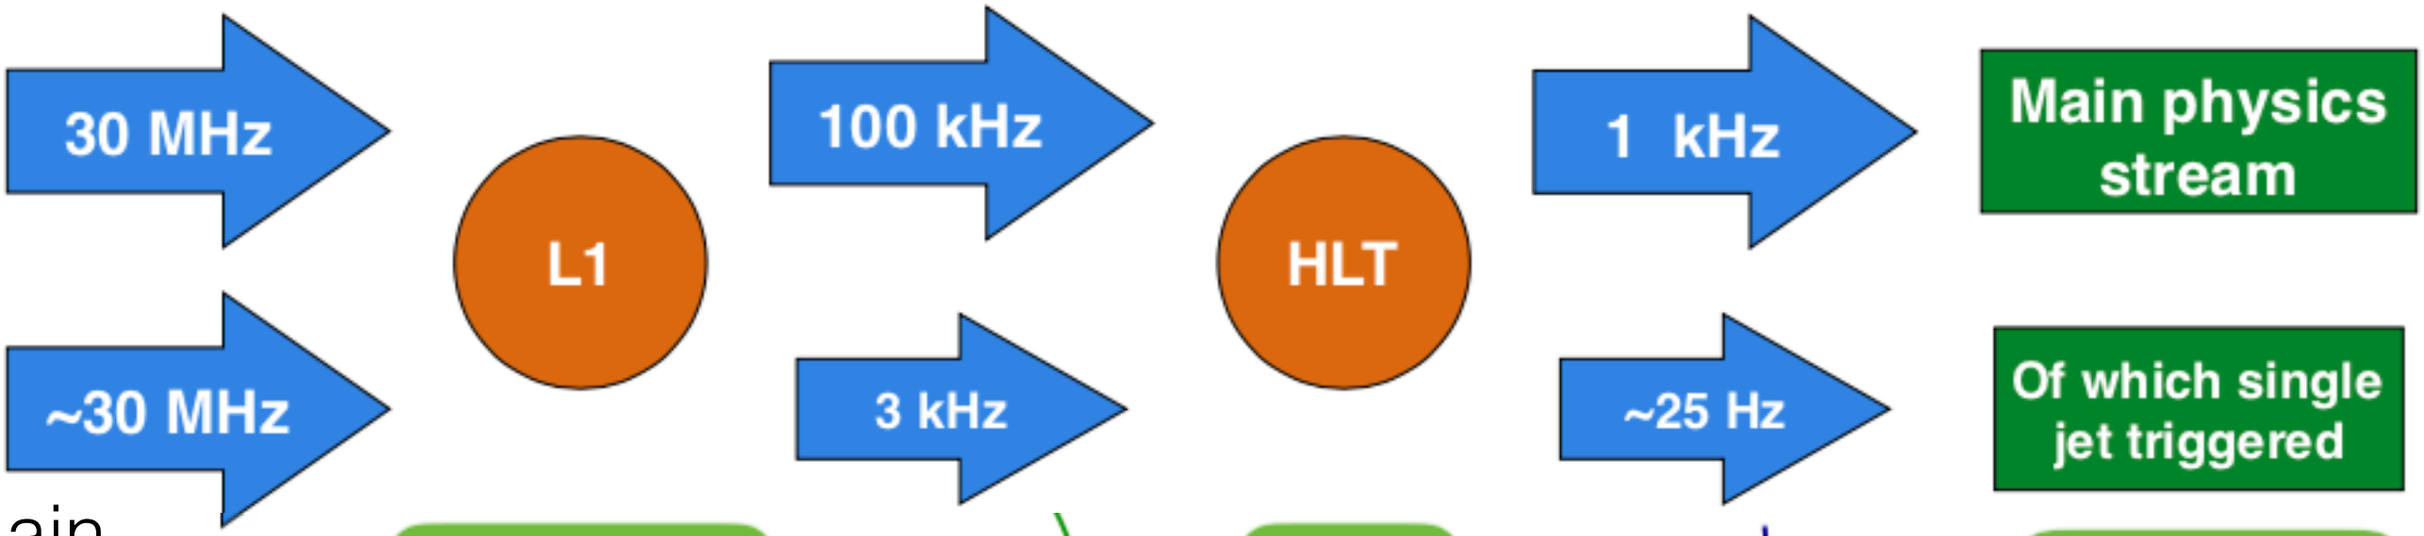

Main Limitations:

[W. Kalderon's talk at HSF / WLCG / OSG workshop 2019]

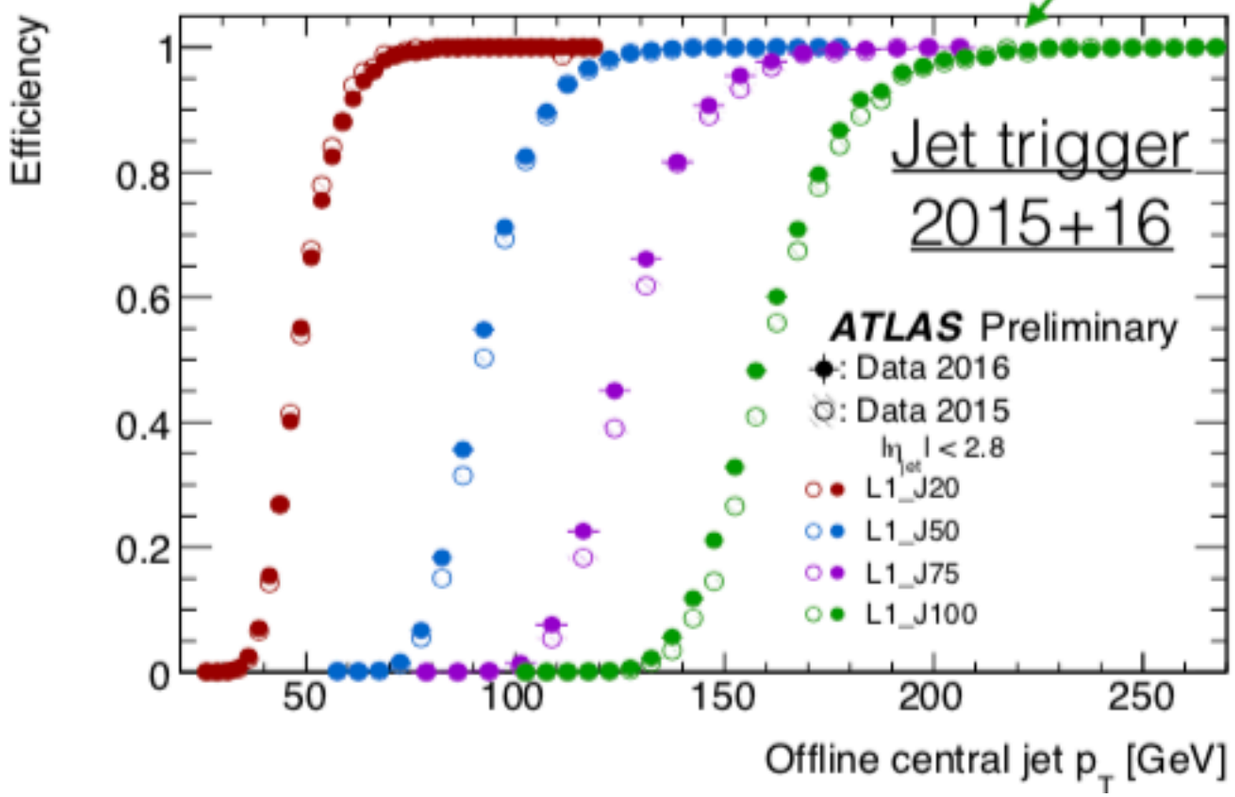

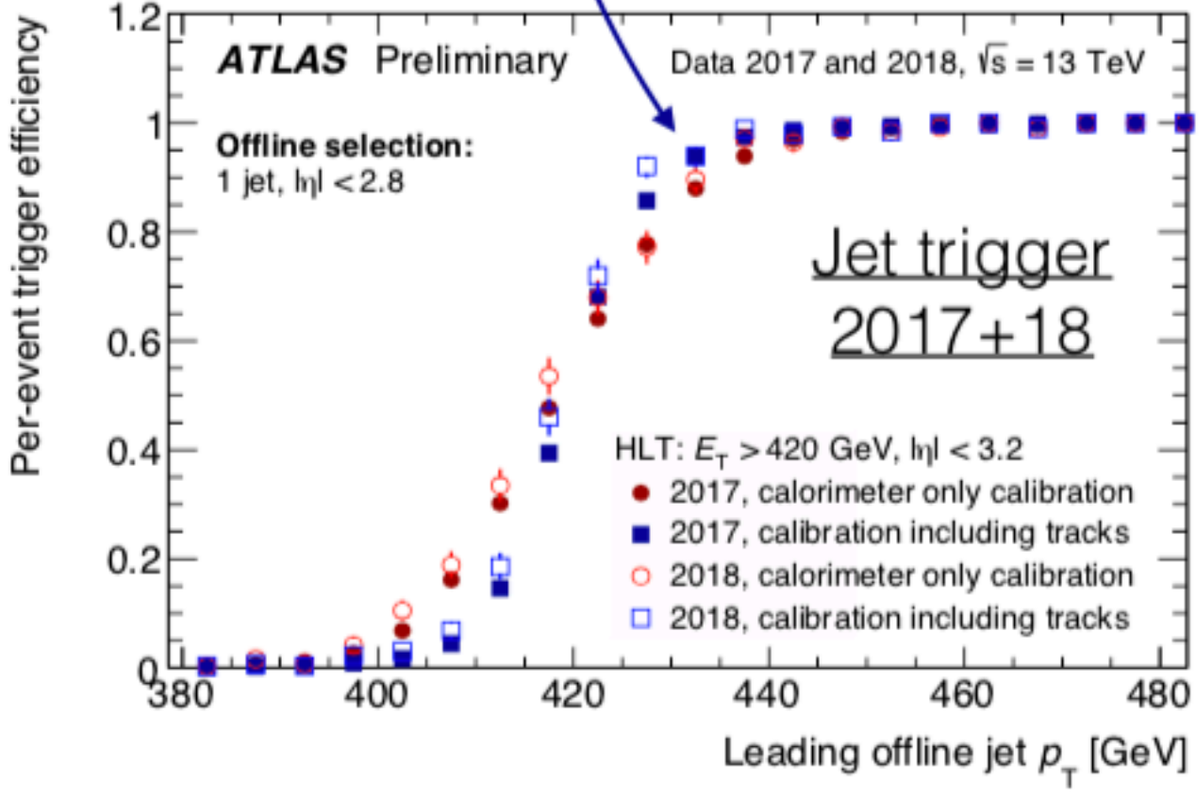

Schematic

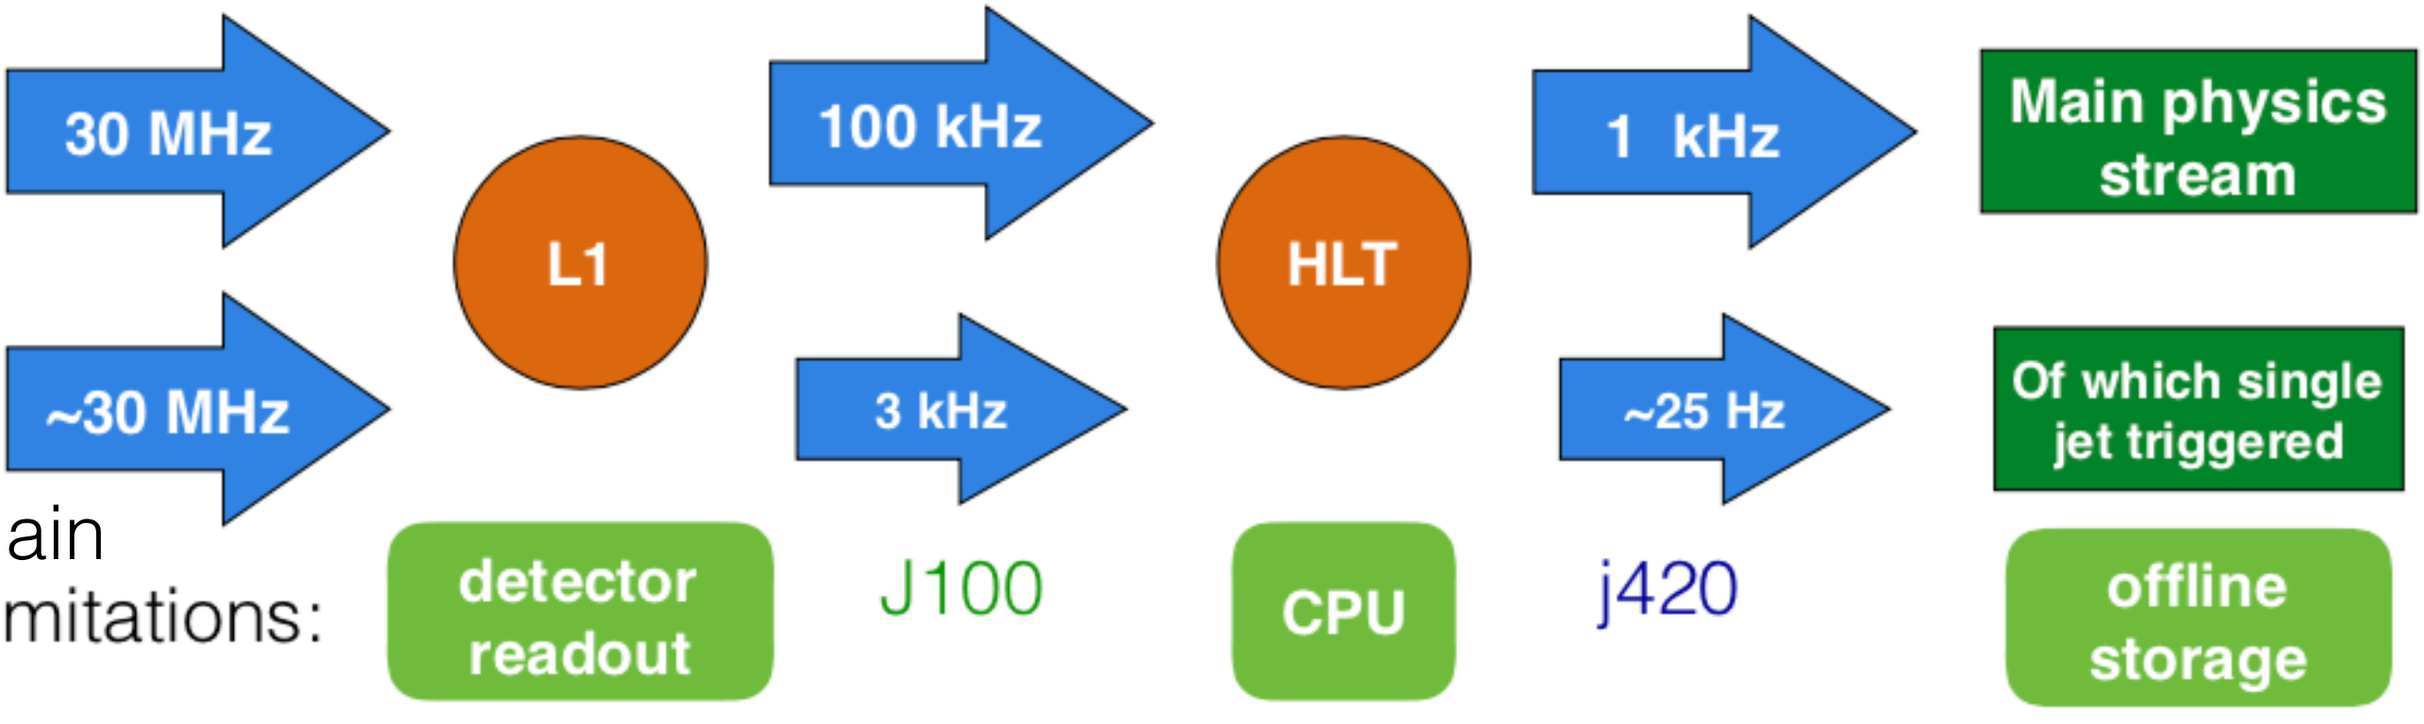

Main Limitations:

[W. Kalderon's talk at HSF / WLCG / OSG workshop 2019]

