DAQ! TRG! PHYSICS! OMG!

O. Hartbrich
(University of Hawaii at Manoa)

B2SS 20201
07/16/2021

Belle II DAQ Components

- The components of the DAQ system:
  - Trigger and Timing Distribution
  - Data Readout and Event Builder
  - Slow Control
- What you need to know about triggers

- This talk heavily relies on materials an input I got (and took) from others, especially
  - Chris Hearty
  - Ewan Hill
  - Michel Hernandez Villanueva

DAQ Is Fun!

- We have all the cables. And all the complicated diagrams.

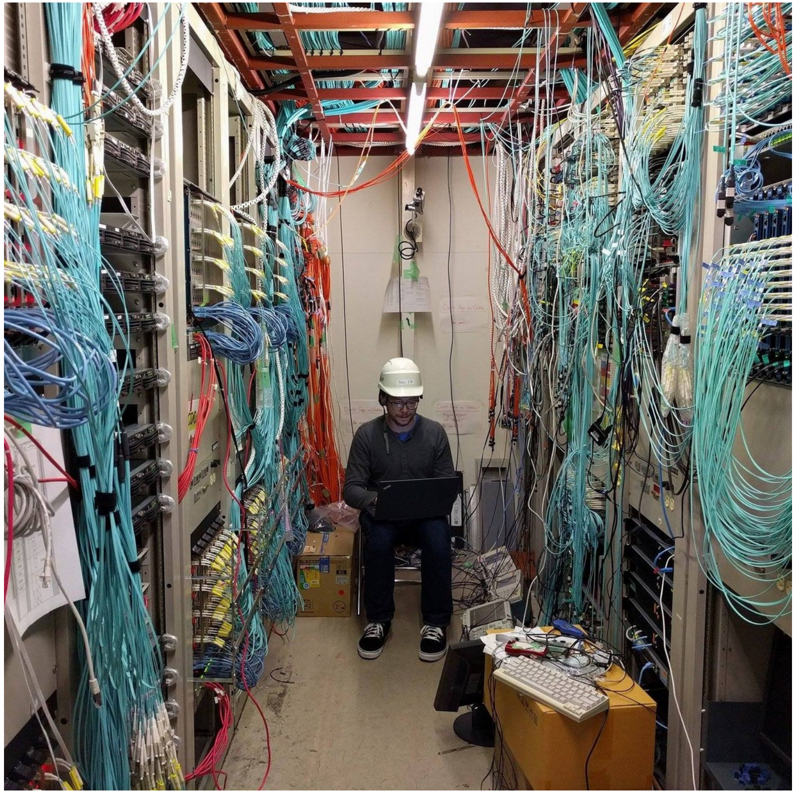
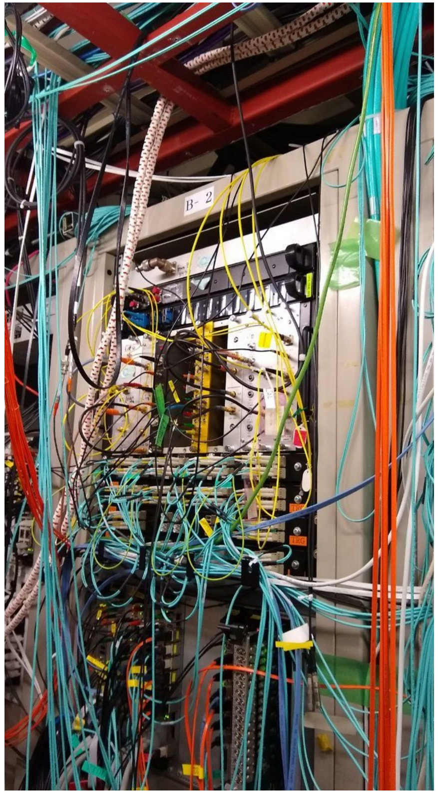

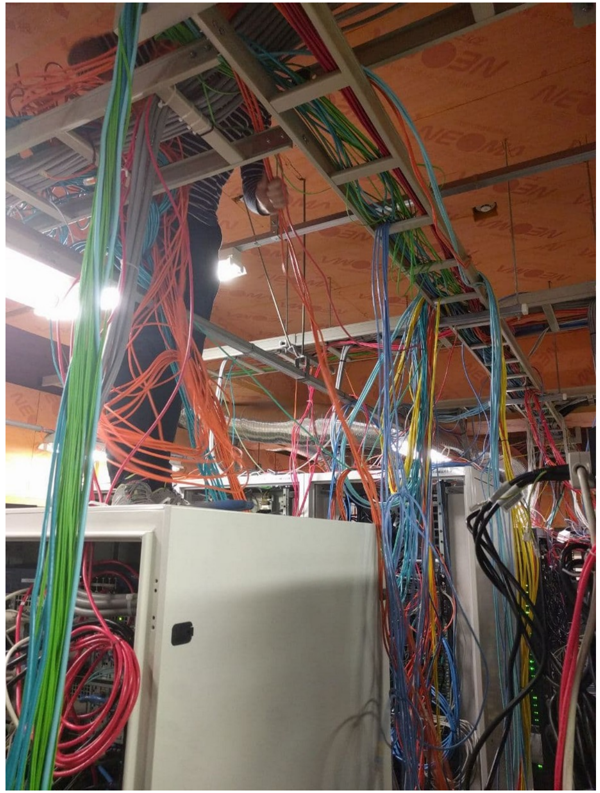

DAQ Is Fun

- We have a … omplicated diagrams.

