

Performance of the Unified Readout System of Belle II

Mikihiko Nakao
(KEK IPNS / SOKENDAI)



mikihiko.nakao@kek.jp


on behalf of Belle II readout experts:

Ryosuke Itoh, Satoru Yamada, Soh Y. Suzuki, Tomoyuki Konno, Qi-Dong Zhou, Takuto Kunigo, Ryohei Sugiura, Seokhee Park, Zhen-An Liu, Jingzhou Zhao, Igor Konorov, Dmytro Levit, Katsuro Nakamura, Hikaru Tanigawa, Nanae Taniguchi, Tomohisa Uchida, Kurtis Nishimura, Oskar Hartbrich, Yun-Tsung Lai, Masayoshi Shoji, Alexander Kuzmin, Vladimir Zhulanov, Brandon Kunkler, Isar Mostafanezhad, Hideyuki Nakazawa, and Yuji Unno

22nd Virtual Real Time Conference 2020, October 23

Belle II at SuperKEKB

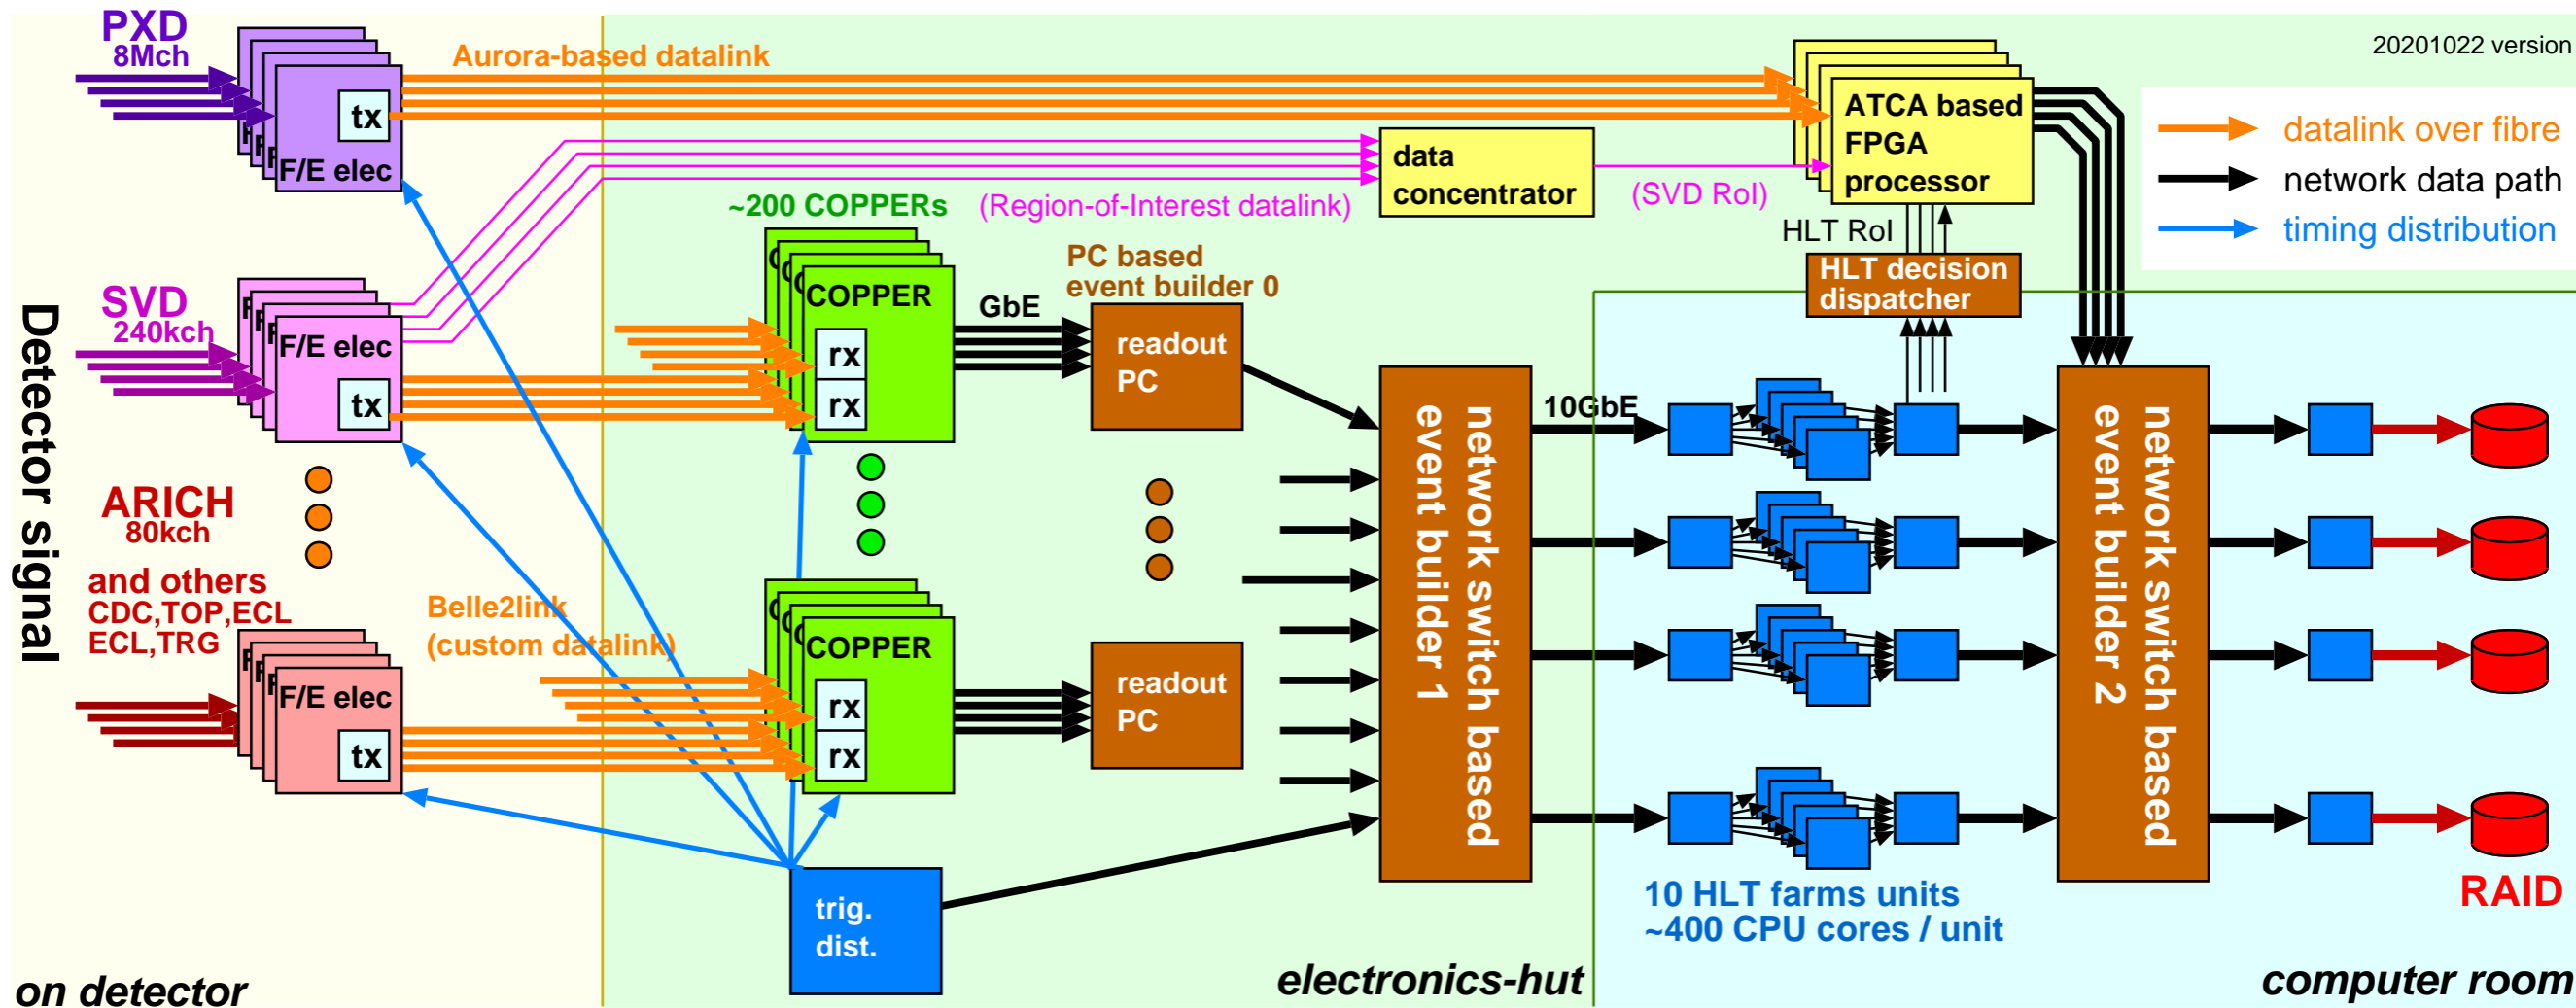
- **Luminosity frontier** for 50 ab^{-1} at KEK, Tsukuba JP
- **Clean e^+e^- collision:** $\sim 1 \text{ kHz}$ each of B , charm, τ
- **Asymmetric-energy:** HER $e^- 7 \text{ GeV}$ \times LER $e^+ 4 \text{ GeV}$
(high energy ring) (low energy ring)
- **7 subdetectors:** PXD SVD CDC TOP ARICH ECL KLM
- **>1000 collaborators** from 26 countries/regions