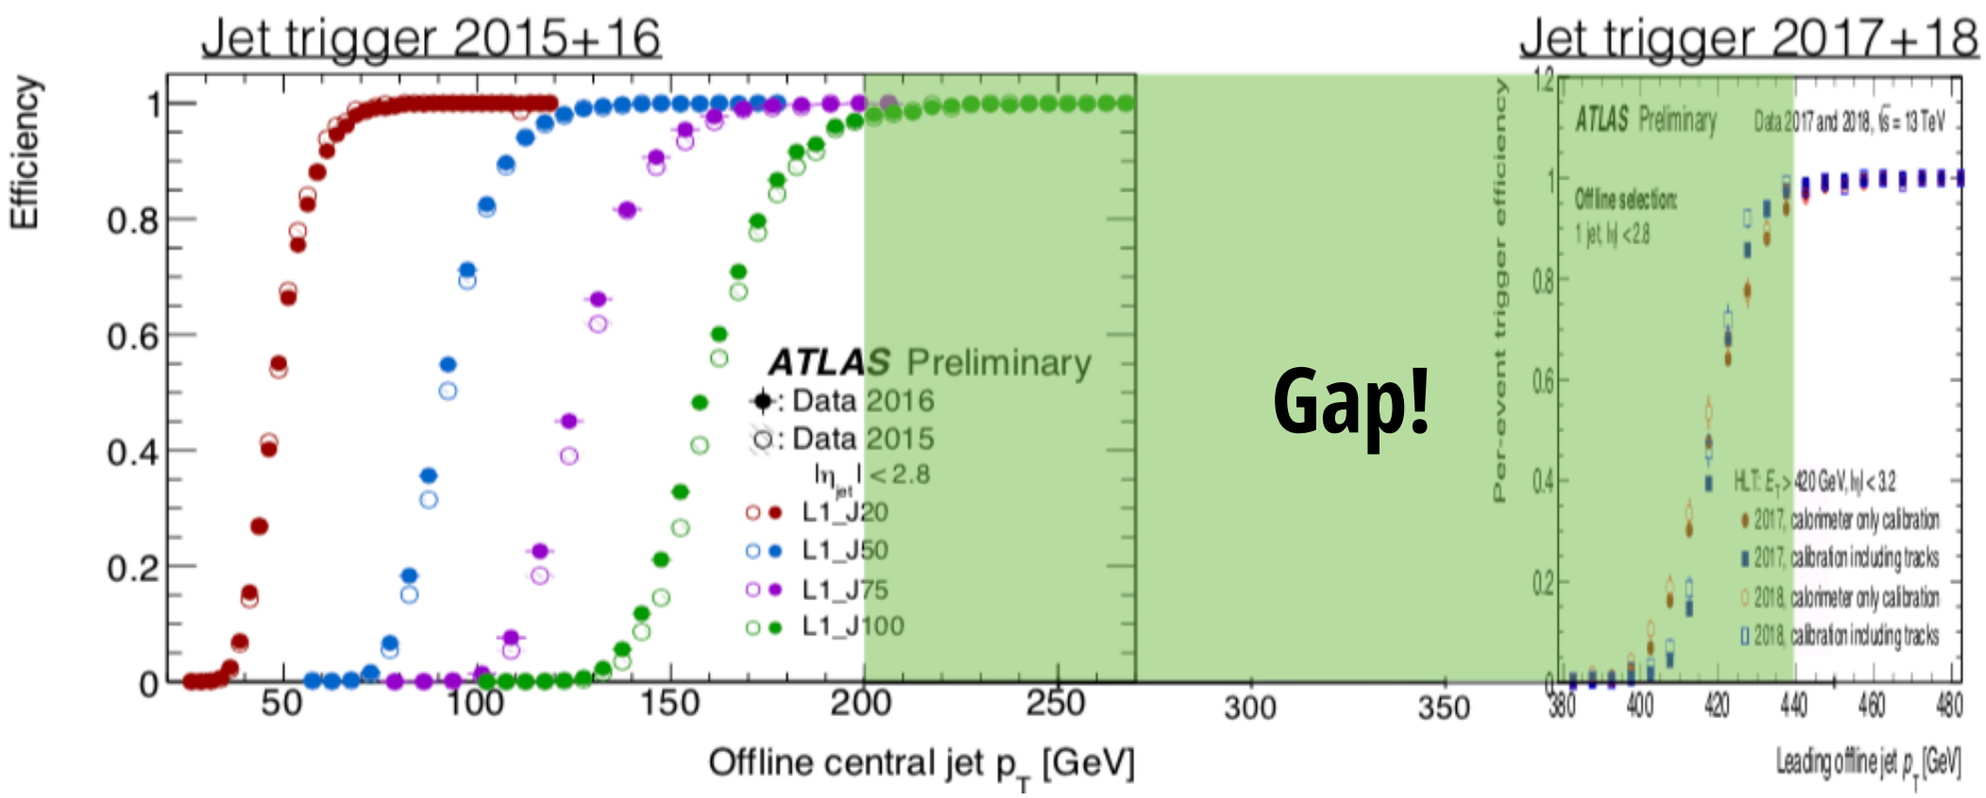

LHCb example: dark photon searches ca 2015

https://arxiv.org/abs/1801.04847

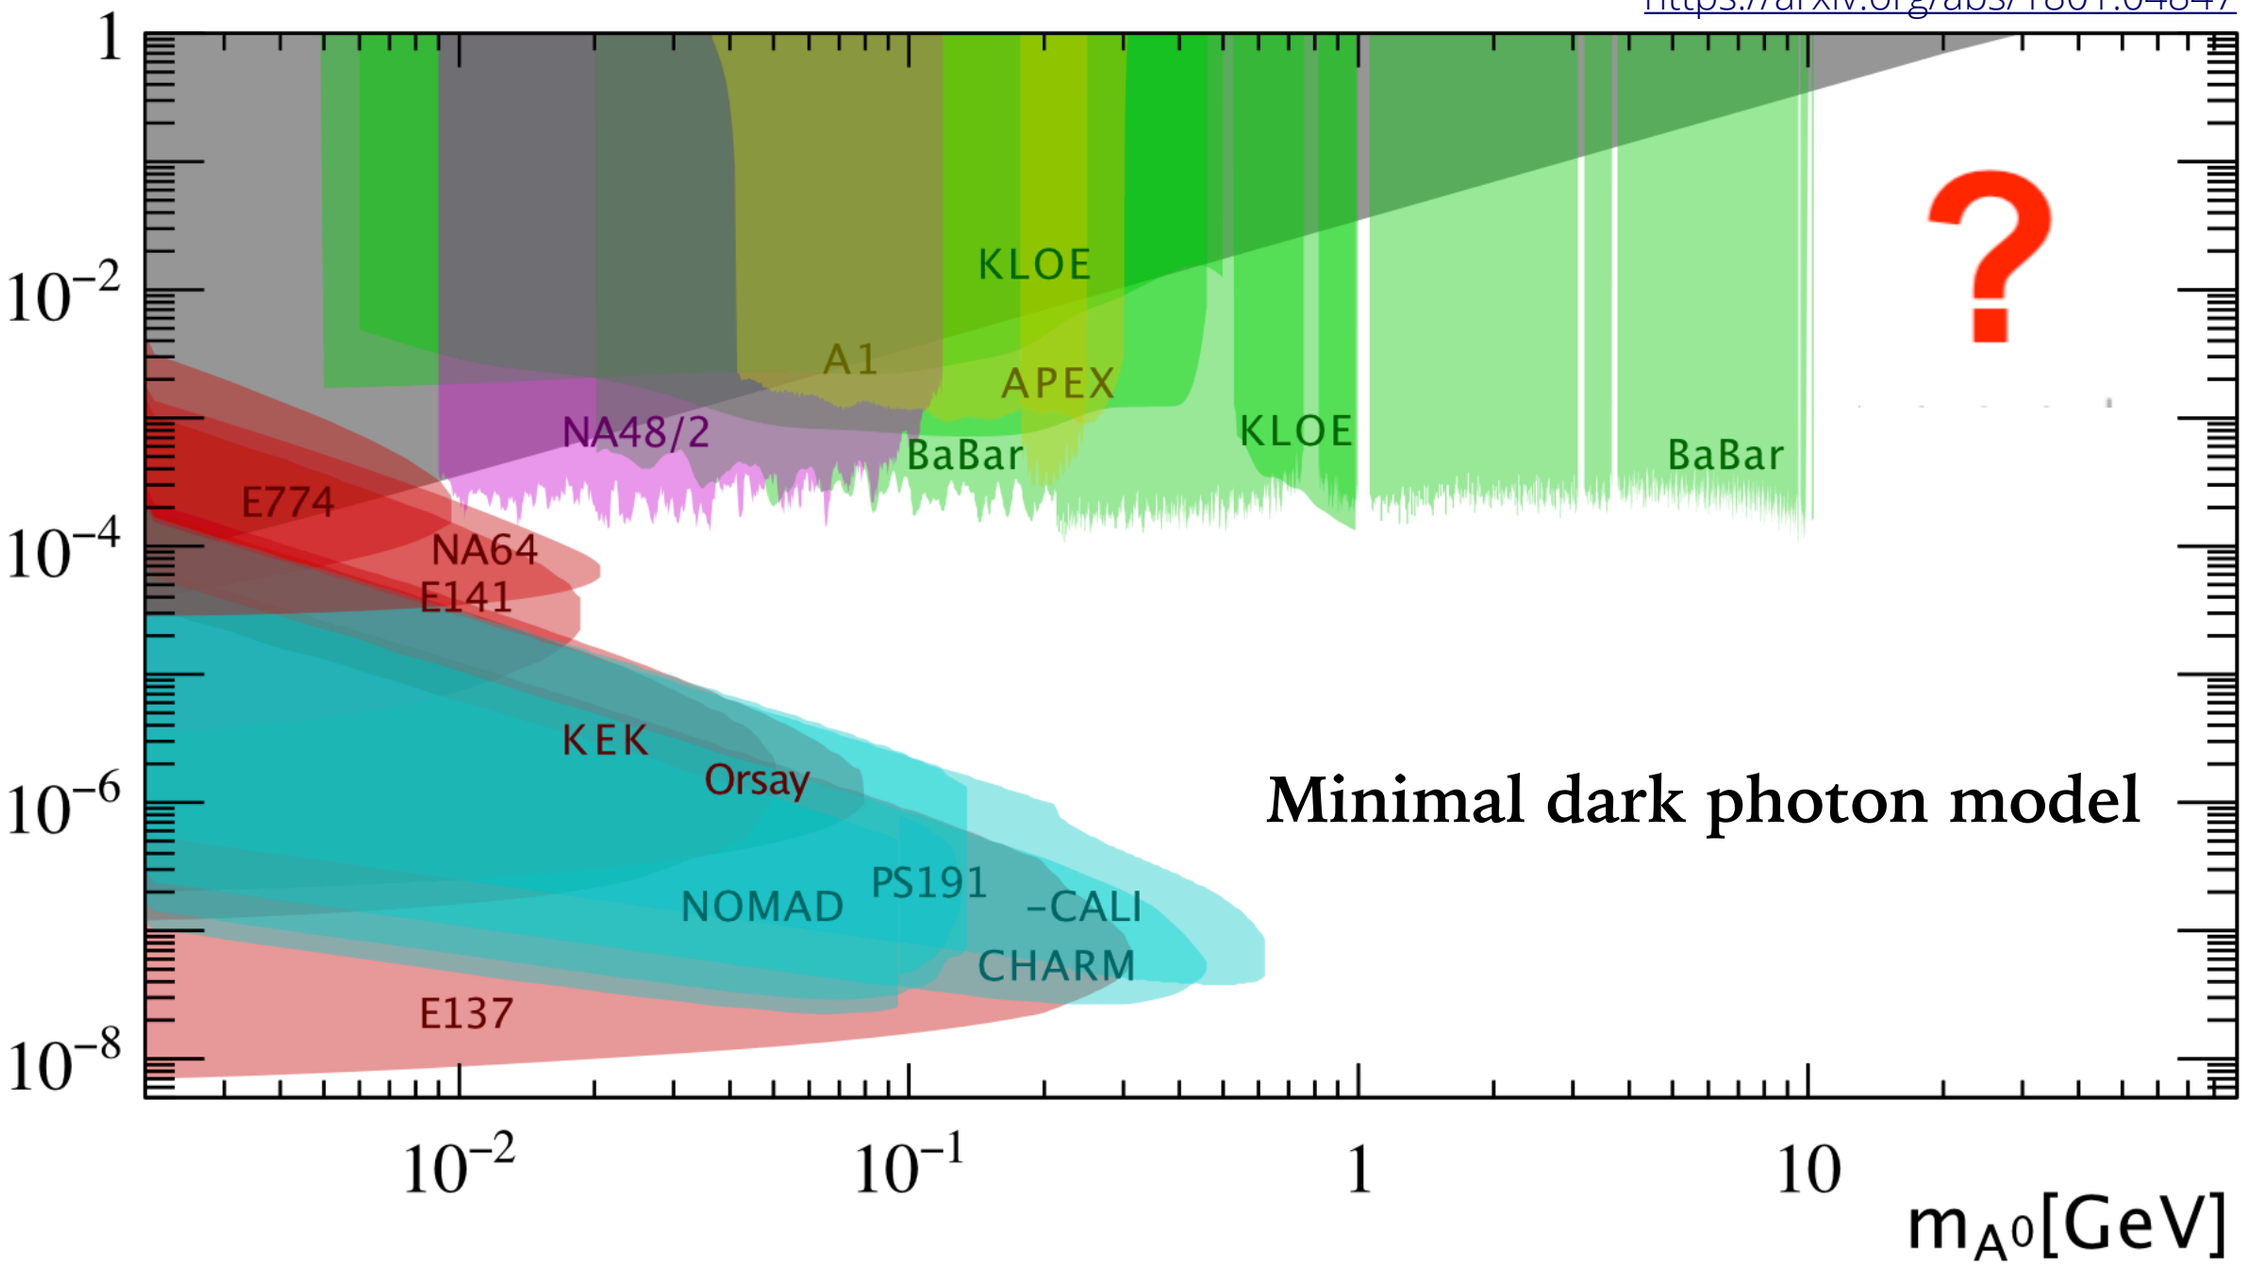

How to get reduced data formats: Real-Time Analysis (RTA)

A paradigm change

Asynchronous data analysis

First record data, then reconstruct/analyze it

Real-time data analysis

Reconstruct/analyze data as soon as it is read out
so that only (**smaller**) final-state objects
or histograms need to be stored

Not to scale

Using high-level trigger data for physics analysis

Continuous readout

Triggerless analysis
(histogramming @ L1…)

Trigger systems

note that

Real-time analysis

can have many meanings, even when only talking about HEP

Not to scale

Using high-level trigger data for physics analysis

Continuous readout

Triggerless analysis
(histogramming @ L1…)

Trigger systems

note that

Real-time analysis

can have many meanings, even when only talking about HEP

(Near-)real-time analysis of LHC data

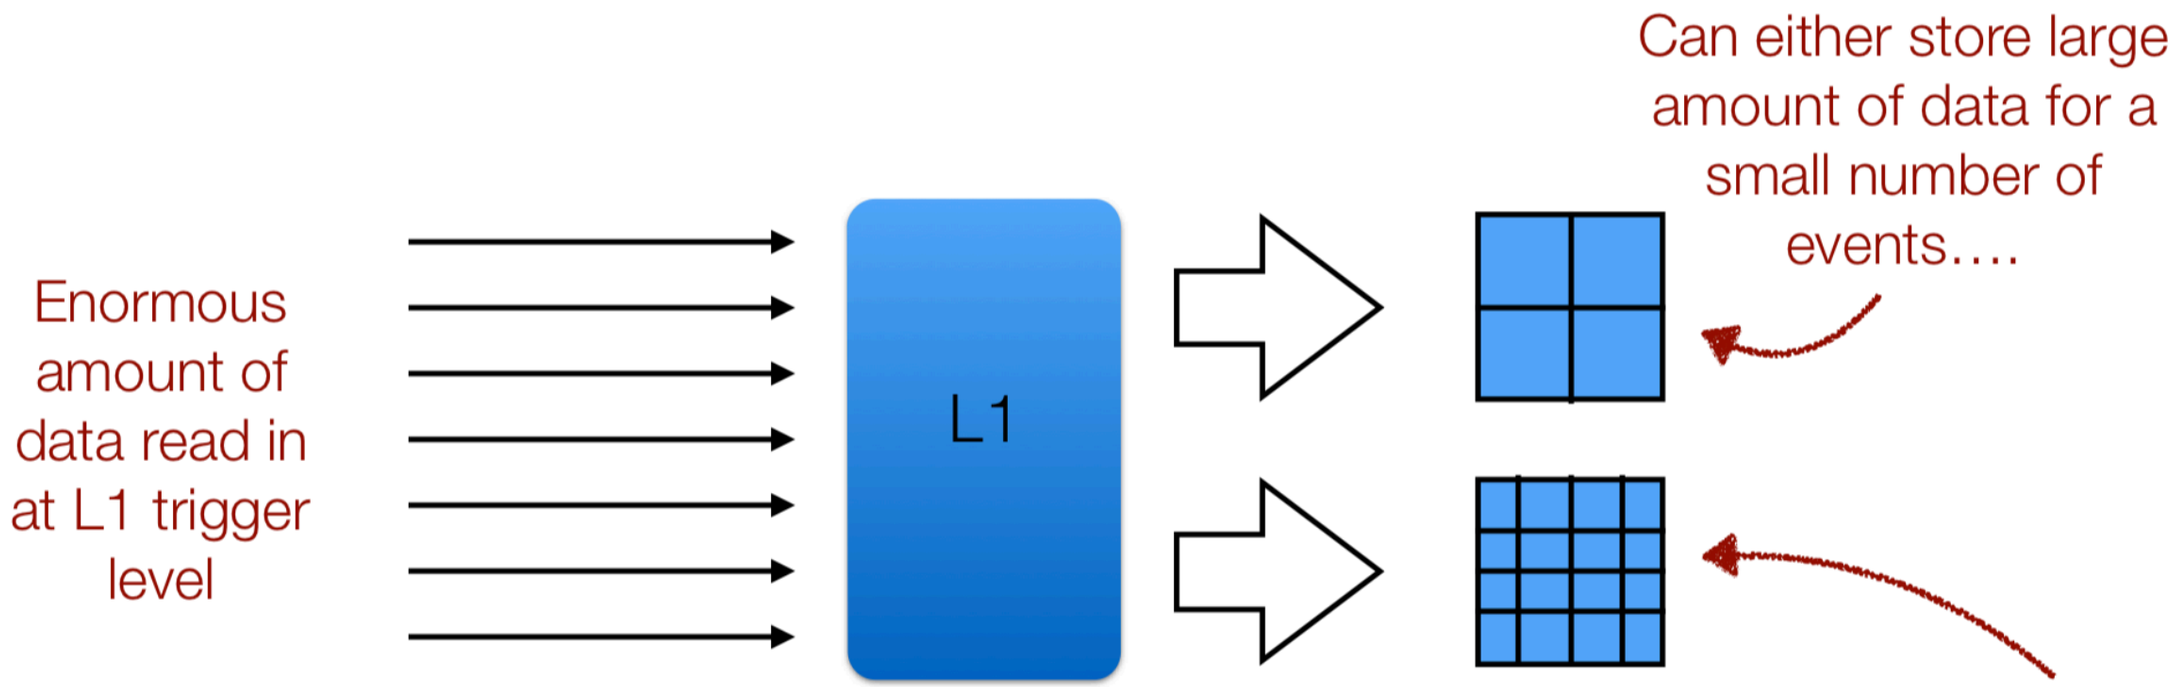

Perform as much "analysis" as possible @ HLT

- Reconstruction & calibration
- First preselection to skim "backgrounds"

Reduced data formats:

- Drop raw data
- Save only "interesting" parts of the detector
- A combination of the two

Turbo/Data Scouting/TLA path

Turbo stream (LHCb),
Data Scouting (CMS),
Trigger-level Analysis (ATLAS),

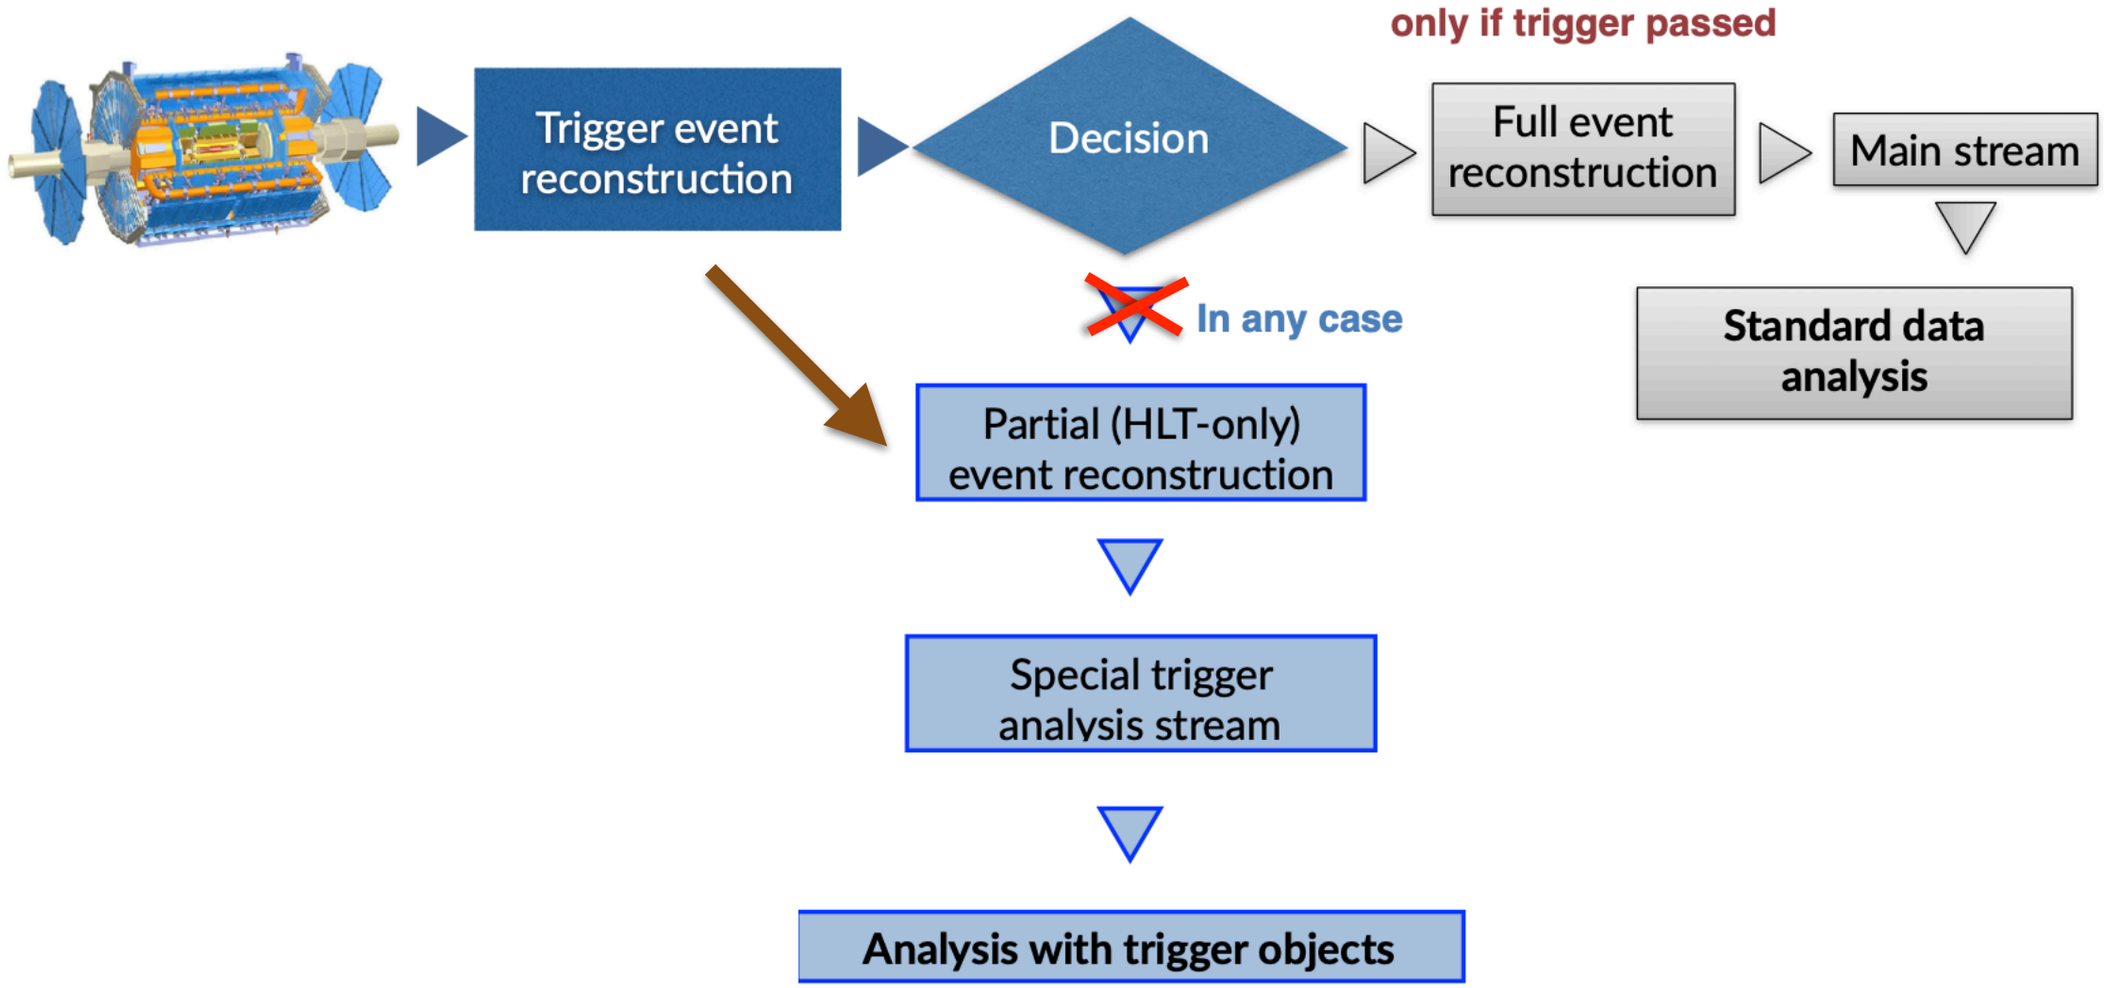

Turbo/Data Scouting/TLA path

Turbo stream (LHCb),
Data Scouting (CMS),
Trigger-level Analysis (ATLAS),

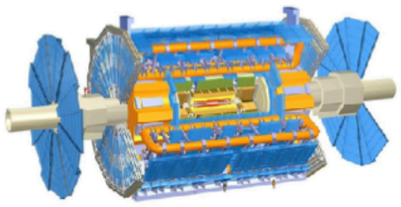

Trigger event reconstruction

Decision

only if trigger passed

Full event reconstruction

Main stream

Standard data analysis

In any case

Partial (HLT-only) event reconstruction

Special trigger analysis stream

Analysis with trigger objects

Customizable output data @ LHCb:

- keep trigger objects only (7 kB)
- keep trigger objects + "on-demand" raw and/or reco in selected regions (< 200 kB)
- keep everything (200 kB)

Data sizes @ ATLAS:

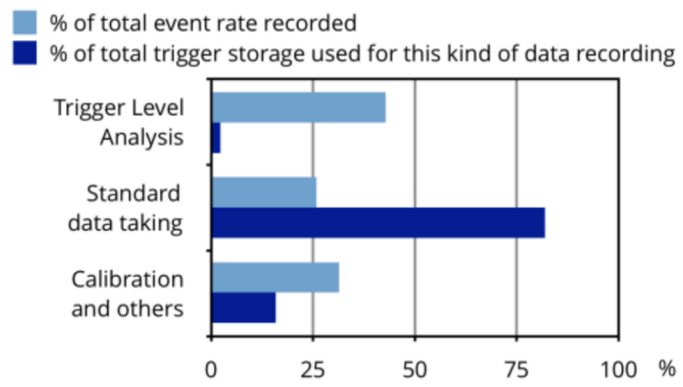

Information from ATLAS Trigger Operation plots, 2017

Objects and data sizes @ CMS:

Stream	Rate (Hz)	Event Size	Bandwidth (MB/s)
PhysicsMuons	420	0.86 MB	360
PhysicsHadronsTaus	345	0.87 MB	300
ScoutingCaloMuon	4580	8.9 KB	40
ScoutingPF	1380	14.8 KB	20

Selected CMS stream rate, event size, and bandwidth at the beginning of LHC Fill 7334 (23 Oct. 2018, L≈1.5×10^{34}cm^{-2}s^{-1})