Data Transfer

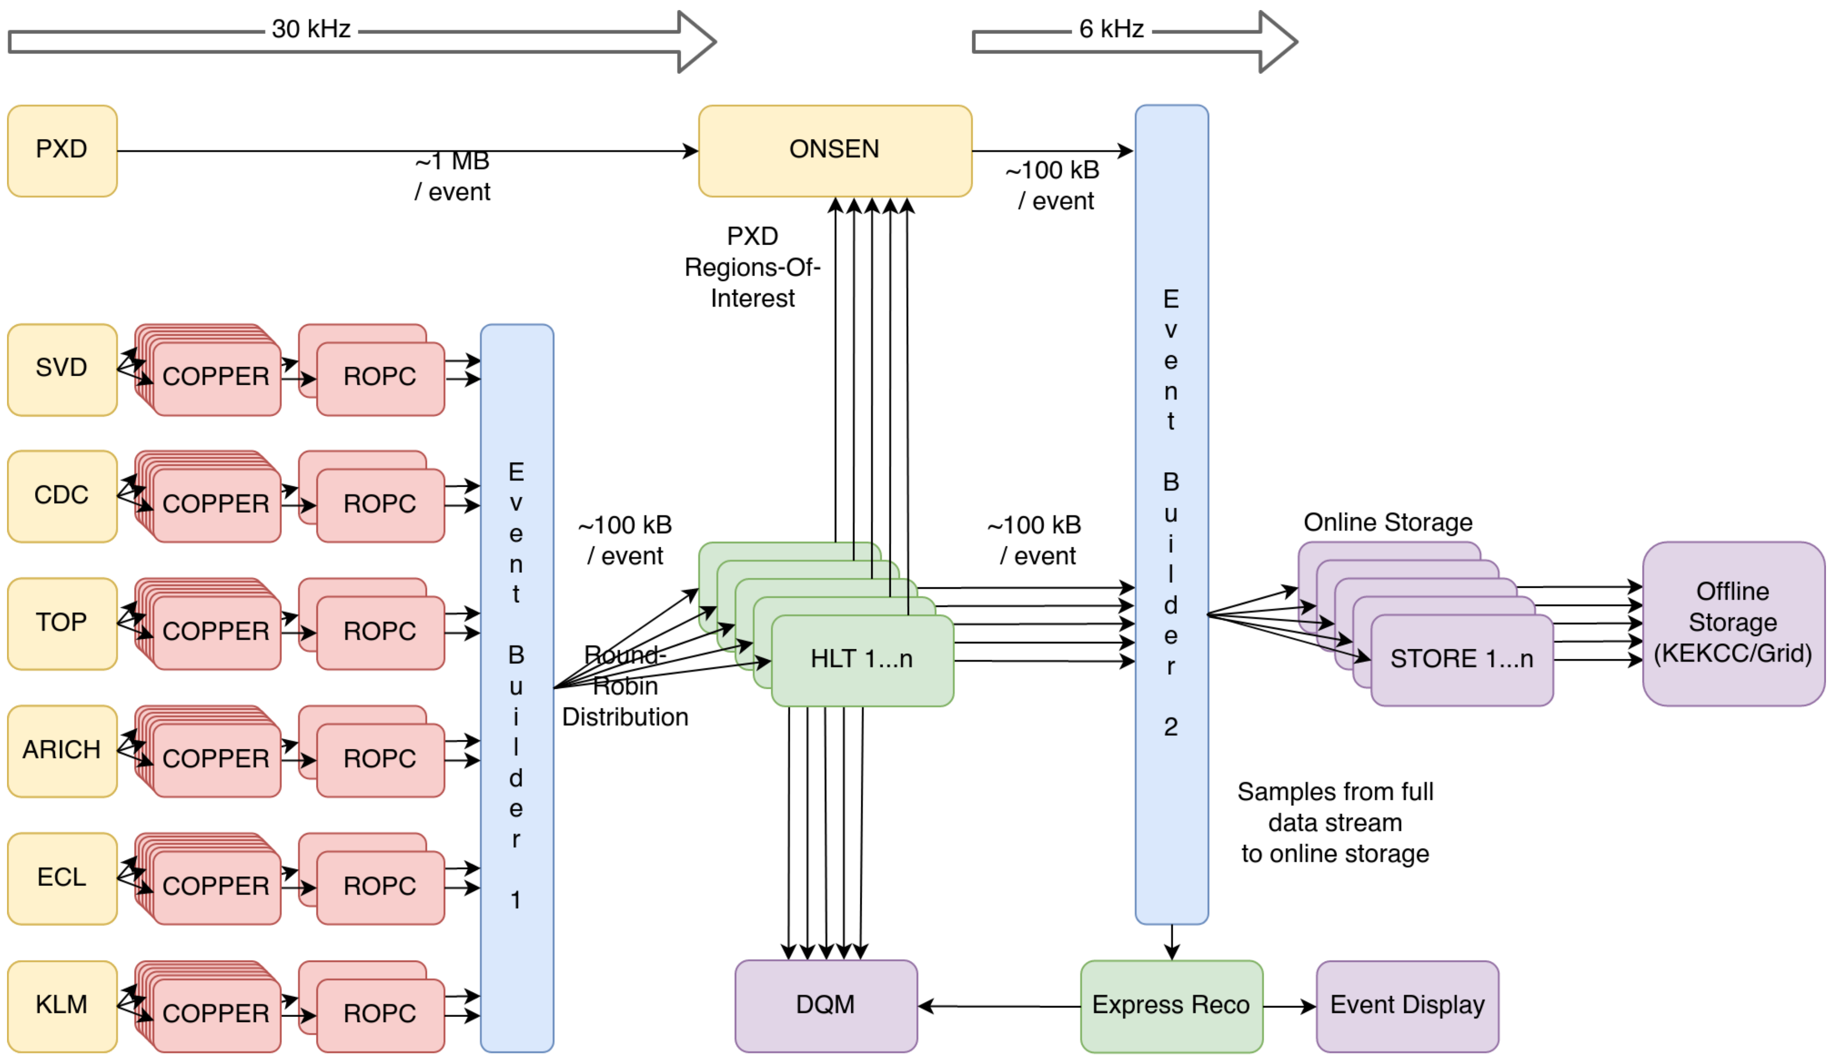

Copper + ROPC

- COPPER (Common Pipelined Platform for Electronics Readout)
  - Receives up to four fibers from front-ends (up to 2.54Gbps each)
  - Integrated Atom CPU board for data packaging, checksums etc.
  - **~200** COPPERs serving **~650** Belle2links
- ROPC (Readout PCs)
  - Receive GbE data from 2-9 COPPERS, forward to event builder via GbE
  - Acts as network boot host for COPPERs
  - **~45** ROPCs for whole Belle II

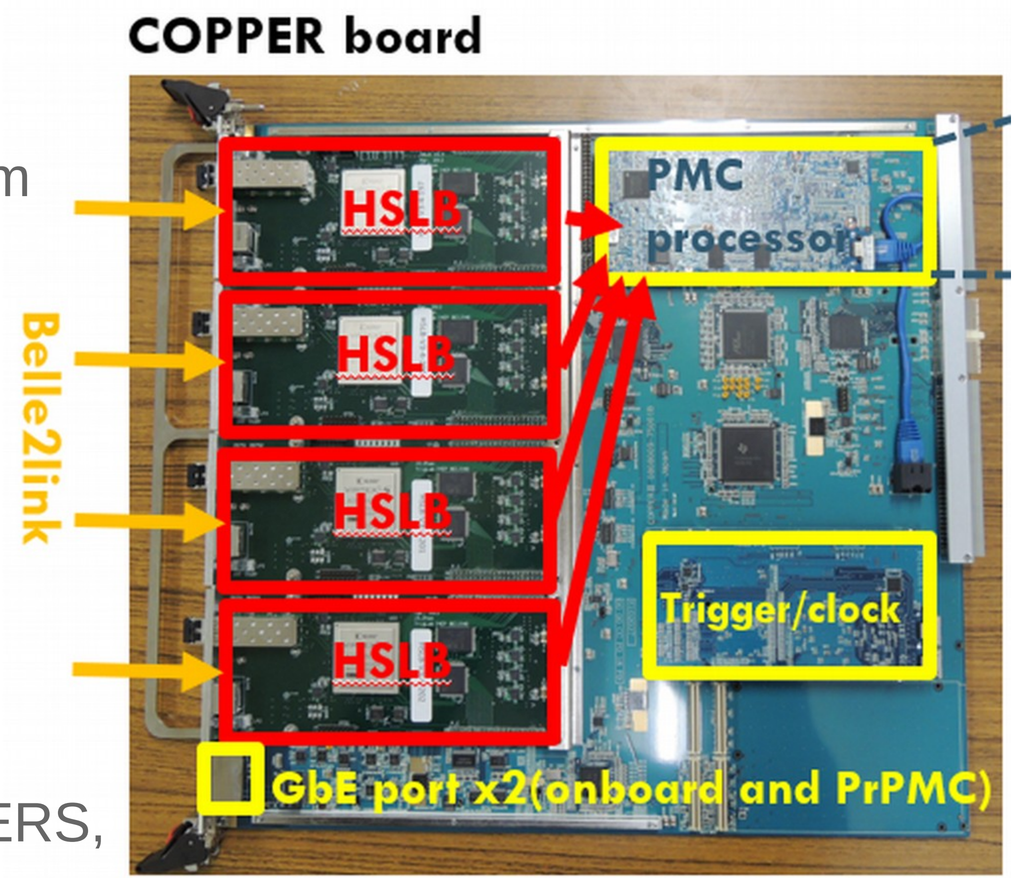

Belle II DAQ Upgrade Project

- COPPER system will be difficult to maintain over the lifetime of Belle II
  - Relatively old Atom CPUs, number of discontinued parts increasing
  - Data rate capabilities are marginal for full luminosity
- Upgrade: “plug-in” replacement for COPPER + ROPCs
  - No changes to front-end links, trigger distribution, HLTs etc.
  - Significant increase in link density, reduction in rack space
- Several technology contenders, selected PCIe40
  - 48 links per PCIe40 card → 19 cards for whole Belle II
  - Whole data readout infrastructure will fit into one rack!
- After some COVID delays and extensive testing,
TOP and KLM will be converted to full PCIe40 this summer

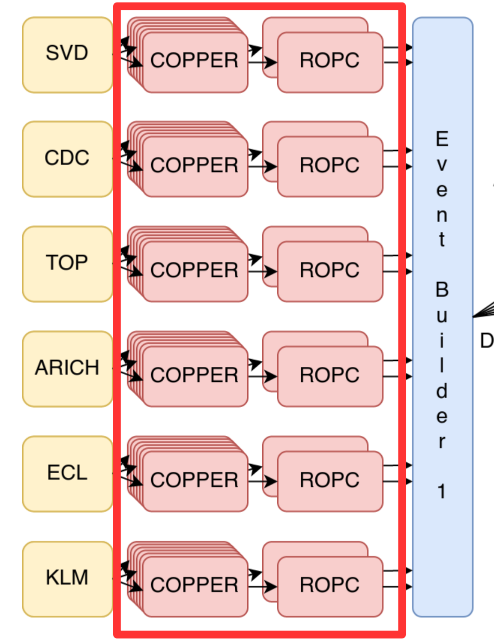

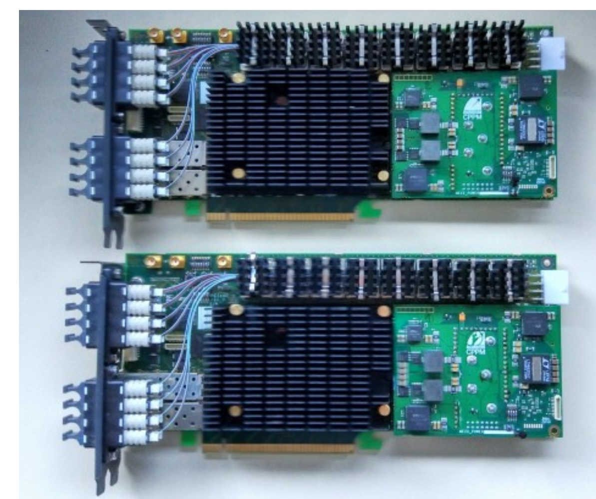

Slow Control

- Non-event process variables (PVs): run control, detector status, machine parameters, etc.
  - If you click “START”, a lot of systems have to react in some way
- Hybrid system for PV distribution: nsm2 and EPICS
- Tens of thousands of Pvs distributed over hundreds of inhomogenous nodes
  - not all of them on the same network
- Introduced continuous integration and (almost) continuous deployment of slow control software

