Status and Performance
of the
Belle II DAQ System

O. Hartbrich
(University of Hawaii at Manoa)
for the Belle II DAQ Group

ICHEP 2020, Prague
07/31/2020

Belle II and SuperKEKB

- SuperKEKB inst. luminosity world record on June 21st 2020
- Accumulated ~74fb^{-1}
- Goal ~50ab^{-1} in 10years
  - Increase inst. lumi by x30

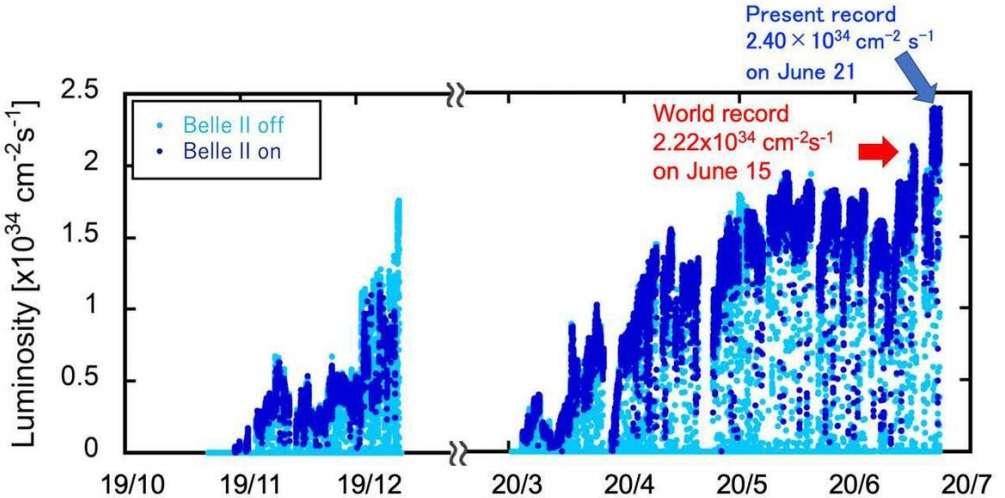

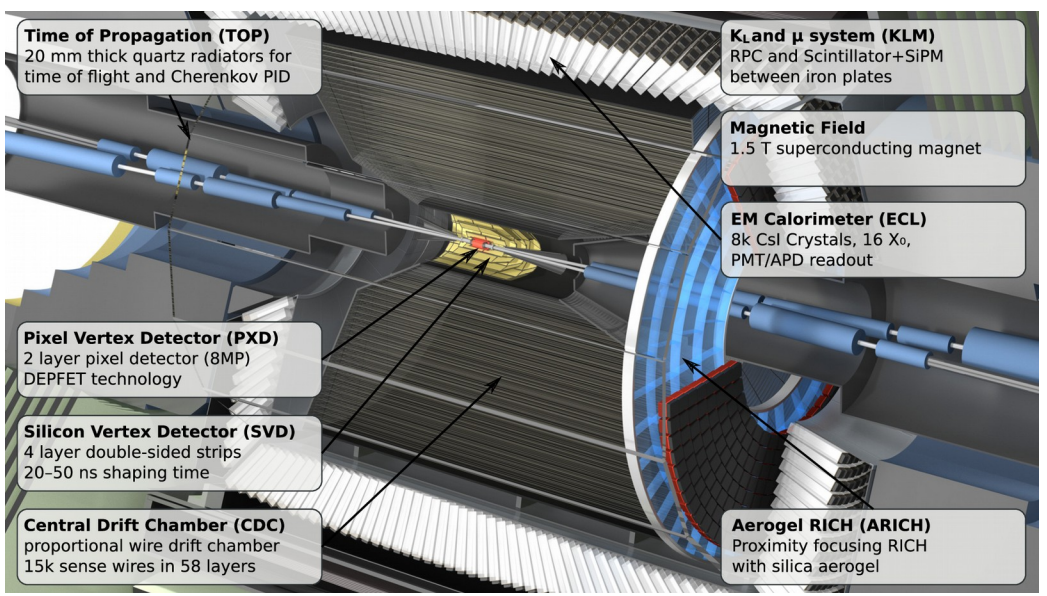

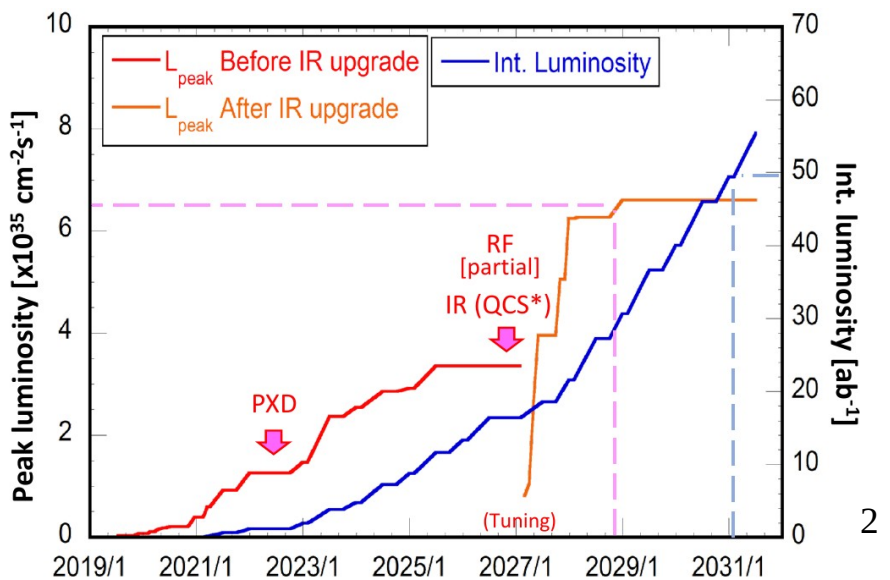

Belle II DAQ Components

- Trigger and Timing Distribution
- Data Readout and Event Builder
- Slow Control
- High Level Trigger Farm: Talk by M. Prim
- Data Storage and Transfer: Talk by M.Barrett

