



Online computing challenges: detector and read-out requirements

Richard Brenner¹, Christos Leonidopoulos^{2,a} 

¹ Uppsala University, Uppsala, Sweden

² University of Edinburgh, Edinburgh, UK

Received: 9 June 2021 / Accepted: 7 November 2021

© The Author(s) 2021, corrected publication 2022

Abstract The operation at the Z-pole of the FCC-ee machine will deliver the highest possible instantaneous luminosities with the goal of collecting the largest Z boson datasets (Tera-Z), and enable a programme of standard model physics studies with unprecedented precision. The data acquisition and trigger systems of the FCC-ee experiments must be designed to be as unbiased and robust as possible, with the goal of containing the systematic uncertainties associated with these datasets at the smallest possible level, in order to not compromise the extremely small statistical uncertainties. In designing these experiments, we are confronted by questions on detector read-out speeds with an extremely tight material and power budget, trigger systems with a first hardware level or implemented exclusively on software, impact of background sources on event sizes, ultimate precision luminosity monitoring (to the 10^{-5} – 10^{-4} level) and sensitivity to a broad range of non-conventional exotic signatures, such as long-lived non-relativistic particles. We will review the various challenges on online selection for the most demanding Tera-Z running scenario and the constraints they pose on the design of FCC-ee detectors.

1 Introduction

The FCC-ee machine is expected to deliver the highest instantaneous luminosities ever achieved, forcing a re-evaluation of the requirements for trigger and data acquisition (DAQ) systems.

The conventional wisdom is that the trigger systems of FCC-ee experiments must rely on simple (low- or minimum-bias¹) triggers with built-in redundancy, e.g. calorimeter-based, muon-based or tracker-based. For example, in the LEP era [1], the online selection was established from calorimeter- and tracker-based triggers. For the ILC studies [2], the assumption has been that the experiments will rely on a ‘triggerless’ DAQ (i.e. no first-level hardware trigger), exploiting the relatively small collision rates. It is worth mentioning that LHCb [3], one of the current experiments, is going to collect all detector data from collisions and feed

¹ In the trigger context, minimum-bias events are inelastic events, typically from both diffractive and non-diffractive processes, which have been selected with as little bias as possible during the online filtering carried out by a trigger system. A low-bias trigger typically applies an additional requirement, like the existence of a low-energy particle candidate such as an electron, muon or jet, in the event.

^a e-mail: Christos.Leonidopoulos@cern.ch (corresponding author)

it into an event selector that will run entirely in software. The experimental environment at FCC-ee is, however, very different from that at LHCb. The event rate is significantly lower than at a hadron collider, but the material budget is much tighter which limits the services and read-out bandwidth. Compared with previous experiments at lepton colliders, the challenge for FCC-ee experiments is the very large data rates (~ 200 kHz when running at the Z-pole), which are orders of magnitude larger than at LEP and are significantly higher than at Belle II.

In this essay, we review studies of hardware and software solutions that will allow FCC-ee experiments to record all of the interesting physics events with very high efficiency and redundancy, leading to minimum uncertainties and biases in the experimental measurements.

2 Event sizes and collision rates at FCC-ee

The parameters of interest in a trigger and DAQ (TDAQ) design are:

- the rate of interesting physics events to record
- the rate of irreducible background events
- the average event size
- the data throughput, which is the product of event size and the interaction rate, corresponding to the data volume per unit time. The data throughput can be derived from the previous three parameters, and is the key element of the TDAQ that determines the read-out and write-out capacity of the system.

At FCC-ee, the nominal instantaneous luminosities per interaction point range from a maximum of $230 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ at the Z-pole, down to 28 (WW), 8.5 (ZH) and $1.8 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ ($t\bar{t}$) for the other running scenarios [4]. In the context of TDAQ design, we concentrate on the most rate-demanding operations at the Z-pole. The total event rate in this scenario is around 200 kHz (Table 1), with the dominant processes including Z-production (100 kHz), Bhabha scattering (50 kHz), $\gamma\gamma \rightarrow$ hadron production (30 kHz) and beam background (20 kHz) [5].

The basic assumption for trigger operation of the FCC-ee experiments is that all interesting physics can be stored for offline analysis with $\sim 100\%$ efficiency, and that the beam background (expected to be $\sim 10\%$ of the total event rate at the Z-pole) is not a major consideration for DAQ. To evaluate average event sizes for each process and their overall contribution to data throughput, several assumptions are made about detector granularities, and that the experiments will be relying on stable zero-suppression² during DAQ. The studies described in Ref. [4] have been performed using early prototypes of the CLD [6] and IDEA [7] detectors and are summarised in Table 2. The main conclusion is that with an appropriate zero-suppression scheme, the major contribution to the average event size for the IDEA detector is from physics, and it should be possible to keep the main backgrounds (e.g. synchrotron radiation) under control at a relatively small fraction of the total event rate.

