ATLAS EventIndex Trigger Counter

una herramienta de análisis del trigger que emplea MapReduce

Carlos García Montoro (IFIC – CSIC / UV)

on behalf of the IFIC EventIndex Team

Jornadas Técnicas del IFIC

13th March 2018

About me @ IFIC

(24) 16-03-2008 – 15-03-2010: I3P

(3,5) 16-03-2010 – 30-06-2010: TS detector ATLAS.

(2,5) 05-07-2010 – 15-09-2010: TS futuros colisionadores.

(57) 16-09-2010 – 09-06-2015: TS CPAN.

(6,5) 10-06-2015 – 31-12-2015: TS ATLAS Tier2.

(36) 01-06-2016 – 31-05-2019: PTA 2014.

10 Years!!!

of service, work, experience and proved value, though instability continues...

What is the trigger?

- A trigger component represents a relevant condition that an event either satisfies or not:
  - L1: Signals with some threshold at the detector, depending on the detector configuration (SMK).
  - HLT: Software conditions, i.e. logical composition of L1 triggers.
- There are several arrays of trigger components for each level:
  - L1: TAV ⊆ TAP ⊆ TBP. Less than 512 possible trigger components per SMK.
  - HLT: PH, PT, RS. Less than 4096 possible trigger components per SMK.
- The trigger decides whether an event is permanently stored on disk/tape or just discarded.

What is AEI Trigger Counter? What does it do?

- AEI Trigger Counter is a **web service** to provide meaningful information about the L1 and HL Trigger of datasets in a “human time” scale.
- It can **COUNT**, **LIST** and **OVERLAP** a particular trigger of a given dataset
  - Can filter by logical expression based on triggers, LB and/or BCID.
    - Expression is converted into **Reverse Polish Notation** and it is **evaluated**.
  - Can group results per LB and BCID.
- A map-reduce task runs and its results are collected, cached, and offered in up to three different formats:
  - An interactive plot is drawn when not listing.
  - The user can always download a CSV file with the results.
  - The raw data in JSON format is always offered to further process it, as it is done by the plot.