JLab HSF workshop session on RTA: LHCb, ALICE, ATLAS, CMS

Overcoming storage (and CPU) bottlenecks

Save many more smaller events

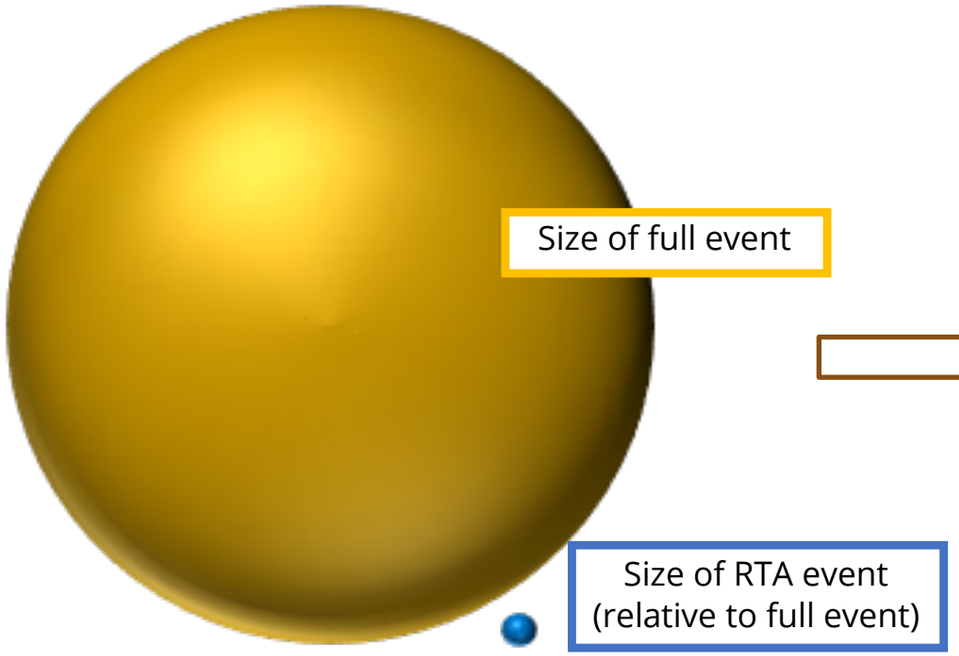

- Allows to record and store much higher event rates

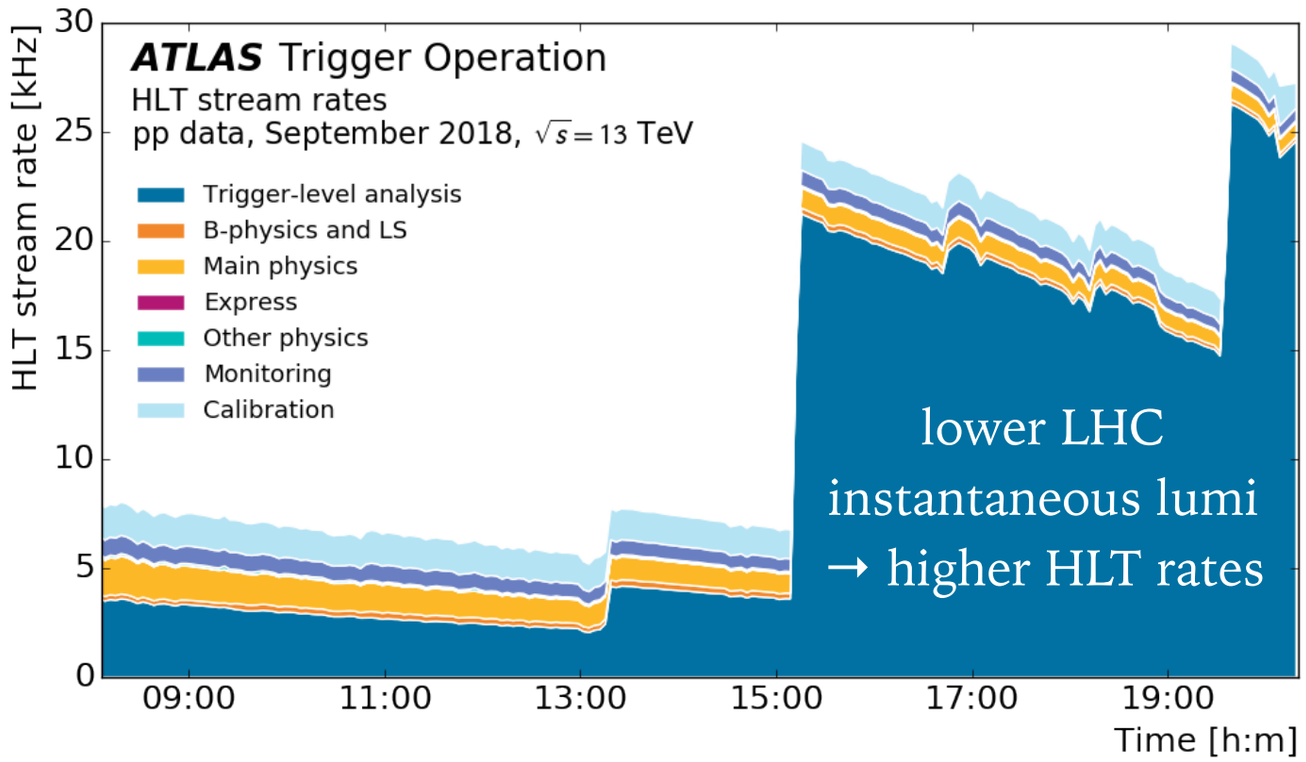

Use all the CPU, all the time

- LHC end-of-fill → unused HLT farm nodes
- Can lower the HLT thresholds to record more RTA(-like) events
- Note: this does not work with lumi-leveling (but GPD could think of getting closer to LHCb buffers for HL-LHC)

https://twiki.cern.ch/twiki/bin/view/AtlasPublic/TriggerOperationPublicResults

ATLAS/CMS results on dijet searches

Low mass "BumpHunt"
new particle → *dijet / multijet* searches
swamped by background
→ reach limited by trigger thresholds,
benefits a lot from RTA as one can
record more events

Phys. Rev. Lett. 121, 081801 (2018)

Dijet+anything

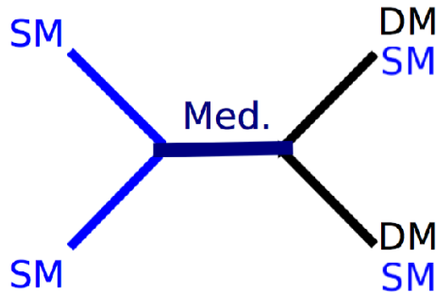

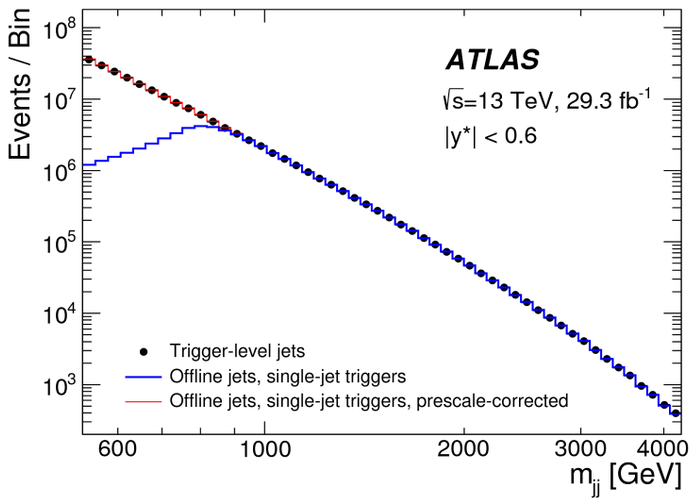

ATLAS/CMS results on dijet searches

Low mass "BumpHunt"
new particle → *dijet / multijet* searches
swamped by background
→ reach limited by trigger thresholds,
benefits a lot from RTA as one can
record more events

Phys. Rev. Lett. 121, 081801 (2018)

Dijet+anything

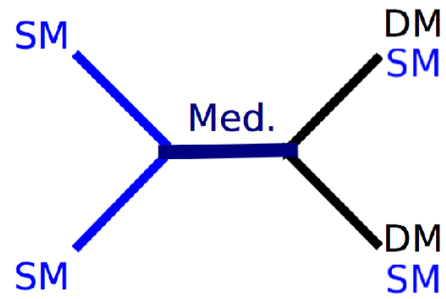

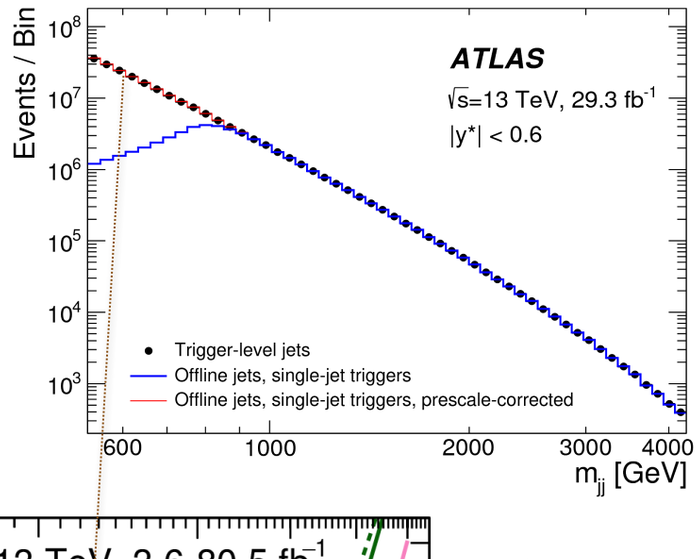

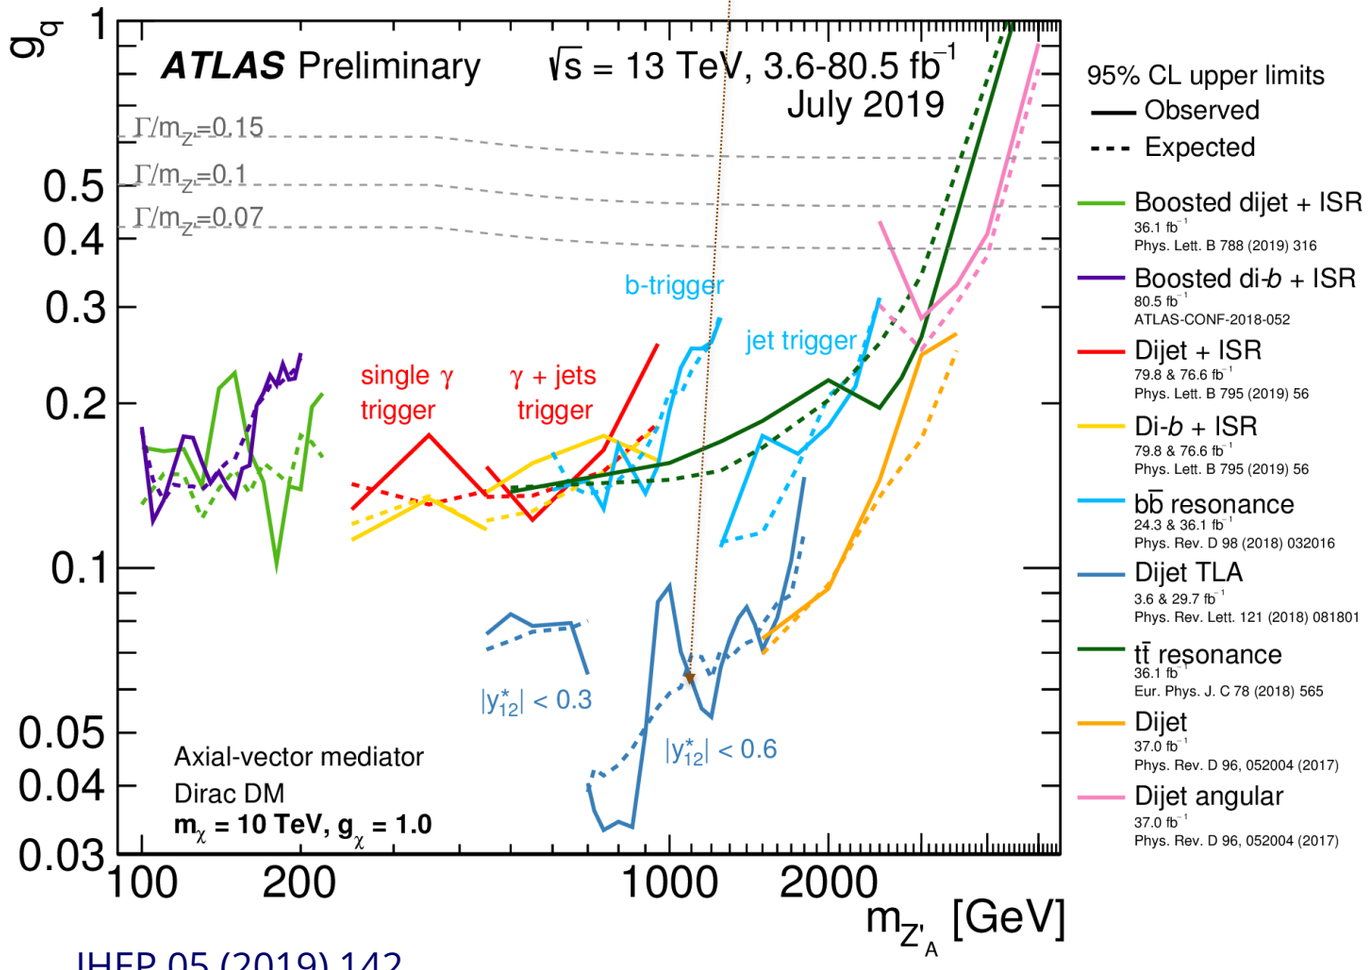

JHEP 05 (2019) 142

ATLAS/CMS results on dijet searches

Low mass "BumpHunt"
new particle → *dijet / multijet* searches
swamped by background
→ reach limited by trigger thresholds,
benefits a lot from RTA as one can
record more events

Dijet+1 other jet

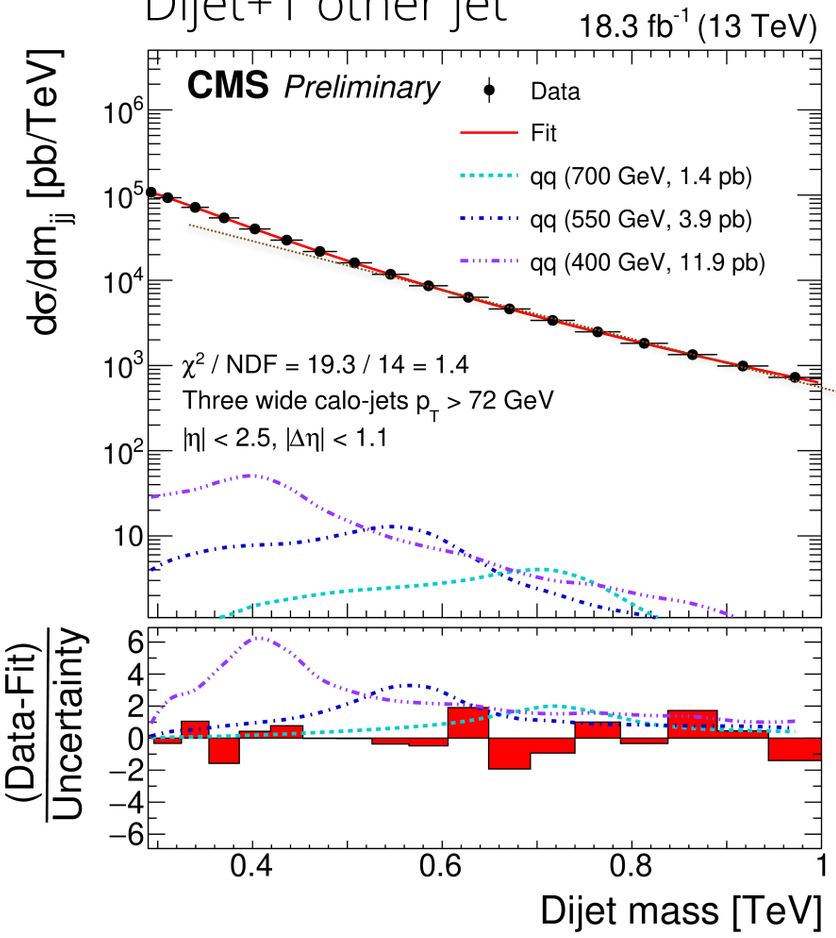

arXiv:1911.03761

Phys. Rev. Lett. 121, 081801 (2018)

Dijet+anything

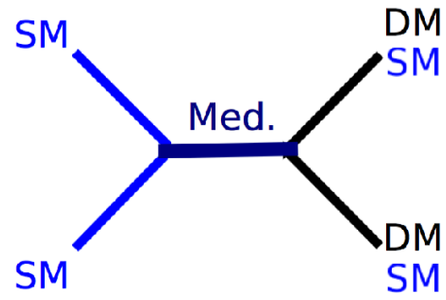

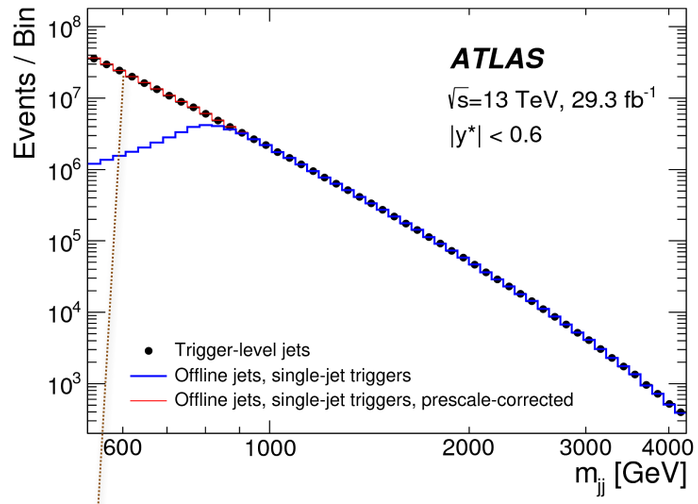

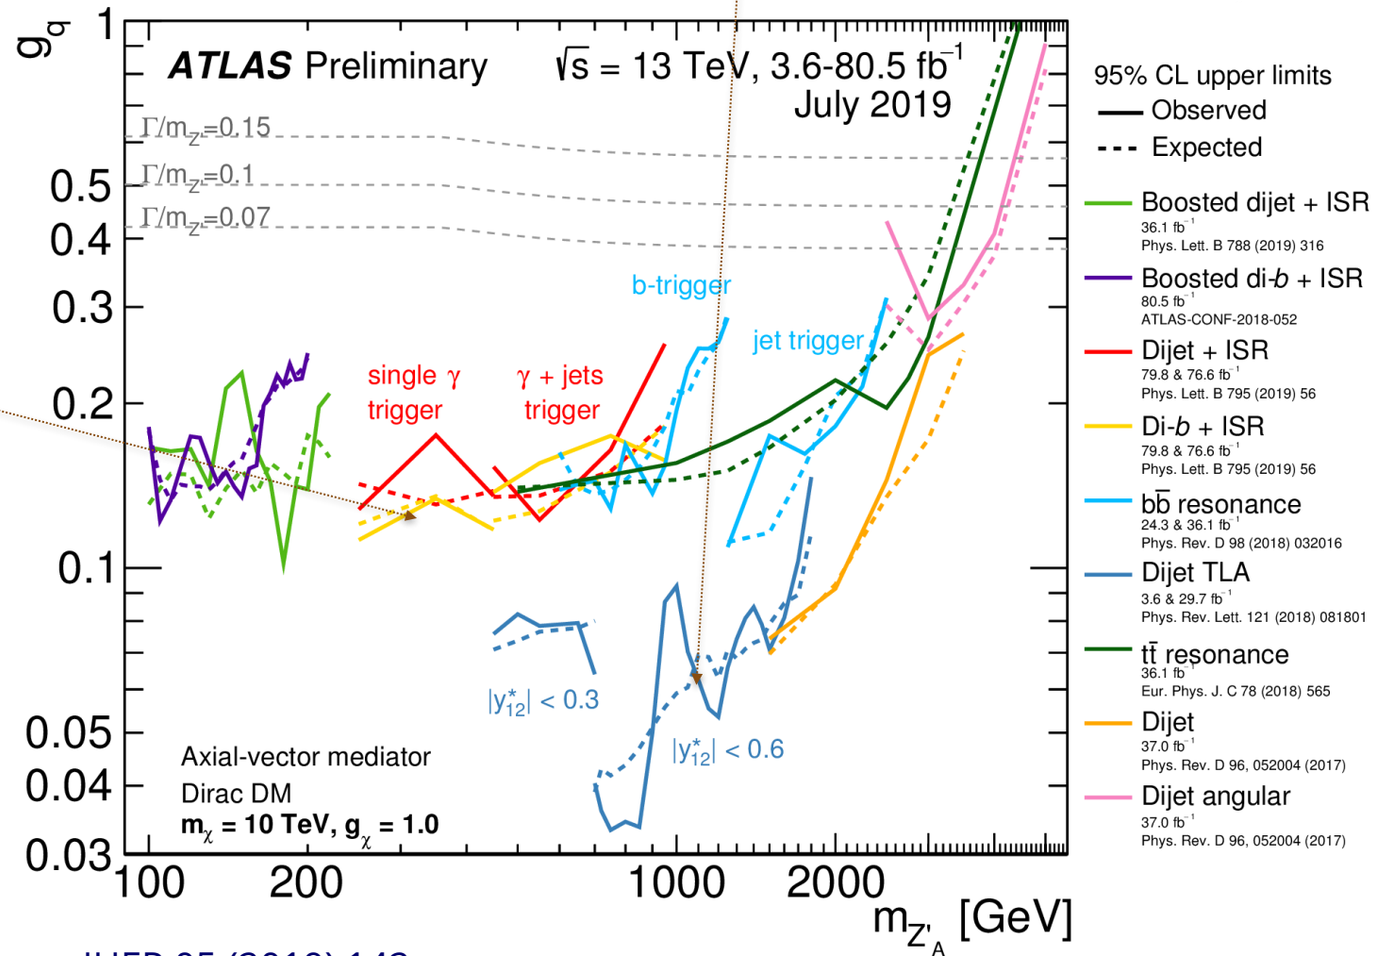

JHEP 05 (2019) 142

Visible decays of dark photon: LHCb and CMS

Dark photon → *dimuon* searches have a more complex background than dijet, but face the same problem at masses below the Z
→ large benefits from RTA

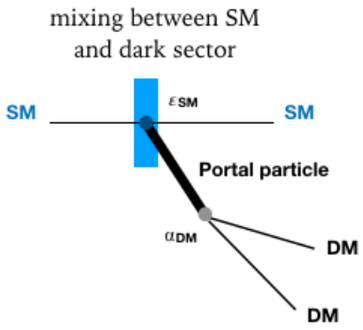

Phys. Rev. Lett. 120, 061801 (2018)