One of the most important ingredients of precision measurements at lepton colliders is the accurate determination of the integrated luminosity. The FCC-ee programme has the very ambitious goal of measuring the absolute luminosity at the 10^{-4} level, with the relative luminosity values between energy-scan points measured 10 times better, at the 10^{-5} level. The challenges for the FCC-ee luminosity monitor design and possible solutions are discussed in a separate essay [8].

² Zero-suppression is the process during which data in the front-end electronics or the read-out system is only recorded above a certain significance level.

Table 1 Event rates expected for various processes at the Z-pole at the FCC-ee [4,5]. The beam background is expected to be $\sim 10\%$ of the total event rate

Physics process	Rate (kHz)
Z decays	100
$\gamma\gamma \rightarrow$ hadrons	30
Bhabha	50
Beam background	20
Total	~ 200

Table 2 Average event data rates expected for the CLD and IDEA subdetectors at the Z-pole for the FCC-ee [4,5]

Subdetector	Physics	Background/noise
CLD vertex detector	150 MB/s	6 GB/s
CLD tracker	160 MB/s	10 GB/s
IDEA drift chamber	60 GB/s	2 GB/s
IDEA Si wrapper	32 MB/s	0.5 GB/s
IDEA DR calorimeter	10 GB/s	1.6 TB/s*
IDEA pre-shower	320 MB/s	820 MB/s
IDEA Muon detector	4 MB/s	67 MB/s

*Assuming no suppression for isolated counts

3 Triggerless design

A software-based online selection (the so-called ‘triggerless’ design) provides a flexibility that cannot be matched by traditional first-level hardware-based filtering systems. Typically, as luminosity increases, a calorimeter- and muon-based filtering system cannot sustain the trigger rates without an increase of threshold values. This means that a trigger inefficiency, which cannot be compensated by a higher integrated luminosity, is incurred. A purely software trigger also allows increased complexity in the online selection by combining information from different systems, provided the algorithmic execution fits within the processing time budget.

For an FCC-ee experiment, the major challenge to a triggerless design is the very high instantaneous luminosity, in particular at the Z-pole ($2.3 \times 10^{36} \text{ cm}^{-2} \text{ s}^{-1}$). R&D studies assume that zero-suppression will be routinely applied at read-out time. However, this requirement necessitates not only a careful calibration (and alignment) of the subdetector systems, but also a technical solution for a trigger that can be deployed online, and updated in semi-real time. In that sense, smooth and stable running conditions and a robust monitoring system are of paramount importance, as has been demonstrated in operation at LEP and the B factories.

Detector choices can, of course, have an impact on the trigger and DAQ design. Consider, for example, the following options:

- Tracking: a time projection chamber (TPC), which is favoured by tracking experts for its lightweight design, cannot be read out every 20 ns. A TPC-based detector would,

therefore, necessitate a hardware-based first filtering level in order to reduce the read-out of the TPC.

- Calorimetry: a fine-granularity but noisy calorimeter, where the application of zero-suppression at the trigger may not be straightforward. For example, the noise data rates for the IDEA DR Calorimeter (Table 2) where the working assumption is no suppression for isolated counts. The noise in the Z-pole running scenario contributes significantly more to the average event data rates (1.6 TB/s) than the physics (10 GB/s), a constraint that would interfere with the optimisation of trigger efficiency of electromagnetic showers.

It is therefore necessary to balance the detector requirements against operational considerations and constraints on the trigger and DAQ systems when designing the FCC-ee experiments.

With TDAQ technology evolving rapidly, and FCC-ee still far into the future, it is perhaps too early to discuss details of concrete TDAQ designs and implementation. However, it is still instructive to review some recent advances in HEP experiments that may be relevant when designing TDAQ systems for FCC-ee experiments. We do this in the following subsections, by reviewing the LHCb TDAQ system, the read-out of ultra-light vertex detectors and tracking with ultra-light TPCs.

3.1 A recent example: the LHCb TDAQ system

The LHCb experiment has opted for a triggerless design for the Run-3 upgrade [3, 9, 10]. Event selection is implemented entirely in software in order to maximise physics yield, increase the amount of data collected, minimise cost and retain flexibility.