AEI TC Web Service

https://ei-atlas-dc.cern.ch/tc/

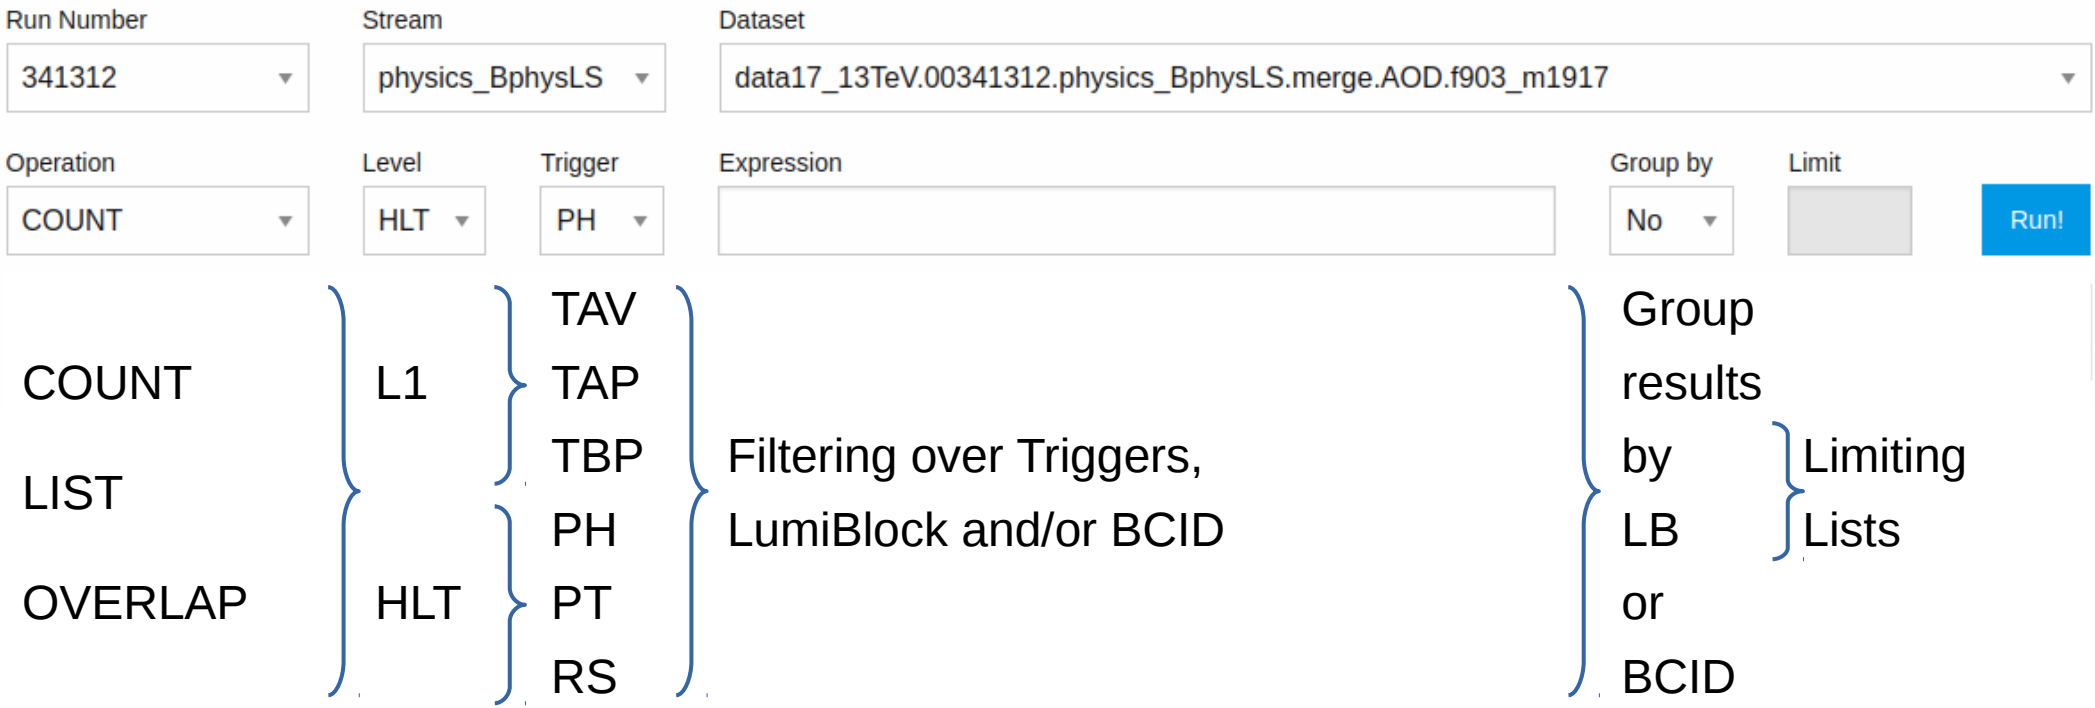

COUNTing Triggers

Results of 8e7f6deb-4180-4d17-884e-0ce058c574b6

COUNT TBP from data17_13TeV.00327103.physics_Main.merge.AOD.f832_m1812 where L1_EM3 or L1_MU21