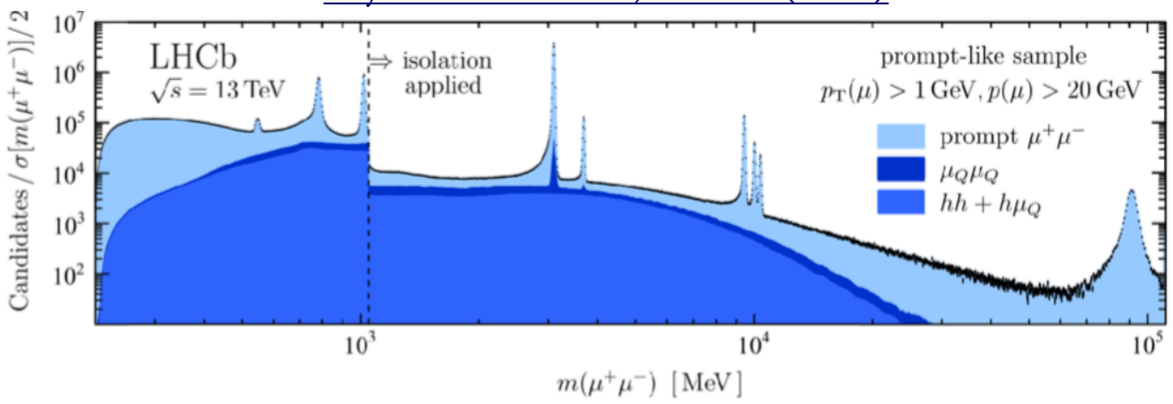

lower rate of events

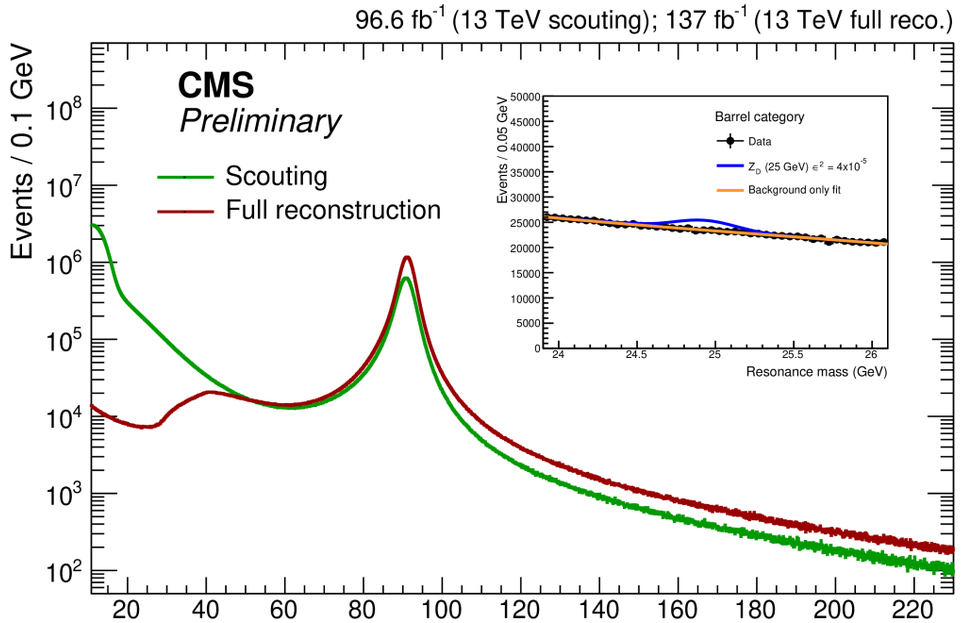

CMS-PAS-EXO-19-018

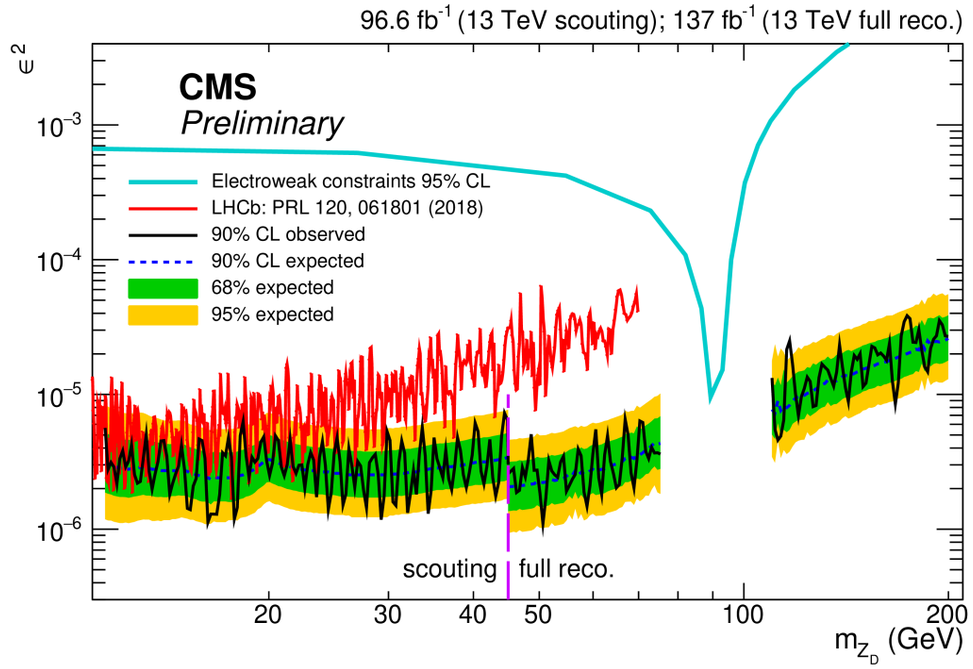

Run-3 prospects in arXiv:1509.06765

More with less: Selective persistency/Partial Event Building

Real-time analysis is necessary for searches
that would otherwise have been impossible due to trigger constraints

Traditional offline analysis still required for a number of searches/final states
where all raw information is needed (but we could do better)

Partial Event Building / Selective Persistency as a middle way:
save raw data && trigger objects only in the regions of interest

H. Russell, EPS-HEP 2019,

Customizable output data @ LHCb:
- keep trigger objects only (7 kB)
- **keep trigger objects + ”on-demand” raw and/or reco in selected regions (< 200 kB)**
- keep everything (200 kB)

HSF Trigger & Reco / Institut Pascal discussion, July 2016:
https://indico.cern.ch/event/835074/

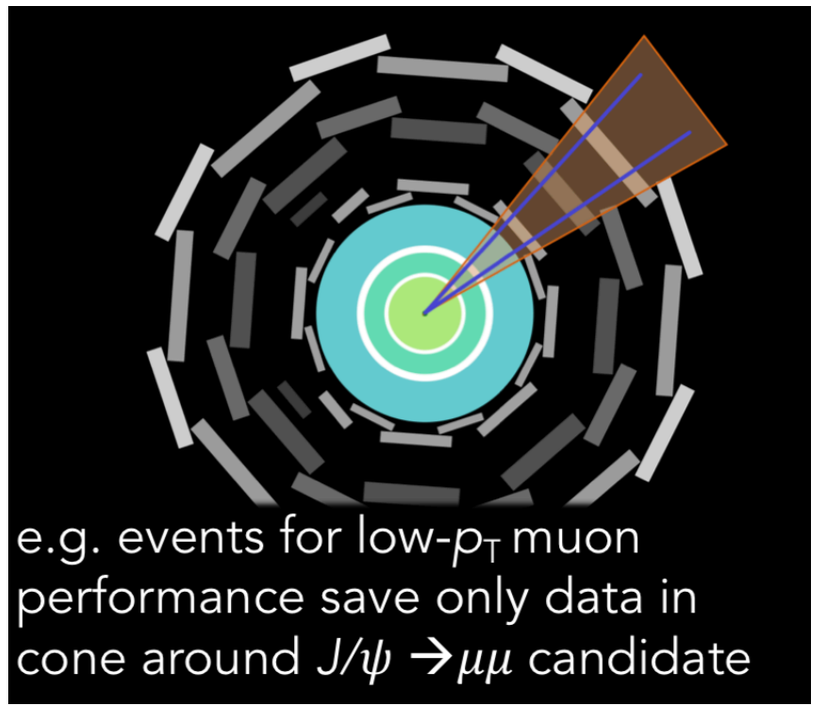

Challenges of real-time analysis

Present and future challenges

See also: Comput Softw Big Sci (2019) 3, 7

- Physics objects used for RTA need to be **identified, reconstructed and calibrated** as well as possible → what does this mean for analysis models?
  - Risk turning storage limitations into CPU limitations?
  - Lack of information may reduce performance / signal discrimination
  - See next slides…

- **MC statistics** will **never** be enough for final analysis
  - Forces to think about alternative solutions!
  - See also C. Gutschow / M. Schoenherr's talks yesterday
  - Background estimation techniques need to be data-driven & robust against signal

- Your event numbers may overflow as some runs collect more events than MAX_INT

LHCb online vs offline reconstruction

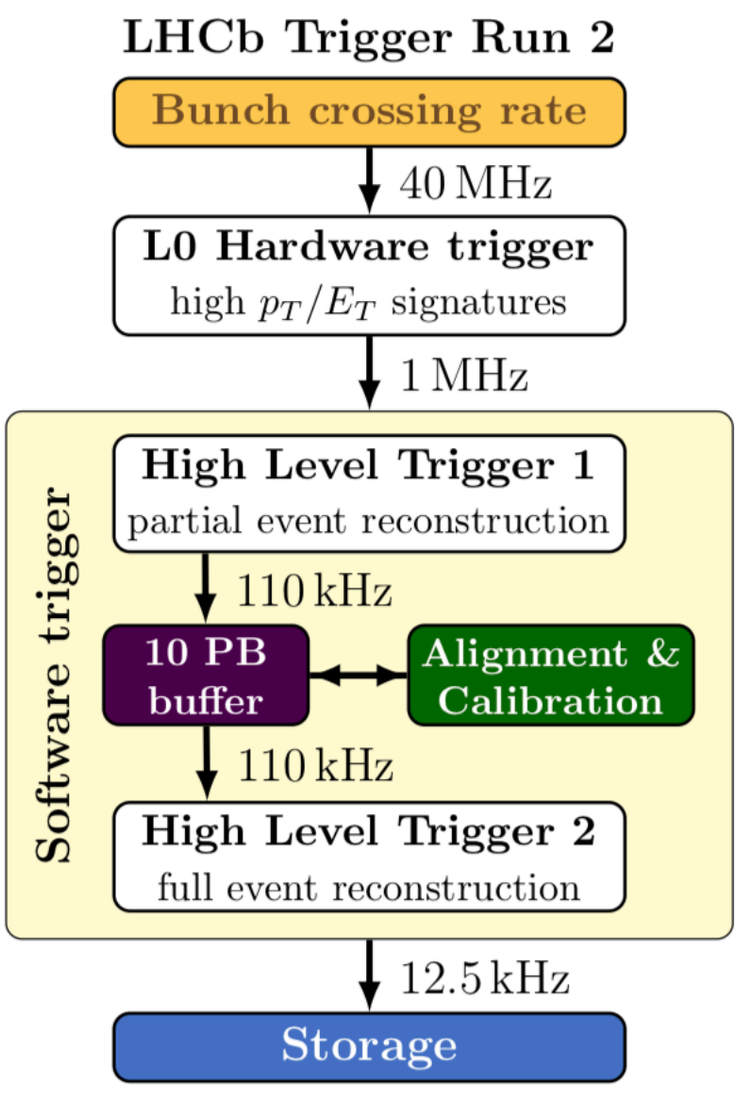

[JINST 14 (2019) P04013](#)
[JINST 14 (2019) P04006](#)

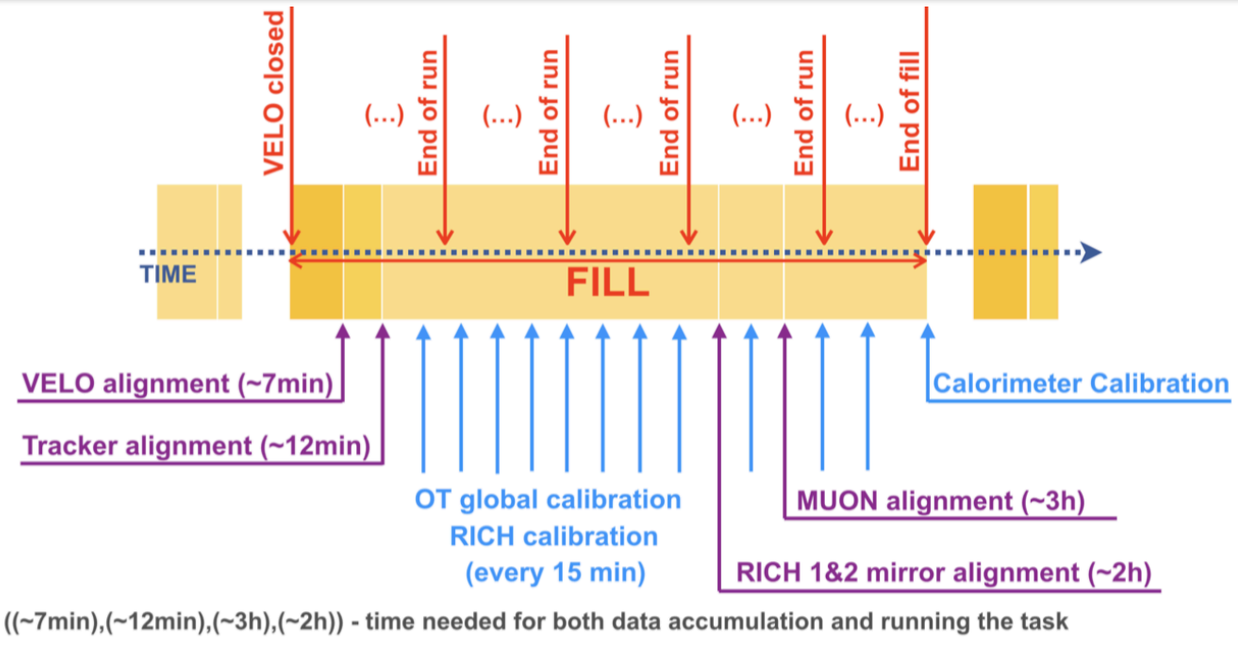

Use of buffer for calibration & alignment

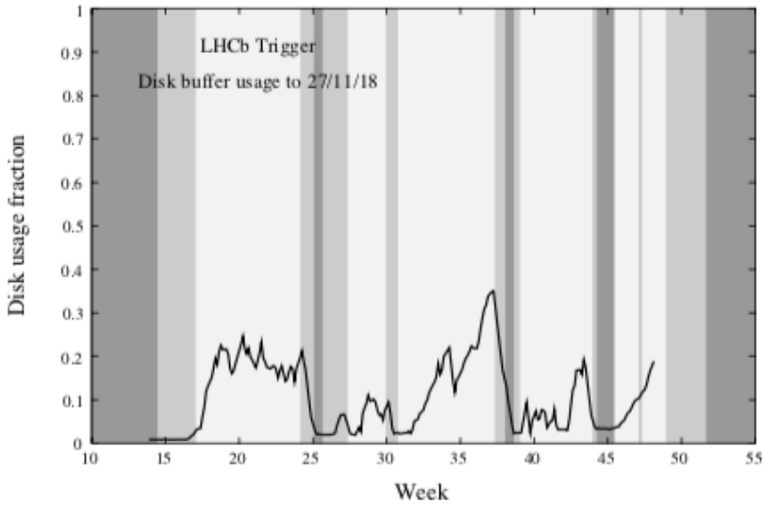

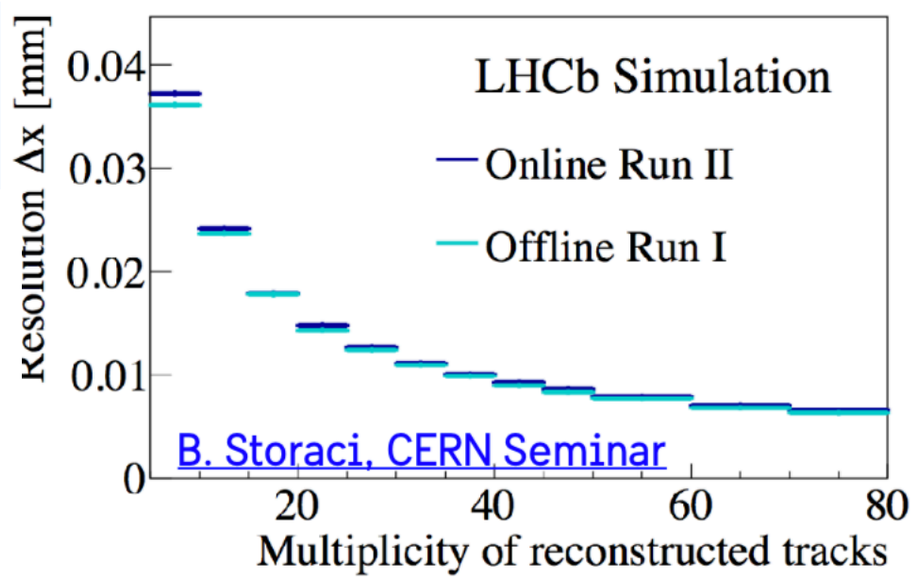

ATLAS/CMS HLT jet calibration not-quite-RTA results

% difference RTA vs offline jets before last calibration step

Used as a calibration factor for RTA jets

JHEP 08 (2018) 130

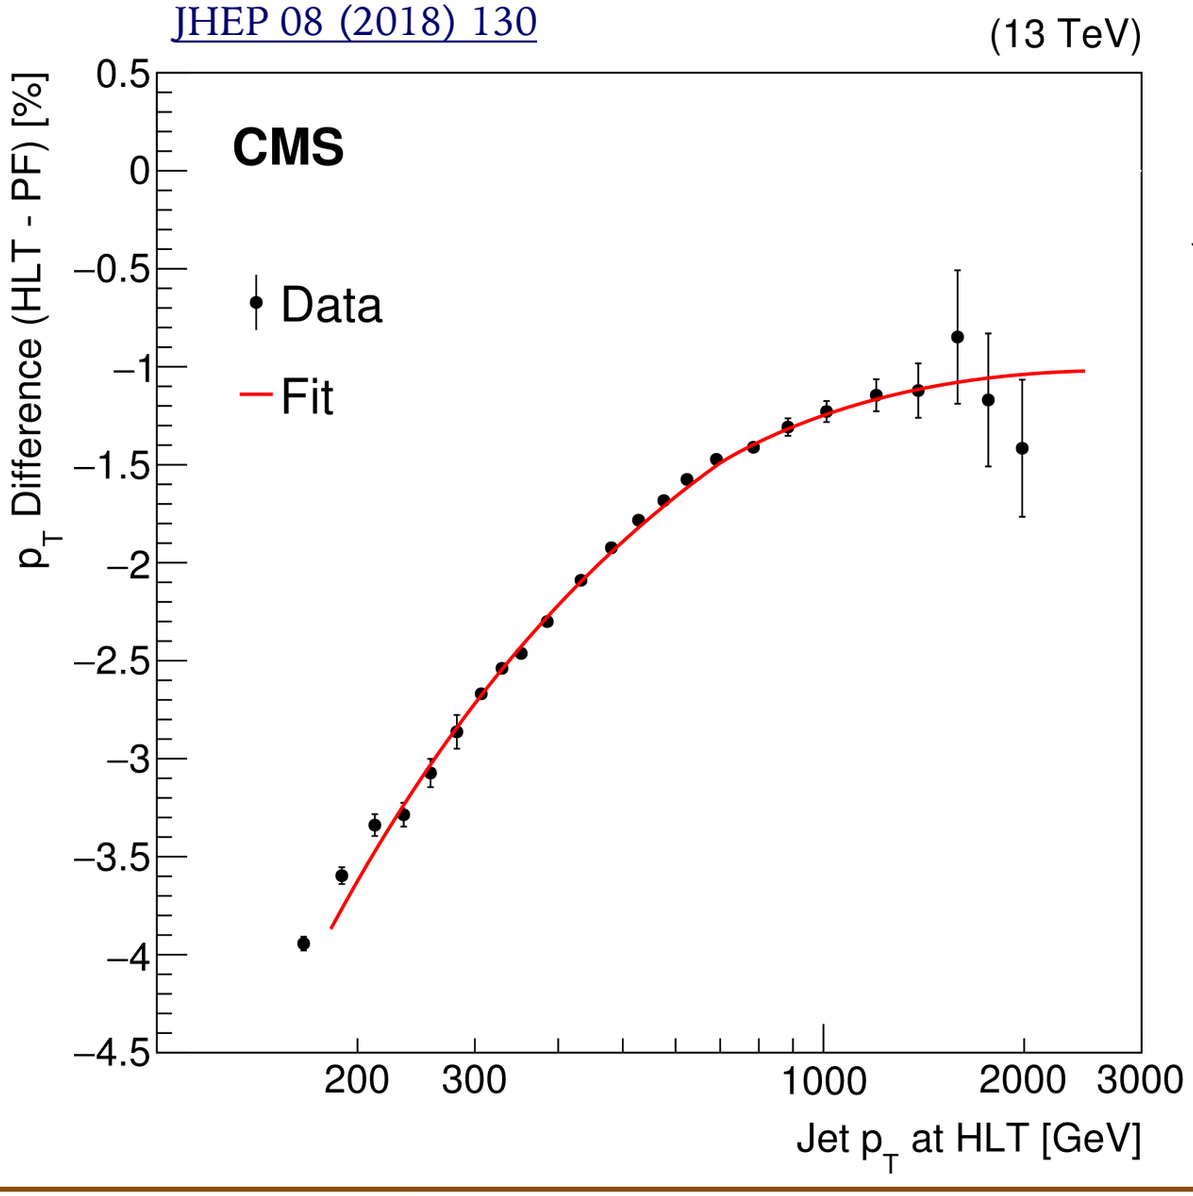

% difference RTA vs offline jets

after all calibrations applied

Phys. Rev. Lett. 121, 081801 (2018)

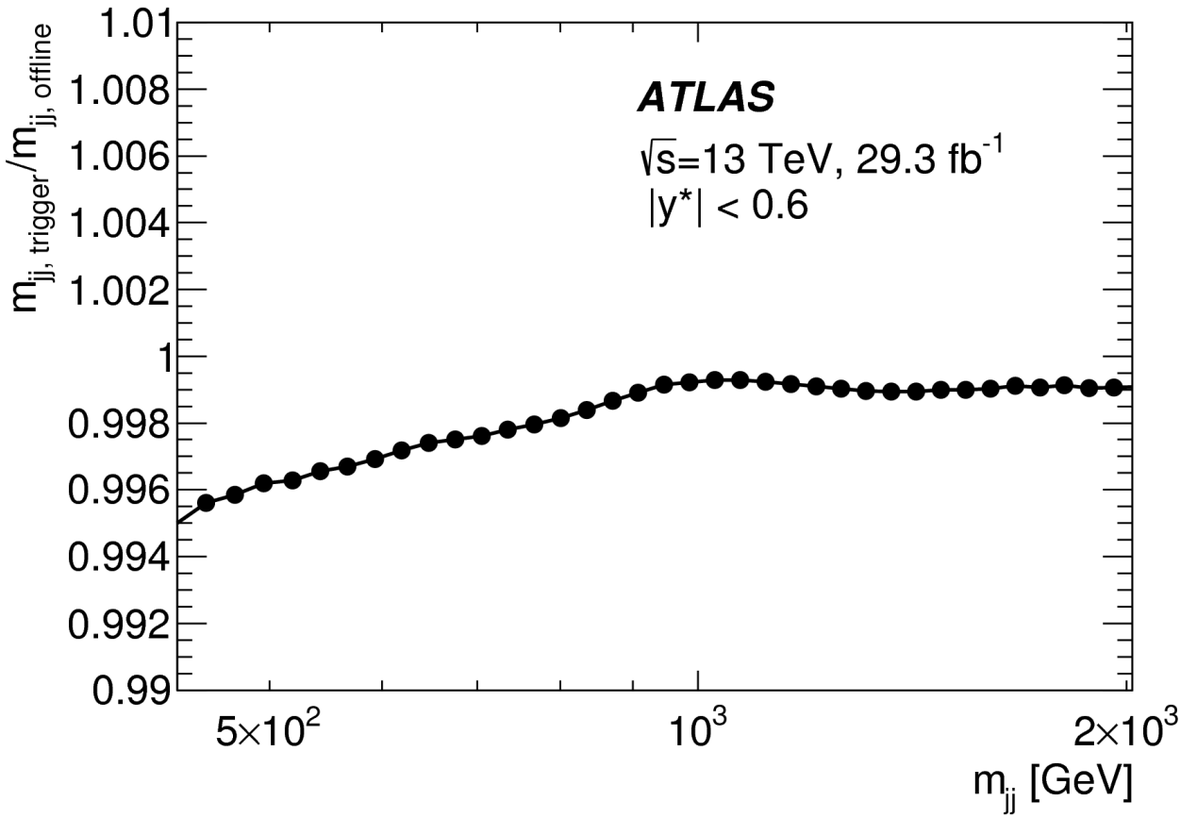

Almost identical performance of jet response for RTA objects (reduced jet resolution, but tolerable)