LHCb employs FPGAs in the ‘middle-layer’ of the DAQ, which is the electronics bridging the read-out of the detectors to the optical fibres that bring the signal to the data centre on the surface. The middle-layer is physically located at the periphery of the detector where the radiation level is moderate, and this allows LHCb to operate FPGAs near the detector hardware. The experiment applies zero-suppression directly at the detector read-out. FPGAs will be tested up to a total ionising dose of 200 krad and a high energy hadron (HEH) fluence of $1.2 \times 10^{12} \text{ cm}^{-2}$ [11]. Given the reduced geometrical acceptance of the detector, the total event size is comparatively small for an LHC experiment and yields ~ 100 kBytes. LHCb is able to perform online selection with offline-like reconstruction, eliminating a significant source of biases.

There are two levels of software filtering: the first one, running on GPUs, reduces the data throughput from 32 Tb/s to 1-2 Tb/s, with the second one bringing it down further to 80 Gb/s. A large storage buffer allows the effective decoupling of the two sub-systems. Calibration and alignment are performed in ‘semi-live’ mode, while the data are being buffered. The challenges to the LHCb TDAQ design are the large memory consumption and the capacity of the network (which needs to collect data in a single physical location from 478 FPGA boards for event selection). High data rates in a fixed bandwidth system can lead to network congestion. To avoid any data loss, LHCb requires a lossless network. The approach followed is to prioritise the network traffic by application, in a technique known as ‘traffic-shaping’ that helps optimise performance and improve the overall latency of the data flow.

The biggest worry that comes with the LHCb design is the scalability and reliability of the network. TDAQ experts’ estimates are that the online system can sustain up to a factor of three increase in network traffic without any substantial changes in its architecture [10]. This scalability is expected to be adequate for the duration of Run-3.

3.2 Read-out of ultra-light vertex detectors

The development of monolithic active pixel sensors (MAPS) has been pioneered by the communities developing detectors for experiments with lepton and heavy ion beams, where low material budget is essential. The most advanced tracker in operation that is using MAPS technology is the inner tracking system of ALICE [12]. It uses a custom made ALPIDE ASIC [13], which has a thick high-resistivity epitaxial layer for sensing charged particles with read-out electronics embedded in a low resistivity n-well. The pixels in the tracker are $29\ \mu\text{m} \times 27\ \mu\text{m}$ in size. Every pixel is equipped with an amplifier, comparator and a 3-hit deep register for storing data. Read-out of the 1024×512 pixel matrix is done with electronics on the periphery that allow up to 1200 Mbps data rate, giving the chip a theoretical maximum hit data transfer capacity of $6\ \text{MHz}/\text{cm}^2$. The ASIC has a continuously active front-end and the read-out is zero suppressed, offering both triggered and continuous read-out with time-stamped hit information. The chip has demonstrated a read-out rate of 100 kHz for Pb-Pb interactions. The performance of ALPIDE already meets the rates presented in Table 2. The material budget of ALPIDE is around $0.3\% X_0/\text{layer}$, which is higher than the target of $0.15\text{--}0.2\% X_0/\text{layer}$ for a vertex detector in a future e^+e^- experiment. Development will be required to optimise the read-out performance with the material and power budgets available.

There are many improvements to the MAPS technology that can be expected for the FCC. The charge collection in the pixel will become significantly faster with HV-CMOS technology, which is based entirely on drifting. Improved time-stamping of hits will allow full 4D tracking. Having 4D tracking capability at FCC-ee detectors opens up the possibility of separating multiple interaction points, and could improve the sensitivity to, for example, slowly propagating exotic particles.

The tracker community is continuously developing ways to reduce the amount of passive material in the tracker. Low power electronics combined with evaporative CO_2 cooling with microchannels etched in silicon are the most recent developments to reach operation in an experiment. Wireless data transmission is being developed by the WADAPT collaboration [14] with the goal of demonstrating the feasibility of this technology for low power, Gbps rate and the short-range data links employed in tracking detectors. The technology will increase the read-out bandwidth without increasing the amount of material in the detector. Local wireless communication inside the tracker can also be used for data reduction inside the tracker prior to read-out.

3.3 Tracking with ultra-light TPC

Time projection chambers are an attractive option for lepton collider experiments since low-mass trackers offer high hit precision, which in turn delivers excellent momentum resolution. As discussed earlier, an important consideration is whether their read-out is compatible with the requirements of a triggerless design. Here, we review methods recently developed by the ALICE collaboration to tackle this problem.

The challenge in developing a TPC with triggerless read-out for ALICE is to master the huge out-of-bunch pile-up during TPC drift time. This is particularly important for searches of rare signals in a massive background. The requirement is to store full minimum bias samples for analysis. The read-out rate for a conventional TPC is limited to about 3 kHz by the gating grid that blocks the back flow of ions. Replacing the multi wire proportional chamber with a gas electron multiplier removes the need for a gating grid since the intrinsic back flow is low. It also opens the possibility for triggerless read-out since the TPC can now be operated in continuous mode. The number of ions entering the TPC region is still large enough to

produce an electric field that distorts the path of the electrons during drifting. The effect of the distortion is rate dependent and must be corrected in order to maintain the intrinsic resolution of the TPC. The correction of space charge distortion will be an important part of the TDAQ.

Two main methods have been developed for correcting the space charge distortions in ALICE:

- Track-based, where tracks are refitted using information from all tracking detectors. It is very accurate but requires large samples of tracks to deliver the necessary statistical precision.
- Integration of currents arriving on the TPC end plates. It requires a full ion history which, in turn, requires continuous read-out. The method is probably less accurate than the track-based method.

In ALICE, the processing of TPC hits is based on 10–20 ms time frames. The calibration intervals need to be fast (~ 5 ms) to be able to correct for short time fluctuations.

Another challenge is the absence of absolute z -coordinates together with the aforementioned space charge distortions. To fully take advantage of the high intrinsic precision of the TPC, the TDAQ will need to have the resources to carry out track reconstruction to calibrate and compensate for any distortions.

Groups developing the new TPC for ALICE have been actively involved in technological innovations for e^+e^- colliders through the AIDA programme [15]. The improvements introduced in the TPC make it a serious alternative for FCC-ee experiments because of the low material budget and continuous tracking capability. The interplay of the TPC- and MAPS-based tracking needs to be optimised to have the best physics performance.

4 The challenge of long-lived non-relativistic particles

The inclusion of the tracking detectors in the fast data handling path will continue to be a big challenge for TDAQ in current and future experiments. The large quantities of data generated by tracking detectors need to be rapidly transferred to off-detector processors. This is not easily done due to the strict requirement of a low material budget for the tracking detectors resulting from the physics goal of high precision measurements. The availability of tracking information through the full TDAQ processing chain will be important for expanding the physics programme to new exotic physics signatures that may otherwise have escaped undiscovered.

Many new models beyond the standard model (SM) postulate long-lived particles (LLP) that have a lifetime long enough for them to travel some measurable distance inside the detector before decaying. These exotic particles differ from the known long-lived particles in the SM in terms of their unusual ionisation and propagation properties. The physics programme at FCC will require reconstruction algorithms in searches for LLP to capture these signatures. Searches for new physics with LLP is a challenge for TDAQ due to the various signatures which have appearing and disappearing tracks that do not point to the primary vertex. The capability to select LLP events in real-time data is usually not a priority in the design phase of an experiment because they are not the primary physics motivation. The complexity of the signature also makes it difficult to find good metrics for the production of design requirements for LLP physics.

A lesson from current experiments which should be considered in future designs is that it is important to have a clear strategy for LLP at an early stage and this may require additional

(and potentially distant) detectors to be integrated in the TDAQ system. If the read-out has to be included in the trigger, the experiment needs a read-out segmentation that lends itself to LLP triggering. This can include hardware track triggering being instrumented in the tracker. The LLP physics programme would, however, benefit from a triggerless read-out including timing information of every hit. This would allow studies of out-of-bunch and out-of-time particles.

5 Machine learning

In recent years, we have witnessed an explosive growth of machine learning techniques in HEP applications. Indeed, there is a strong R&D activity concerning the deployment of machine learning technologies in both off-the-shelf commercial processors and FPGAs with a more limited computing footprint [16, 17]. Some examples include: front-end data compression, particle identification with multivariate classifiers, pattern recognition, tracking and reconstruction with neural networks and regression for improved resolution. It is widely anticipated that some of these innovations will soon find their way into TDAQ systems [18, 19]. We expect that in the next five years these developments will mature and form the basis for proto-TDAQ design contributions to the FCC feasibility study.

6 Conclusions

The biggest challenge for FCC-ee experiments is the all-time record high instantaneous luminosity at the Z-pole. TDAQ systems sustaining similar data throughput rates will already be operational at the LHC. The baseline assumption for future lepton collider experiments is that they will rely on a software-only triggering system with $\sim 100\%$ efficiency and built-in redundancy. The possibility of using an ultra-light tracking detector, such as TPC, in a triggerless system requires the management of the very large out-of-bunch pile-up and its operation in continuous mode. Full timing information of detector hits is expected to be fully exploited in the TDAQ of FCC-ee experiments because it will be needed for calibration, reconstruction and in searches for new exotic particles. Of course, when making projections on rapidly developing technologies and how these may affect the design of future experiments, it should be remembered that they come with a degree of uncertainty.

Funding Funding was provided by Science and Technology Facilities Council (Grant number ST/S000828/1). This project is co-funded from the European Union's Horizon 2020 research and innovation programme under grant agreement No 95175.

Data Availability Statement This manuscript has no associated data or the data will not be deposited. [Author's comment: There are no associated data available.]

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. R. Tenchini, LEP trigger strategies. Presentation given at the 1st FCC-ee Workshop on Detector Requirements . <https://indico.cern.ch/event/393093/contributions/1830018>
2. ILC Collaboration, The international linear collider. Technical Design Report, Vol. 1. <https://arxiv.org/abs/1306.6327> arXiv:1306.6327
3. LHCb Collaboration, LHCb trigger and online upgrade technical design report, CERN-LHCC-2014-016 ; LHCb-TDR-016 <https://cds.cern.ch/record/1701361>
4. A. Abada et al., FCC-ee: The Lepton Collider, future circular collider conceptual design report volume 2. Eur. Phys. J. Spec. Top. **228**, 261 (2019). <https://doi.org/10.1140/epjst/e2019-900045-4>
5. Fr. Grancagnolo, Event sizes at the Z-pole. Presentation Given at the 4th FCC Physics and Experiments Workshop. <https://indico.cern.ch/event/932973/timetable/#day-2020-11-12>
6. N. Bacchetta, et al., CLD—a detector concept for the FCC-ee (2019). arXiv:1911.12230 [physics.ins-det]
7. F. Bedeschi, et al., IDEA detector letter of intent, Snowmass 2021. https://www.snowmass21.org/docs/files/summaries/EF/SNOWMASS21-EF1_EF4-IF3_IF6-096.pdf
8. M. Dam, Challenges for FCC-ee luminosity monitor design (2021). arXiv:2107.12837 [physics.ins-det]
9. R. Aaij, et al., A comparison of CPU and GPU implementations for the LHCb experiment run 3 trigger, LHCb-DP-2021-003 (2021). arXiv:2105.04031 [physics.ins-det]
10. T. Colombo, LHCb Run3 DAQ & event filter, presentation given at the 4th FCC physics and experiments workshop. <https://indico.cern.ch/event/932973/timetable/#day-2020-11-12>
11. V.M. Placinta, et al., Proton-induced radiation effects in the I/O blocks of an SRAM-based FPGA. JINST **14** T10001. <https://iopscience.iop.org/article/10.1088/1748-0221/14/10/T10001>
12. S. Beolè, for the ALICE Collaboration, The ALICE inner tracking system: performance with proton and lead beams. Phys. Procedia, Vol. 37 (2012). ISSN 1875-3892. <https://doi.org/10.1016/j.phpro.2012.02.443>
13. M. Mager, for the ALICE Collaboration, ALPIDE, the monolithic active pixel sensor for the ALICE ITS upgrade, Nucl. Inst. Meth. A, **824** (2016) pp. 434–438. <https://www.sciencedirect.com/science/article/pii/S0168900215011122>
14. R. Brenner, et al., Development of wireless techniques in data and power transmission - application for particle physics detectors (2015). arXiv:1511.05807 [physics.ins-det]
15. AIDA 2020, Advanced European Infrastructures for Detectors at Accelerators, Project funded by European Union's Horizon 2020 Research and Innovation programme under Grant Agreement no. 654168. <https://aida2020.web.cern.ch/aida2020>
16. S. Giagu, for the ATLAS Collaboration, Fast and resource-efficient Deep Neural Network on FPGA for the Phase-II Level-0 muon barrel trigger of the ATLAS experiment, EPJ Web of Conferences **245**, 01021 (2020). <https://doi.org/10.1051/epjconf/202024501021>
17. J. Duarte, J.R. Vlimant, Graph neural networks for particle tracking and reconstruction (2020). arXiv:2012.01249 [hep-ph]
18. ATLAS Collaboration, Technical design report for the phase-II upgrade of the ATLAS TDAQ System. CERN-LHCC-2017-020 ; ATLAS-TDR-029 (2017). <https://cds.cern.ch/record/2285584>
19. CMS Collaboration, The Phase-2 Upgrade of the CMS Level-1 Trigger, CERN-LHCC-2020-004; CMS-TDR-021 (2020). <http://cds.cern.ch/record/2714892>