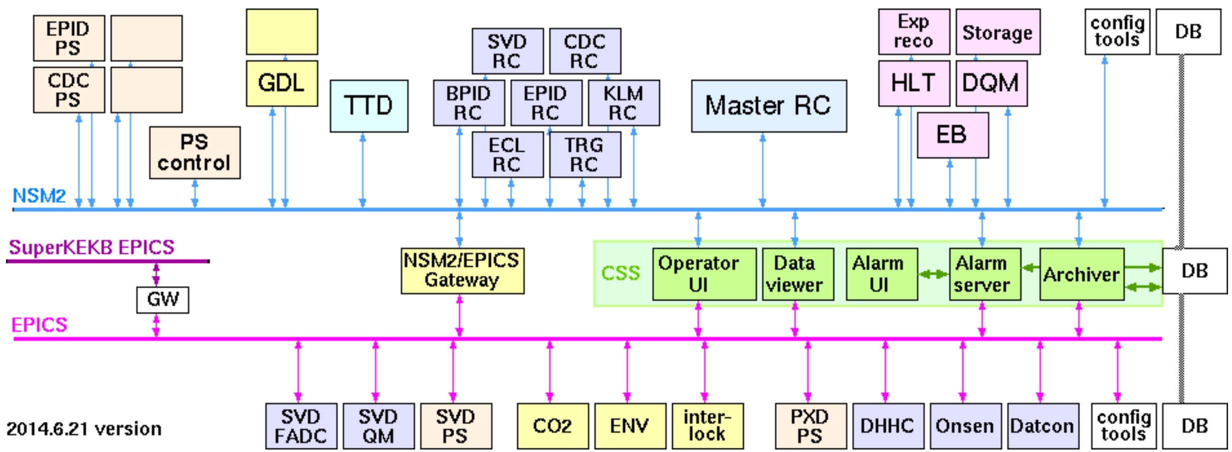

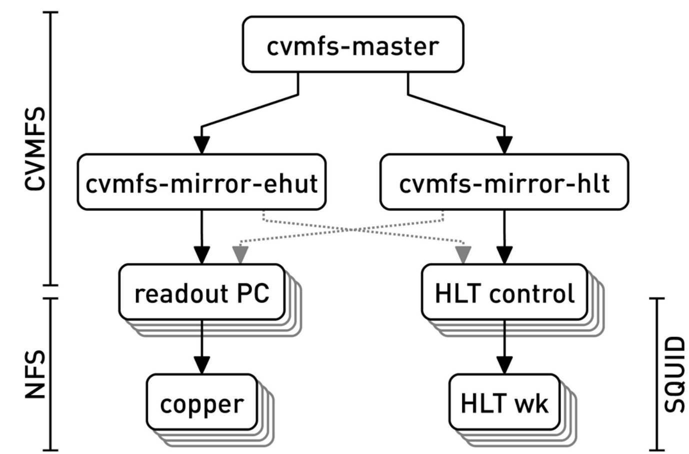

Trigger and Timing Distribution

- Custom Fast Timing Switch (FTSW) infrastructure
  - Distributes global clock, trigger information, injection signals, JTAG
  - Gathers FEE readout/busy status
- Cascaded tree distribution
  - Up to four levels deep, up to 30m cat7, 20m multimode fiber, 620ns latency
  - ≤25ps clock distribution jitter (clock routed through FPGAs)
- Some operational instabilities with individual front-ends investigated this summer