Need to plan for availability of validation data with matching RTA/offline objects

ATLAS/CMS (jet) calibration

- Important note: **difficult** (not impossible) to calibrate ATLAS & CMS physics objects ”in real time”:
  - Quasi-online calibration already happens in ATLAS/CMS for certain things (e.g. luminosity, beamspot, cells - see backup slide)
  - Some steps of jet calibration depend on use of ”clean” objects balancing the jets
    - ”clean” == rare → need to collect whole dataset before calibrating
    - other downside: limited statistics of calibration sample

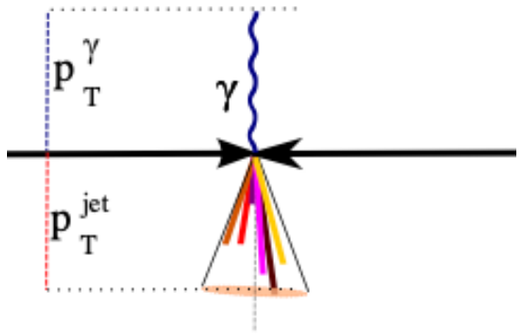

- Studies ongoing of what calibration steps we could move online for HL-LHC: see e.g. ATL-DAQ-PUB-2017-003

RTA-driven data-taking & calibration improvements

An **HL-LHC wish:** get RTA objects to look as close as possible as regular ones
Good for RTA, good for the experiment

The earlier the calibration is applied, the better
(it sharpens the trigger efficiency turn-on curve)

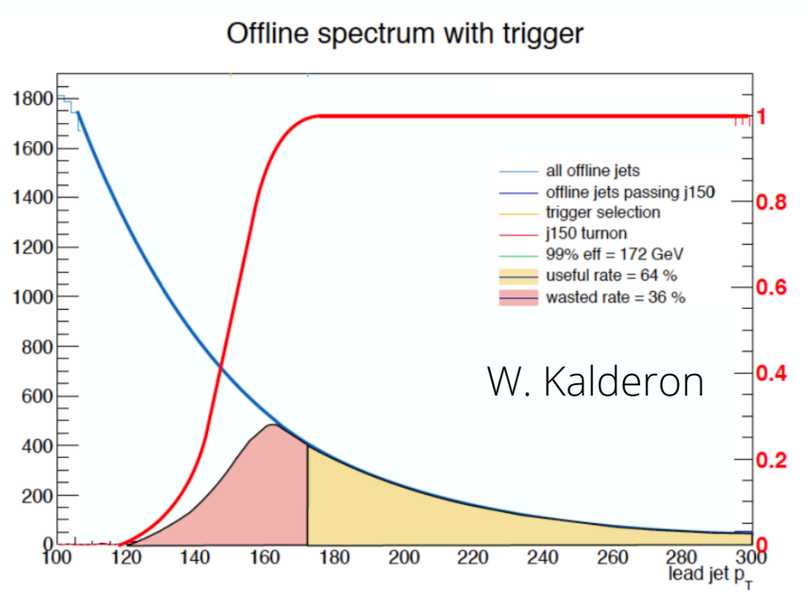

- This benefits the whole experiment: less wasted data
  - Some early ideas in ATL-DAQ-PUB-2017-003 (tests ongoing on whether ideas work)
  - CMS calorimeter calibration with ML data challenge at Institut Pascal real-time analysis workshop

Requires thinking of **HLT and offline software** as **the same thing**
(and when you think about it from a coding perspective, it is desirable anyway)

Where we need to work for more RTA

Tracking is important to extend RTA to other objects / lower pT (pile-up)

- *"Tracking is expensive"* → need to make it cheaper
  Just some examples from CHEP (+ next talk):
  - Hardware: CHEP talk *Level-1 track finding with an all-FPGA system at CMS for the HL-LHC*, Thomas Owen James
  - Accelerators: CHEP talk *Physics performance and event throughput of the GPU High Level Trigger 1 of LHCb, Dorothea Von Bruch*
  - Software: ATL-PHYS-PUB-2019-041, *M. Elsing, talk at the joint IRIS-HEP / HSF Trigger&Reco meeting, minutes*

ATL-PHYS-PUB-2019-041

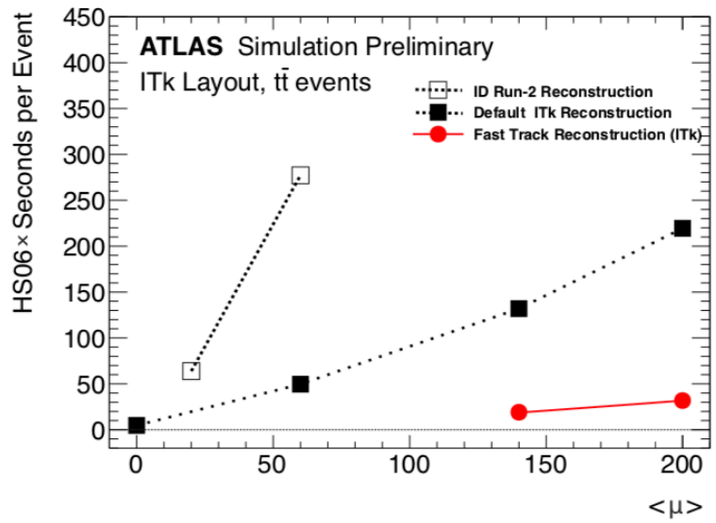

Note that the physics performance needs have to be tailored case-by-case, e.g.

- tracking for **pile-up** suppression
- tracking for **precision** object ID / measurements
- tracking for l**ong-lived particles** (large-radius tracking)

Much work ongoing on efficient solutions…but equivalent work needed from the physics side: **connecting the two worlds is necessary**

Discussion point: can we use physics benchmarks to drive efficient software & computing improvements? Which ones?

Conclusions

How LHC collaborations make the most of the data

Interesting time for high energy accelerator physics:
we don't know what to expect from DM/new physics
(but we have a prior: it should be *somewhere*)
we have the **LHC running now / HL-LHC running soon**,
and the data we discard is **gone**

1. Analyze as much data as possible, **as fast as possible → RTA**
2. Save data for **further reconstruction, later** (possible, but not covered here)
3. Implement more refined algorithms to **look for the unexpected**
   - Including unsupervised searches / novelty detection
     - E.g. CHEP talk New-Physics agnostic searches for New Physics, Kinga Anna Wozniak

Real-time analysis, in different contexts

C. Fitzpatrick, **LHCb**

E. Bellm, **Large Synoptic Survey Telescope**

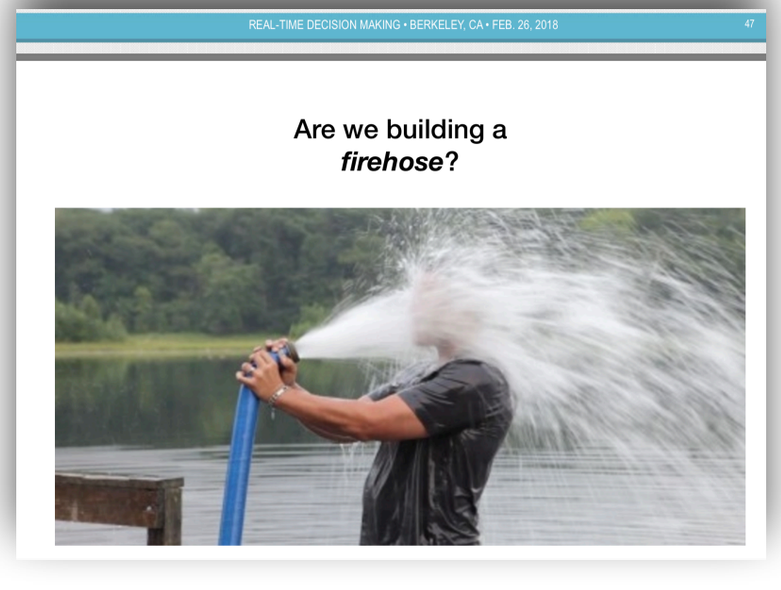

The LHC and modern astrophysics surveys are **data firehoses → RTA can help**

"Real-time analysis will be a keyword for multimessenger astronomy"
E. Cuoco, Bohr Seminar @ U of Manchester

REALTIME Study Group, Pufendorf Institute of Advanced Studies @ Lund University
connecting HEP, psychology, law, astronomy, maths, engineering

Foundations of physics strategies, and synergies

APPEC Astroparticle

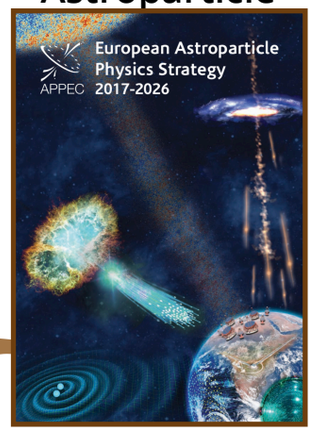

ECFA Particle

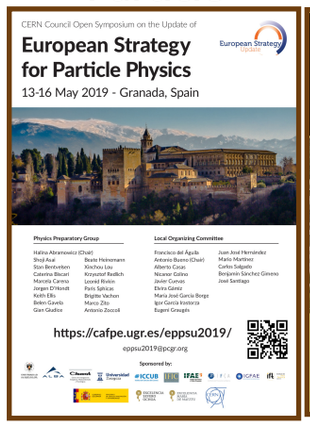

NuPECC Nuclear

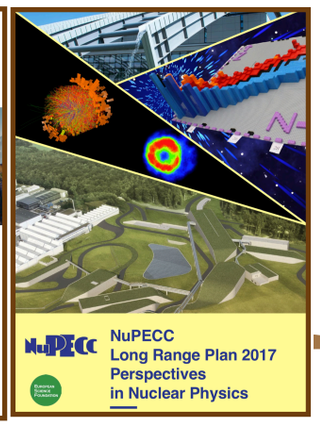

Common theory ground

instrumentation (accelerators, beams, detectors, vacuum & cryogenics, control & automation...)

data acquisition, computing, data sharing & open science

HSF HEP Software Foundation

ESCAPE European Science Cluster of Astronomy & Particle physics ESFRI research Infrastructures

iris hep

& more...

Talk at EPS-HEP / ECFA session 2019, CERN EP Newsletter
HSF meeting @ JENAS, 2019

Discussion point from an "external observer": understand if/when/how to connect the SoI in-progress to non-UK / inter-experiment initiatives towards efficient software & computing (see also G. Stewart / H. Schreiner's talks)

Join us for the HSF/WLCG workshop!

Jointly organised between the HSF and WLCG, the focus of this workshop is the challenge of **adapting our software and computing infrastructures to increased data rates, new computing technologies and facility evolution**. All of this is targeted to maximise the physics opportunities from future upgrades and new facilities.

The workshop will take a forward look at key topics for software and computing, reviewing progress, looking at new approaches, and discussing opportunities and challenges. There will be **plenty of time for discussion** and the development of R&D ideas that should be explored.

The workshop is **open to everyone in the field**, from LHC experiments to the intensity frontier, dark matter, astroparticle and other data intensive sciences. **Participation of Early Career Researchers is particularly welcome.**

http://indico.cern.ch/e/HSFWLCG2020

Backup slides

LHCb event sizes

Turbo

- Turbo is LHCb's Real-Time Analysis paradigm for reduced event format data[5]

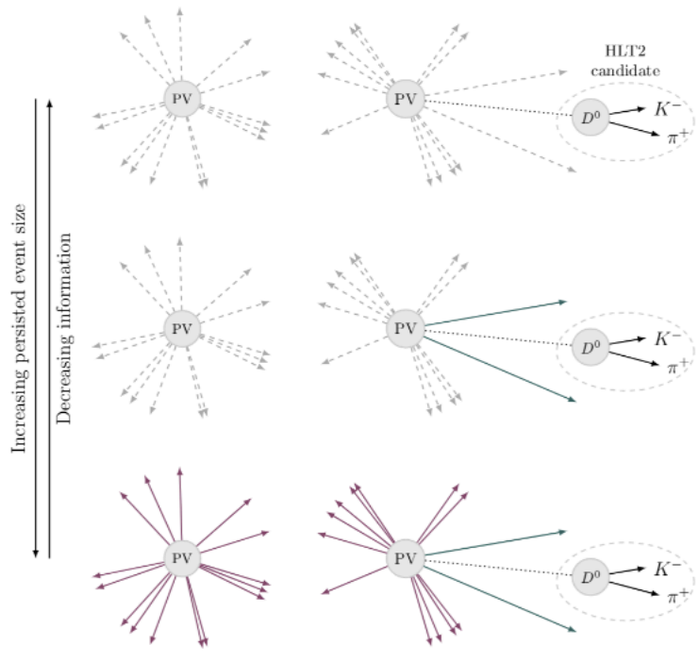

- High degree of flexibility: Save only as much of the event as is needed for analysis
  - Keep all reconstructed objects, drop the raw event: < 100kB
  - Keep only objects used to trigger: 7kB
  - 'Selective Persistence' objects used to trigger + user-defined selection: $7 \rightarrow 100$kB