Unified Readout System for Belle II

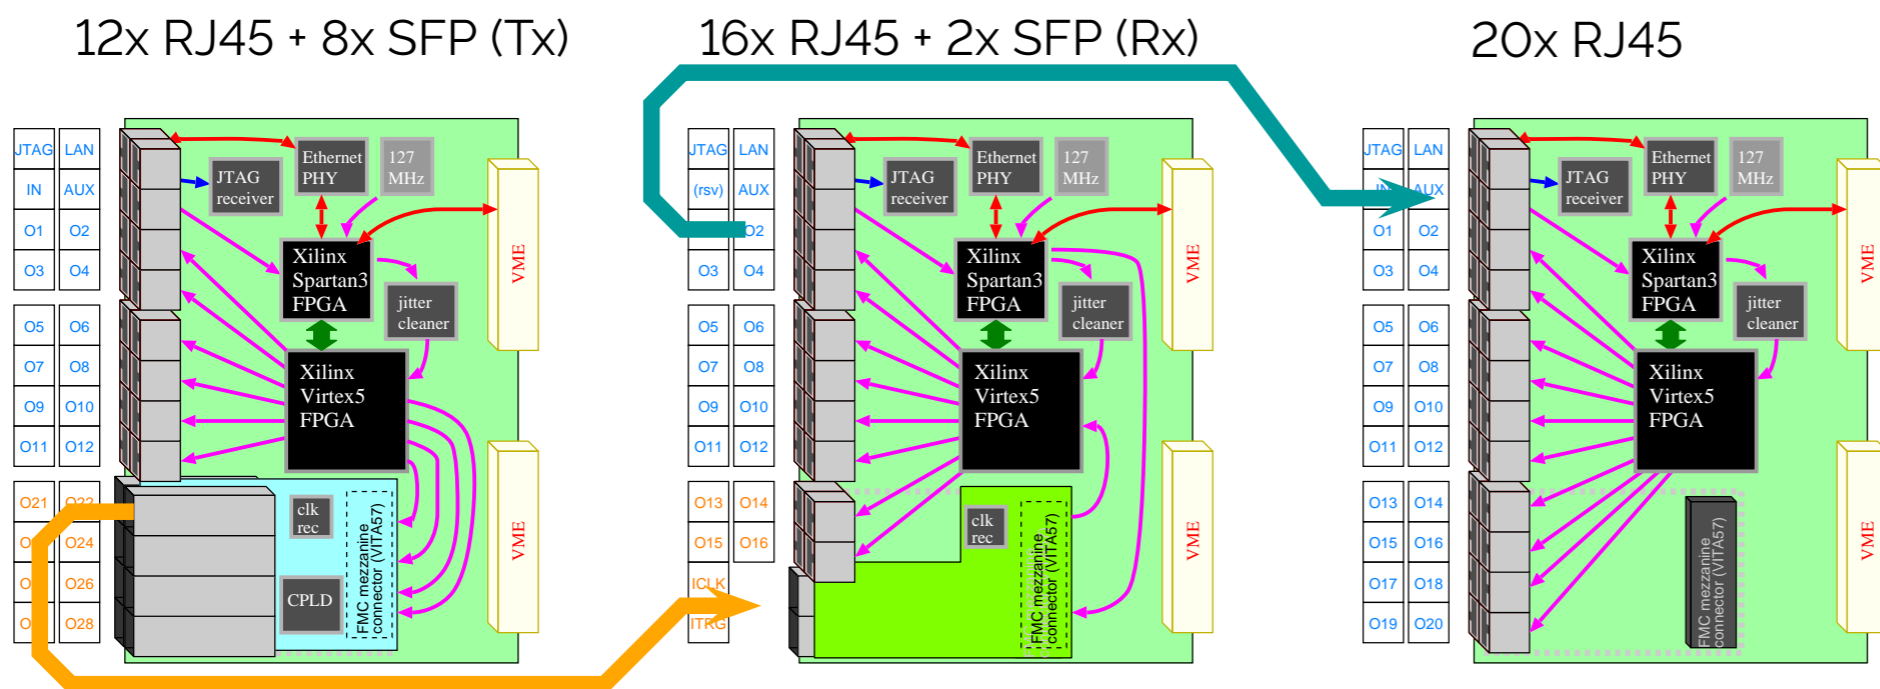
Belle II DAQ

- Up to 30 kHz level-1 trigger
- **Unified timing distribution system**
- **Unified data transport** (except PXD)
- Separate readout chain for PXD
- HLT for filtering and PXD data reduction
- 2-stage event building