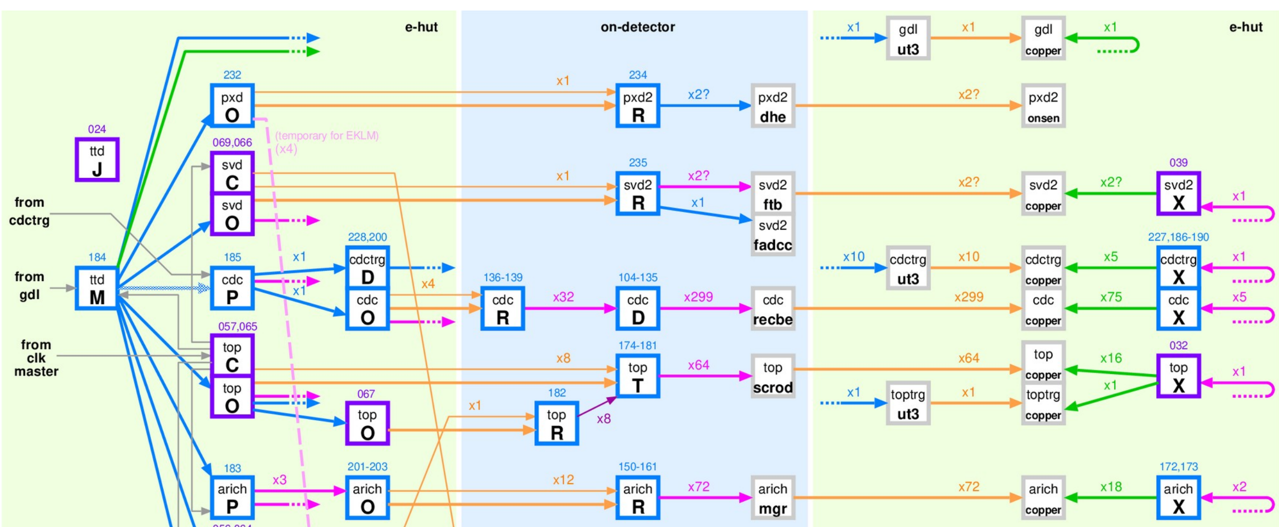

Luminosity & Event rates

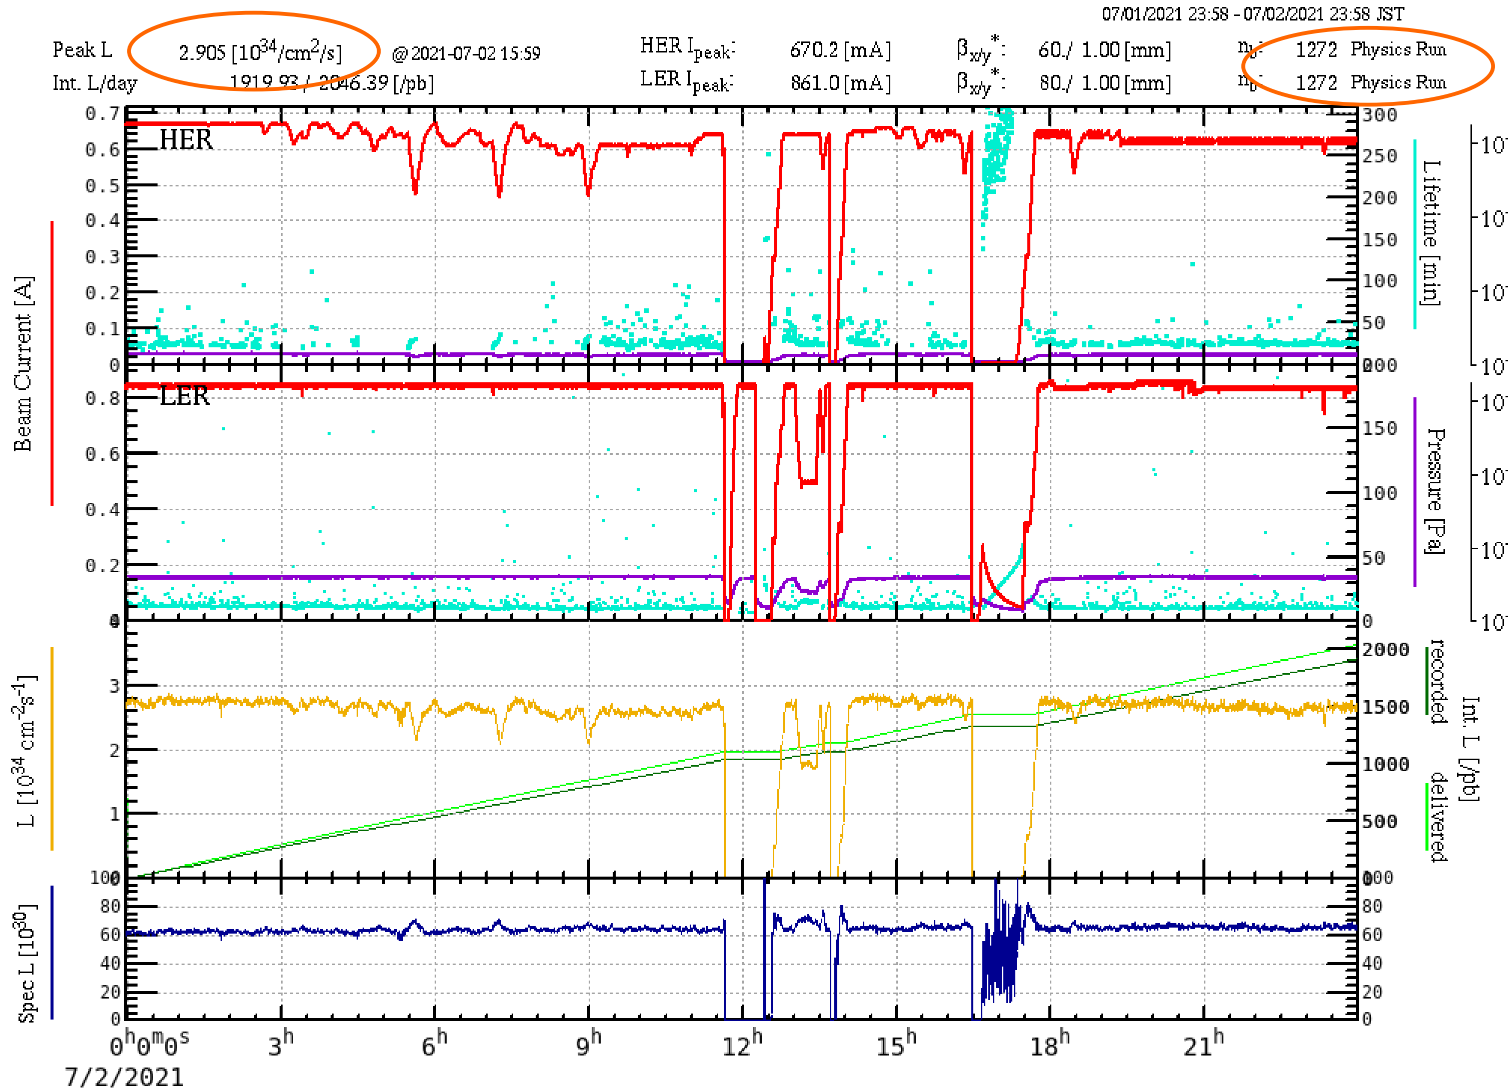

Event Rates

- 1278 bunches cross IP every 10us (~3000m ring circumference): ~130MHz collision rate
  - We cannot possibly collect all of these
- $B\overline{B}$ production cross section ~1nb, instantaneous luminosity ~30nb/s
  → only 30Hz of $B\overline{B}$ events
- Goal of trigger system: out of 13MHz collision rate, pick out the interesting events.
- Uninteresting events are not empty:
  - Bhabha scattering diverges at low angles!
    e+e- → e+e- with both tracks > 0.5deg is ~122,000 nb (=5MHz)
  - Beam backgrounds, ...

We Can Only Keep So Much…!

- Two main limitations:
  - Trigger/data rate from detectors to DAQ system (design limits, $$$)
  - Amount of data written to disk ($$$)
- Current goal: store not more than 20nb: 130MHz → 600Hz
  - Should be enough for all interesting physics
    + calibration channels etc

Physics process	Cross section [nb]	Cuts
$\Upsilon(4S)$	1.05 ± 0.10	-
$u\bar{u}(\gamma)$	1.61	-
$d\bar{d}(\gamma)$	0.40	-
$s\bar{s}(\gamma)$	0.38	-
$c\bar{c}(\gamma)$	1.30	-
$e^+e^-(\gamma)$	300 ± 3 (MC stat.)	$10° < \theta^*_{e's} < 170°$, $E^*_{e's} > 0.15$ GeV
$e^+e^-(\gamma)$	74.4	e's (p >0.5GeV) in ECL
$\gamma\gamma(\gamma)$	4.99 ± 0.05 (MC stat.)	$10° < \theta^*_{\gamma's} < 170°$, $E^*_{\gamma's} > 0.15$ GeV
$\gamma\gamma(\gamma)$	3.30	γ's (p >0.5GeV) in ECL
$\mu^+\mu^-(\gamma)$	1.148	-
$\mu^+\mu^-(\gamma)$	0.831	μ's (p >0.5GeV) in CDC
$\mu^+\mu^-\gamma(\gamma)$	0.242	μ's (p >0.5GeV) in CDC, ≥ 1 γ (E_γ >0.5GeV) in ECL
$\tau^+\tau^-(\gamma)$	0.919	-
$\nu\bar{\nu}(\gamma)$	0.25×10^{-3}	-
$e^+e^-e^+e^-$	39.7 ± 0.1 (MC stat.)	$W_{\ell\ell} > 0.5$GeV
$e^+e^-\mu^+\mu^-$	18.9 ± 0.1 (MC stat.)	$W_{\ell\ell} > 0.5$GeV

Mode	σ (nb)
BB	1.05
cc	1.30
qq	2.13
τ τ	0.80
μ μ	0.78
γ γ	2.01
bhabha	8.35
eeee	1.14
eeμμ	0.82
TOTAL	18.35

The Belle II Trigger System

- Belle II Trigger system ultimately decides which collision events are written to disk.
  - If something misbehaves and we only realise afterwards, important data is simply lost! There is no second try with triggering!
- Two stage process:
  Level 1 (L1) and “High Level Trigger” (HLT)
- Level 1 trigger:
  - Primarily based on CDC and ECL energy/clusters
  - Receives a stream of raw, coarse detector data (not the readout data!)
  - L1 correlates streamed information based on various “conditions” and issues an event trigger.
    - Processing in advanced FPGA logic
    - Individual processes might be prescaled (only every Nth occurrence of trigger X is let through)
    - This is what defines “an event” in Belle II
- High Level Trigger: Server Farm next to Belle II
  - Receives event readout data for each issued, decides based on running “full” basf2 reconstruction on the fly
  - ~5000 CPU cores, can only cope with a limited input rate (~20kHz max)

The Level 1 Trigger Menu

- L1 has a multitude of trigger conditions with exotic names

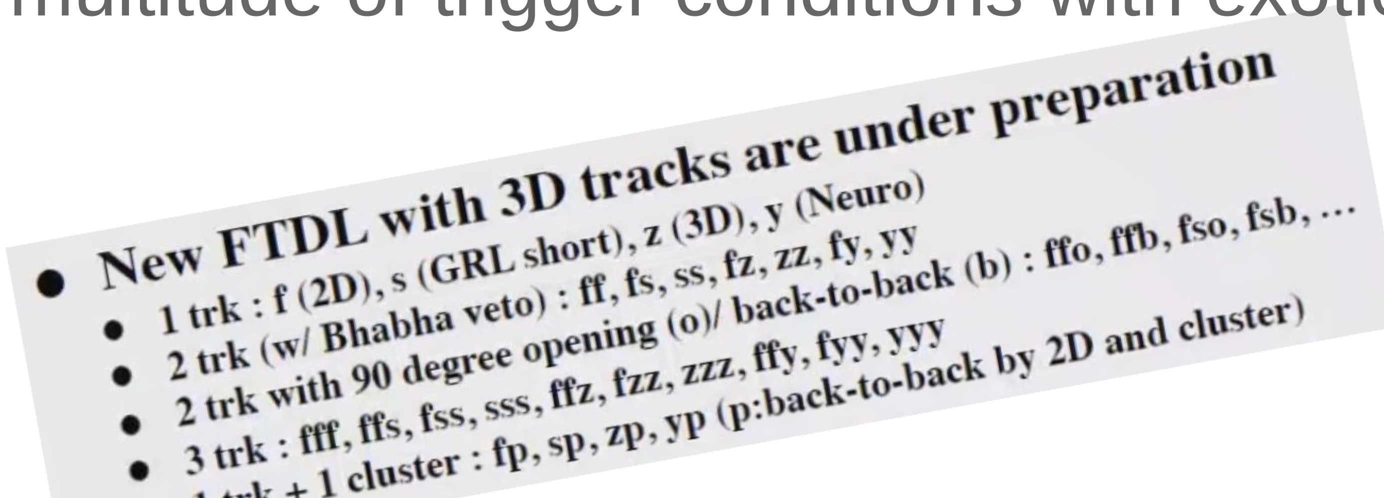

- New FTDL with 3D tracks are under preparation
  - 1 trk : f (2D), s (GRL short), z (3D), y (Neuro)
  - 2 trk (w/ Bhabha veto) : ff, fs, ss, fz, zz, fy, yy
  - 2 trk with 90 degree opening (o)/ back-to-back (b) : ffo, ffb, fso, fsb, …
  - 3 trk : fff, ffs, fss, sss, ffz, fzz, zzz, ffy, fyy, yyy
  - 1 trk + 1 cluster : fp, sp, zp, yp (p:back-to-back by 2D and cluster)