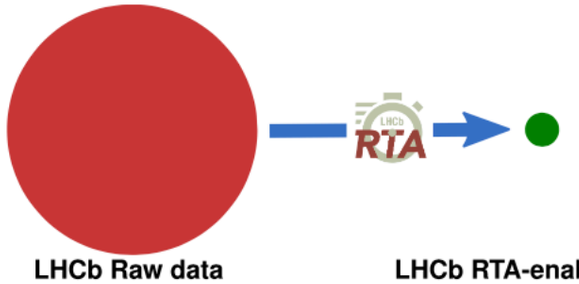

LHCb Raw data
100 kB per collision

LHCb RTA-enabled data
7kB per collision

[5] arXiv:1604.05596, JINST 14 (2019) P04006

C. Fitzpatrick

February 17, 2020

UK Research and Innovation

Data parking in CMS

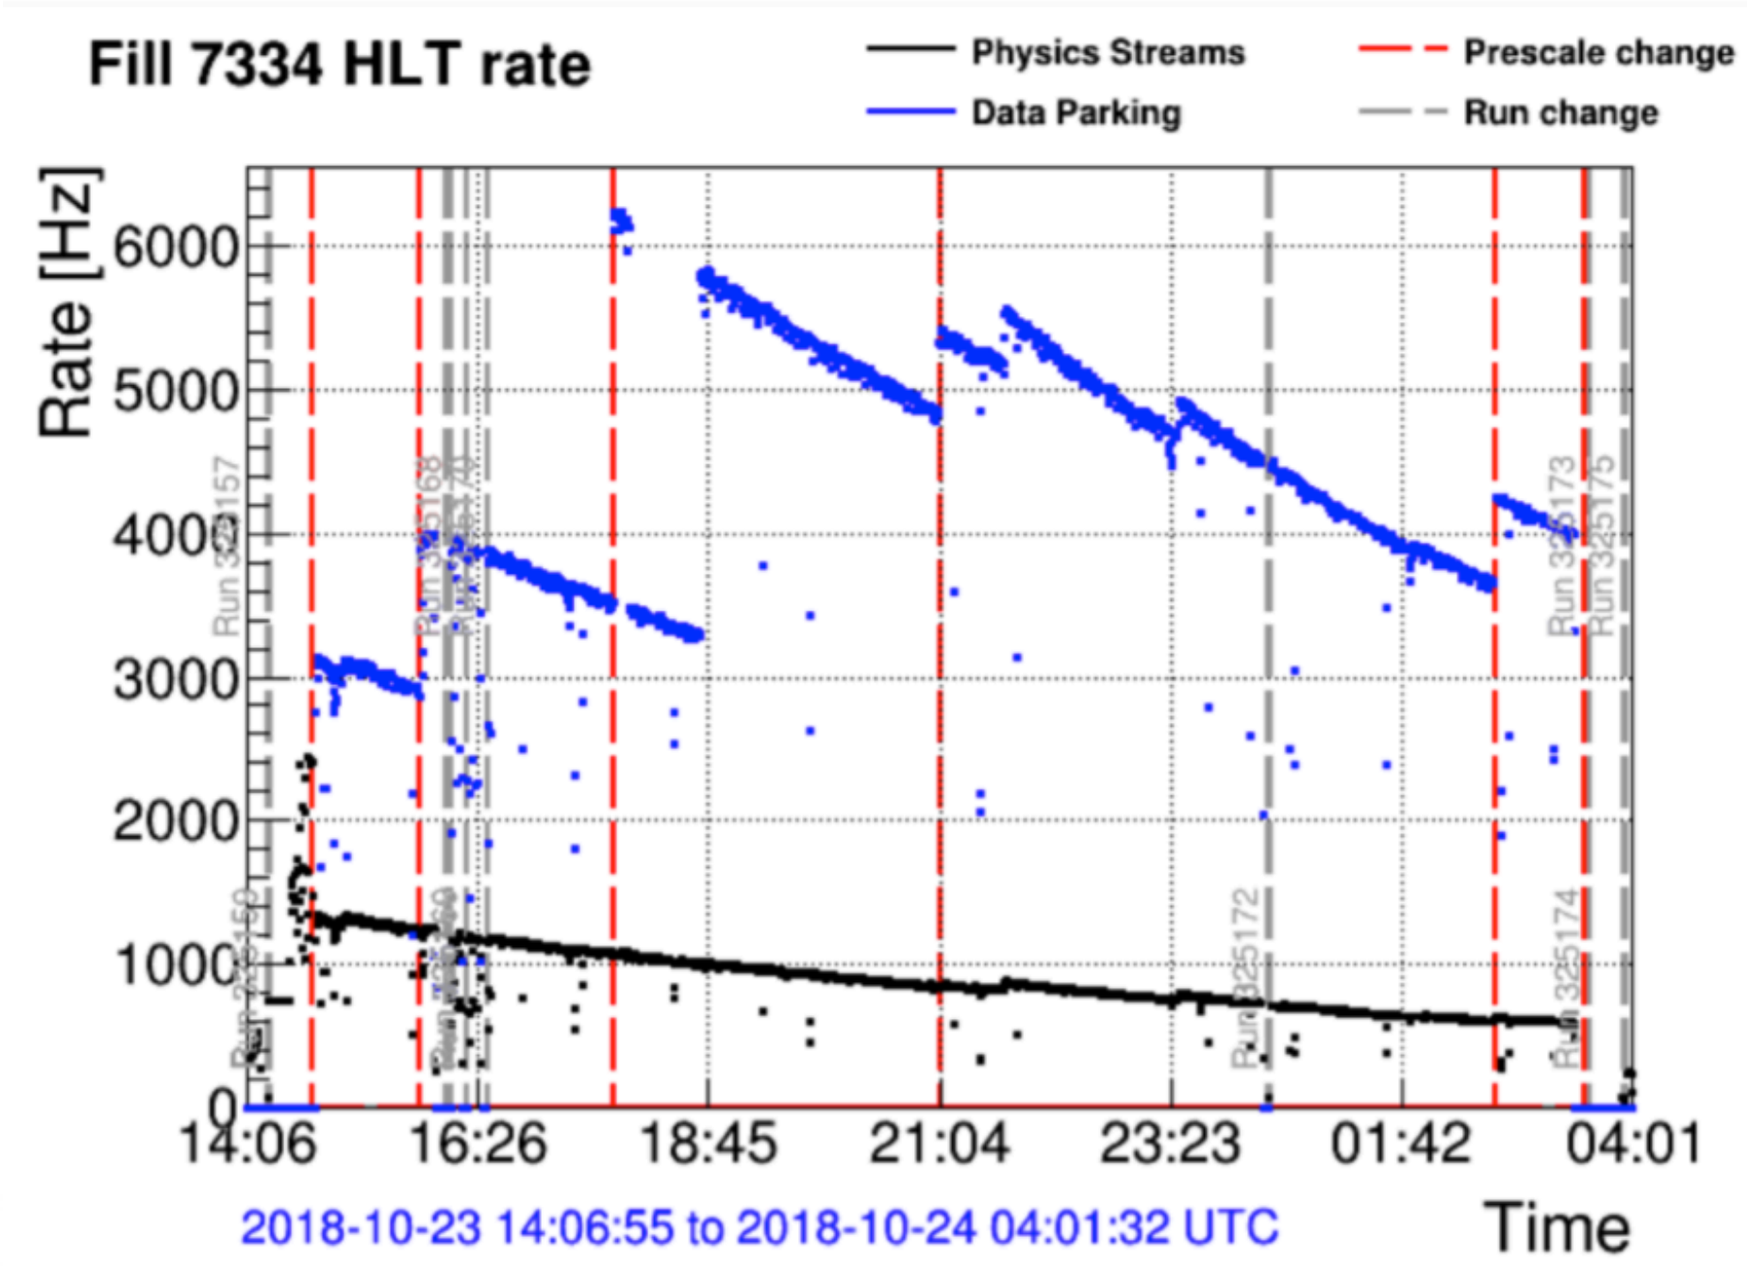

End-of-fill in ATLAS

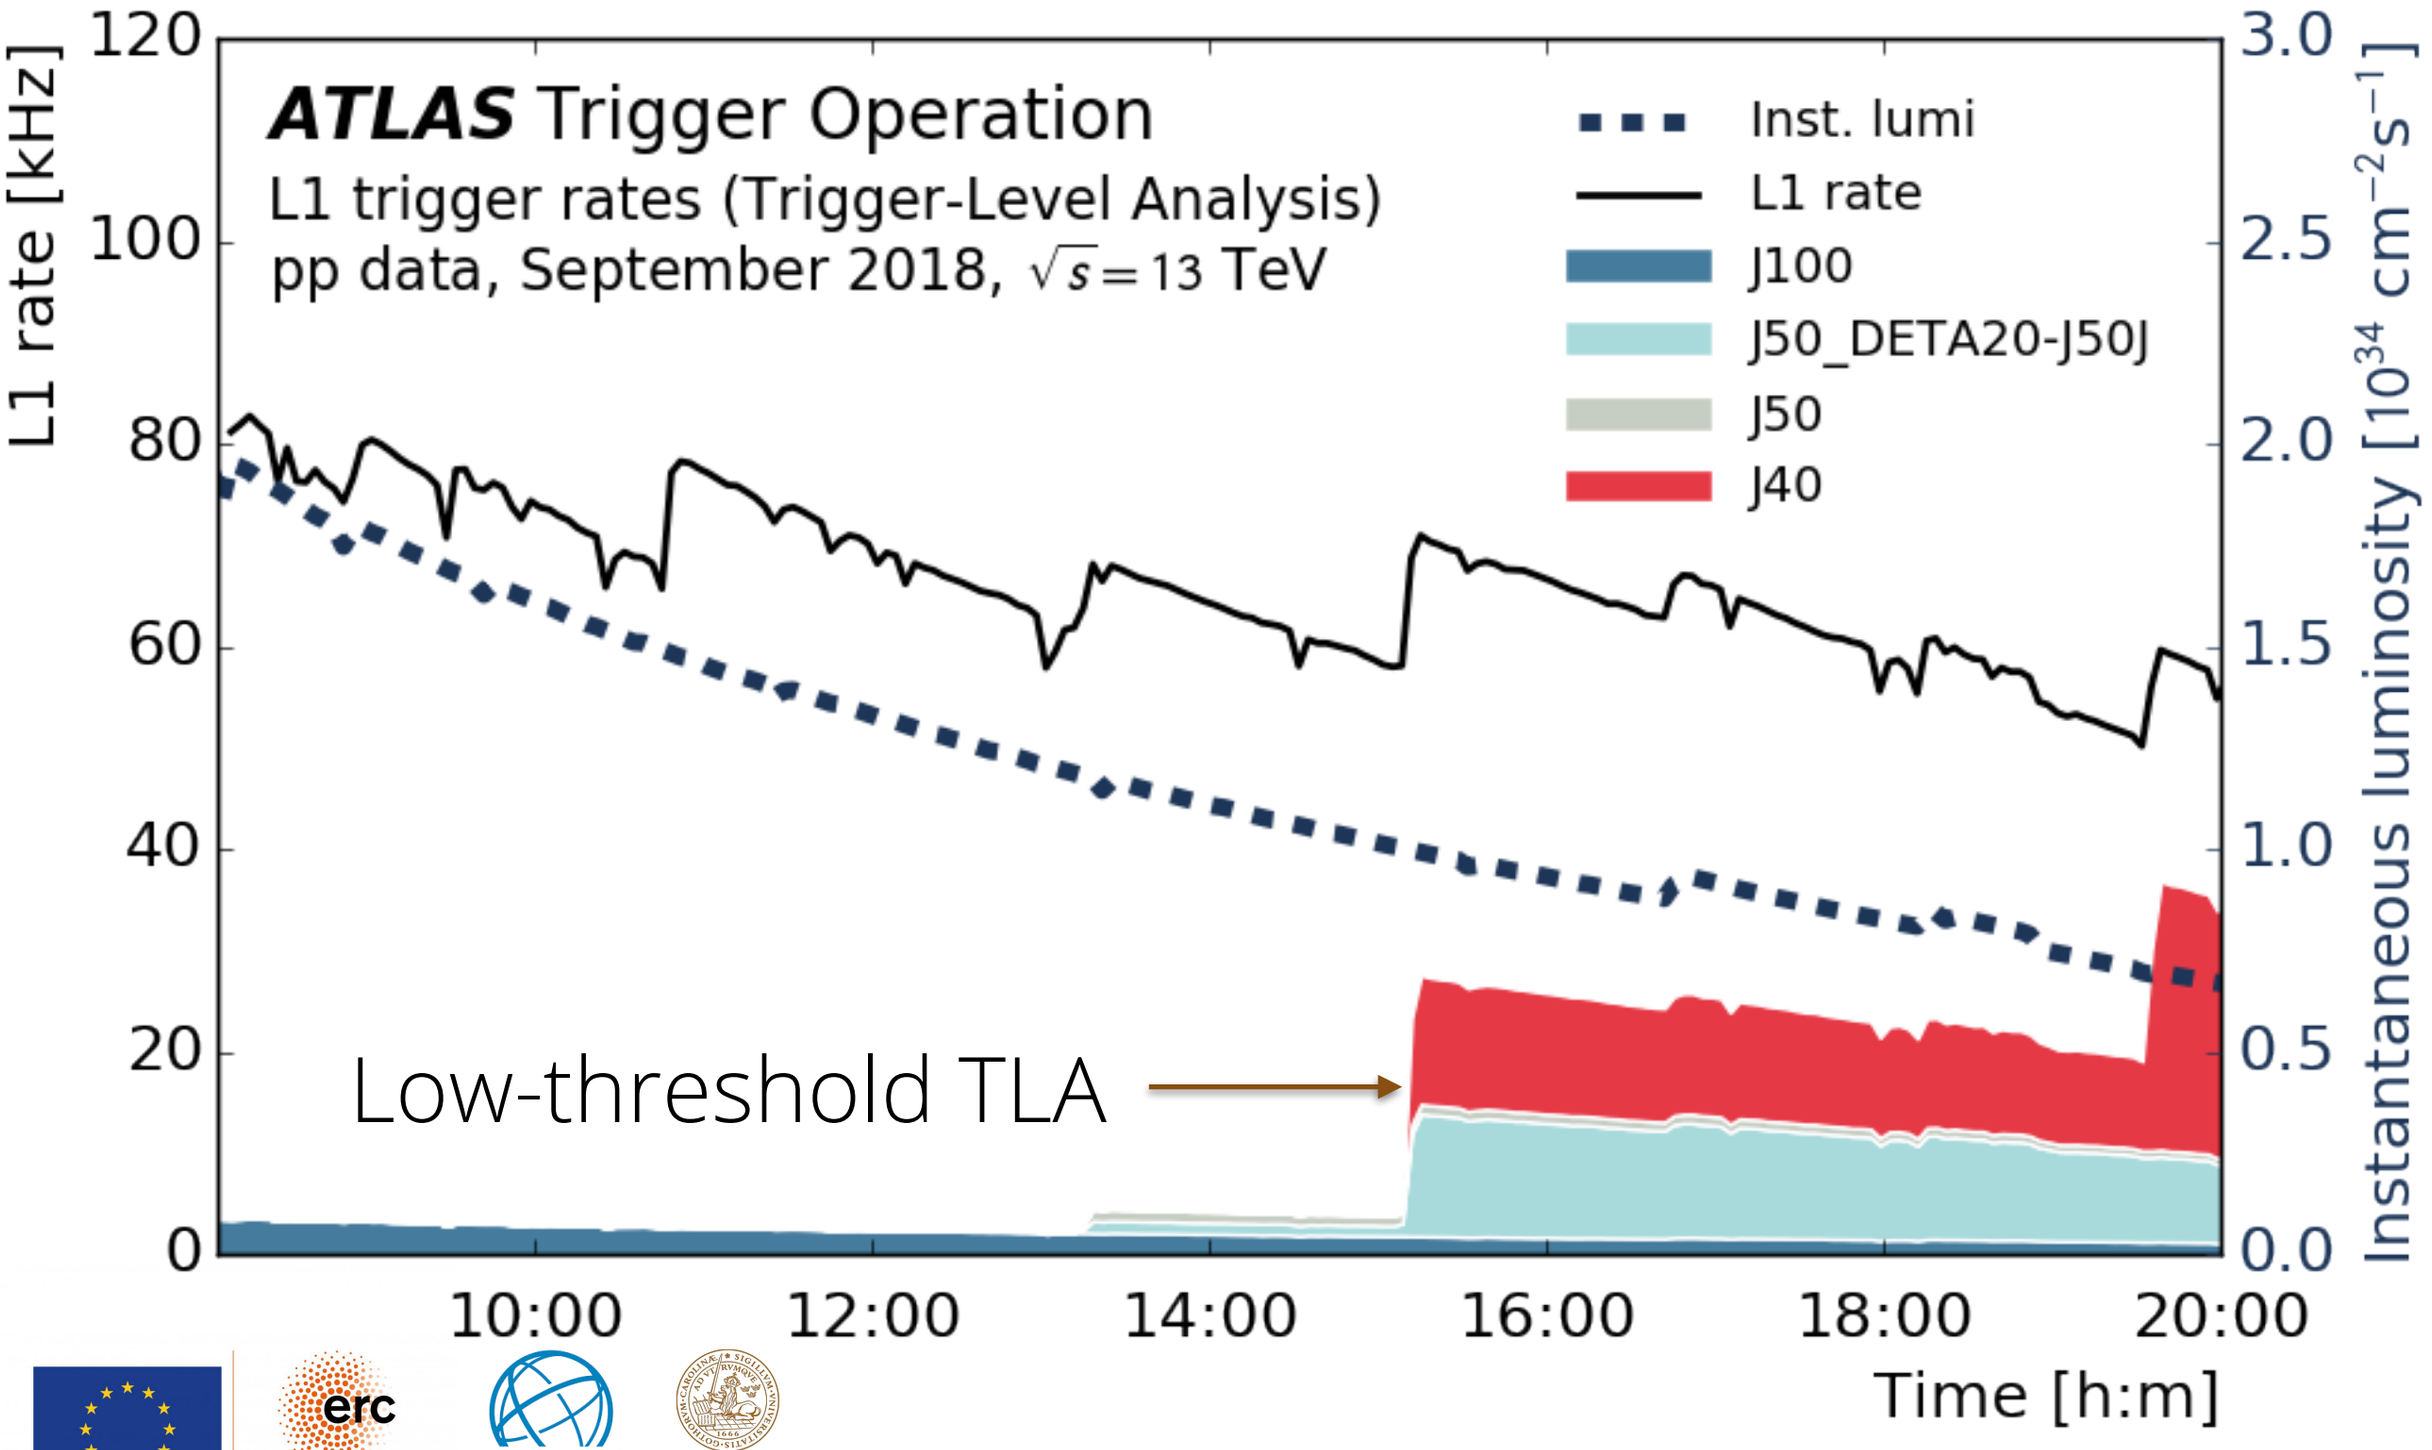

Complementarity of visible/invisible searches

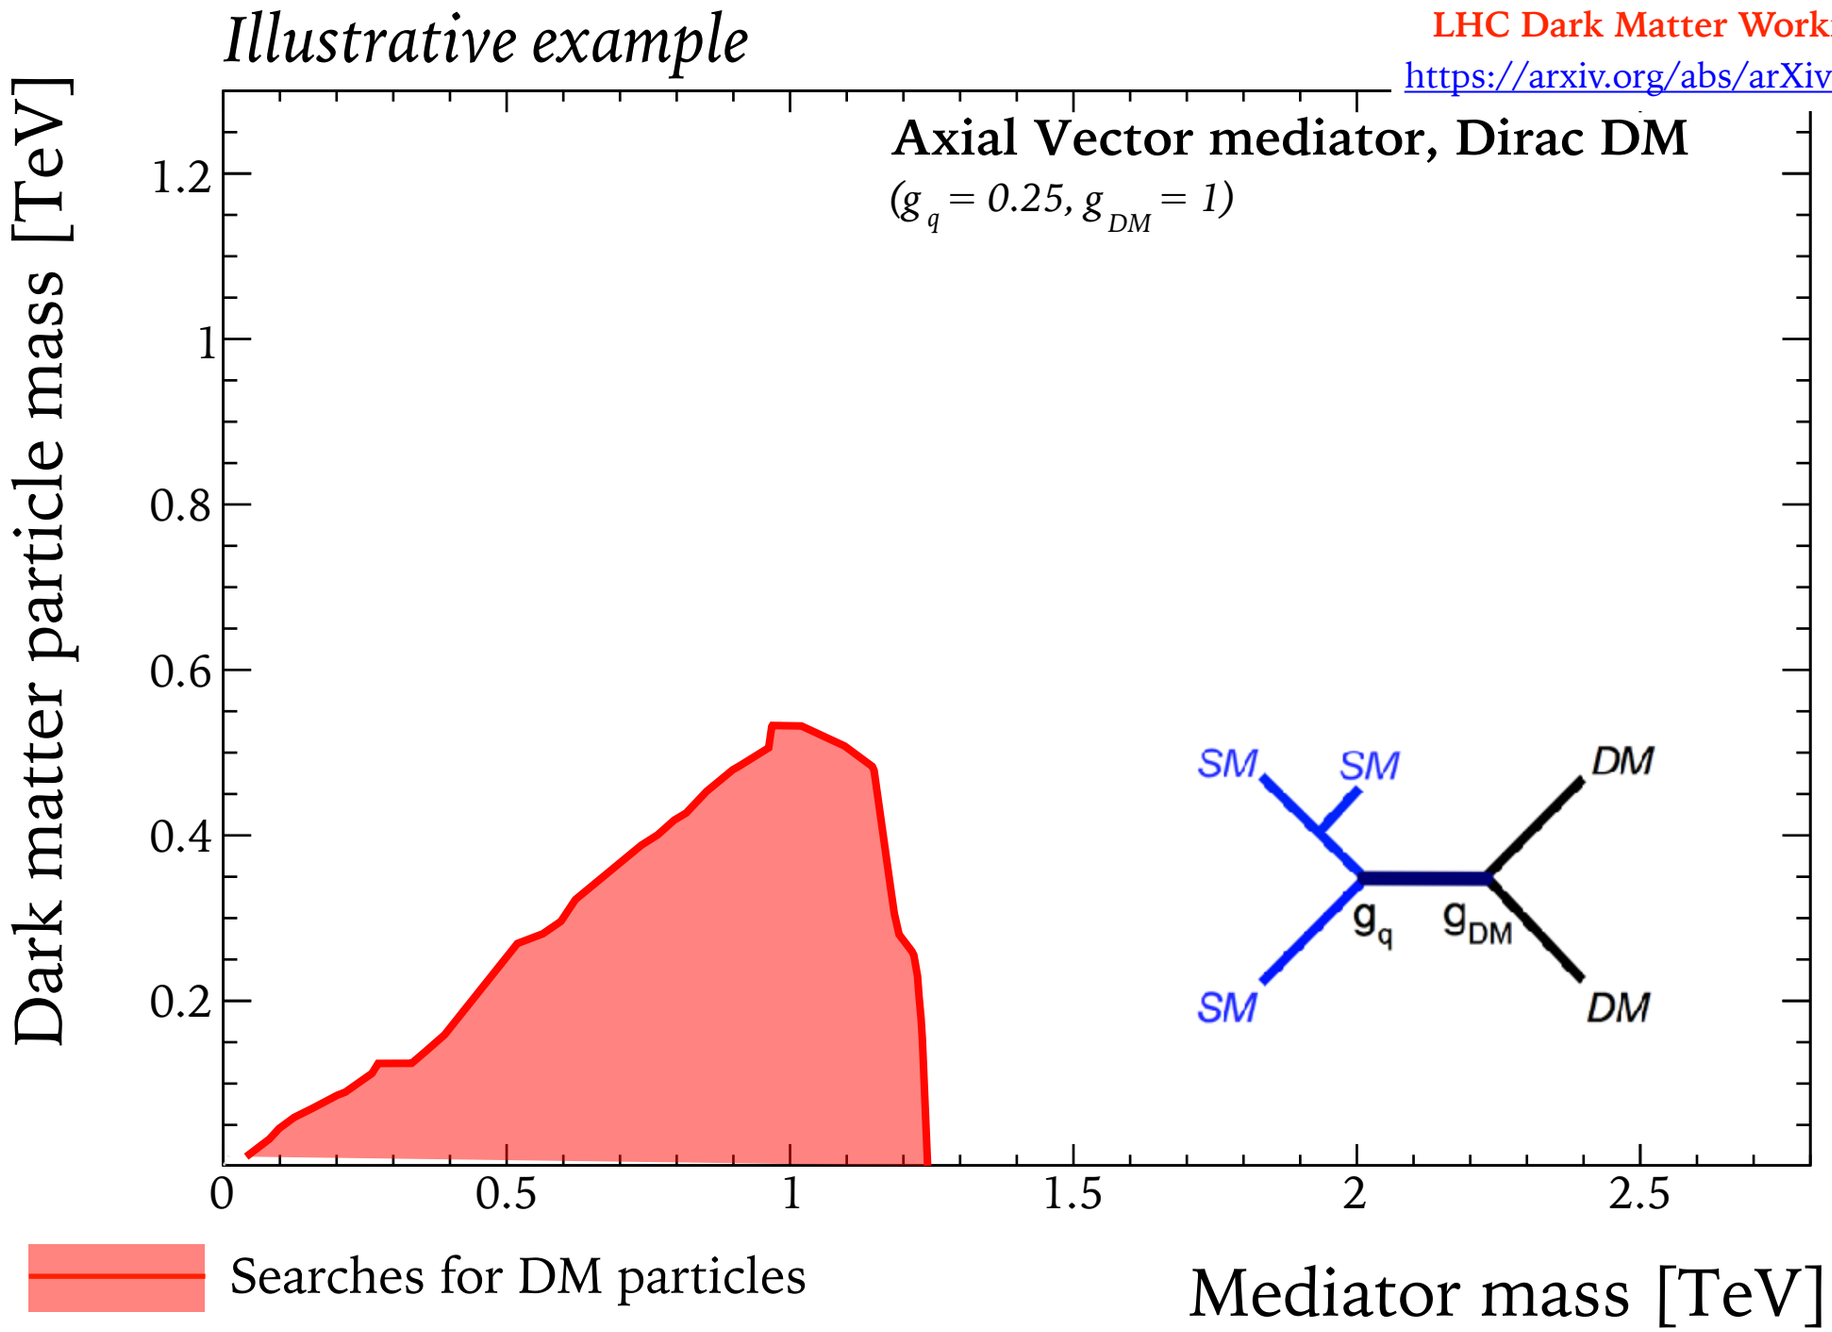

ATLAS EXPERIMENT
European Research Council
Established by the European Commission
European Commission
Horizon 2020
European Union funding
for Research & Innovation
LUND UNIVERSITY

Complementarity of visible/invisible searches

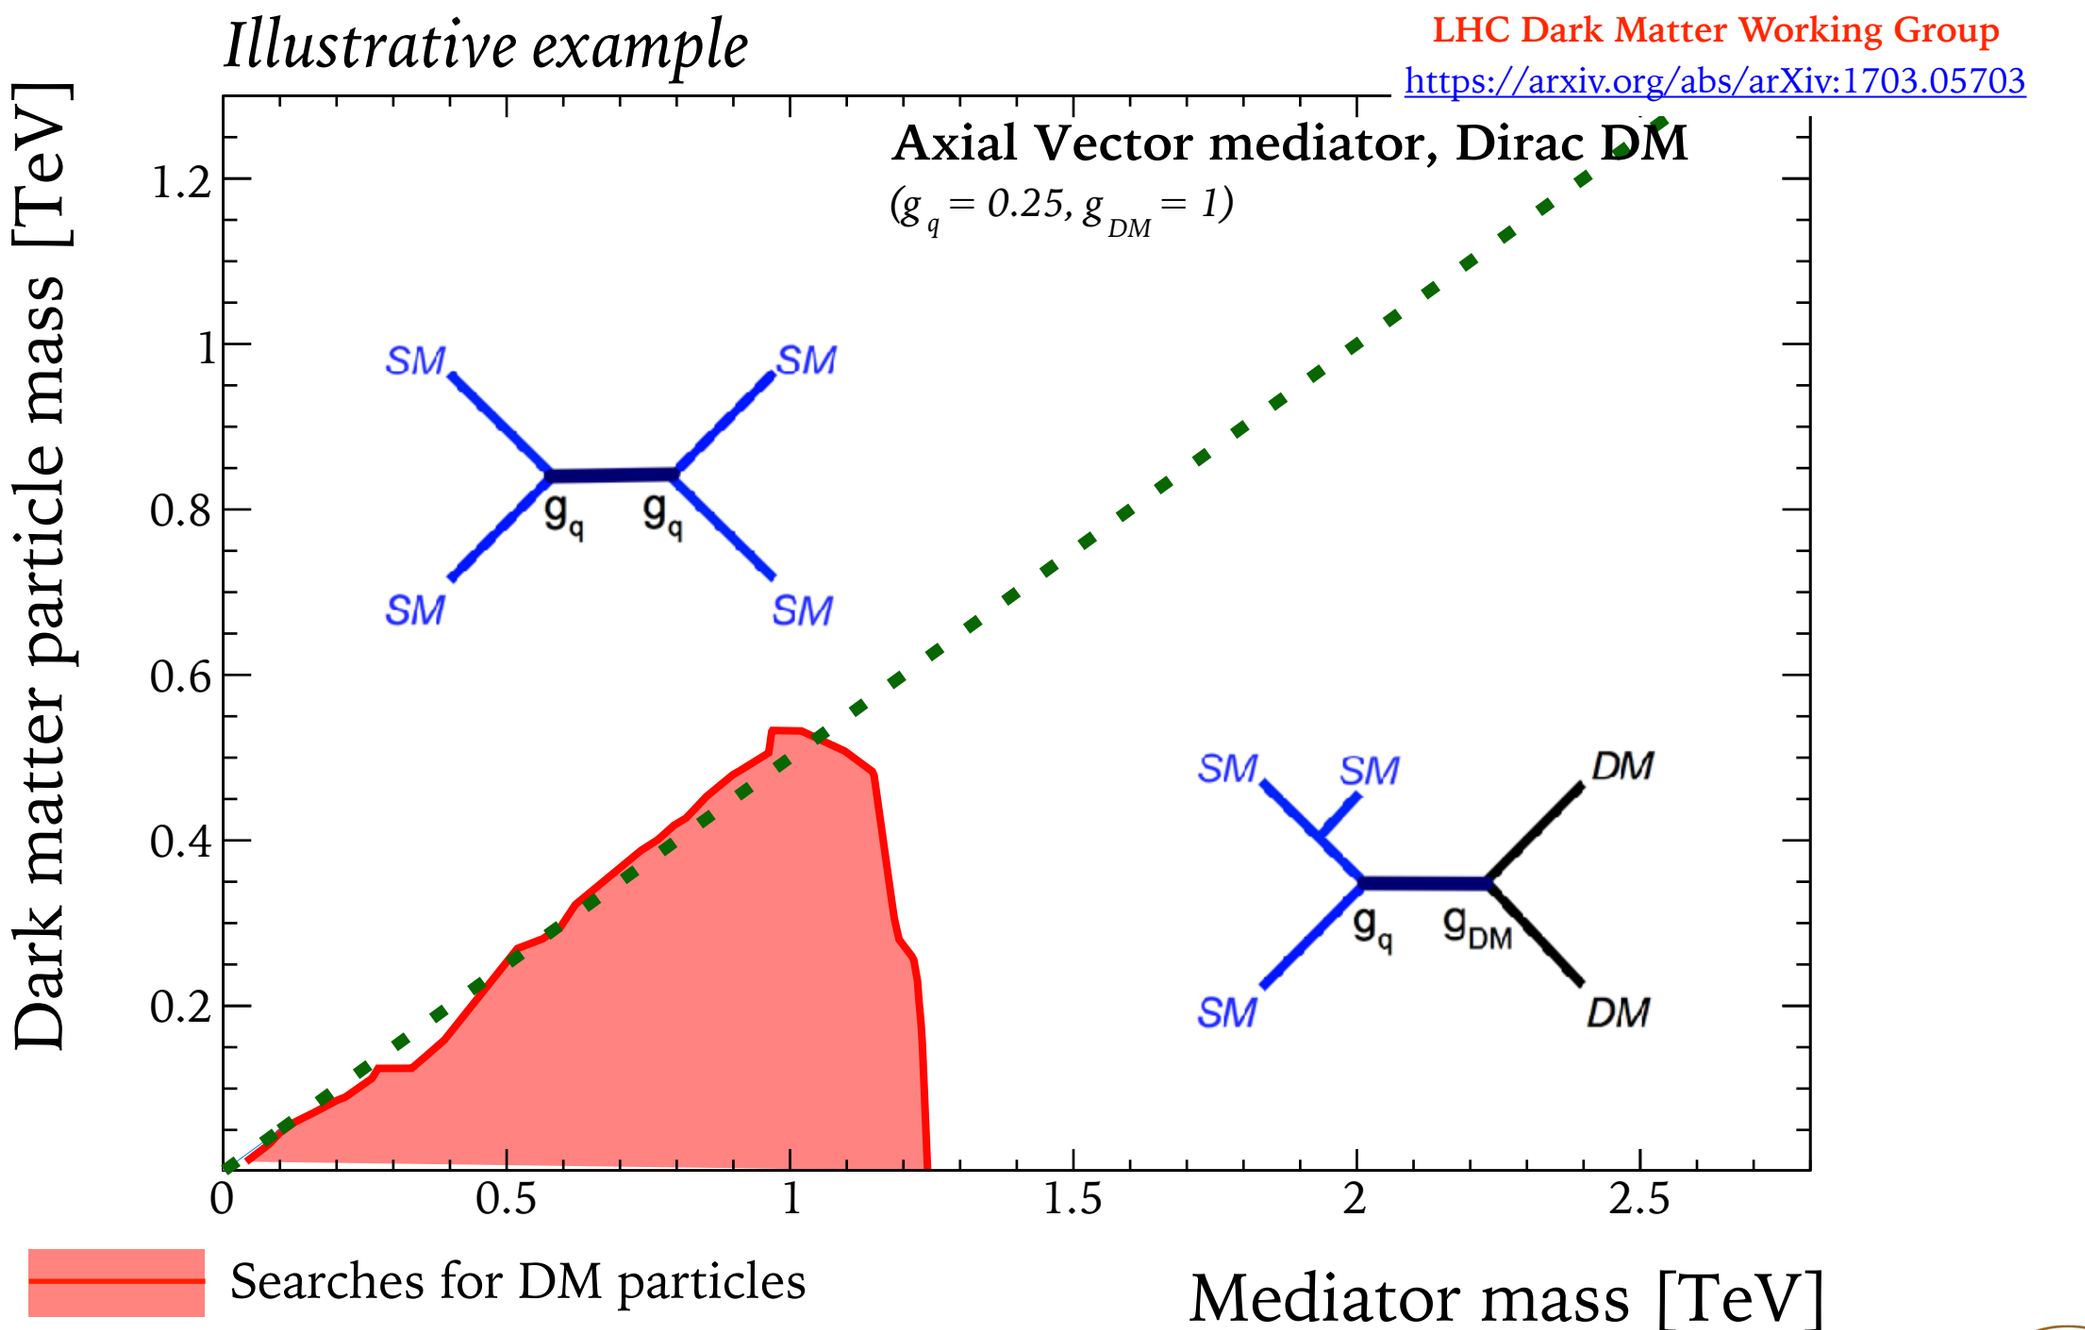

LHC Dark Matter Working Group
https://arxiv.org/abs/arXiv:1703.05703

ATLAS EXPERIMENT
European Research Council
Established by the European Commission
European Commission
Horizon 2020
European Union funding
for Research & Innovation
LUND UNIVERSITY