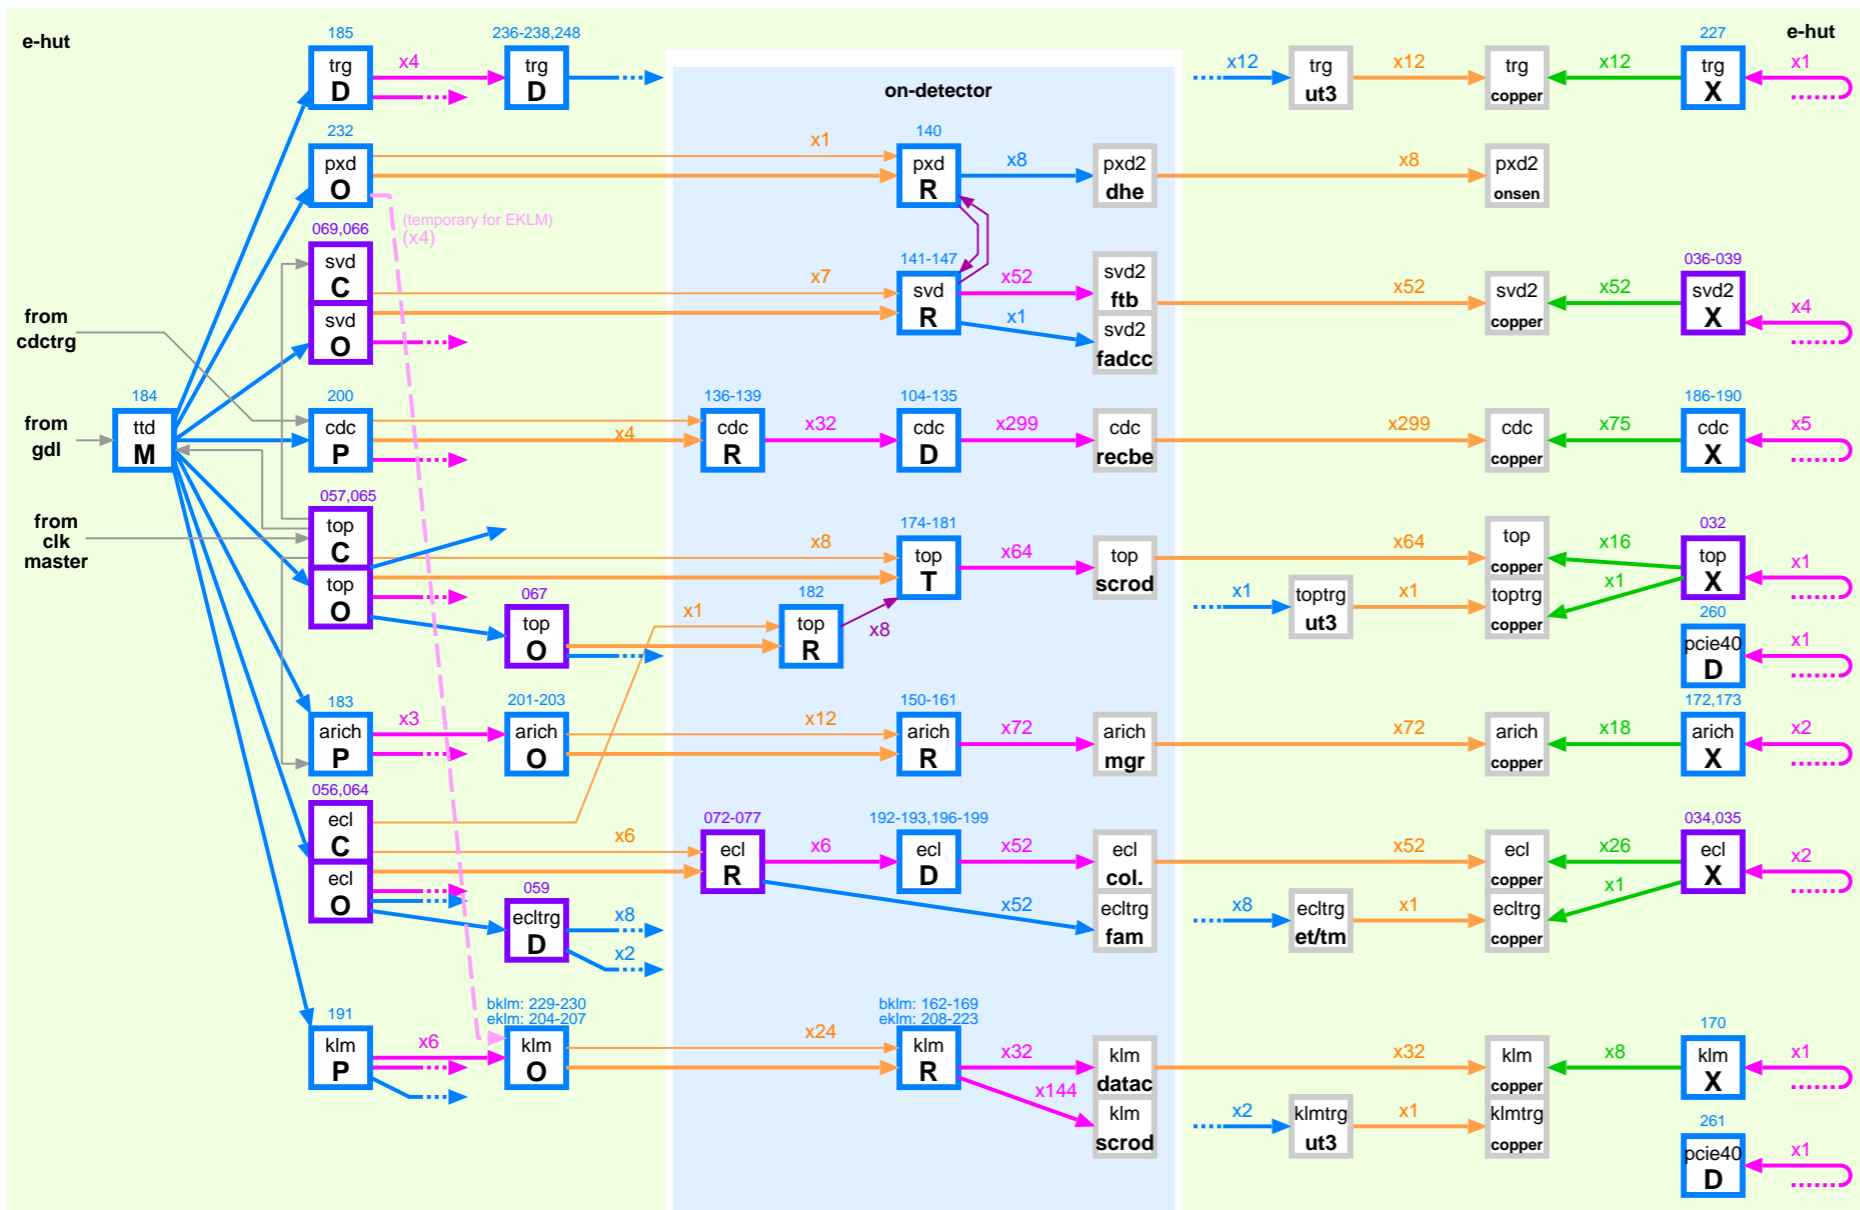
Unified Trigger Timing Distribution (TTD)

- 127 MHz system clock (1/4 of SuperKEKB RF)
- **“b2tt”** – custom 254 Mbps bidirectional serial protocol
- **“FTSW”** (frontend timing switch) – single PC-board for multiple TTD functions
- Up to 20 RJ-45 / 8 optical output from 6U double-width VME
- Common receiver firmware everywhere for Xilinx/Altera FPGAs
- JTAG connections to frontend using RJ-45 ports