- Find details, prescales and their changes over runs/experiments on confluence: https://confluence.desy.de/display/BI/TriggerBitTable
  - Prepare for the worst

trigger bit PSNM

trigger name

run number

prescale; 0 means trigger is not used

exp12 Physics runs

Output Bits

	797-	psv	1660-	psv	1743-	psv	1908-	psv	2134-	psv	2309-	psv	2335-	psv	2720-	psv	5888-	psv	6373-	psv
0	fff	1	fff	1	fff	1	fff	1	fff	1	fff	1	fff	1	fff	1	fff	1	fff	1
1	ffs	100	ffs	100	ffs	100	ffs	100	ffs	100	ffs	100	ffs	100	ffs	100	ffs	100	ffs	100
2	fss	0	fss	0	fss	0	fss	0	fss	0	fss	0	fss	0	fss	0	fss	0	fss	0
3	sss	100	sss	100	sss	100	sss	100	sss	100	sss	100	sss	100	sss	100	sss	100	sss	100
4	ffz	0	ffz	0	ffz	0	ffz	0	ffz	0	ffz	0	ffz	0	ffz	0	ffz	0	ffz	0
5	fzz	0	fzz	0	fzz	0	fzz	0	fzz	0	fzz	0	fzz	0	fzz	0	fzz	0	fzz	0
6	zzz	0	zzz	0	zzz	0	zzz	0	zzz	0	zzz	0	zzz	0	zzz	0	zzz	0	zzz	0
7	ffy	1	ffy	1	ffy	1	ffy	1	ffy	1	ffy	1	ffy	1	ffy	1	ffy	1	ffy	1
8	fyy	0	fyy	0	fyy	0	fyy	0	fyy	0	fyy	0	fyy	0	fyy	0	fyy	0	fyy	0
9	yyy	0	yyy	0	yyy	0	yyy	0	yyy	0	yyy	0	yyy	0	yyy	0	yyy	0	yyy	0
10	ff	20	ff	20	ff	20	ff	20	ff	20	ff	20	ff	20	ff	20	ff	20	ff	20
11	fs	0	fs	0	fs	0	fs	0	fs	0	fs	0	fs	0	fs	0	fs	0	fs	0
12	ss	400	ss	400	ss	400	ss	400	ss	400	ss	400	ss	400	ss	400	ss	400	ss	400
13	fz	0	fz	0	fz	0	fz	0	fz	0	fz	0	fz	0	fz	0	fz	0	fz	0
14	zz	0	zz	0	zz	0	zz	0	zz	0	zz	0	zz	0	zz	0	zz	0	zz	0
15	fy	0	fy	0	fy	0	fy	0	fy	0	fy	0	fy	0	fy	0	fy	0	fy	0
16	yy	0	yy	0	yy	0	yy	0	yy	0	yy	0	yy	0	yy	0	yy	0	yy	0
17	ffo	1	ffo	1	ffo	1	ffo	1	ffo	1	ffo	1	ffo	1	ffo	1	ffo	1	ffo	1
18	fso	0	fso	0	fso	1	fso	1	fso	1	fso	1	fso	1	fso	1	fso	1	fso	1

Rules of Thumb

- If your analysis is based on $B\bar{B}$ or $c\bar{c}$ processes, the trigger efficiency for your events is most likely excellent
- If you are interested in low-multiplicity final states, you need to be careful.
  - Especially if they tend to look anything like Bhabhas

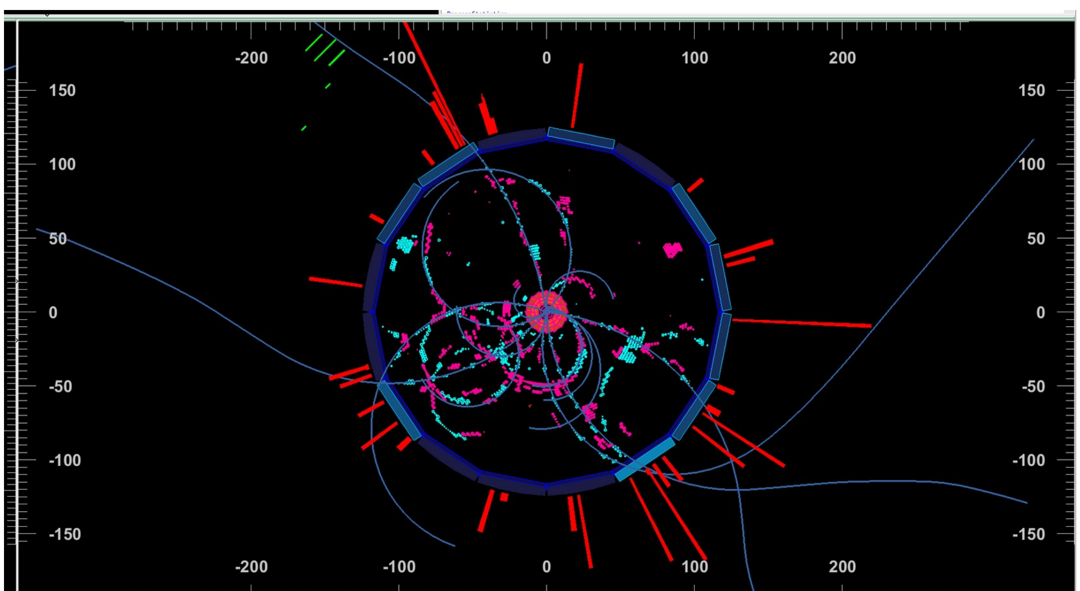

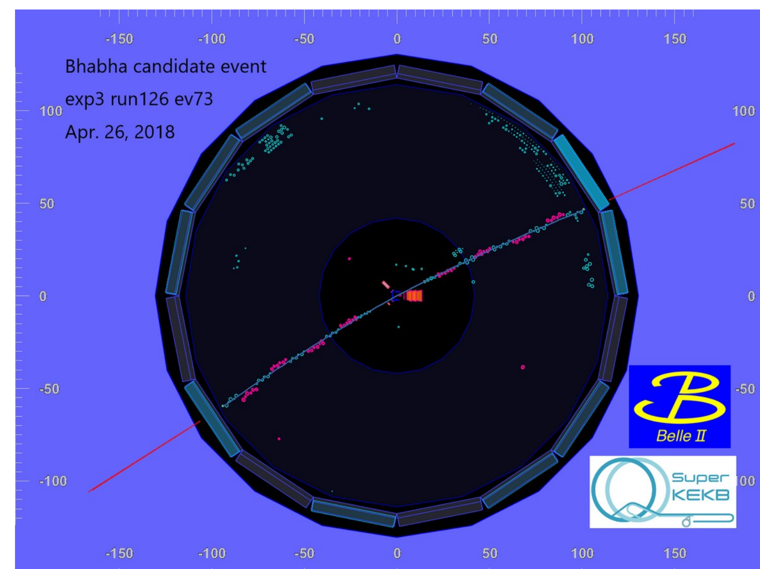

Example 1: Mono Photon Searches

- Look for nothing but a single photon
- Dedicated single photon L1 triggers:
  - **lml6**: exactly one cluster with E* > 1 GeV in the ECL barrel and no other cluster with E > 300 MeV anywhere
  - **lml13**: exactly one cluster with E* > 0.5 GeV with 44.2° < θ < 94.8° and no other cluster > 300 MeV.
  - There are also triggers for events that have a cluster with E*>2 GeV (and any number of other clusters), and "hie", which requires 1 GeV in the barrel or part of the forward endcap.