Complementarity of visible/invisible searches

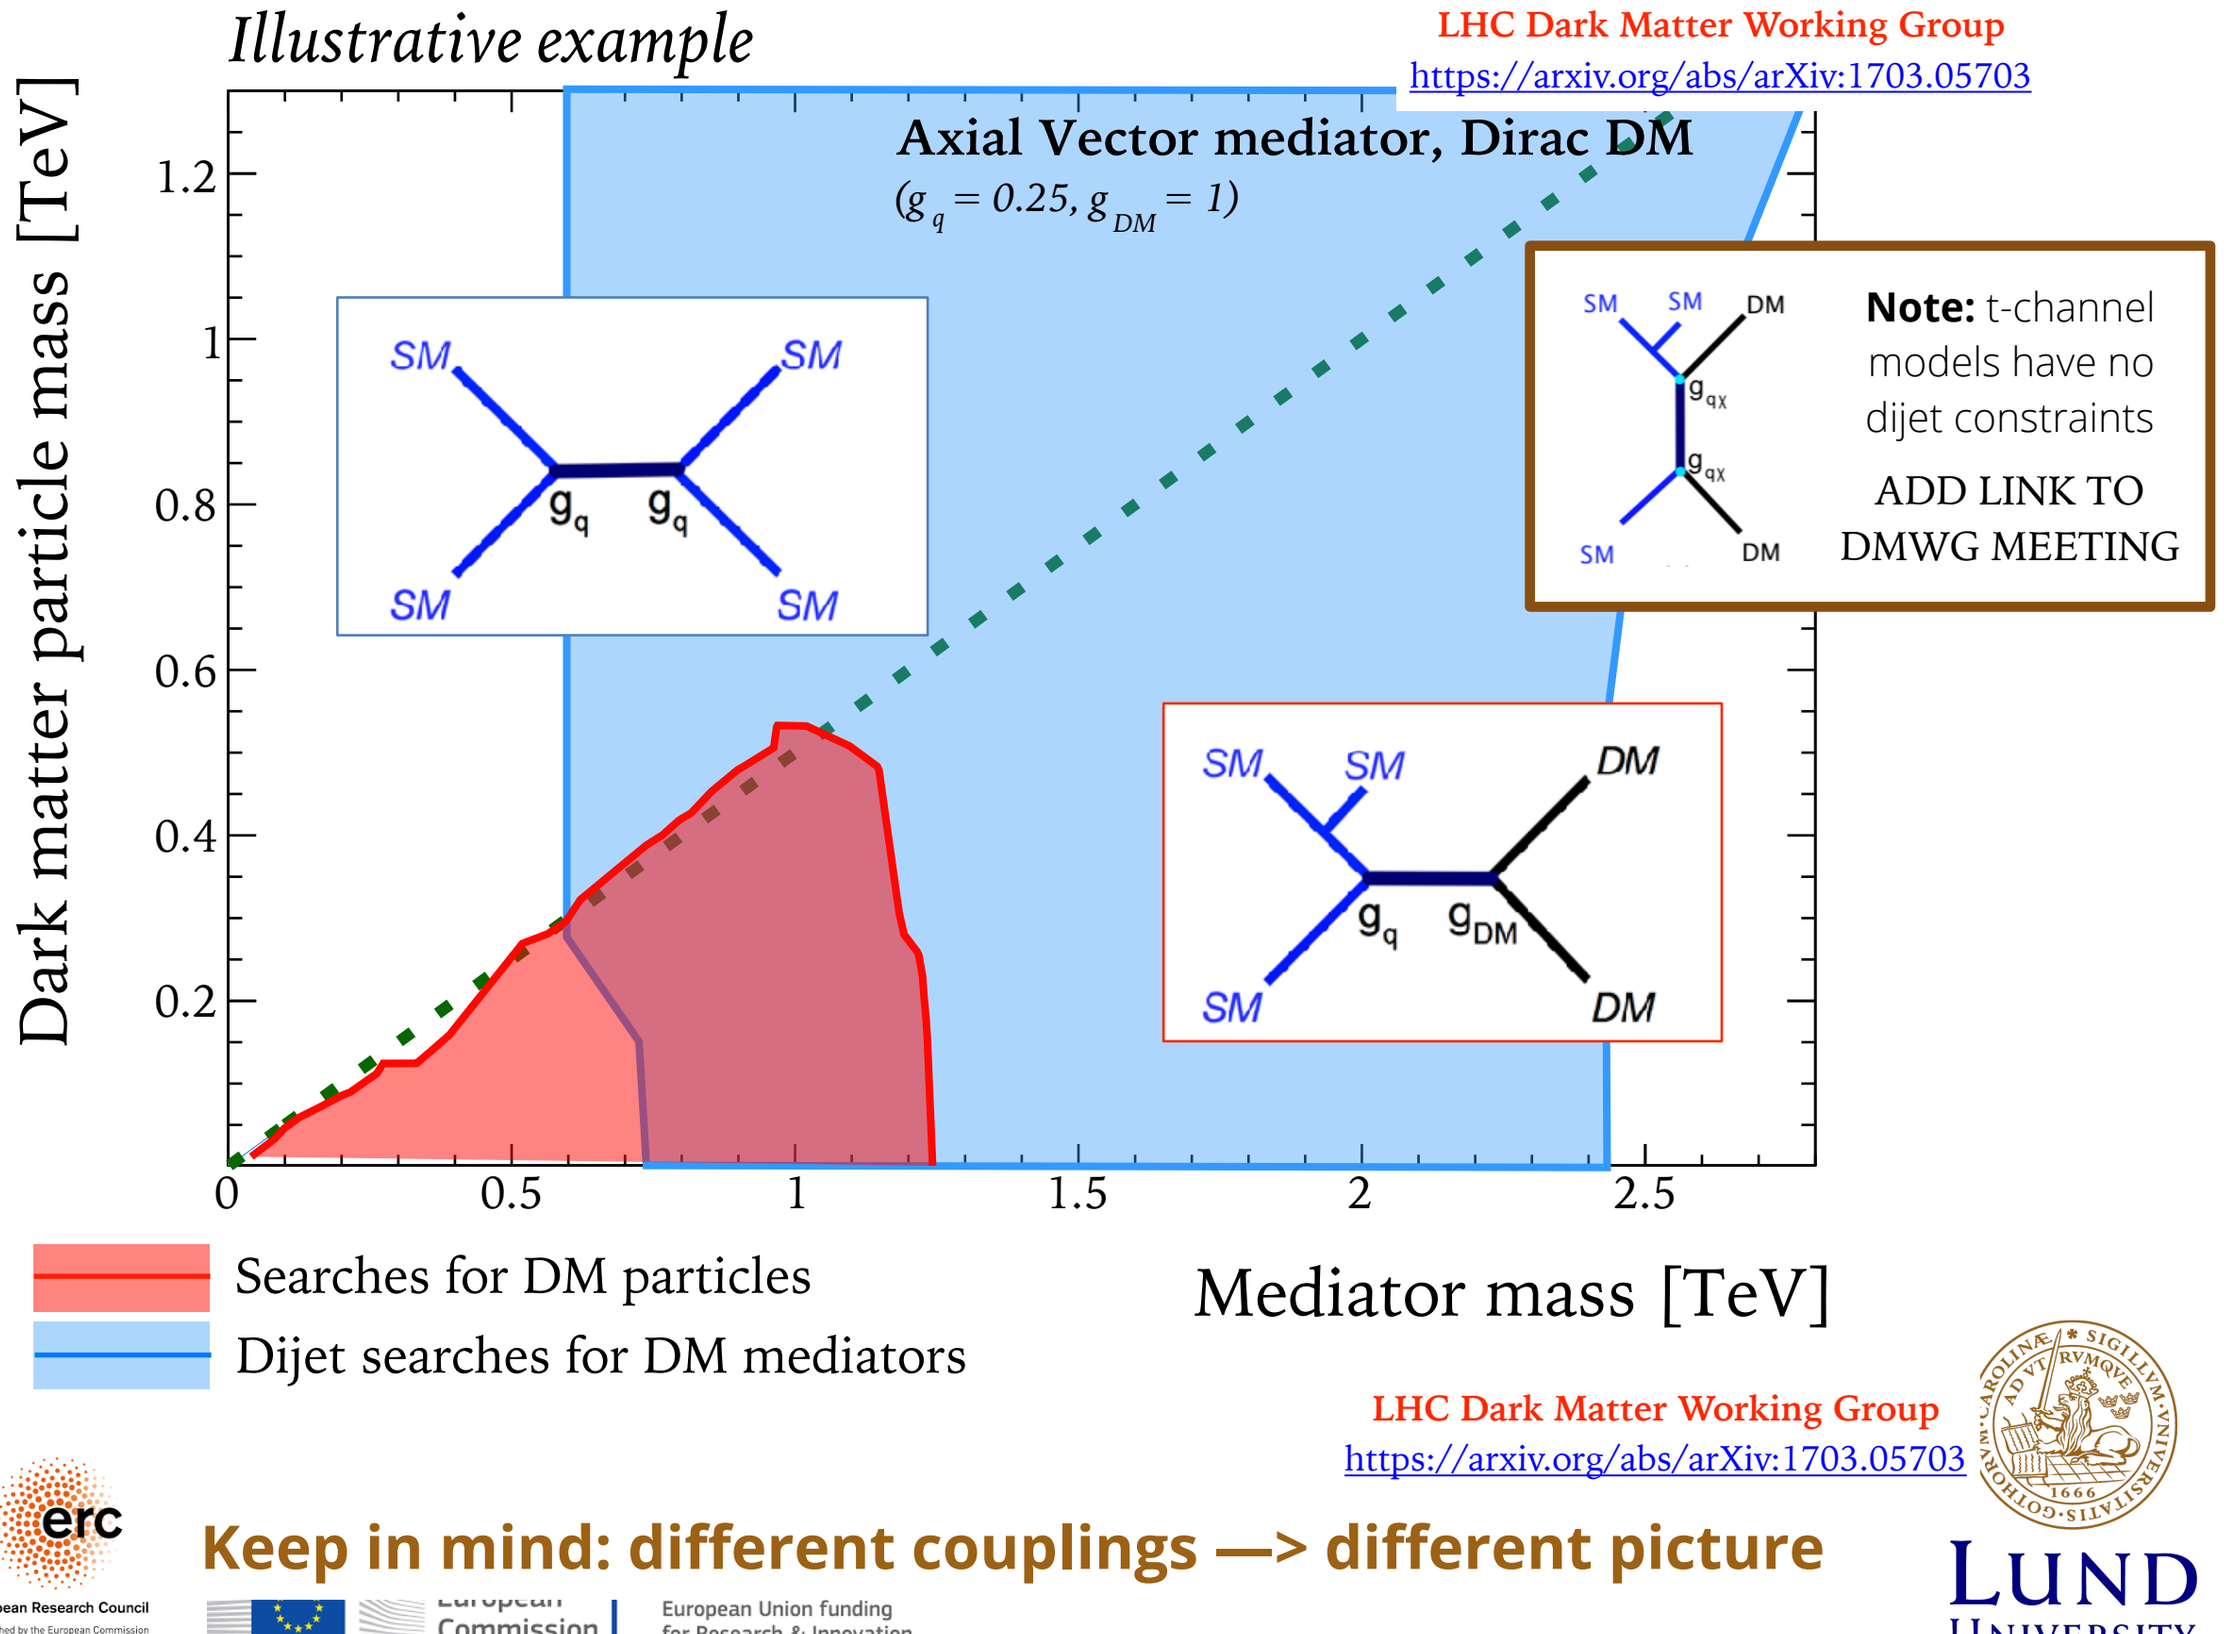

LHC Dark Matter Working Group
https://arxiv.org/abs/arXiv:1703.05703

Keep in mind: different couplings —> different picture

ATLAS EXPERIMENT
European Research Council
Established by the European Commission
European Commission
European Union funding for Research & Innovation
LUND UNIVERSITY

Complementarity of visible/invisible searches

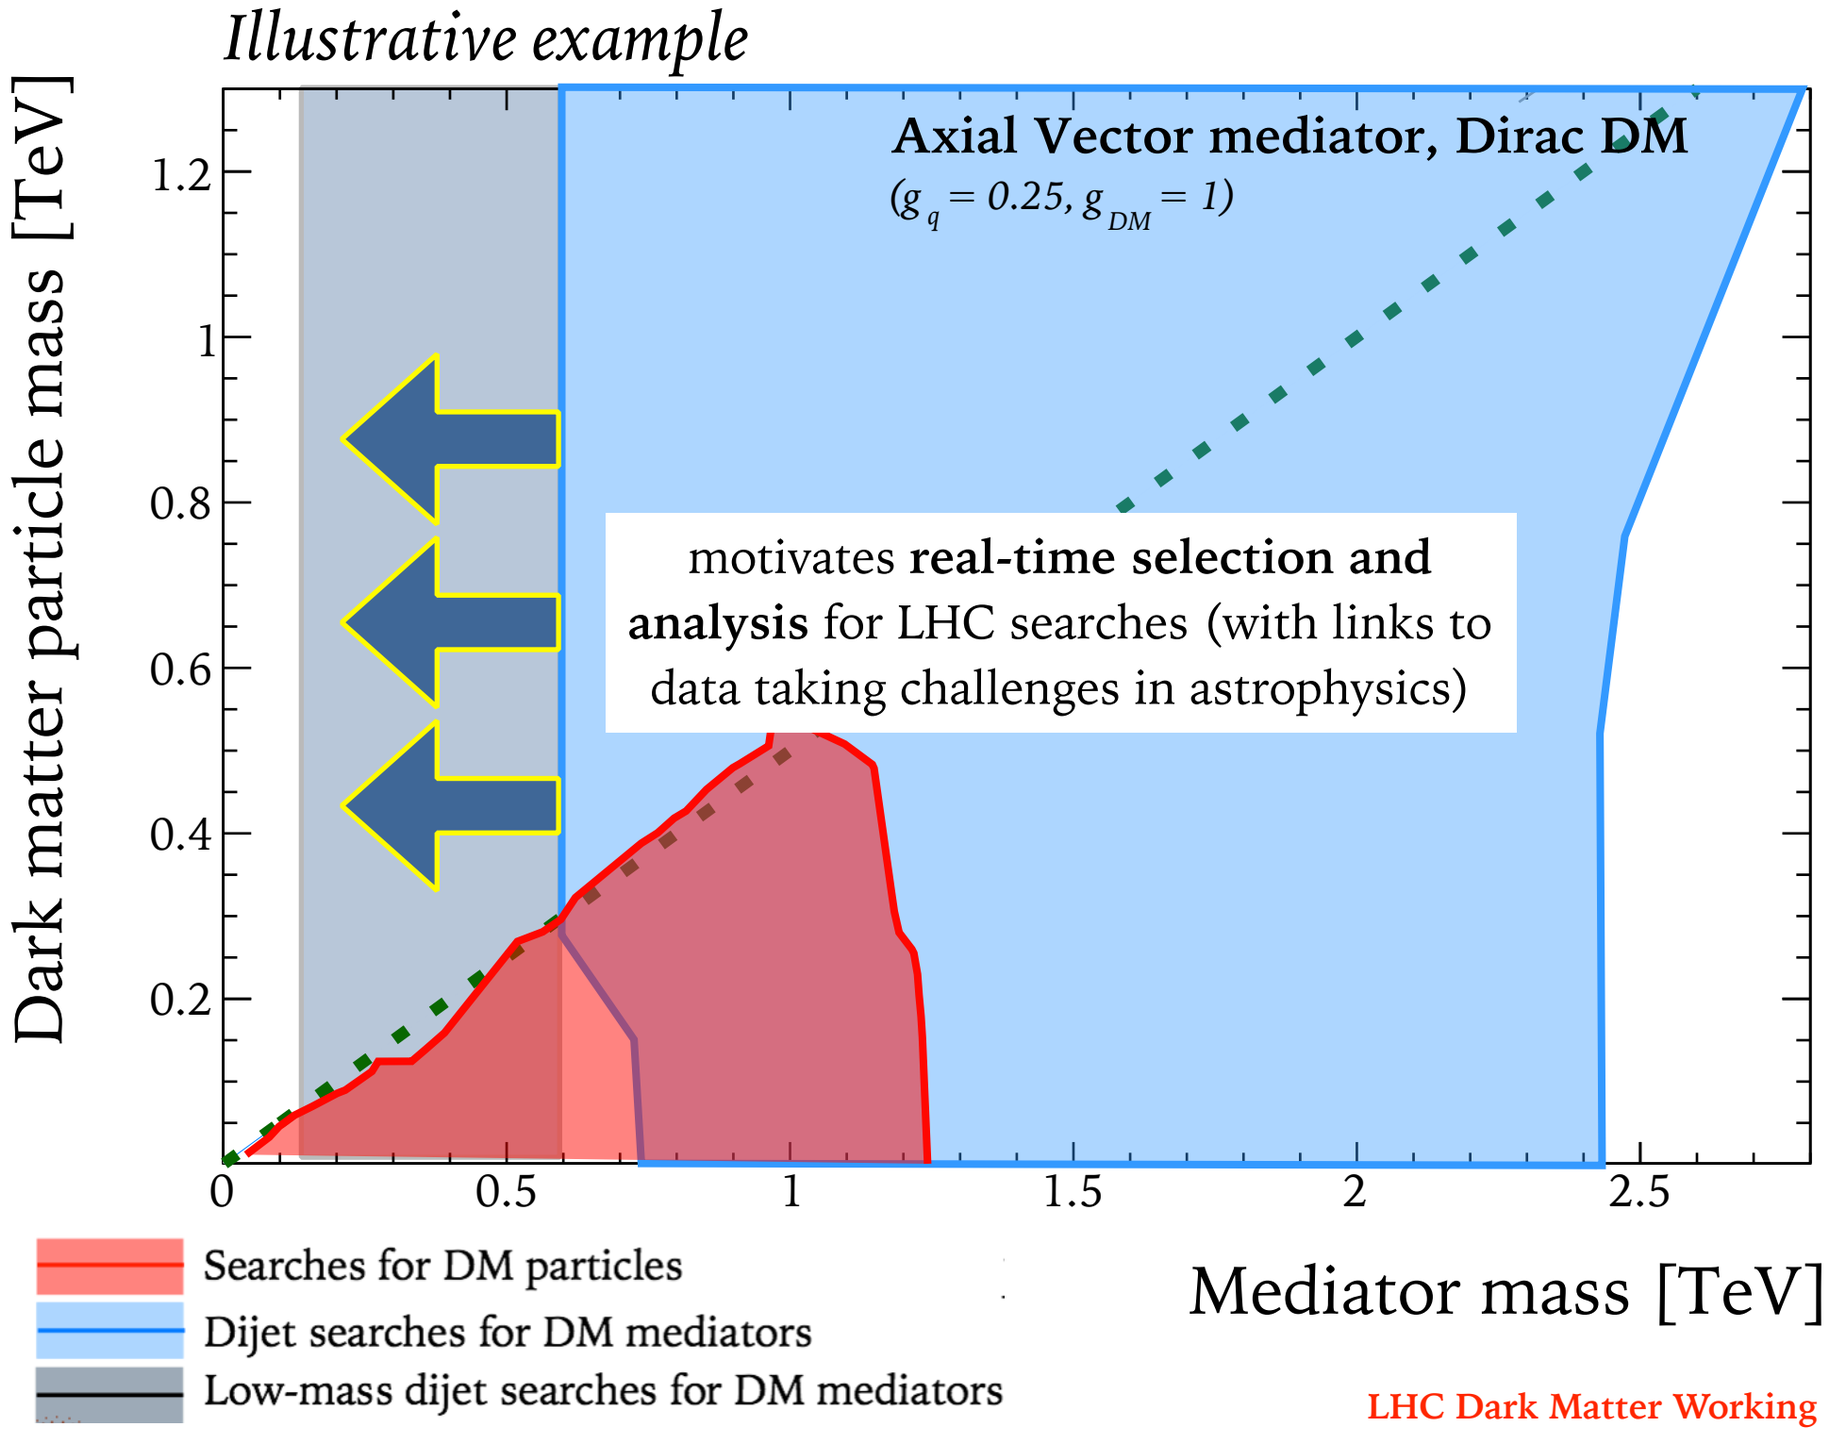

LHC Dark Matter Working Group
https://arxiv.org/abs/arXiv:1703.05703

ATLAS EXPERIMENT
European Research Council
Established by the European Commission
European Commission
Horizon 2020
European Union funding
for Research & Innovation
LUND UNIVERSITY