Main TTD rack



TTD tree



Glossary

FTSW nodes
M: main-global
C: clock
P: primary
O: optical
R: receiver
D: distributor
T: top-special
X: copper/trrx

→ b2tt on CAT7
→ b2tt/B2L optical
→ clock optical
→ b2tt+JTAG
→ JTAG
→ ttrxlink
→ clock

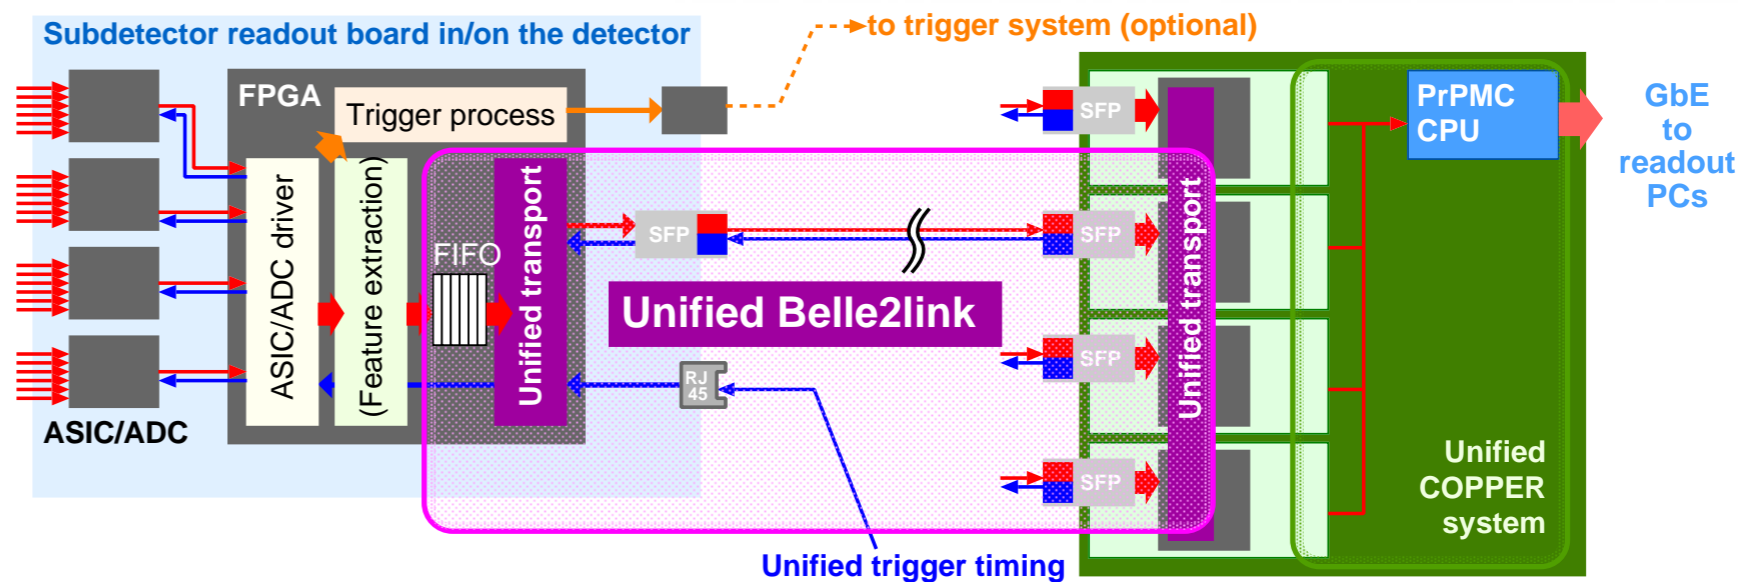



- Network of Cat-7 cables + fibers to >1000 frontend+backend
- Electronics-hut and detector are electrically isolated by fibers

Unified Belle2link and COPPER

- **2.54 Gbps custom protocol for data transport:**
unified firmware components for various Xilinx families
- **Slow control path to the frontend:**
16-bit address space for 32-bit registers in FEE
- **HSLB receiver + 1 MB FIFO on COPPER per link:**
Hardware event building on COPPER
- **209 COPPERs and 42 readout-PCs:**
75 COPPERs for CDC,
up to 13 COPPER data to one readout-PC


COPPER board



Frontend Electronics (FEE)



Simplest: CDC, all-in-one board solution

- pre-amp./shaper, 32 MHz sampling flash ADC
- Xilinx Virtex 5 FPGA
- 1ns LSB TDC in the FPGA logic fabric, with 254 MHz \times 4 clock phases



Tough: TOP, stack of small boards

- Custom sampling ADC chip
- 2-steps of Xilinx ZYNQ processors
- Feature extraction in a software
- In a very small corner inside the detector, ultra tight space constraint



Inside the detector: CDC, TOP, ARICH

Outside: outer subdetectors and (unified part of) SVD/PXD (processing on the sensors)

Data Flow Control

- Various trigger throttles in the main-global TTD
 - Minimum trigger interval of $0.5 \mu s$ (limited by trigger logic)
 - Programmable number of trigger limit for a given interval
 - **Main throttle** using emulation of deterministic SVD buffer usage in the main-global TTD firmware (tightest front-end buffer constraint among subdetectors)

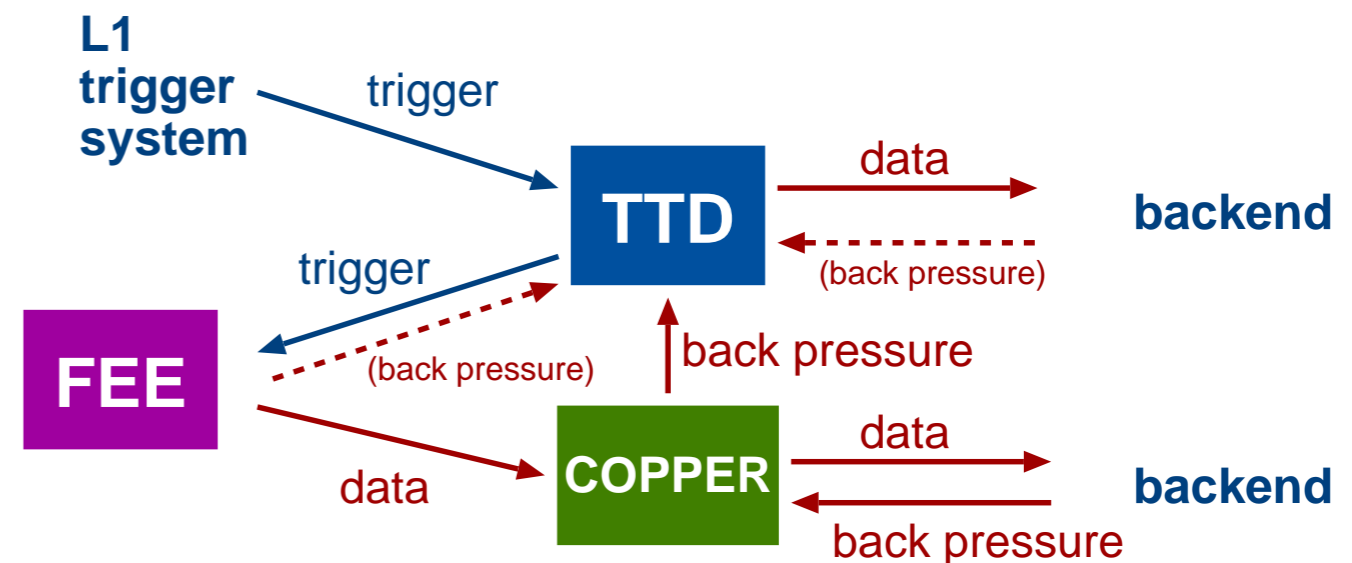
- **Back pressure from backend**

- Backend slow-down \Rightarrow COPPER FIFO threshold \Rightarrow **busy signal generated**
- TTD receives the back pressure well ahead of time before COPPER buffer gets full