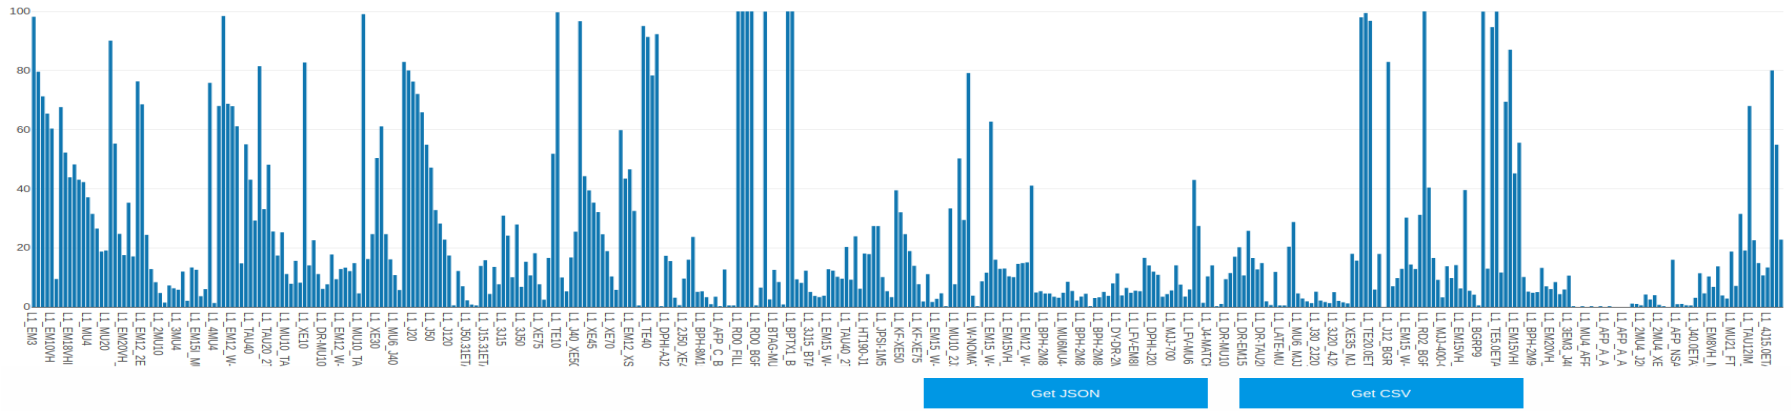

Results of 84807e84-55db-469d-87af-5cc9595d9b59

COUNT TAV from data17_13TeV.00327103.physics_Main.merge.AOD.f832_m1812 where L1_EM3 or L1_MU21 group by LumiBlock

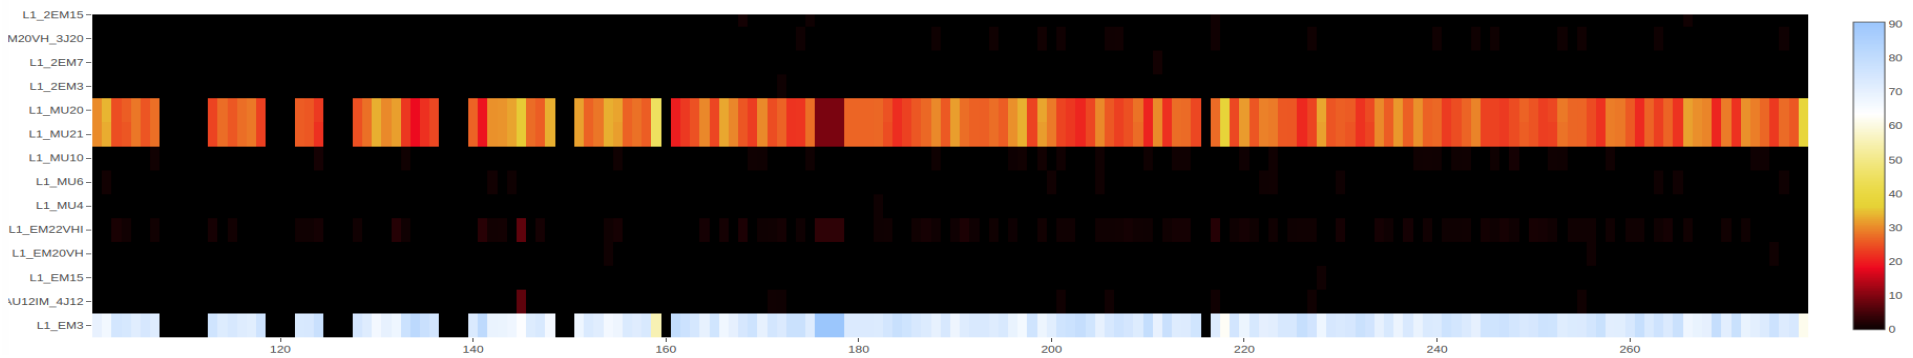

OVERLAPing Triggers

Results of 58668579-dc70-46bd-b685-1dfd87a1b800

OVERLAP TAV from data17_13TeV.00327103.physics_Main.merge.AOD.f832_m1812 where not(L1_EM3 or L1_MU21)