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

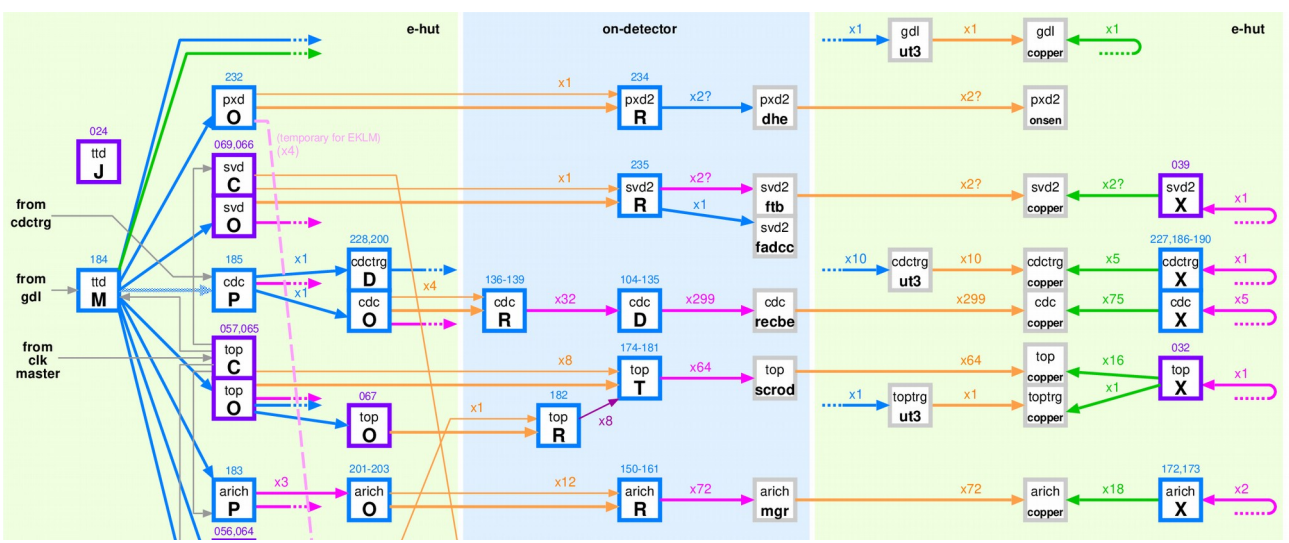

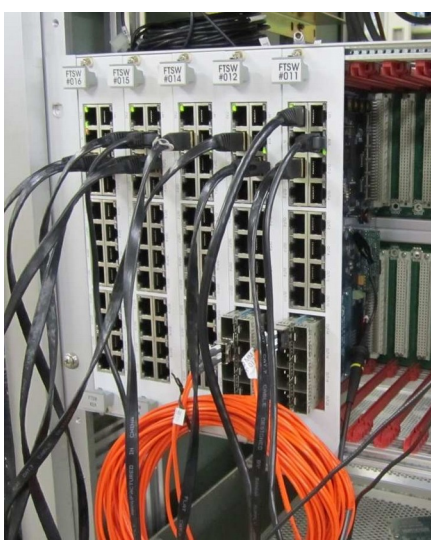

Data Transfer

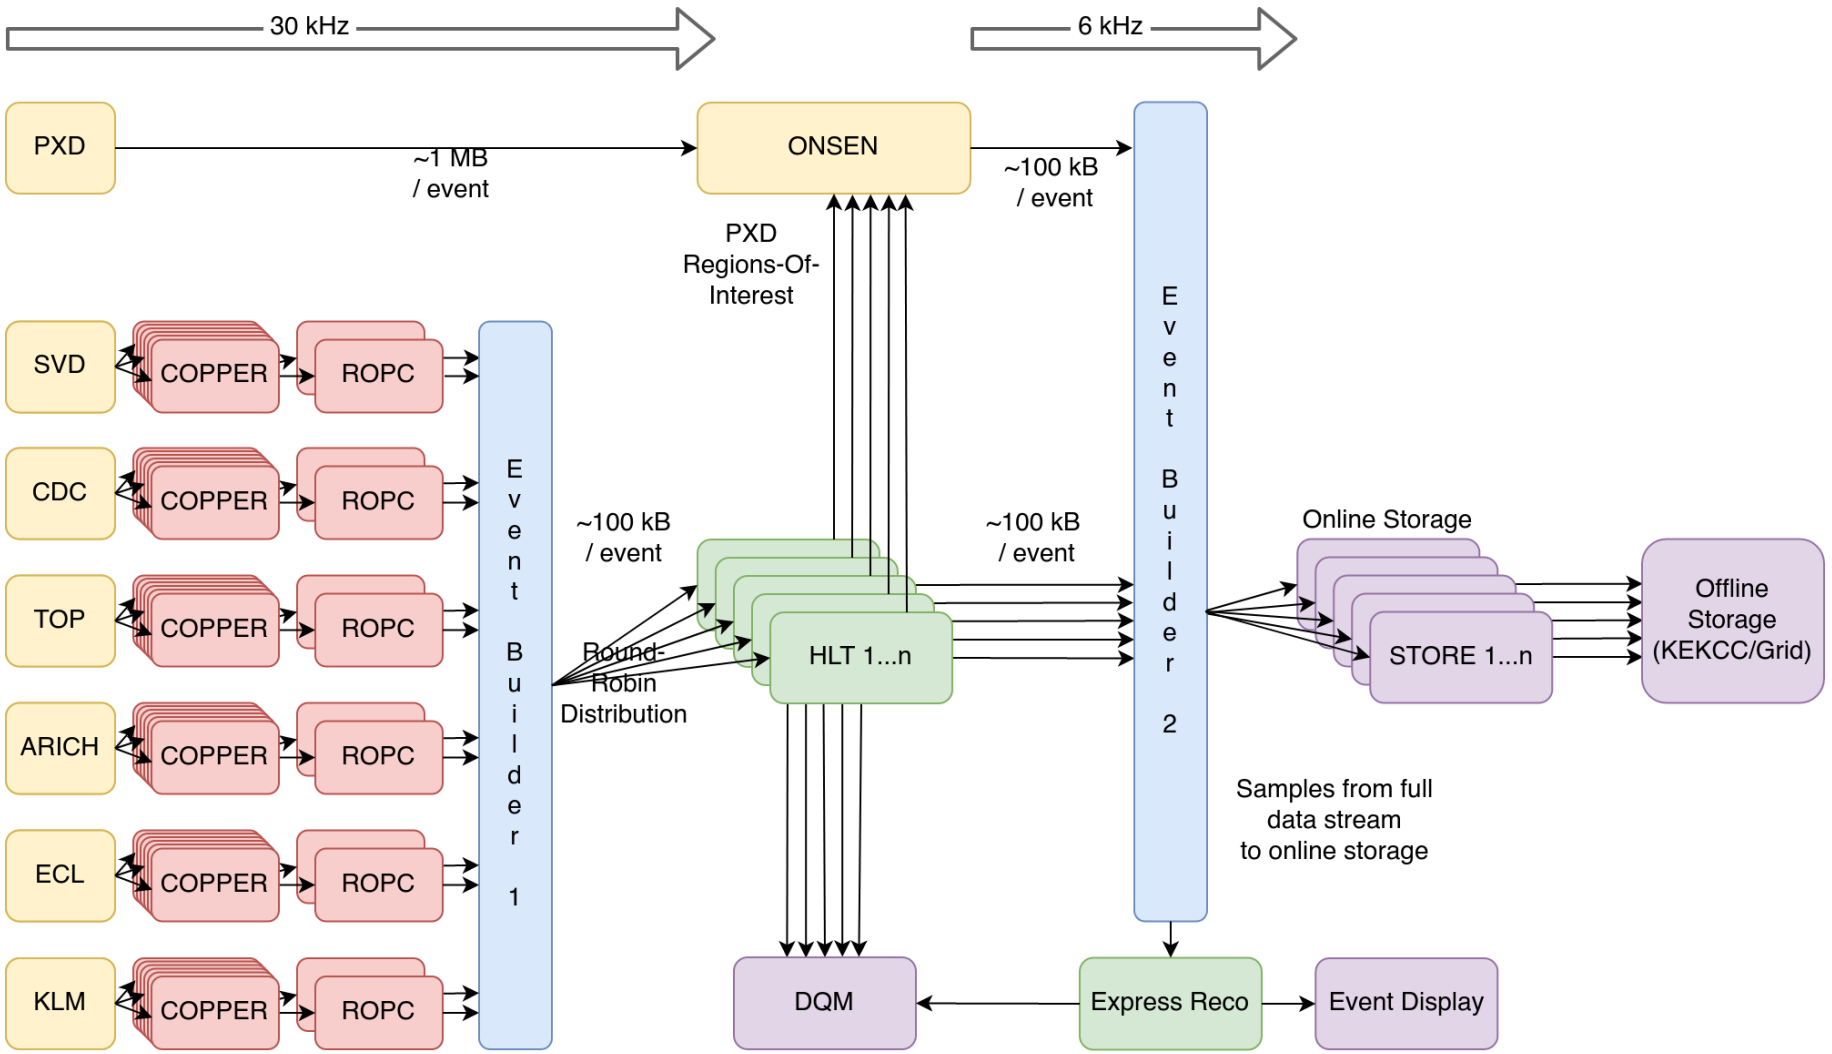

Copper + ROPC

- COPPER (Common Pipelined Platform for Electronics Readout)
  - Receives up to four fibers from front-ends via custom Belle2link protocol (up to 2.54Gbps)
  - Receives central trigger and clock information
  - Integrated Atom CPU board for data packaging, checksums etc.
  - GbE transfer to ROPCs
  - ~200 COPPERs serving ~650 Belle2links in Belle II
- 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