- **No back pressure from COPPER to FEE**

- **Back pressure from FEE**

- For elaborated buffer management by TOP and SVD
- b2tt round-trip time $\sim 1 \mu s$



Performance and Lessons: 2019–2020 Runs

2019-2020 Runs

- Belle II run history: Mar–Jul 2019, Oct–Dec 2019, Feb–Jul 2020

- Pilot run without PXD/SVD in 2018, but debugging was far from complete...

- 2019 runs were still quite unstable for initial troubles in software / firmware / hardware

- Improved stability in 2020 runs, no loss due to COVID-19 (most of the experts could work remotely)

- Now starting up for the Oct–Dec 2020 run


- 24h / 7d non-stop operation, 1+1 local+remote shift

- Typical run lifetime of about a few hours by various reasons (beam loss, DAQ error)

- Biweekly maintenance day for accelerator and detector access


- Physics / accelerator study time sharing

- Accelerator study time used to run DAQ with 30 kHz dummy trigger + noise data with lower threshold (No high-voltage applied to the detector)



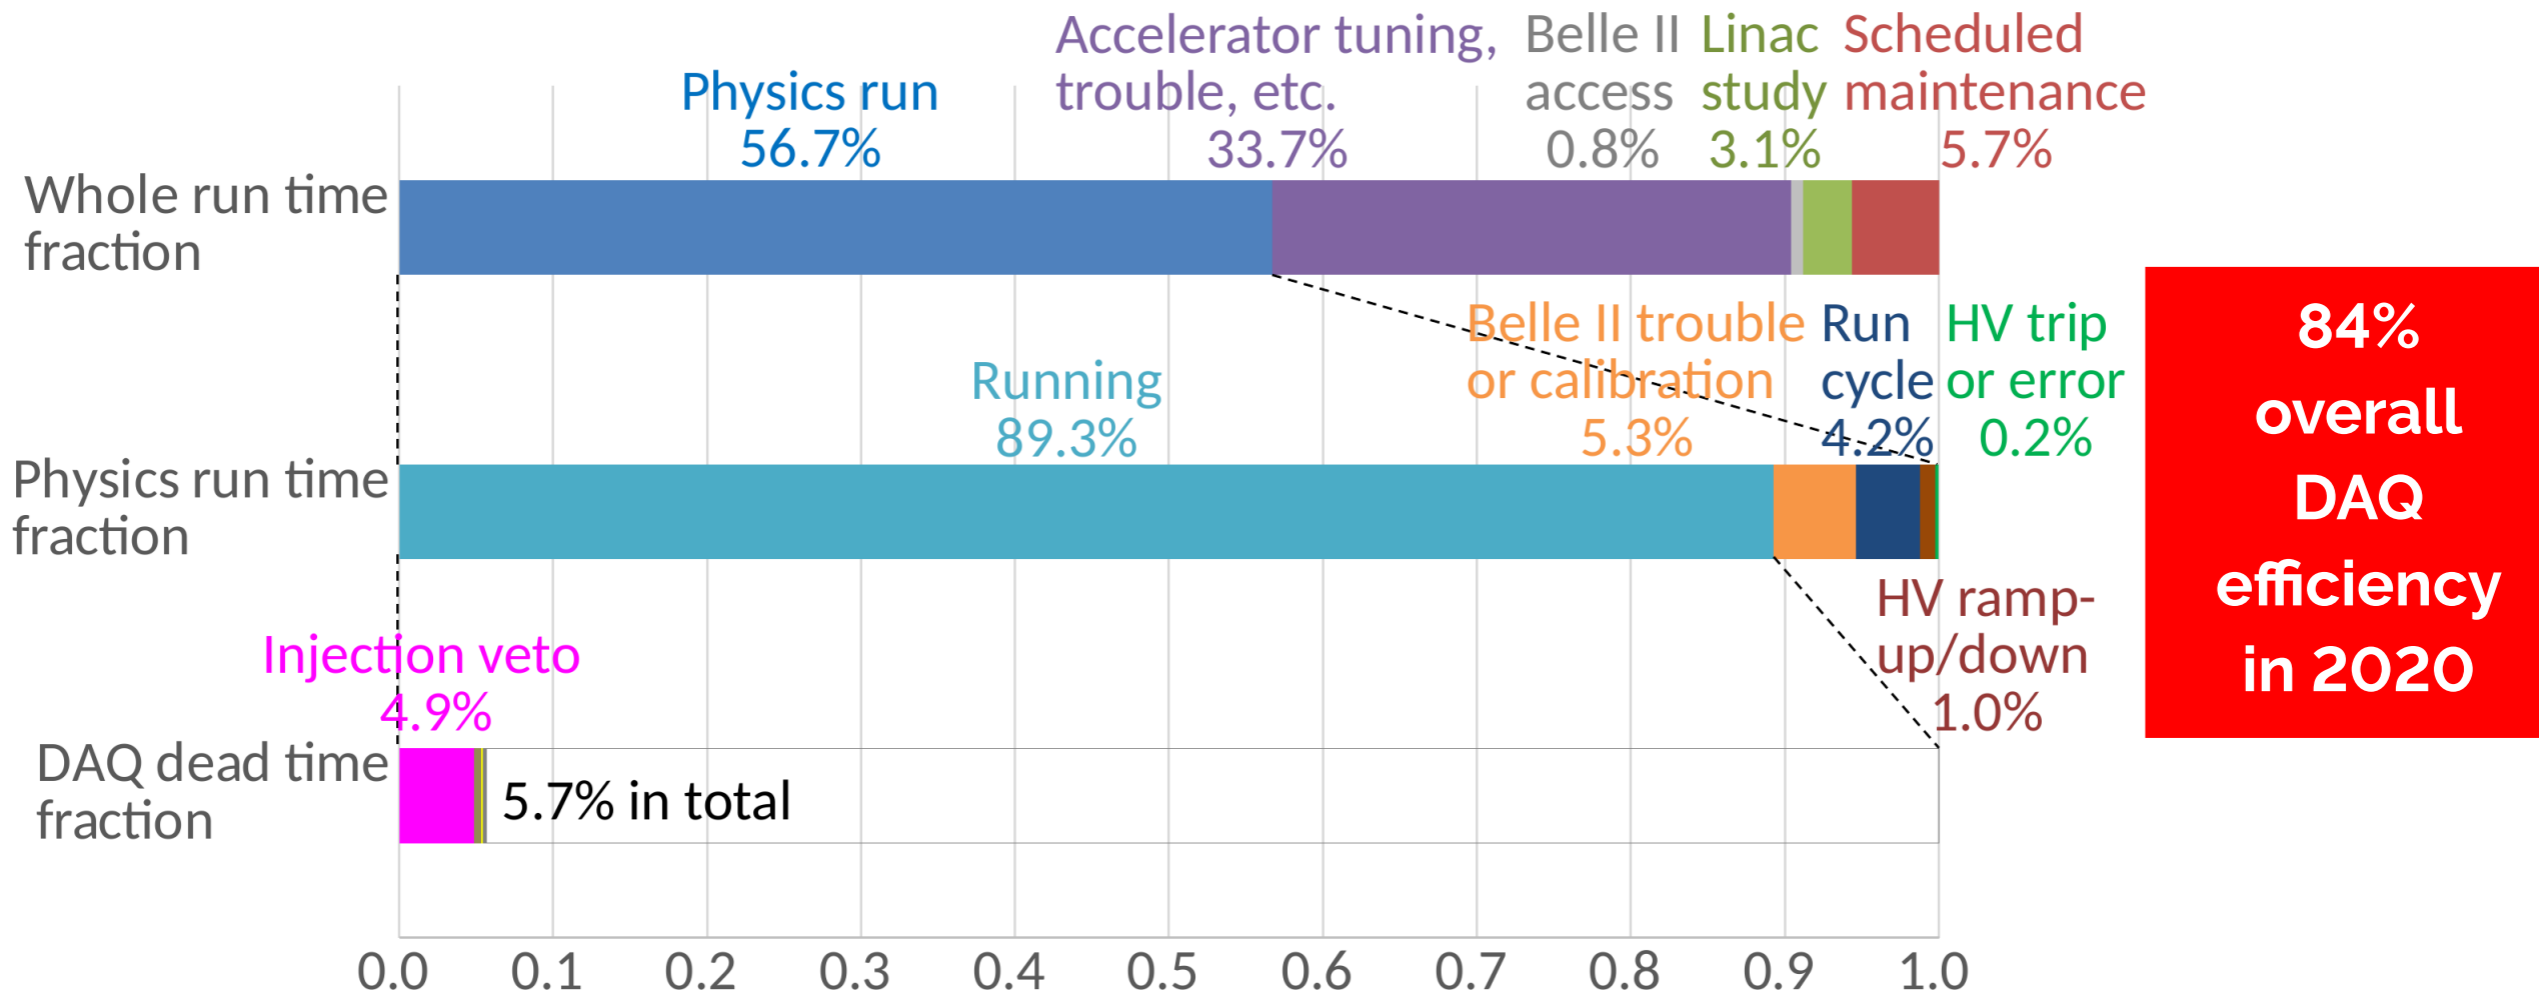