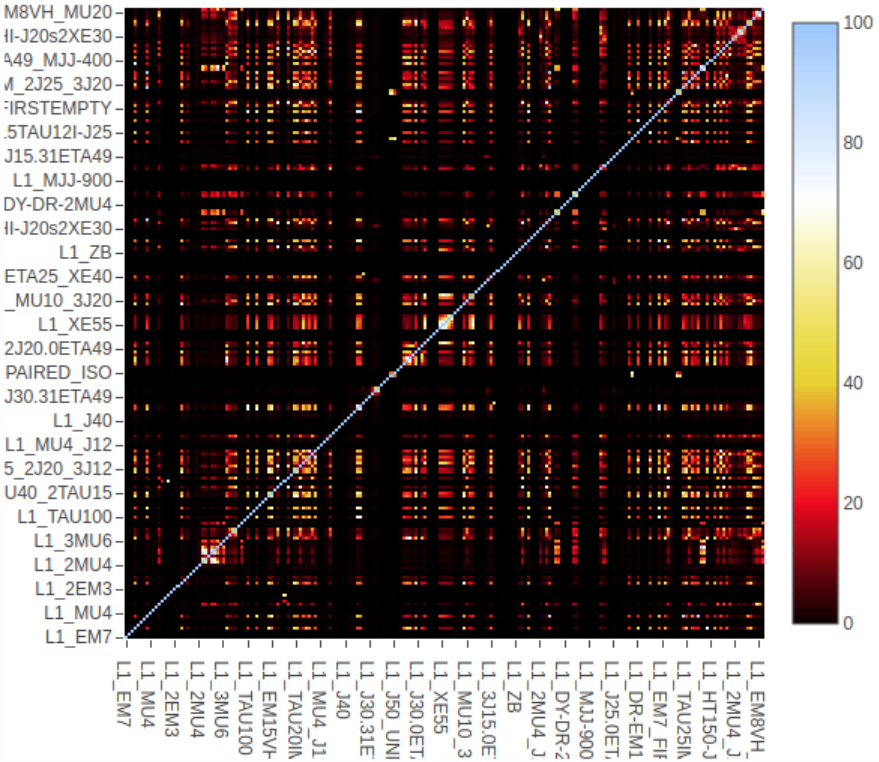

Get JSON

Get CSV

Behind the Scene

- The amount of data to store is huge and the way it is processed should not take long, but be responsive.
- We store trigger information in compressed Sequence Files using some tricks to further reduce their size.
  - **Sequence Files**: containers that contain data in key and value format.
  - The files contain minimal information of the trigger based on its properties, creating a file format which is very compact.
  - Their process can be **highly parallelizable**.
- A **MapReduce** task runs over a Sequence File.
  - Process information in parallel.
  - The proccess occurs **where the information is stored**:
    - No transfers.

What does “big data” means for us?

- Above 10 billion real events.
- Up to 175 million events into a single dataset.
- 1725 datasets that belong to 2017 real data
  - HLT: 367 GB
  - L1: 518 GB
- 375 Runs for several streams:
  - physics_Background
  - physics_BphysLS
  - physics_CosmicCalo
  - physics_Main
  - physics_MinBias
  - physics_ZeroBias