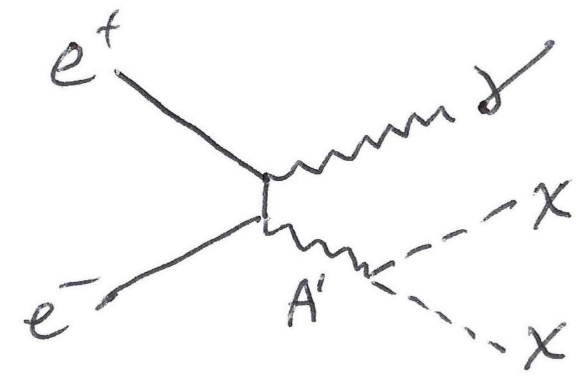

Mono Photon L1 Trigger: 1GeV

- Efficiency evaluation: find orthogonal trigger that acts as reference
  - Caveat: even fundamentally unrelated triggers might be correlated
- L1 calibration and resolution are not perfect

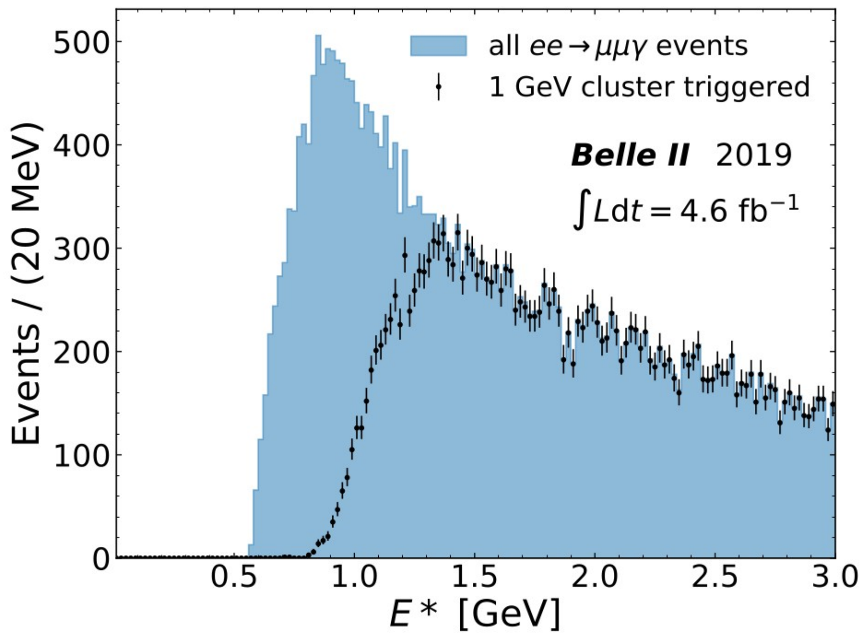

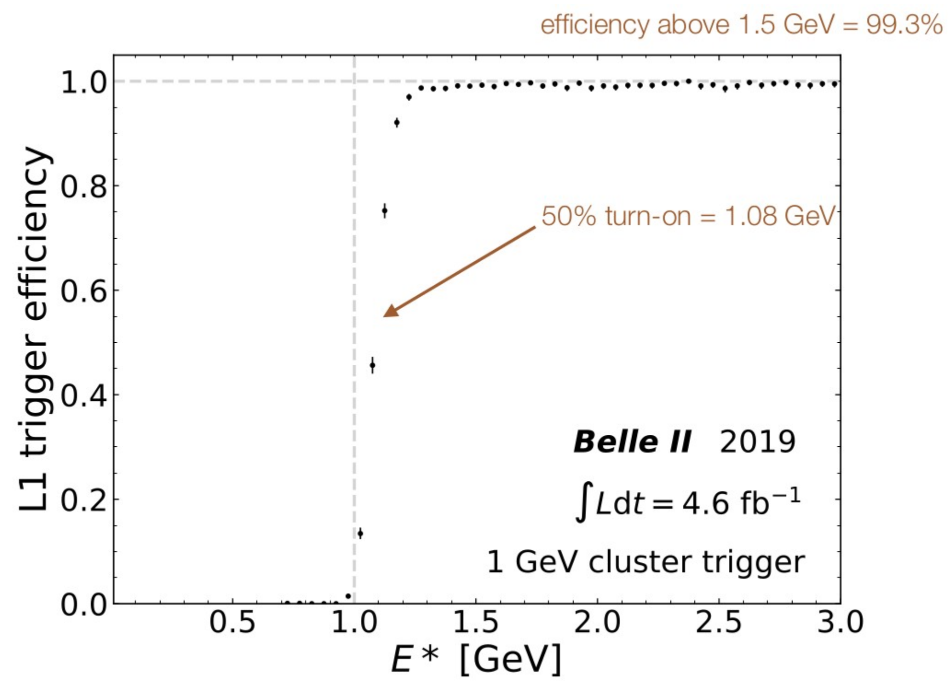

Mono Photon HLT Trigger: 1GeV

- HLT Trigger decision almost perfect
  - HLT uses “online” calibration, which might lag behind “latest greatest”

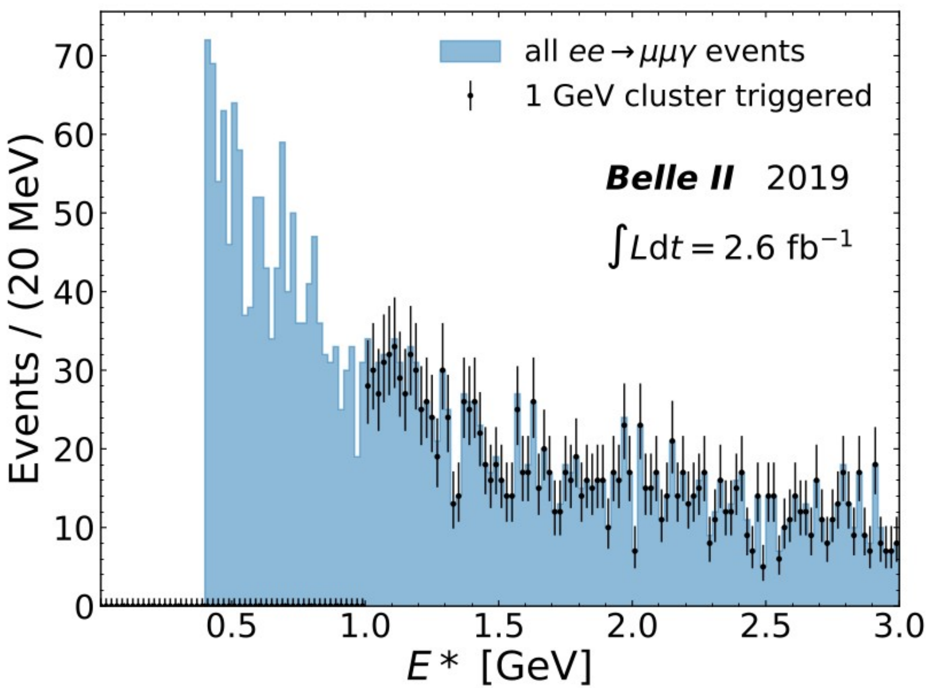

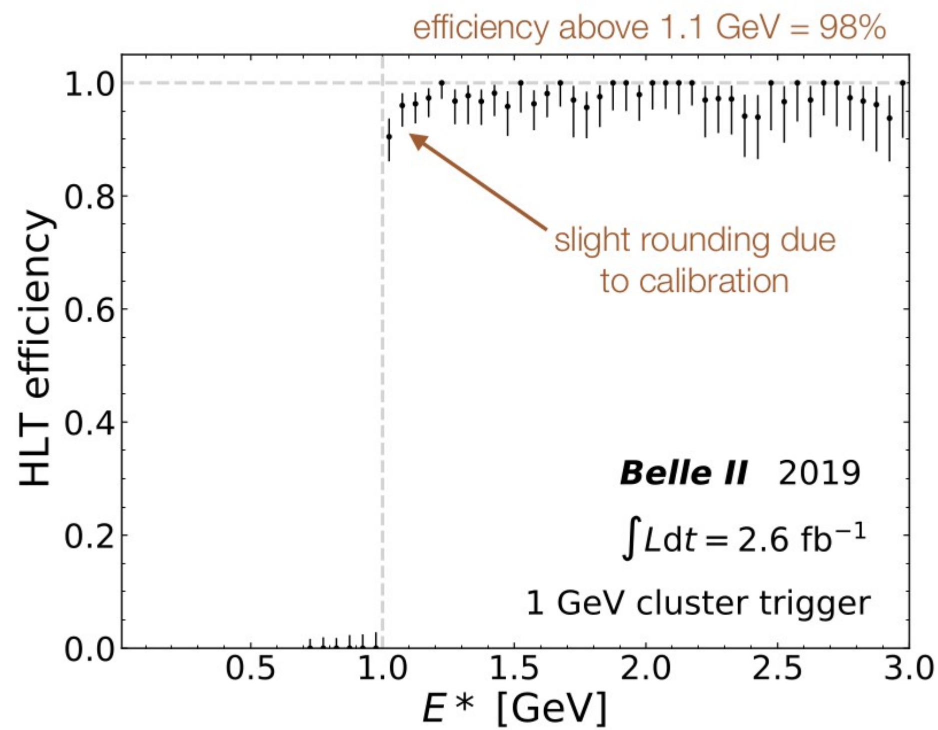

Example 2: Taus

- B factories are tau factories!
- Tau events are low multiplicity with “random” momentum tracks