Luminosity delivered / recorded

Peak L	2.394 [$10^{34}/\text{cm}^2/\text{s}$]	@ 2020-06-22 20:53	HER I_{peak} :	640.2 [mA]	$\beta_{x/y}^*$:	60./ 1.00 [mm]	n_b :	978	Physics Run
Int. L/day	1345.95 / 1498.13 [fb]		LER I_{peak} :	740.2 [mA]	$\beta_{x/y}^*$:	80./ 1.00 [mm]	n_b :	978	Physics Run



- Almost constant beam current / luminosity \Leftarrow continuous injection
- World best peak luminosity of $2.4 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ — still $\times 30$ below the design
- Best day integrated 1.35 fb^{-1} — so far 74 fb^{-1} collected, $\times 675$ more to go
- Steadily recording — but not at a 100% efficiency \diamond \Rightarrow **this talk**

Run-time sharing and efficiency



84% overall DAQ efficiency in 2020

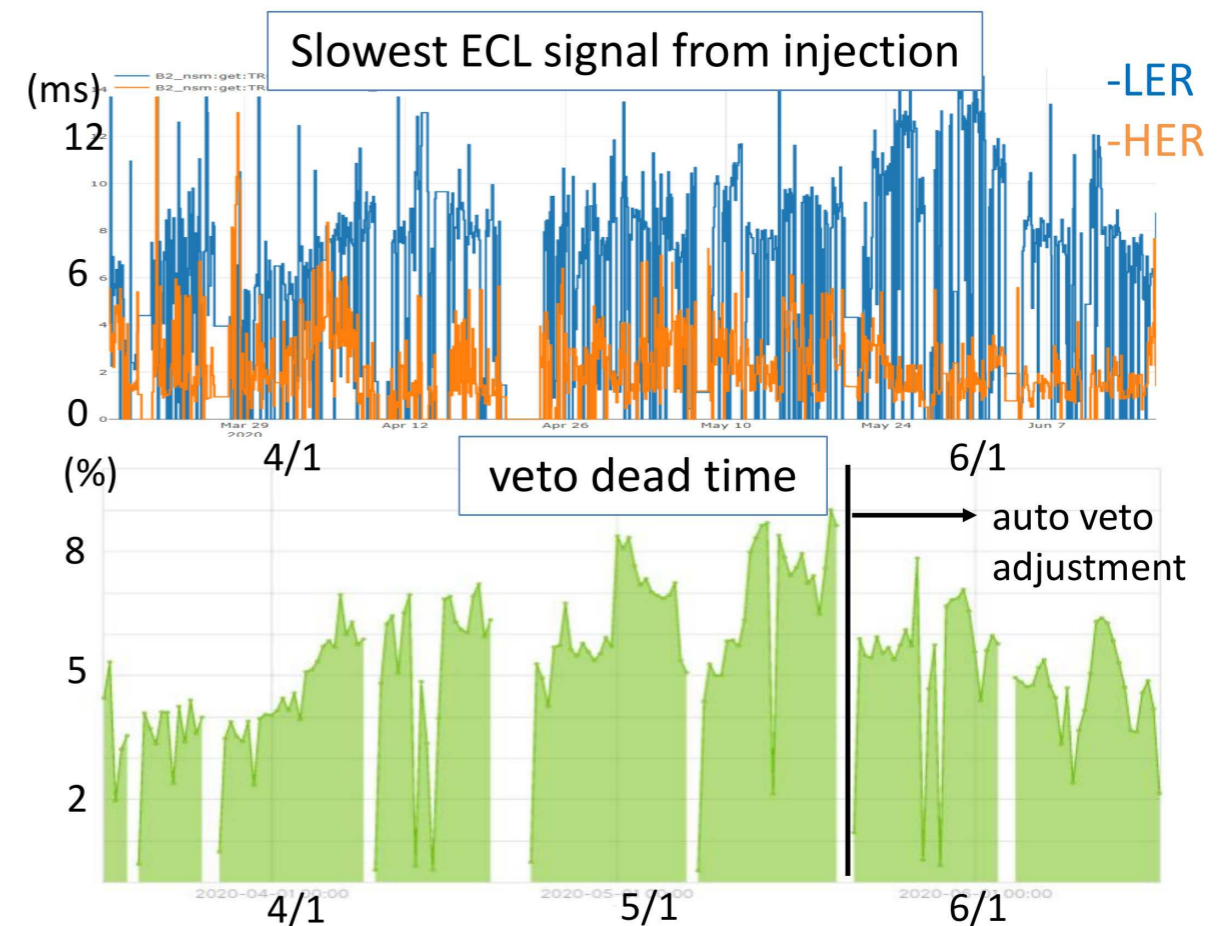
- Large fraction of run-time devoted to accelerator study / tuning
- ~10% of presumably reducible dead-time fraction in physics run-time
- DAQ dead time during a stable run is less than 1% (thanks to the lower trigger rate than the design...)