Workflow & Tools

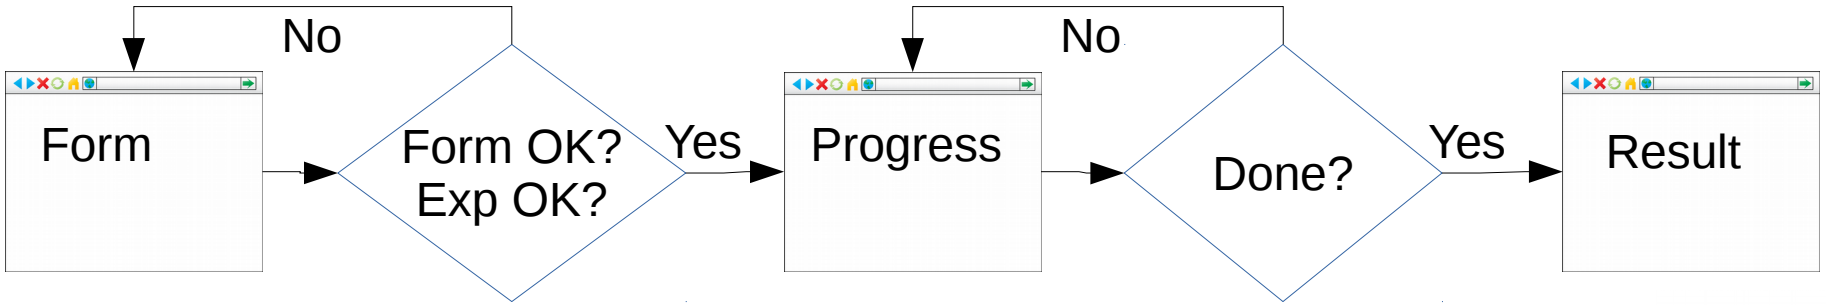

Front-end

Back-end

Flask
web development, one drop at a time

Submit
Progress
Result
SH
SH
SH

hadoop
Map Reduce

Java

SQLAlchemy
SQLite

MapReduce

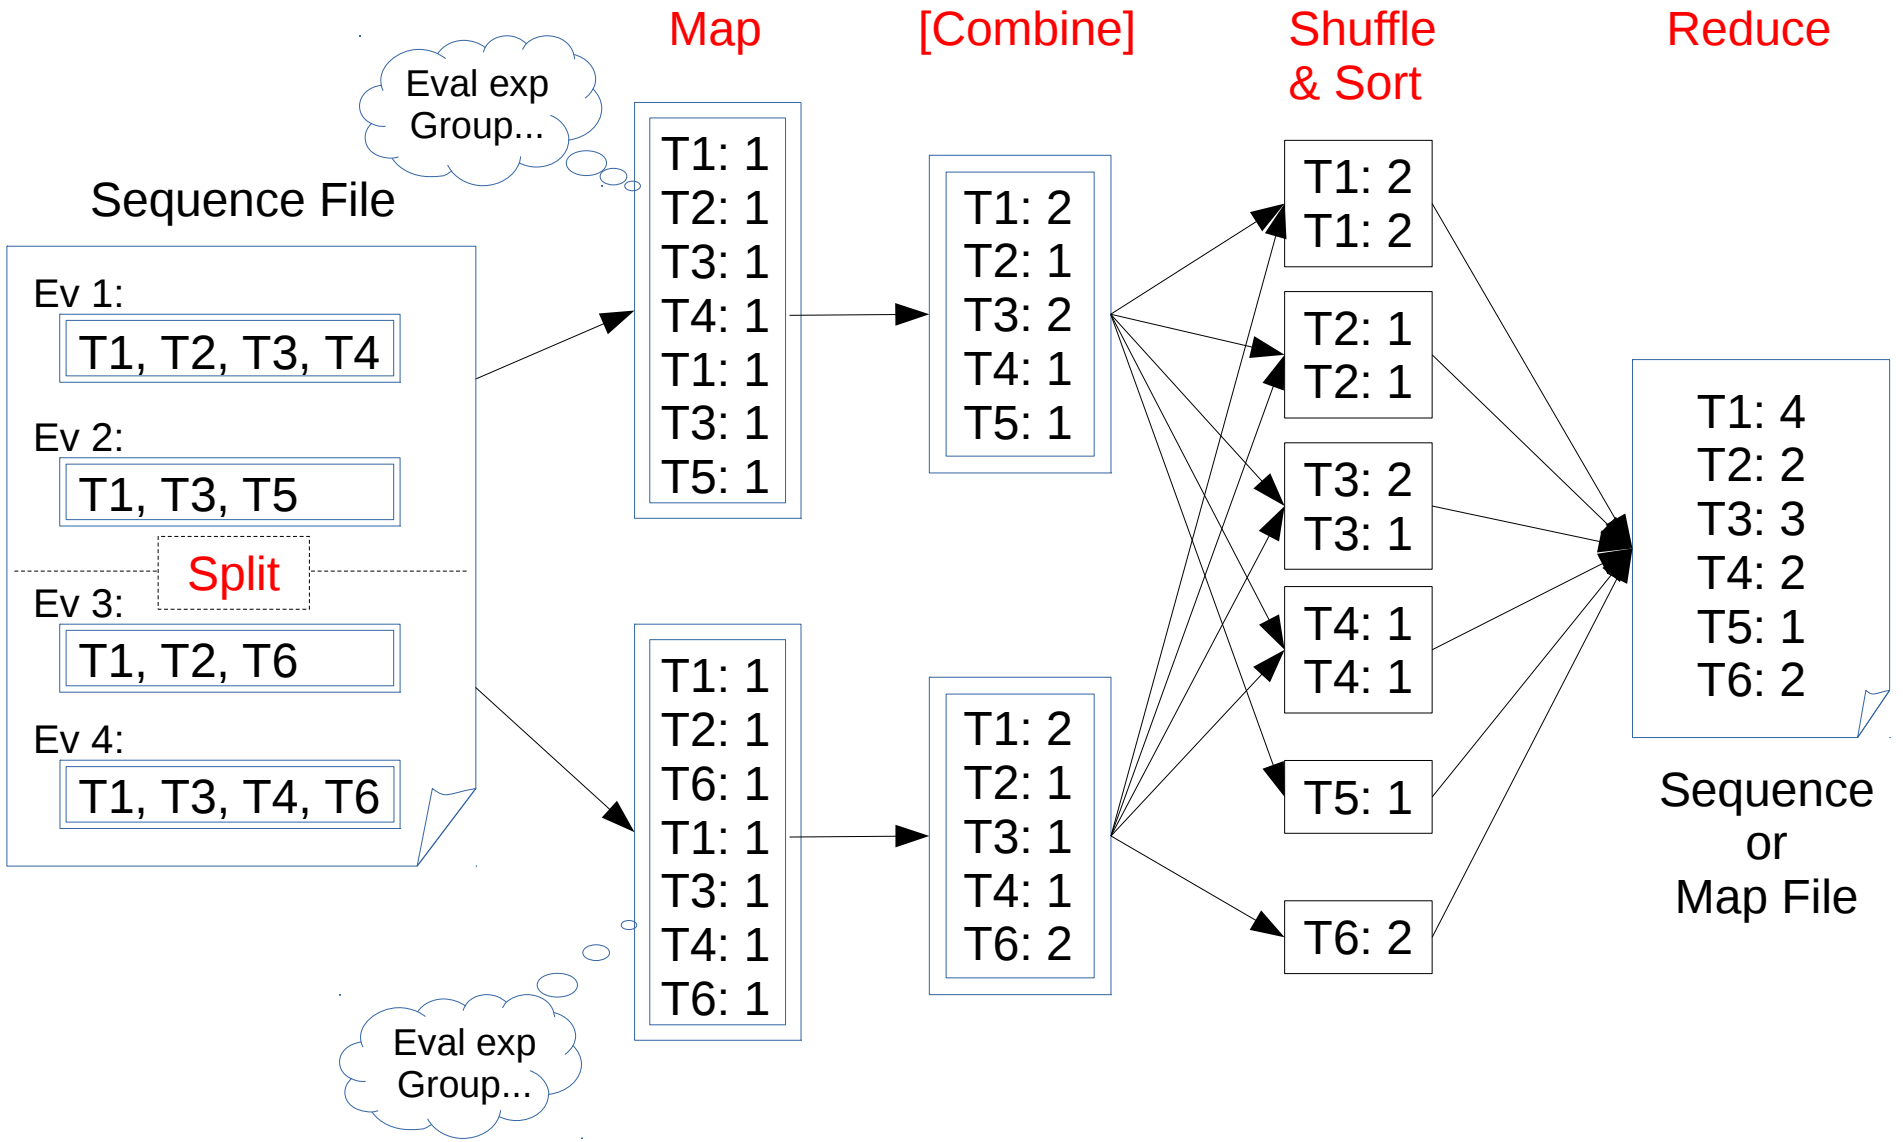
Map
[Combine]
Shuffle & Sort
Reduce
Eval exp Group...
Sequence File
Ev 1:
T1, T2, T3, T4
Ev 2:
T1, T3, T5
Split
Ev 3:
T1, T2, T6
Ev 4:
T1, T3, T4, T6
T1: 1
T2: 1
T3: 1
T4: 1
T1: 1
T3: 1
T5: 1
T1: 1
T2: 1
T6: 1
T1: 1
T3: 1
T4: 1
T6: 1
T1: 2
T2: 1
T3: 2
T4: 1
T5: 1
T1: 2
T2: 1
T3: 1
T4: 1
T6: 2
T1: 2
T1: 2
T2: 1
T2: 1
T3: 2
T3: 1
T4: 1
T4: 1
T5: 1
T6: 2
T1: 4
T2: 2
T3: 3
T4: 2
T5: 1
T6: 2
Sequence or Map File
Eval exp Group...