S-channel A-V mediator: coupling to quarks = 0.25

http://cds.cern.ch/record/2684864?ln=en

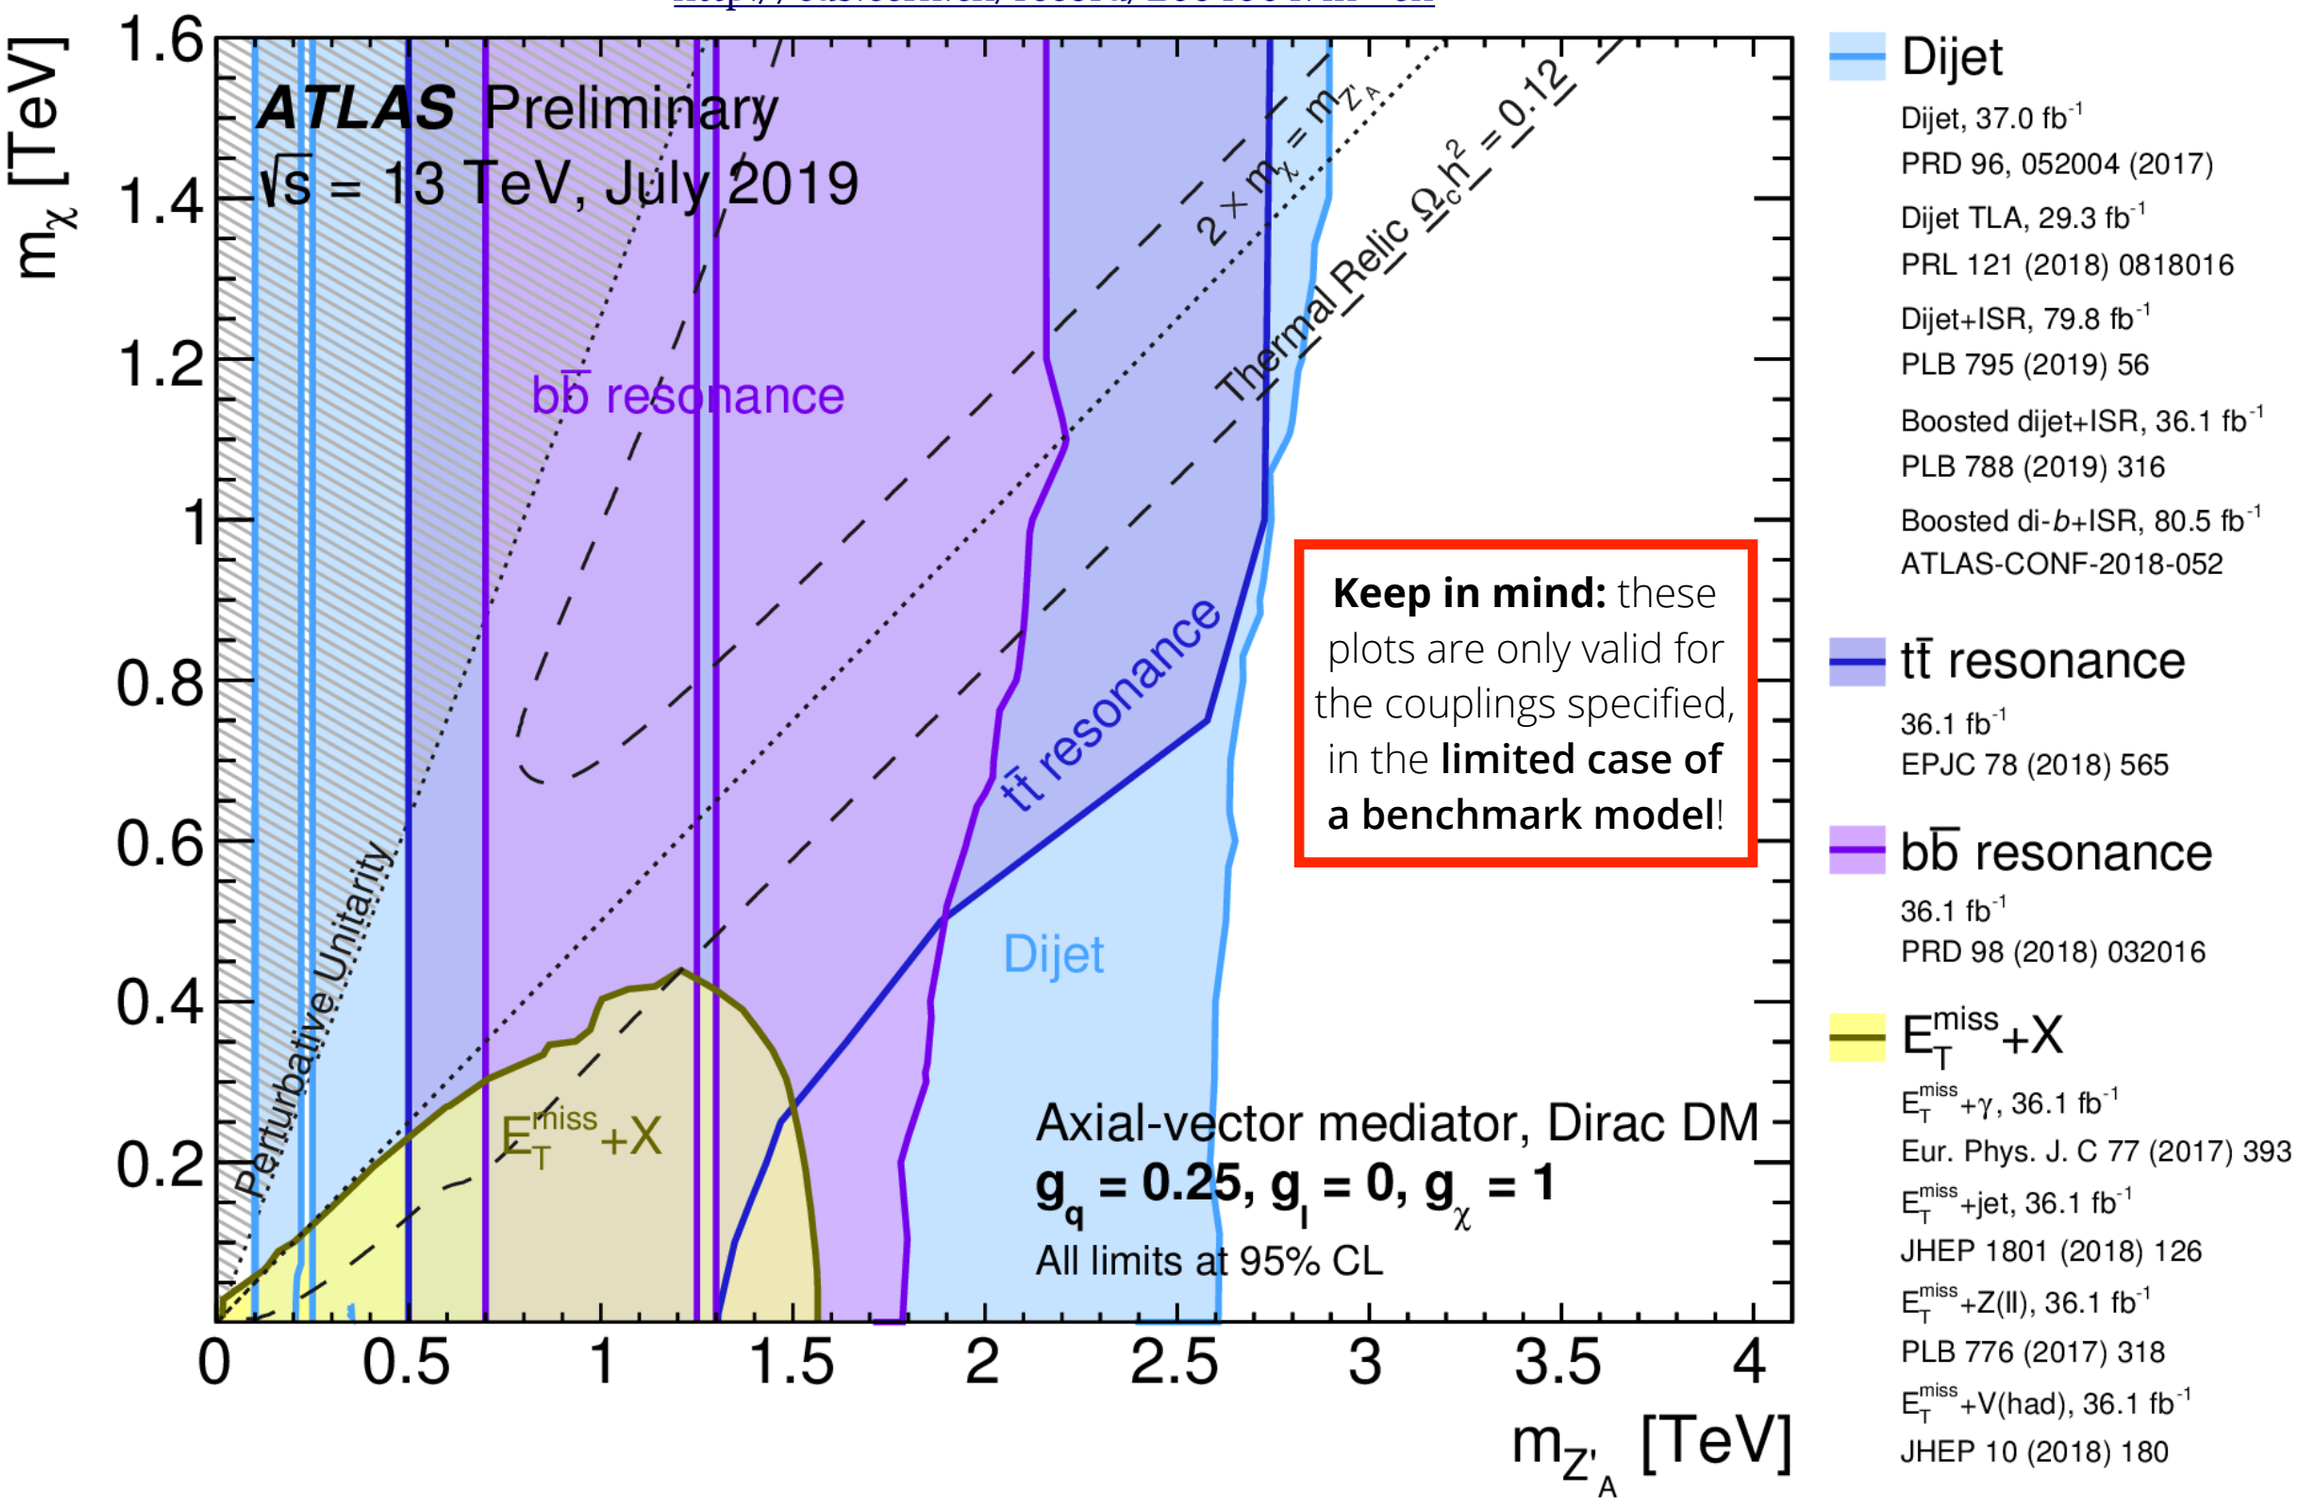

S-channel A-V mediator: coupling to quarks = 0.1

http://cds.cern.ch/record/2684864?ln=en

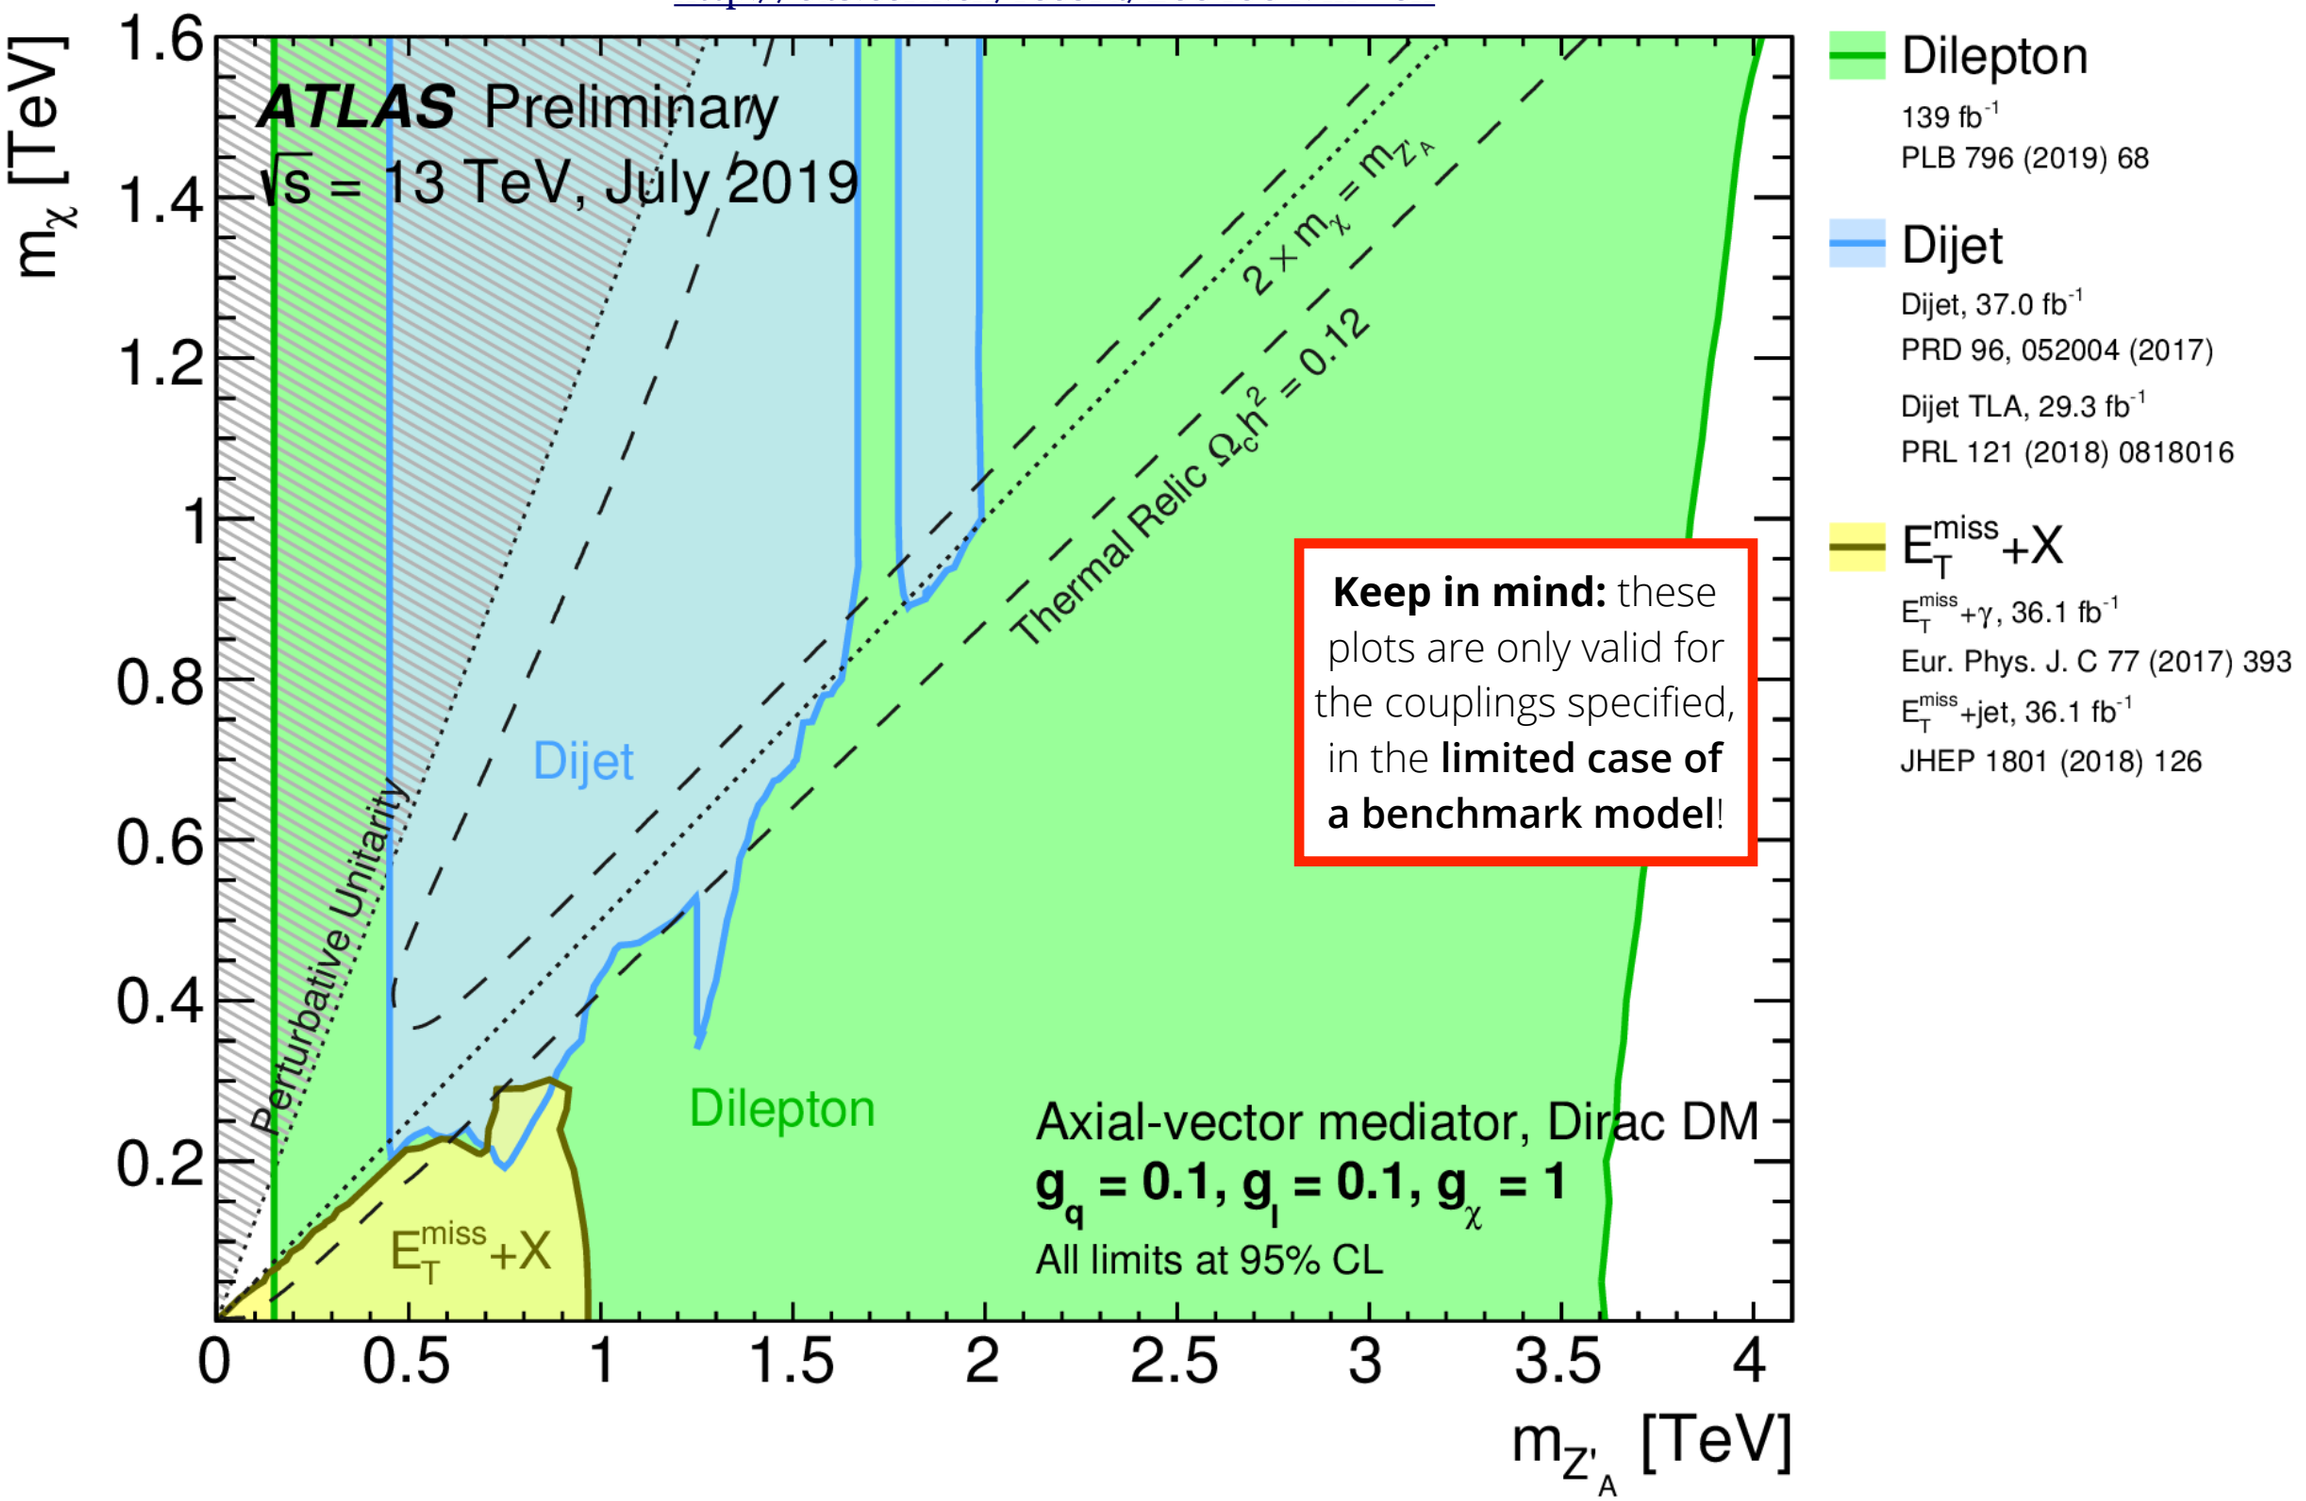

ATLAS/CMS jet calibration (big picture)

What runs where

- The quasi online calibrations
  - For high-level Trigger & express stream calibration
    - Beam spot: based on track- and pixel-only vertexing: one measurement every ~2 min

quasi-real time: 2 min

- Prompt alignment and calibration loop
  - To feed the prompt reconstruction; relies on pre-defined data streams obtained from 0-latency processing ('express')
    - Beam spot, ECAL crystal response -> measured
    - Problematic cells, pixel alignment -> monitored (updated upon need)

PCL: <48 h

- Offline alignment and calibration
  - To feed data reprocessing with ultimate accuracy; full treatment of detectors' alignment inter-dependencies
  - Relies on calibration data streams with dedicated event selection and reduced event content to optimize statistics and resource usage

Offline O(months)

M. Pierini, Institut Pascal RTA workshop 2019