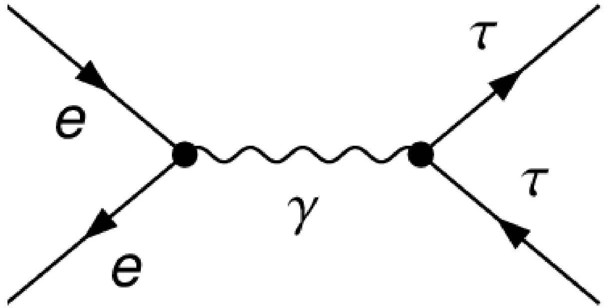

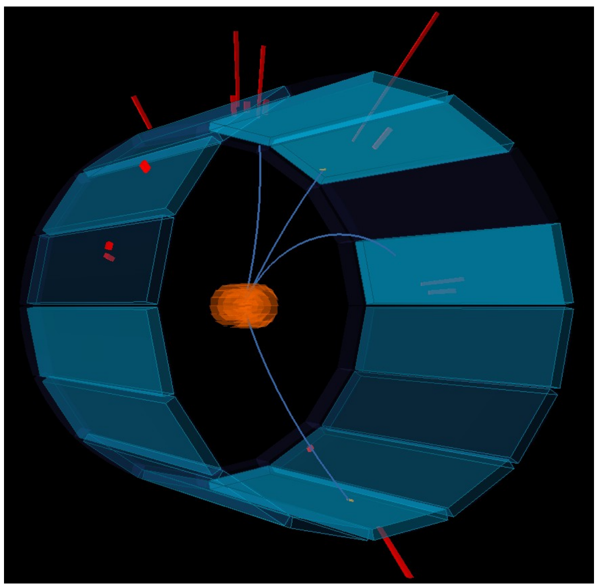

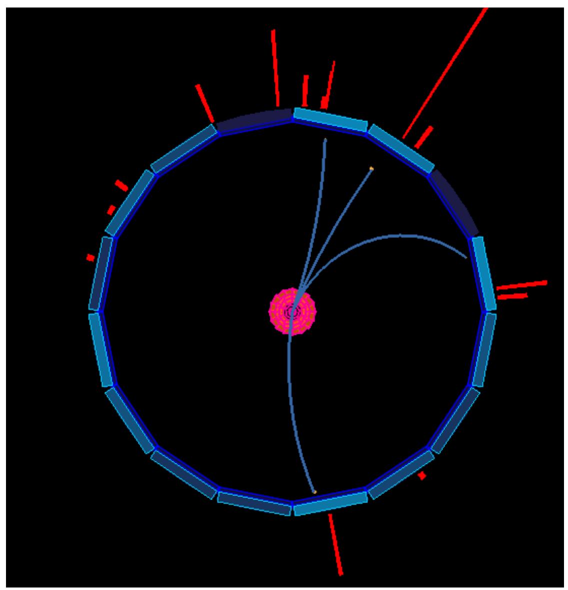

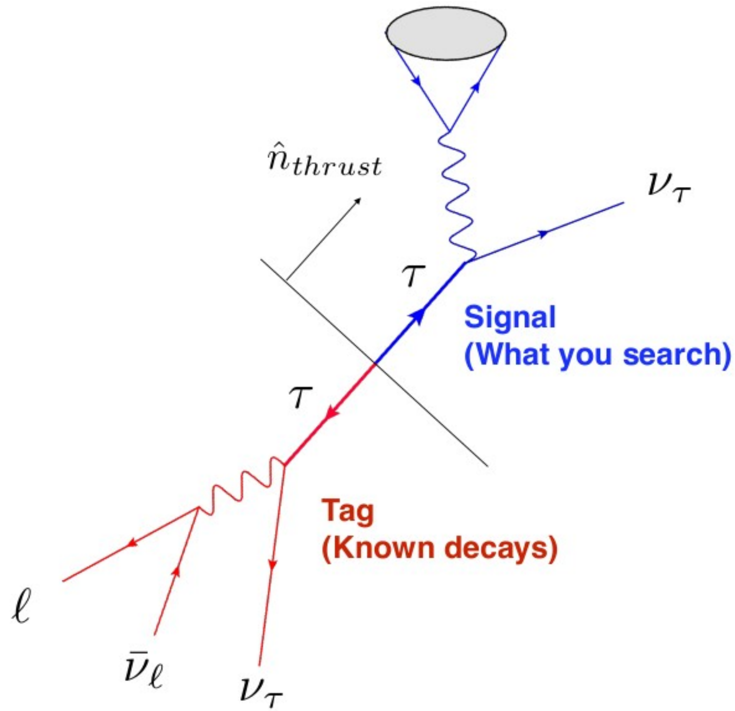

Low Multiplicity Track Triggering

- The “standard” CDC tracking trigger does not know about the Z-coordinate
  - Harder to distinguish backgrounds from interesting events

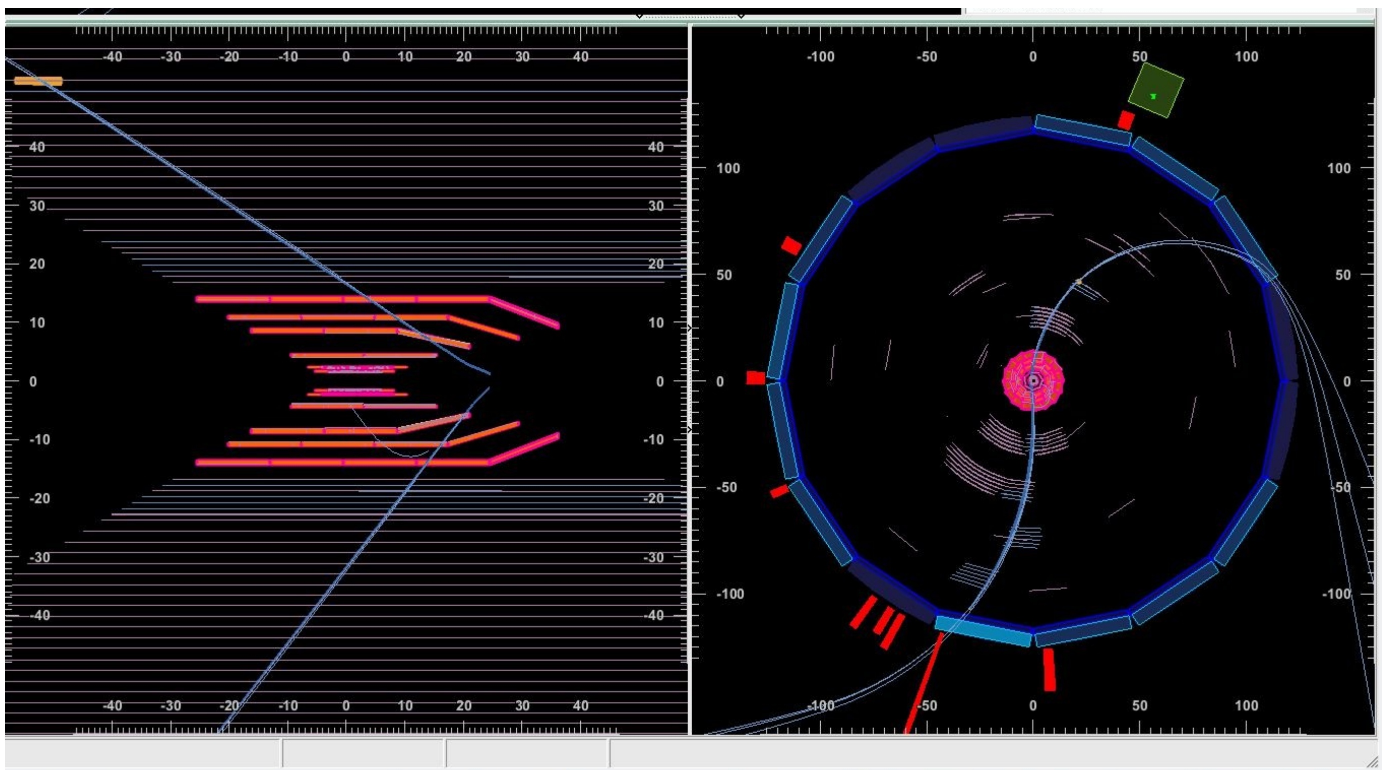

Neuro-Z Trigger

- Use Neural Networks in trigger FPGAs to reconstruct Z-origin of track in L1 trigger
- Still a little bit experimental, but default now (I believe)