HDFS L1Trigger File Format

SEQ file with:

- Key: <LongWritable> Value: <BytesWritable>

EventNumber	LB	BCID	L1Trigger

Less than 512 possible triggers, i.e. enumerating triggers, a trigger can be represented as 9 bits

- LB and BCID are uint32
- L1Trigger is a collection of packed 10 bits triggerNumbers into 32 bit words:

#TBP'	#TAP'	#TAV'	trig3	trig2	trig1	…		trigN	trig(N-1)

- Since TAV ⊆ TAP ⊆ TBP :

TAV' = TAV
TAP' = TAP - TAV
TBP' = TBP - TAP

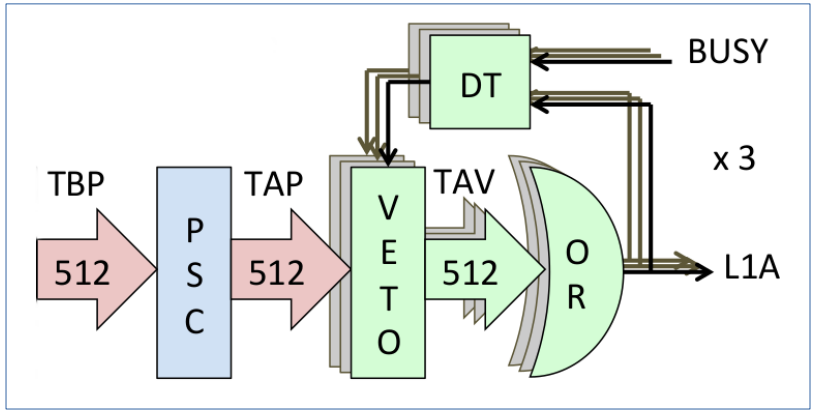

HDFS HLTrigger File Format

SEQ file with:

- Key: <LongWritable> Value: <BytesWritable>

Less than 4096 possible triggers, i.e. enumerating triggers, a trigger can be represented as 12 bits

- LB and BCID are uint32
- HLTrigger is a collection of packed 12 bits triggerNumbers into 64 bit words:

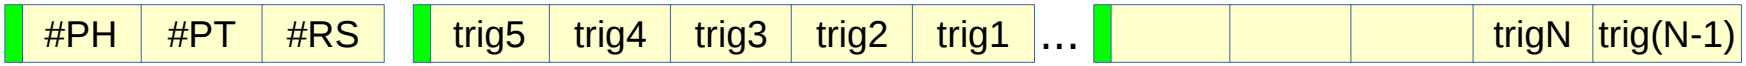

  - There is no inclusion between HLT sets, therefore all triggers have to be stored

Backup Slides

Trigger Counter Web Interface

Run Number: 334955 | Stream: physics_Backgrour | Dataset: data17_13TeV.00334955.physics_Background.merge.AOD.f863_m1865

Operation: COUNT | Trigger: TAV | Expression: | Group by: Nothing | Limit: | Run!

Show/Hide Triggers

0: L1_EM3	1: L1_EM7	2: L1_EM12	3: L1_EM8VH
4: L1_EM10VH	5: L1_EM15VHI_2TAU12IM_4J12	6: L1_EM15	7: L1_EM15VH
8: L1_EM18VHI	9: L1_EM20VH	10: L1_EM20VHI	11: L1_EM22VHI
12: L1_EM3_EMPTY	13: L1_EM7_EMPTY	14: L1_MU4	15: L1_MU6
16: L1_MU10	17: L1_MU21	18: L1_MU20	19: L1_MU4_EMPTY
20: L1_MU4_FIRSTEMPTY	21: L1_MU11_EMPTY	22: L1_MU4_UNPAIRED_ISO	23: L1_2EM3
24: L1_2EM7	25: L1_2EM10VH	26: L1_EM20VH_3J20	27: L1_2EM15
28: L1_2EM15VH	29: L1_EM7_2EM3	30: L1_EM12_2EM3	31: L1_EM15VH_3EM7
32: L1_2MU4	33: L1_2MU6	34: L1_2MU10	35: L1_2MU20_OVERLAY
36: L1_MU10_2MU6	37: L1_MU11_2MU6	38: L1_3MU4	39: L1_MU6_2MU4
40: L1_3MU6	41: L1_4J15.0ETA25	42: L1_EM15I_MU4	43: L1_2EM8VH_MU10
44: L1_EM15VH_MU10	45: L1_TAU12	46: L1_4MU4	47: L1_TAU12IM
48: L1_TE15.0ETA24	49: L1_TAU20	50: L1_EM12_W-MT25	51: L1_TAU20IM
52: L1_TAU100	53: L1_TAU30	54: L1_TAU40	55: L1_TAU60
56: L1_TAU8	57: L1_TAU8_EMPTY	58: L1_TAU20IM_2TAU12IM	59: L1_TAU20_2TAU12
60: L1_EM15VHI_2TAU12IM	61: L1_EM15VHI_2TAU12IM_J25_3J12	62: L1_EM15VHI_TAU40_2TAU15	63: L1_MU10_TAU12IM
64: L1_MU10_TAU12IM_J25_2J12	65: L1_EM7_MU10	66: L1_MU6_EMPTY	67: L1_MU10_TAU20IM
68: L1_XE10	69: L1_EM20VHI_TAU20IM_2TAU20_J25_3J20	70: L1_TAU20IM_2TAU12IM_J25_2J20_3J12	71: L1_J25_2J20_3J12_BOX-TAU20ITAU12I
72: L1_DR-MU10TAU12I_TAU12I-J25	73: L1_MU10_TAU12I-J25	74: L1_TAU20IM_2J20_XE45	75: L1_J15.31ETA49_UNPAIRED_ISO
76: L1_TAU60_DR-TAU20ITAU12I	77: L1_EM12_W-MT35_XS30_W-15DPHI-JXE-0_W-15DPHI-EMXE	78: L1_EM15VHI_2TAU12IM_XE35	79: L1_EM12_W-MT35_XS40_W-05DPHI-JXE-0_W-05DPHI-EMXE
80: L1_EM12_W-MT35_W-90RO2-XEHT-0_W-15DPHI-JXE-0_W-15DPHI-EMXE	81: L1_MU10_TAU12IM_XE35	82: L1_TE50	83: L1_TAU20IM_2TAU12IM_XE35