DAQ Dead-time Breakdown

- Injection veto
- Readout dead time
- Single event upset
- Link errors and following run restart cycle
- Hardware failures, other troubles

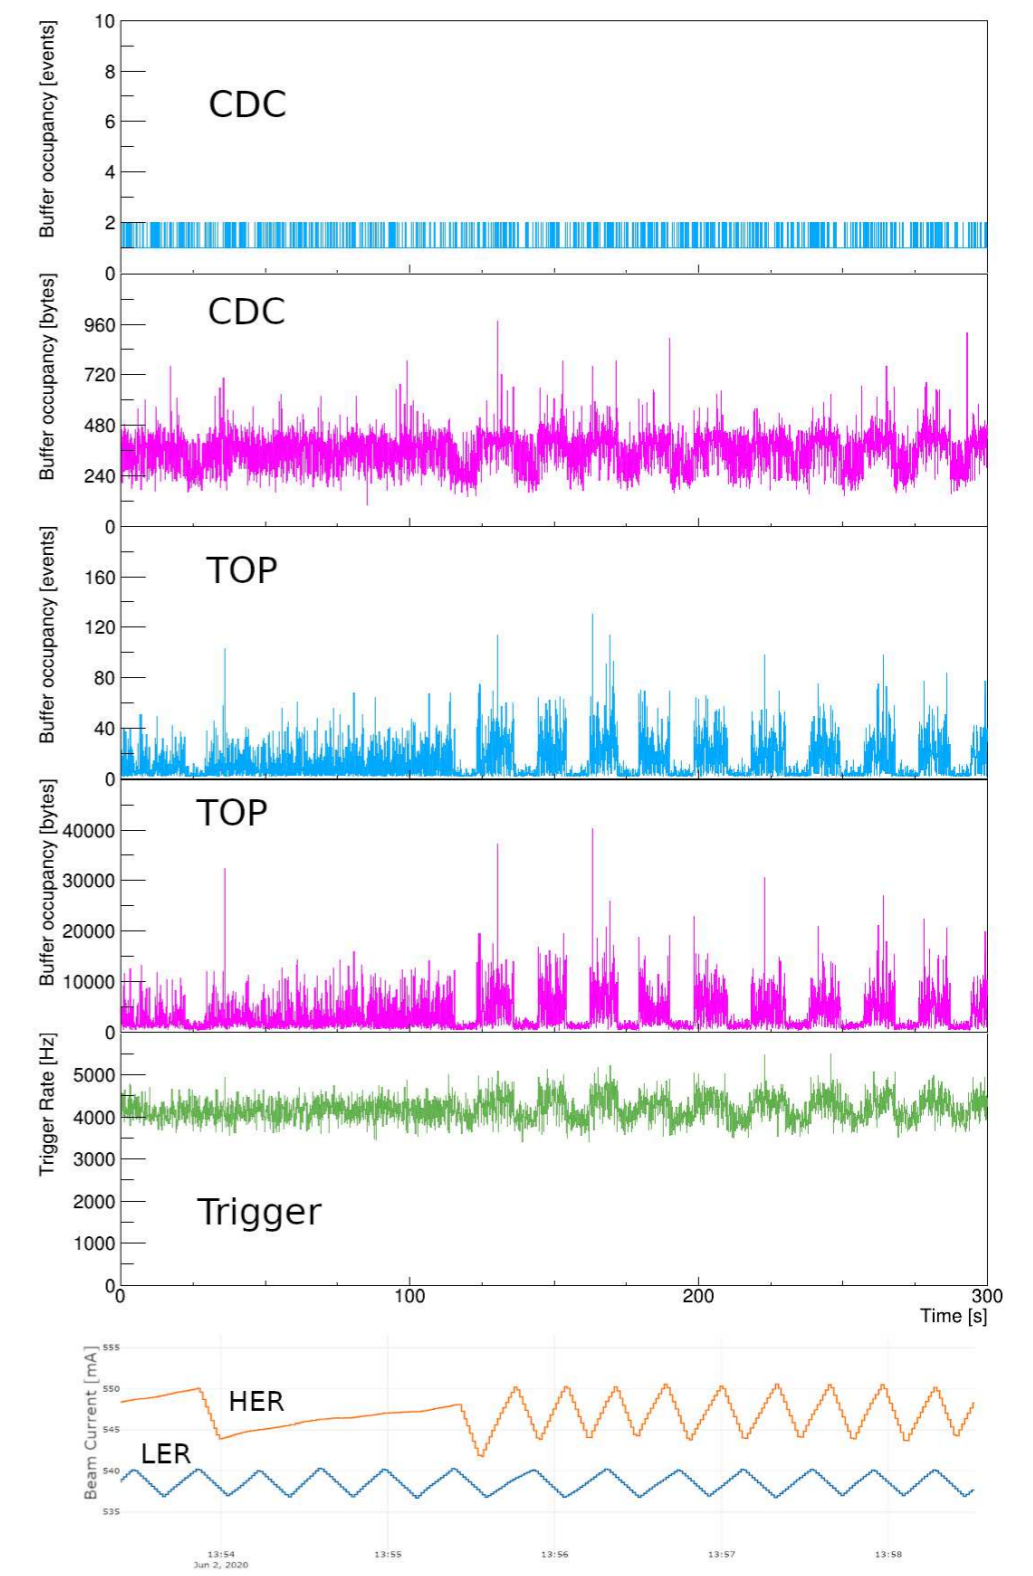
Injection Veto

- Continuous injection disturbs the already stored beam
 - Background over the whole ring — **short dumping time < 1 ms**
 - Background in $\sim 1/4$ of the ring after the injected bunch — medium dumping time
 - Background of the injected bunch train — **long dumping time ~ 10 ms**
 - LER (low energy ring, e^+) has a longer dumping time than HER (high energy ring, e^-)
- Programmable veto to the level-1 trigger
 - Full veto and gated veto
 - Tuned based on ECL trigger distribution
 - **$\sim 5\%$ dead time** for 25 Hz injection
- Further improvements foreseen
 - Narrow and wide gated veto
 - Better accelerator injection condition




Readout Characteristics

- **COPPER buffer occupancy estimated**
 - From the difference of timestamps between trigger and Belle2link transmission
 - COPPER has to wait for events from all 4 links are aligned
 - No latency variation in CDC, **large latency variation in TOP**
- **Effect of injection**
 - Trigger rate is kept to be **4-5 kHz** even during injection
 - HER (e^+) injection has a larger effect to the TOP buffer occupancy
- **Expectation for a higher trigger rate**
 - COPPER buffer is large enough for 30 kHz trigger rate
 - TOP FEE back pressure (at 150 events) will be asserted more frequently