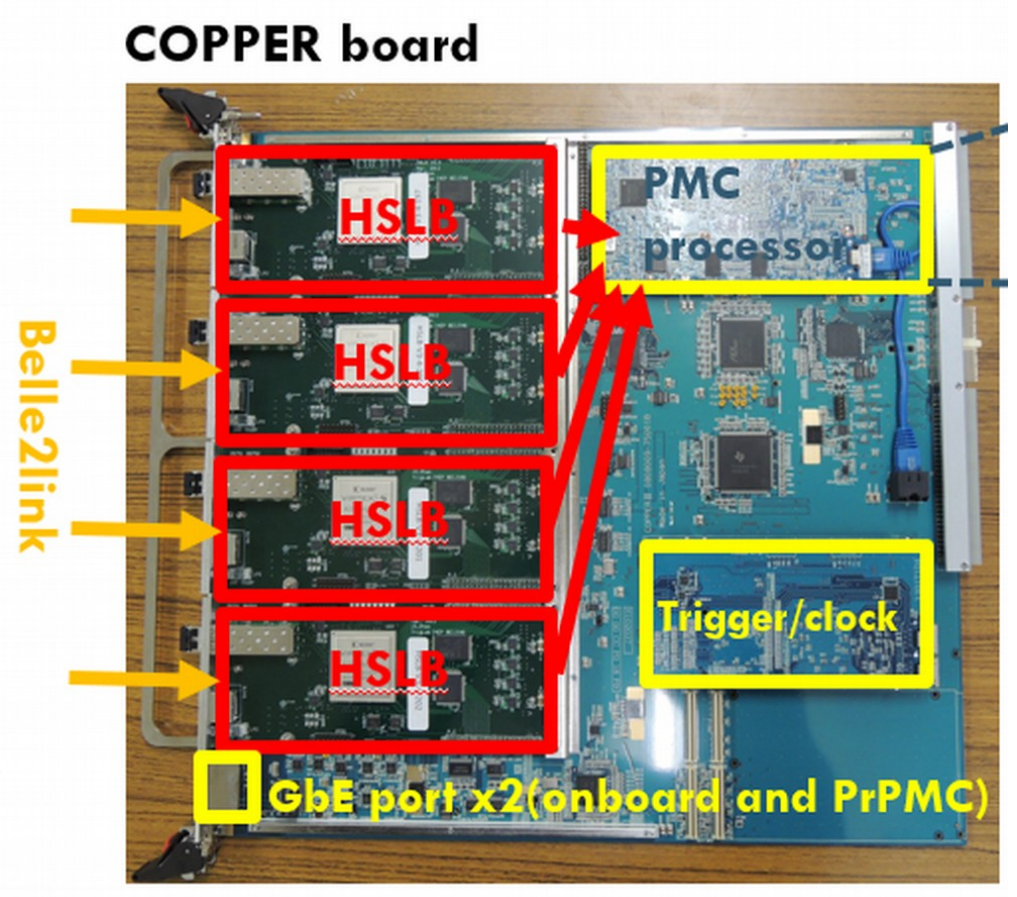

Slow Control

- Non-event process variables (PVs): run control, detector status, machine parameters, etc.
- 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
  - Largely automated build system for all target distributions in Docker containers
  - Central software distribution through local CVMFS installation
  - Working on integrating all non-DAQ hosts into centralised deployment