Trigger Counter Web Interface: Progress of a Task

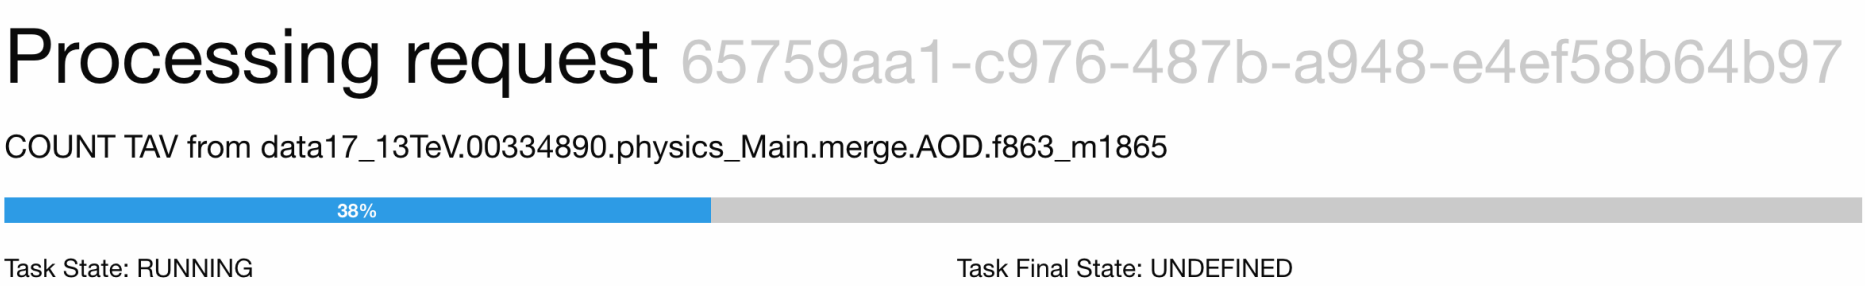

- User has output of the progress of its request, which refresh every 10 seconds.
- Map-reduce ‘state’ and ‘finalStatus’ are shown.
- If the task succeeds, the page redirects to the results.
- Trigger Counter caches the tasks, avoiding multiple runs of the same task.

Trigger Sizes and Compression

Typical trigger counts and final Lz4 size in bytes per event:

dsname	TBP/e	TAP/e	TAV/e	TBP'/e	TAP'/e	TAV'/e	totTrig	lz4
data17_13TeV.00331951.physics_Background.merge.AOD.f848_m1844	13.1	10.6	1.1	2.4	9.6	1.1	13.1	15.6
data17_13TeV.00333994.physics_BphysLS.merge.AOD.f859_m1860	62.3	45.3	6.4	17.1	38.8	6.4	62.3	53.6
data17_13TeV.00324839.physics_CosmicCalo.merge.AOD.f823_m1799	7.2	4.6	1	2.6	3.6	1	7.2	11.2
data17_13TeV.00335302.physics_EnhancedBias.merge.AOD.f869_m1870	49.5	34	3	15.5	31	3	49.5	46.4
data17_13TeV.00326923.physics_Late.merge.AOD.f832_m1812	11.4	7.8	3.8	3.7	4	3.8	11.5	12.4
data17_13TeV.00327761.physics_Main.merge.AOD.f832_m1812	84.6	57.8	8.8	26.8	49	8.8	84.6	67.6
data17_13TeV.00334191.physics_MinBias.merge.AOD.f859_m1860	9.4	5.4	1	4	4.4	1	9.4	12
data17_13TeV.00324857.physics_ZeroBias.merge.AOD.f823_m1799	21.9	20.4	1	1.4	19.4	1	21.8	23.2

Compression ratios:

- Deflate 3.0
- Bzip2 4.5
- Lz4 1.9
- Snappy 1.8

⇨ LZ4 has been chosen because:
- It offers a good compression ratio.
- It is fast on decompression.