Single Event Upset

- **FEE boards inside the detector are affected by neutrons**
 - Xilinx SEU mitigation code for CDC (Virtex 5) and TOP (ZYNQ)
 - Custom SEU mitigation code for ARICH (Spartan 6) (see talk by R.Giordano)
- **Unrecoverable SEU is not negligible**
 - Either detected by the mitigation code, inconsistent data, or by the malfunctioning of the firmware
 - **Once per day in average** for CDC, need FPGA reprogramming (this rate is almost as expected)
- **Reducing dead time**
 - Reprogramming FPGA takes only ~ 10 s for CDC, but decision takes time if it is from data error
 - Automated reprogramming procedure in preparation



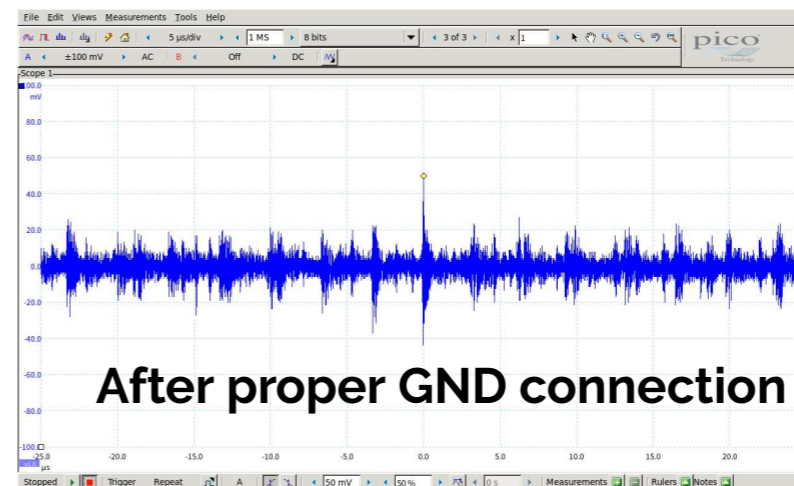
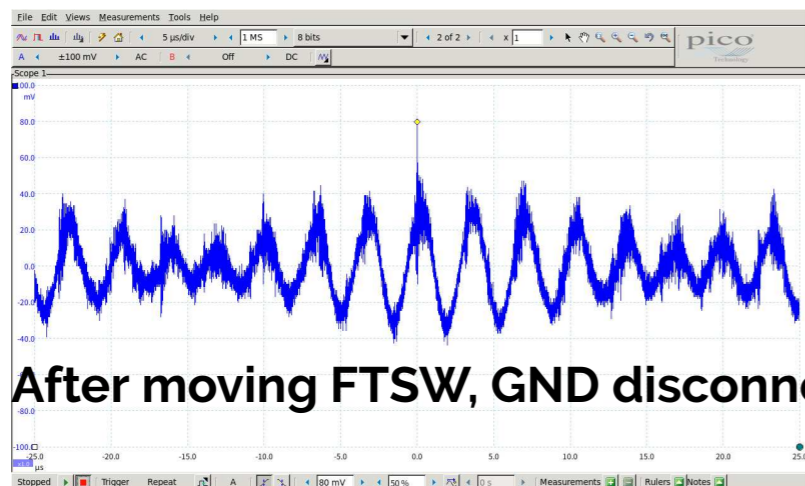
history of last 50 days of 2020 run

● detected by SEU mitigation code

● others

Link Errors and Mitigation (1)

- Both b2tt and Belle2link link errors stop the run
 - **Most of errors caused by the clock or serial data glitches on a long ($\geq 10\text{m}$) Cat-7 cable**
 - Link error can be reset in a few seconds, but run restart cycle takes 2.5 min (see HLT talk by M.Prim)
- **KLM (barrel part) case:**
 - GND of FEE board was not connected to anywhere and picking up external noise through a long Cat-7 cable (somehow nobody noticed and running for more than 1 year)
 - FTSW to new locations with new VME crates to reduce the Cat-7 cable length from 20m to 10m
 - **Very stable after these fixes during summer 2020**



~50mV noise from power supply (?) is still visible

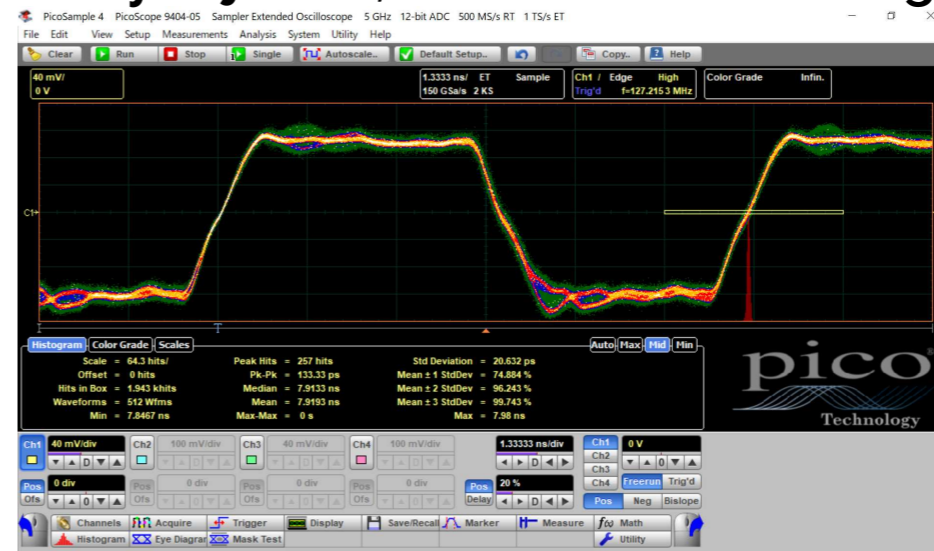
Link Errors and Mitigation (2)

- **CDC case** (also effective for ARICH)
 - Crosstalk from the edge of serial link to clock was found in some FTSW ports:
⇒ **making a large clock jitter**
 - Combination of lower LVDS swing of a particular FTSW port (by a few%) and a larger voltage drop in a particular FEE channel made the link very unstable
 - Cured by adding delay to the serial link output (IODELAY of Virtex-5)
 - **Updating the FTSW firmware made 10/3 unstable CDC/ARICH FEEs to back online**

No delay, crosstalk around the clock edge



Delay adjusted, crosstalk after the edge



Other Errors and Troubles

- **More link failures** (not very often, very little impact to dead time)
 - **TOP:** crosstalk between other b2tt signals, cannot be cured by IODELAY
 - Links in **electronics hut:** VME access to FTSW generates noise on b2tt lines, work in progress
- **Broken / unstable FEE boards**
 - **KLM:** LVDS driver broken, both at FEE and FTSW, most likely because of the GND problem
 - **ARICH and TOP:** faulty FEE boards, no access until next major shutdown in 2022 (causing small loss in particle id)
- **Software / PCs / slow control troubles**
 - Outside of the scope of this talk, but steadily improving (see ELK poster by T. Kunigo)

Conclusions

Conclusions

- **Unified readout system of Belle II** (timing distribution and data link) has been successfully integrated and “basically running stable”
- Dead time due to various reasons has been analyzed and steadily improved: **some of the problems were cured during the shutdown of summer 2020**
- **84% overall efficiency now, improvement to >90% is expected** (last 10% includes injection veto and HV/run cycle and will be tough to reduce)
- **Currently running at 4–5 kHz trigger rate, and 30 kHz design rate is regularly tested**
- **Backend upgrade project (to replace COPPER) is on-going** (next talk by Q.-D. Zhou)