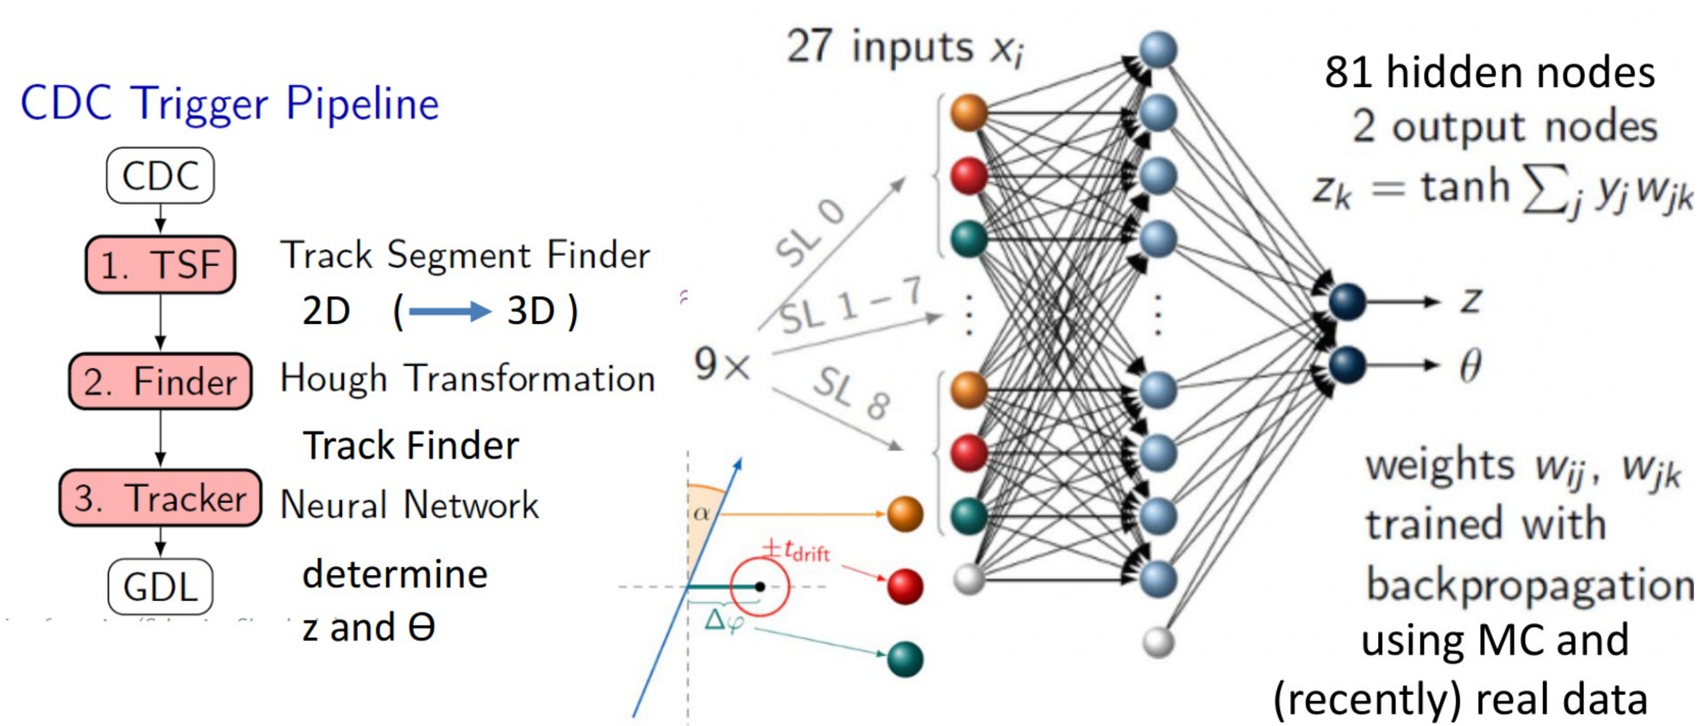

Single Track Neuro Trigger for Taus

- Triggering on single neuro track

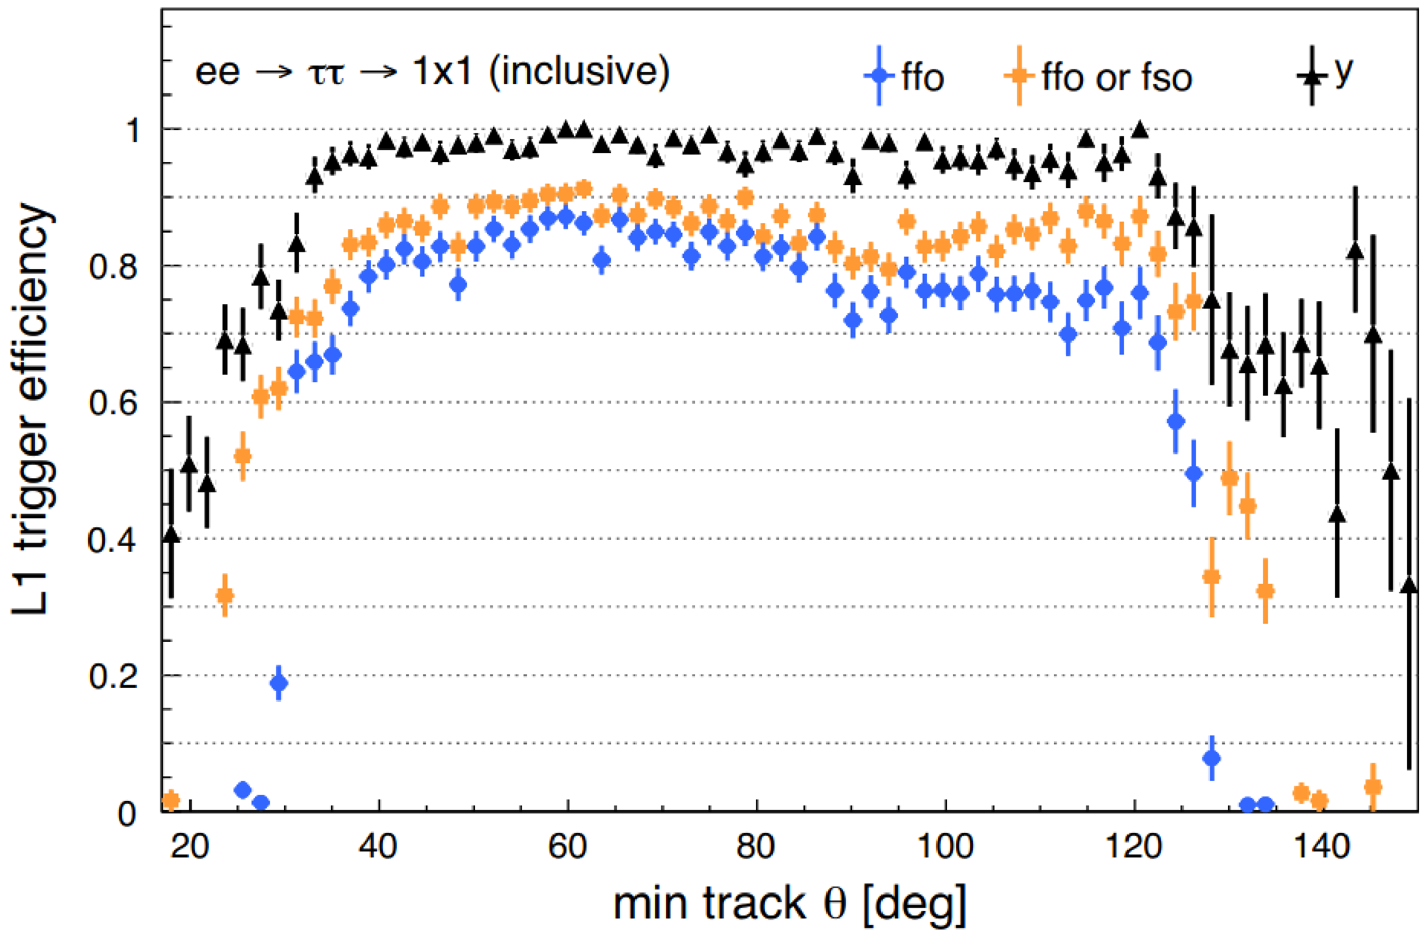

Summary

- More examples and instructions how to access trigger information and how to perform trigger efficiency analyses in Chris Hearty's talk at B2SS 2020: https://indico.belle2.org/event/1501/contributions/11211/attachments/6418/9962/Trigger_8-Jul-2020.pdf

- Trigger reduces the ~130 MHz collision rate to ~few hundred Hz to disk.
  - Especially if you are working with a low multiplicity final state: make sure your events are kept by the trigger system!
- Interact with your local trigger specialists. They can help you find what you need.

Summary

- More examples and instructions how to access trigger information and how to perform trigger efficiency analyses: Chris Hearty’s talk at B2SS 2020: