Online Computing Challenges

Detector & readout requirements for future e^+e^- colliders

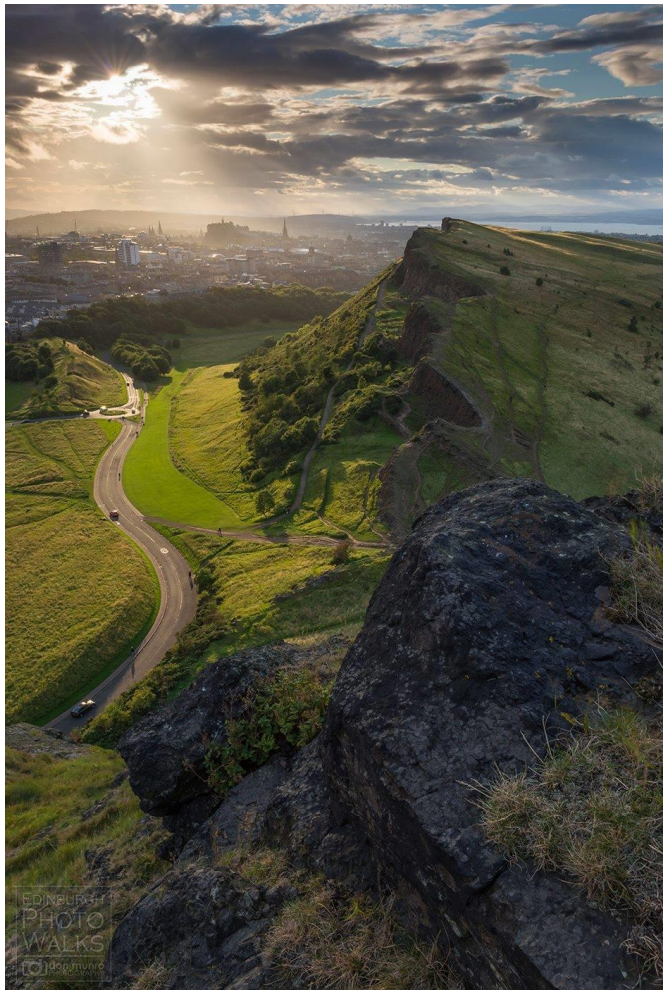

Christos Leonidopoulos

THE UNIVERSITY of EDINBURGH

2021 International Workshop
on the high-energy CepC
10 November 2021

Online Computing Challenges

Detector & readout requirements for future e^+e^- colliders

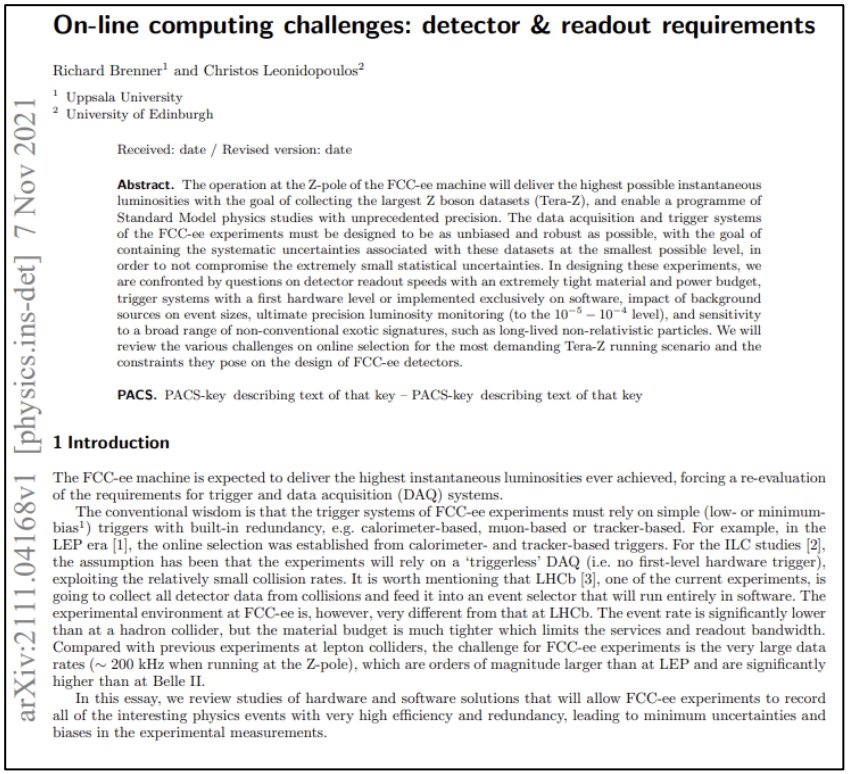

On-line computing challenges: detector & readout requirements

Richard Brenner[1] and Christos Leonidopoulos[2]

[1] Uppsala University
[2] University of Edinburgh

Received: date / Revised version: date

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 readout 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.

PACS. PACS-key describing text of that key – PACS-key describing text of that key

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-bias[1]) 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 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 readout 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.

Talk based on recent invited FCC essay on online computing challenges for future e^+e^- colliders submitted to EPJ+ (jointly with Richard Brenner)

"Focus Point on A Future Higgs & Electroweak Factory (FCC): Challenges towards Discovery"

- Accepted for publication 3 days ago
- Emphasis is on FCCee, but general discussion is machine-agnostic

https://arxiv.org/abs/2111.04168

Online challenges for lepton colliders

- High rates, high precision
- Detector with extremely tight material & power budget
- Event sizes & production rates
- Luminosity measurement (FCCee: at 10^{-4} – 10^{-5} level!)
- First triggering level: custom hardware or software-only?
- Recent detector R&D
- Non-conventional signatures: challenge of triggering on long-lived, non-relativistic particles
- Machine Learning

Online challenges: what do others do?

Conventional wisdom: rely on simple triggers with built-in redundancy

- LEP: when life was simple. Calo-, muon- or tracker-based selection
- ILC: “trigger-less” DAQ (aka: no custom hardware for Level-1 filtering)
- LHCb: collector all detector data from all collisions, and feed into event selection (run entirely on software)
  - But: material budget at future e^+e^- colliders limits readout bandwidth & services

“Good artists copy; Great artists steal”

Instantaneous luminosities: FCCee

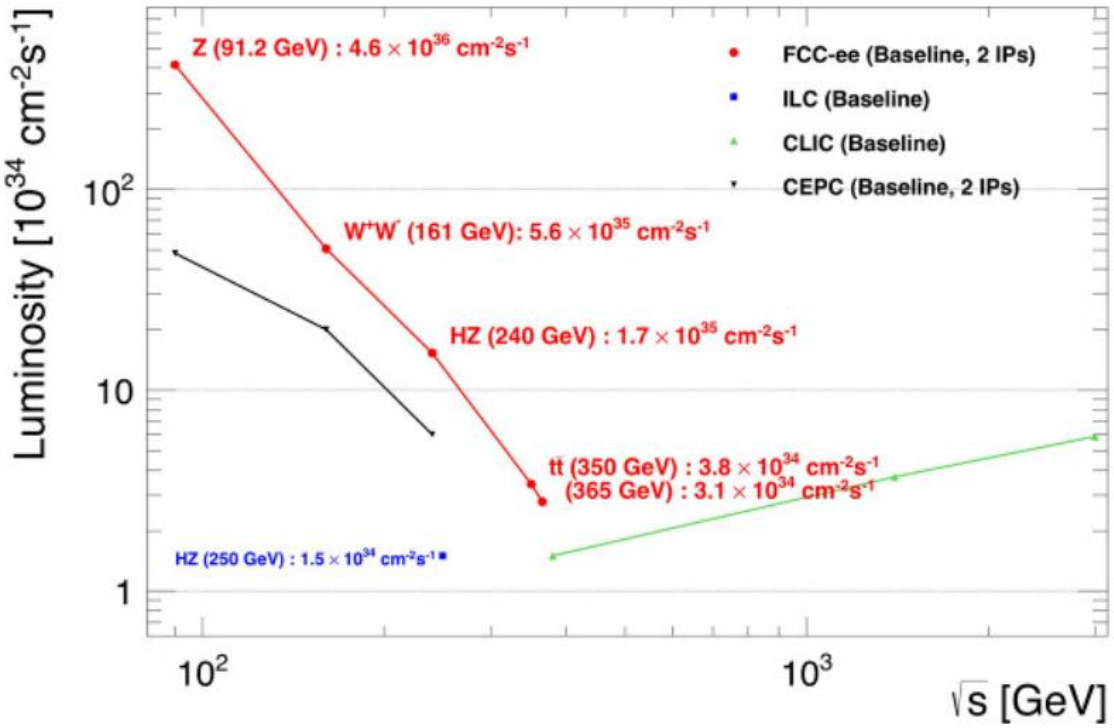

Fig. 2. Baseline luminosities expected to be delivered (summed over all interaction points) as a function of the centre-of-mass energy $\sqrt{s}$, at each of the four worldwide e^+e^- collider projects: ILC (blue square), CLIC (green upward triangles), CEPC (black downward triangles), and FCC-ee (red dots), drawn with a 10% safety margin. The FCC-ee performance data are taken from this volume, the latest incarnation of the CEPC parameters is inferred from [20], and the linear collider luminosities are taken from [15,17].

FCC CDR vol.2: `https://link.springer.com/article/10.1140%2Fepjst%2Fe2019-900045-4`

Instantaneous luminosities: FCCee

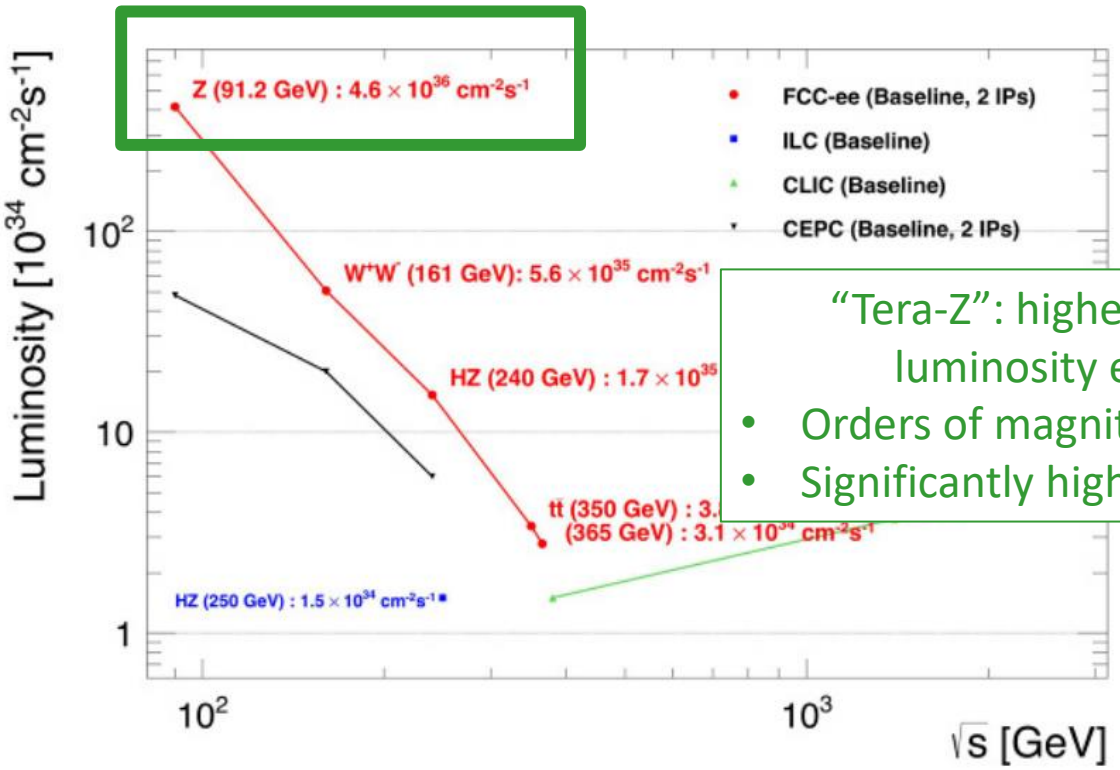

"Tera-Z": highest instantaneous luminosity ever achieved

- Orders of magnitude higher than LEP
- Significantly higher than Belle-II

Fig. 2. Baseline luminosities expected to be delivered (summed over all interaction points) as a function of the centre-of-mass energy $\sqrt{s}$, at each of the four worldwide e^+e^- collider projects: ILC (blue square), CLIC (green upward triangles), CEPC (black downward triangles), and FCC-ee (red dots), drawn with a 10% safety margin. The FCC-ee performance data are taken from this volume, the latest incarnation of the CEPC parameters is inferred from [20], and the linear collider luminosities are taken from [15,17].

https://link.springer.com/article/10.1140%2Fepjst%2Fe2019-900045-4

Instantaneous luminosities: CepC

	ttbar	Higgs	W	Z
Number of Ips		2		
Circumference [km]		100.0		
SR power per beam [MW]		30		
Half crossing angle at IP [mrad]		16.5		
Bending radius [km]		10.7		
Energy [GeV]	180	120	80	45.5
Energy loss per turn [GeV]	9.1	1.8	0.357	0.037
Piwinski angle	1.21	5.94	6.08	24.68
Bunch number	35	249	1297	11951
Bunch population [10^10]	20	14	13.5	14
Beam current [mA]	3.3	16.7	84.1	803.5
Momentum compaction [10^-5]	0.71	0.71	1.43	1.43
Beta functions at IP (bx/by) [m/mm]	1.04/2.7	0.33/1	0.21/1	0.13/0.9
Emittance (ex/ey) [nm/pm]	1.4/4.7	0.64/1.3	0.87/1.7	27/1.4
Beam size at IP (sigx/sigy) [um/nm]	39/113	15/36		/35
Bunch length (SR/total) [mm]	2.2/2.9	2.3/		2.5/8.7
Energy spread (SR/total) [%]	0.15/0.20		0.07/0.14	0.04/0.13
Energy acceptance (DA/RF) [%]	2.3	/2.2	1.2/2.5	1.3/1.7
Beam-beam parameters (ksix/ksiy)	0.071	0.015/0.11	0.012/0.113	0.004/0.127
RF voltage [GV]	10	2.2	0.7	0.12
RF frequency [MHz]	650	650	650	650
HOM power per cavity (5/2/1cell)[kw]	0.4/0.2/0.1	1/0.4/0.2	-/1.8/0.9	-/-/5.8
Qx/Qy/Qs	0.12/0.22/0.078	0.12/0.22/0.049	0.12/0.22/	0.12/0.22/
Beam lifetime (bb/bs)[min]	81/23	39/18	60/717	80/182202
Beam lifetime [min]	18	12.3	55	80
Hour glass Factor	0.89	0.9	0.9	0.97
Luminosity per IP[1e34/cm^2/s]	0.5	5.0	16	115

2021 Improved Design

67%⇑

259%⇑

Rates & event sizes at colliders

- Three (or four) parameters here
  - Rate of interesting physics to record
  - Event size
  - Data throughput (ie. Read-out & write-out data volume/time)

- Key TDAQ parameter: data throughput, not rate!
  - Capacity: data volume per unit time =

    (event size) × (interesting physics rate)
  - Determining readout & write-out capacity of system

Rates & event sizes at LHC

Experiment	Rate	Event size	Throughput
Detector Readout			
ATLAS/CMS Run 1/2	100 kHz	1 MB	100 GB/s
LHCb Run 1/2	1 MHz	100 kB	100 GB/s
ATLAS/CMS Run 4 –	O(500 kHz)	4 MB (PU = 200)	2 TB/s
LHCb Run 4 –	40 MHz	100 kB	4 TB/s
Throughput to disk			
ATLAS/CMS Run 1/2	1-2 kHz	1 MB	1-2 GB/s
LHCb Run 1/2	10 kHz	100 kB	1 GB/s
ATLAS/CMS Run 4 –	5 kHz	4 MB (PU = 200)	20 GB/s
LHCb Run 4 –	20 kHz – ?	100 kB	2 GB/s

Notes:

- Figures refer to order-of-magnitude estimates
- Generally, disk space capacity is the actual bottleneck here, not trigger rate or output to disk

Rates & event sizes at FCCee (Z-pole)

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 ~10% of the total event rate.

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

Basic assumptions

- Store all interesting physics with ~100% efficiency
- Beam background: not a major consideration for DAQ

https://link.springer.com/article/10.1140%2Fepjst%2Fe2019-900045-4

Rates & event sizes at FCCee (Z-pole)

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 ∼10% of the total event rate.

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

https://link.springer.com/article/10.1140%2Fepjst%2Fe2019-900045-4

Rates & event sizes at FCCee (Z-pole)

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 ~10

- 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
- Zero-suppression requires continuous calibration in semi-real time, smooth/stable running conditions, robust monitoring

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

https://link.springer.com/article/10.1140%2Fepjst%2Fe2019-900045-4

Luminosity monitoring

Challenges for FCC-ee Luminosity Monitor Design

Mogens Dam

Niels Bohr Institute, Copenhagen University, Denmark

Received: July 28, 2021/ Revised version: July 28, 2021

Abstract. For cross section measurements, an accurate knowledge of the integrated luminosity is required. The FCC-ee Z lineshape programme sets the ambitious precision goal of 10^{-4} on the *absolute* luminosity measurement and one order of magnitude better on the *relative* measurement between energy-scan points. The luminosity is determined from the rate of small-angle Bhabha scattering, $e^+e^- \to e^+e^-$, where the final state electrons and positrons are detected in dedicated monitors covering small angles from the outgoing beam directions. The constraints on the luminosity monitors are multiple: *i*) they are placed inside the main detector volume only about 1 m from the interaction point; *ii*) they are centred around the outgoing beam lines and do not satisfy the normal axial detector symmetry; *iii*) their coverage is limited by the beam pipe, on the one hand, and the requirement to stay clear of the main detector acceptance, on the other; *iv*) the steep angular dependence of the Bhabha scattering process imposes a geometrical precision on the acceptance limits at about 1 μrad, corresponding to geometrical precisions of $\mathcal{O}(1\,\mu m)$; and *v*) the very high bunch crossing rate of 50 MHz during the Z-pole operation calls for fast readout electronics. Inspired by second-generation LEP luminosity monitors, a proposed ultra-compact solution is based on a sandwich of tungsten-silicon layers. A vigorous R&D programme is needed in order to ensure that such a solution satisfies the more challenging FCC-ee requirements.

https://arxiv.org/abs/2107.12837

Trigger-less design?

- A software-based solution provides flexibility that cannot be matched by traditional first-level hardware-based filtering systems
- For a future e^+e^- collider, the major challenge is the very high luminosity (especially at the Z-pole). R&D studies assume zero-suppression will be routinely applied at read-out. However, this necessitates not only careful calibration (&alignment), but also a technical solution that can be deployed online and updated in semi-real time.
- Smooth & stable running conditions and robust monitoring system are of paramount importance
- Detector choices can have a major impact on TDAQ design. It is important to balance detector requirements against operational considerations & constraints on TDAQ when designing future experiments.

Trigger-less design? #2

- Tracking: Time-Projection Chambers (TPC) which is favoured by tracking experts for lightweight design cannot be read out every 20 ns. A TPC-based detector would require hardware-based filtering system
- Calorimetry: a fine-granularity but noisy calorimeter may lead to non-straightforward zero-suppression (see IDEA example). A high-noise calorimeter that contributes significantly to average event data rates would interfere with optimisation of trigger efficiency of electromagnetic showers.

Detector R&D

With TDAQ technology evolving rapidly, and e^+e^- colliders 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 these experiments.

LHCb TDAQ in Run-3

- FPGA: middle-layer between detector readout & optical fibres brining signal to data centre, located at periphery of detector
- Zero-suppression directly at detector readout
- Event-size relatively small: 100 kB
- Online selection with offline-like reconstruction
- Two-levels of s/w filtering: GPUs 32 → 1-2 Tb/s, CPUs: 80 Gb/s
- Calibration & alignment: semi-live mode, while data is being buffered
- Challenges: large memory consumption, network capacity (data from 478 FPGA boards transferred into single physical location)
- Prioritising network traffic using “traffic-shaping”, optimising performance & improving latency of data flow
- Worries: scalability & reliability

Ultra-light vertex detectors

- Monolithic Active Pixel Sensors (MAPS) technology: being developed for detectors operating with lepton & heavy-ion beams (e.g. Inner Tracking System of ALICE)
- Custom ALPIDE ASIC with theoretical maximum hit data transfer capacity of 6 MHz/cm^2 $\longrightarrow$ 100 kHz for Pb-Pb collisions
- Material budget: 0.3% X_0/layer (need ~ 0.1-0.2% X_0/layer for future e^+e^- experiments)

Future improvements:

- Faster charge collection with HV-CMOS technology
- Time-stamping of hits allows for full 4D tracking $\longrightarrow$ possibility of separating multiple interaction points $\longrightarrow$ slow propagating exotic particles
- Lower-power electronics with evaporative CO_2 cooling
- Wireless transmission (WADAPT collaboration) $\longrightarrow$ increase readout bandwidth without increasing material budget

Ultra-light TPCs

- Low-mass trackers offer high-hit precision
- Challenge: huge out-of-bunch pile-up during TPC drift time
- Conventional TPC: readout ~ 3kHz (gating grid that blocks the back flow of ions: drift velocity ~ 1 m/s)
- Replacing multi-wire proportional chamber with gas electron multiplier removes need for gating grid (since intrinsic back flow is low) $\longrightarrow$ TPC can operated in continuous mode

- Number of ions entering TPC region still large enough to produce electric field that distorts path of electrons during drifting. Effect of charge distortion is rate-dependent and must be corrected in order to maintain intrinsic TPC resolution.
- TPC: potentially serious alternative for e^+e^- tracking
- Interplay of TPC & MAPS needs to be optimised

Long-Lived Particles

- Dark Sector models give rise to long-lived signatures
- Challenge for TDAQ with appearing/ disappearing tracks that do not point to primary vertex
- Selection of LLP events in real-time usually not a priority in design phase of experiments. Complexity of signature makes it harder to find good metrics for design specs

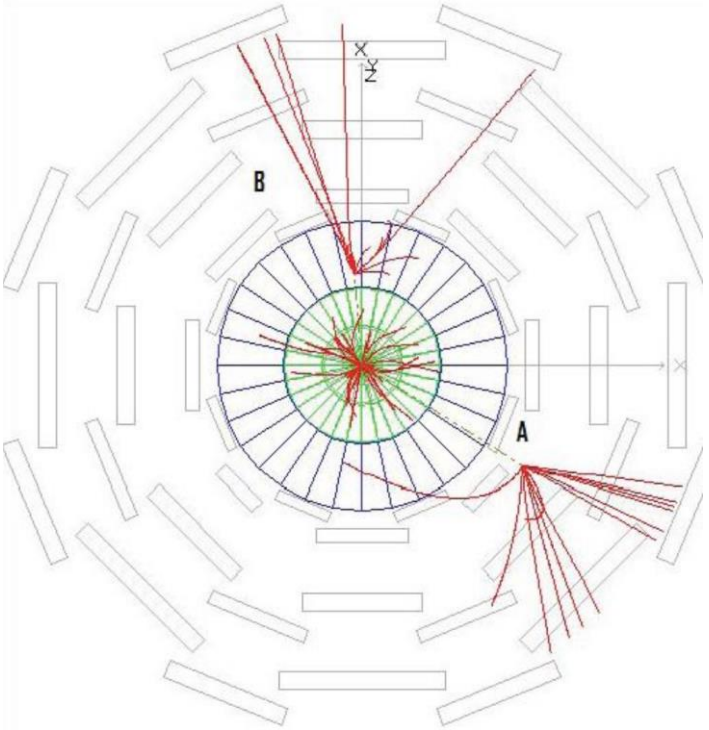

- Timing info of every hit would allow studies of out-of-bunch/out-of-time particles.
- Hardware track triggering requires instrumentation on tracker
- Important to have clear strategy for LLP searches. Require distant detectors? integrate in TDAQ?

Machine Learning

- ML on HEP applications (including TDAQ) already here:
  - ➢ front-end data compression
  - ➢ particle ID with multivariate classifiers
  - ➢ pattern recognition
  - ➢ tracking & reconstruction with NN
  - ➢ regression for improved resolution
- Some of these developments will form basis for proto-TDAQ design in next few years

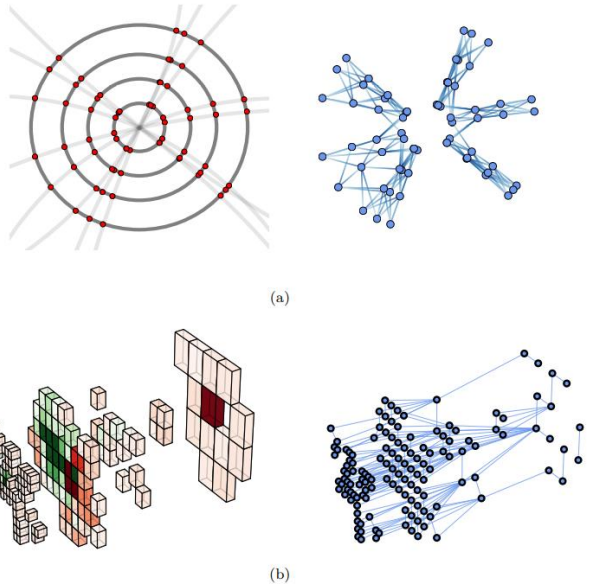

Summary

- Biggest challenge for future e^+e^- is all-time high instantaneous luminosity (but: TDAQ systems will sustain similar data throughput rates already at LHC)
- Baseline assumption: software-only triggering system; some detector choices (e.g. TPC) will challenge this assumption & require management of very large out-of-bunch pile-up and operation in continuous mode
- Full timing info on detector hits will be a game-changer for calibration, reconstruction and exotic searches
- Devil is in details: careful planning needed in detector R&D
- ML is expected to be everywhere: calibration, monitoring, and (yes) event selection