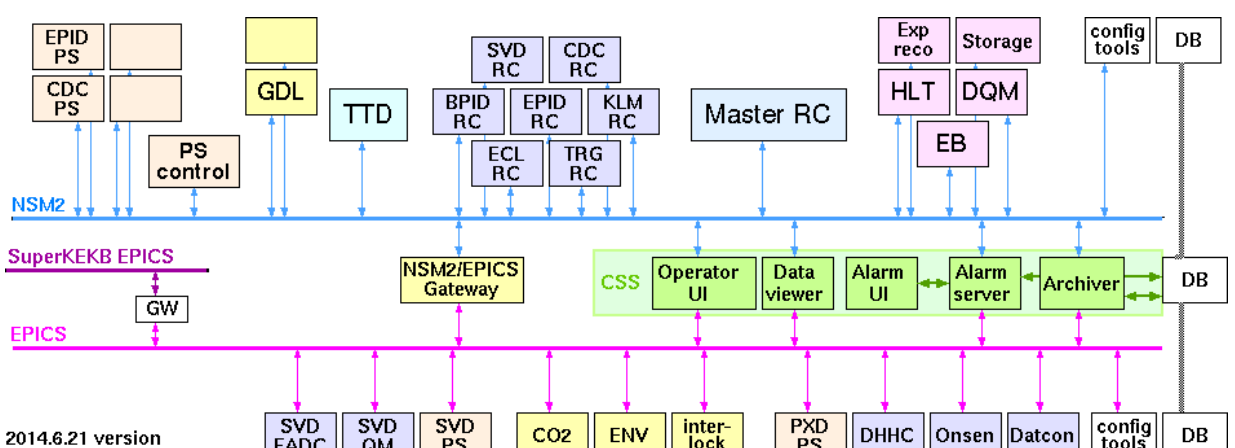

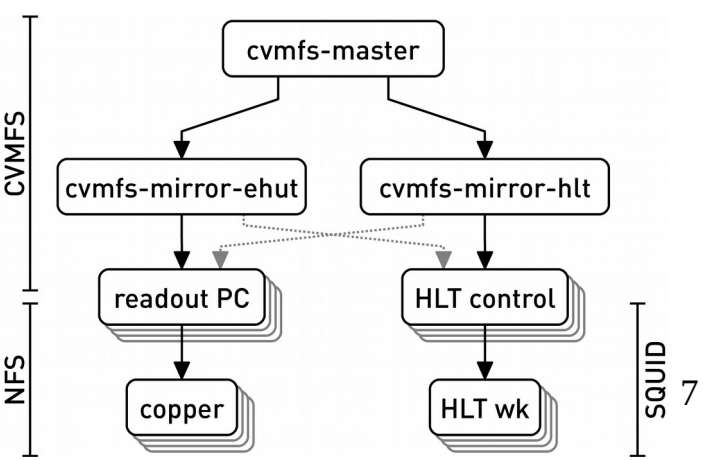

Monitoring

- Zabbix for monitoring individual hosts
  - crucial processes, software versions etc.
- Monitoring machines and processes in the Belle II DAQ network using Elastic Stack:
  - Elasticsearch: database/search engine
  - Logstash: ingest pipeline
  - Kibana: web-based visualisation
  - Extremely scalable, open source, huge community
- Predefined dashboards for subdetectors, DAQ experts, run coordinators etc.
- Integrated alarm messages to central Belle II operations RocketChat channel
- Significant step forward for operations and post-mortem error analysis

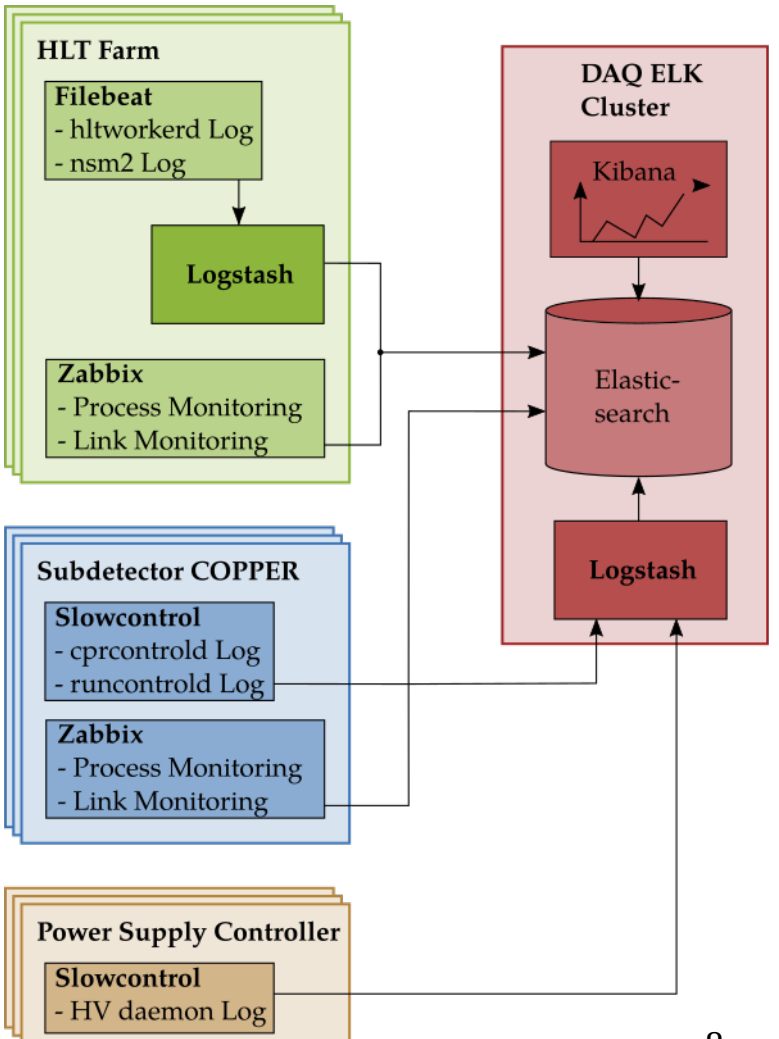

DAQ Operation in 2020

- Belle II and SuperKEKB operated through the pandemic, see talk by K. Matsuoka
- Overall data taking efficiency ~85%
  - ~5% lost to issues requiring expert intervention
  - ~5% lost to issues fixed by run restart
  - ~5% lost to injection trigger veto
  - Many days without any significant loss due to DAQ/detector troubles
    → major improvement over 2019 run periods
- Sufficient readout bandwidth to keep most trigger lines wide open
  - Running at ~3% of design lumi,
    but ~20% of design trigger rate

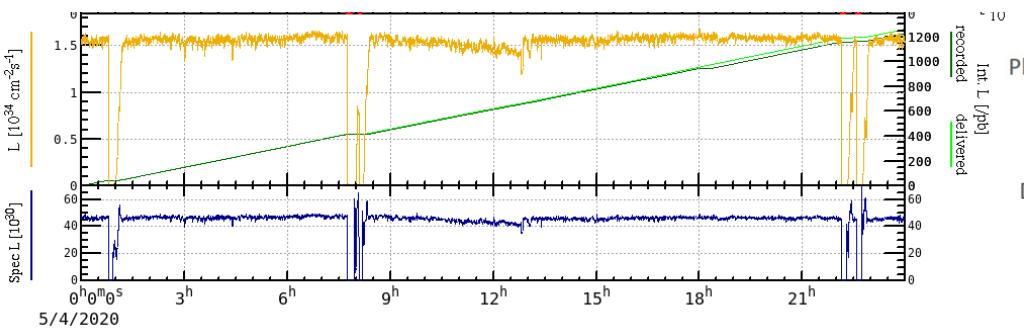

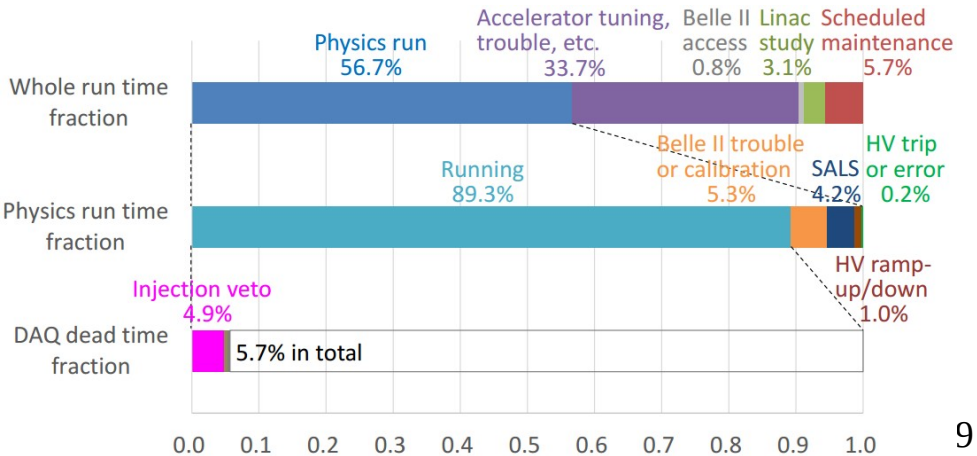

Belle II DAQ Upgrade Project

- 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
- Selection process: DAQ boards used/developed for other HEP applications
  - FELIX: ATLAS (trigger/daq upgrade), BNL, USA
  - CPPF: CMS (trigger), IHEP, China
  - PCIe40: ALICE + LHCb, CPPM + IJCLab, France
- All contenders shown to fullfil requirements, selected PCIe40

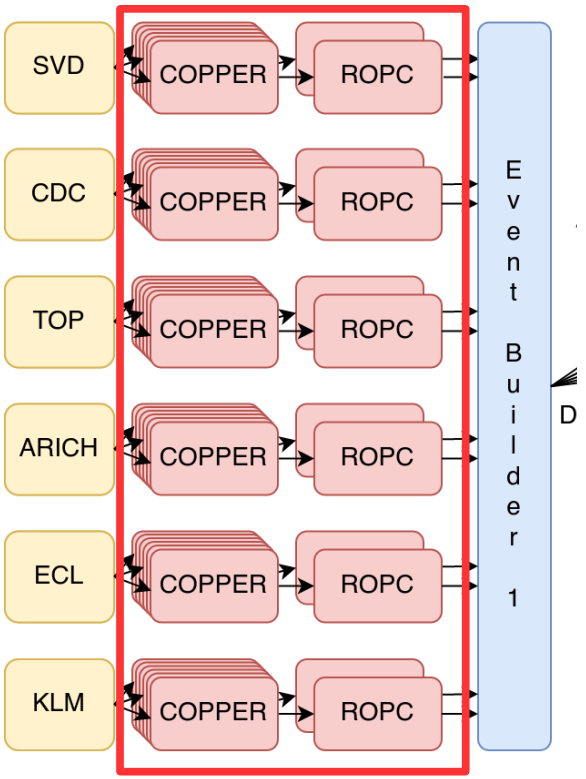

PCIe40 + Layout

- PCIe card (2x8 lanes) mounted in host server
  - Based on Intel Arria10 FPGA
  - 48 bidirectional fiber links
  - Direct Memory Access data transfer to host
  - Large parts of firmware already implemented from ALICE/LHCb projects
- 19 PCIe40 cards needed to convert whole Belle II readout
  - 31 cards produced, tested and on their way to KEK now
- PCIe40 host servers connect to event builder via 10GbE on new switches
- Plan for full conversion of TOP and KLM systems to PCIe40 readout by September, to be used in physics runs starting in October

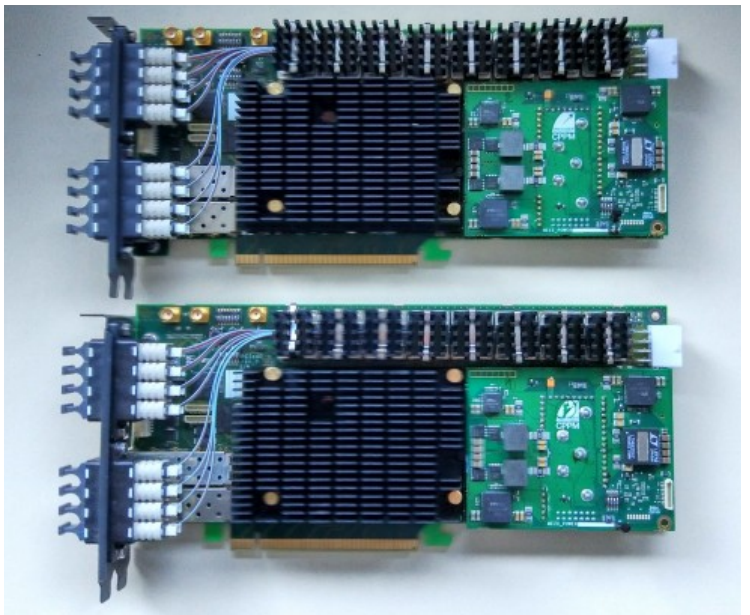

Upgrade Status

- Test benches available at KEK and Hawaii
  - KEK: CDC, ARICH front-ends, dummy data sources
  - Hawaii: KLM + TOP front-ends
- Firmware utilising all 48 links works
  - Front-end register access for slow control validated at full readout speed
- Initial versions of modified slow control framework software tools available, integration into test benches ongoing
- Expect full DAQ integration tests in August
  - So far little delay due to COVID travel restrictions, but integration tests might suffer

Summary

- Belle II DAQ system operated smoothly in 2020 period
  - Significant improvements in stability since 2019

- PCIe40 chosen as platform for Belle II DAQ upgrade
  - Development well underway, integration tests starting ~now
  - Plan to convert two Belle II subdetectors to upgraded PCIe40 infrastructure for next beam period in October