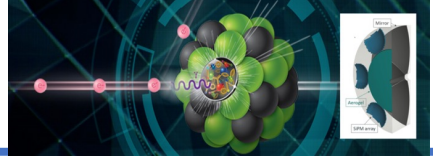

dRICH DAQ – design status

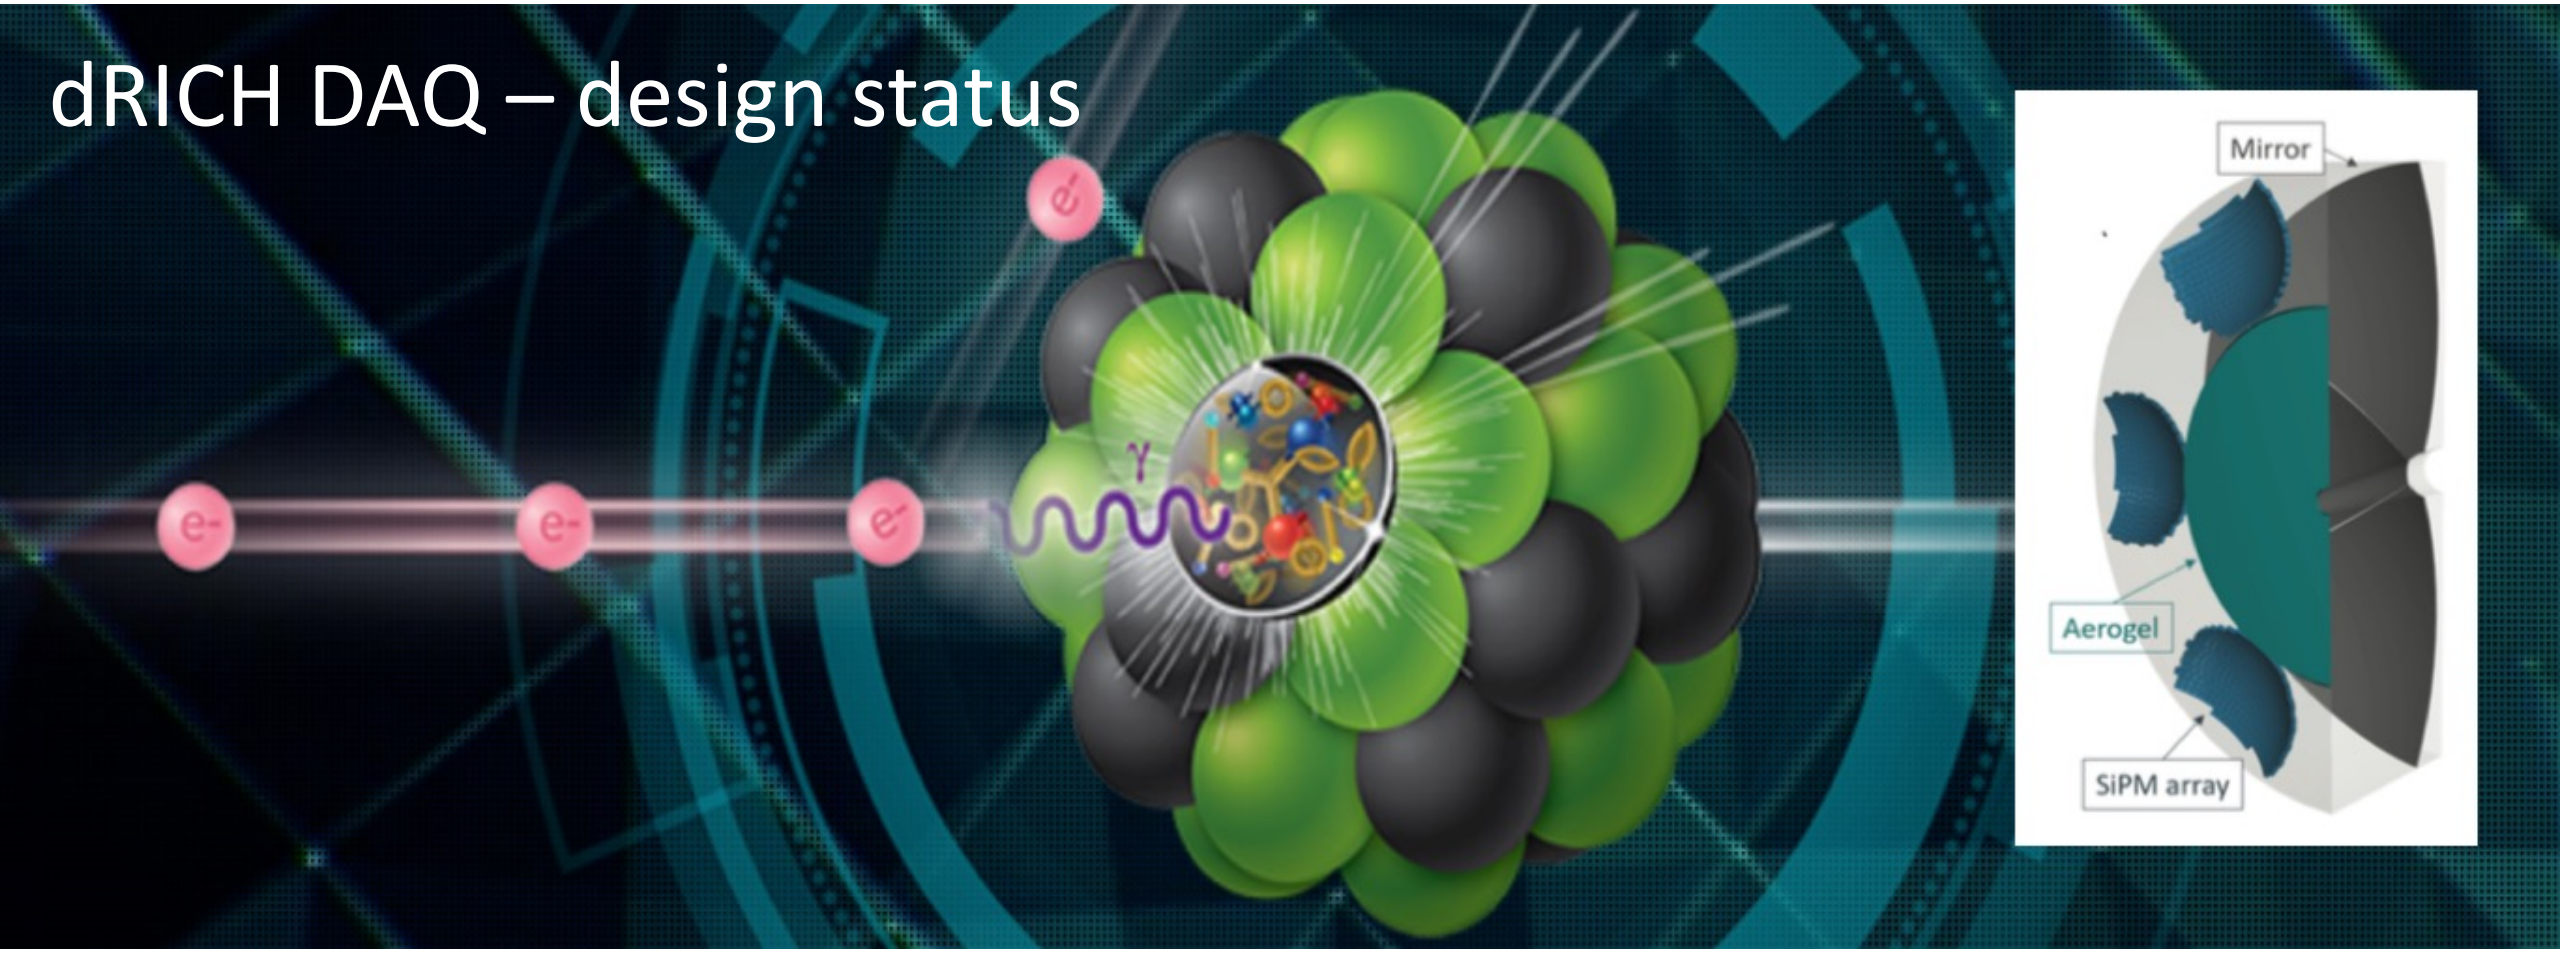

P. Antonioli – INFN Bologna

Outline

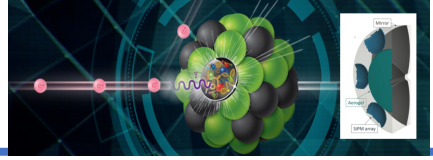

- dRICH baseline (as per ATHENA & ECCE proposals)
- DPAP report random excerpts and the R&D to study high throughput mitigation strategies
- dRICH DAQ status (@ATHENA proposal time) + the ASIC
- radiation, SiPM cooling and annealing: why DAQ matters?
- miscellanea thoughts on DAQ "handles" (brainstorming)
- next steps

dRICH baseline (ECCE & ATHENA)

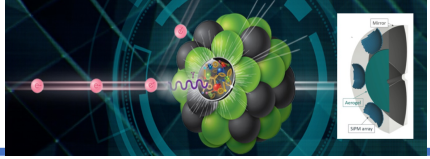

- SiPM readout
- O(3 m^2) of sensors
- ~ 300000 channels (ATHENA: 317440, ECCE: 344046)
- 3 mm x 3 mm pixels
- high throughput foreseen given DCR increase due to SiPM radiation damage
- both proposals selects as "working example" ATLAS FELIX card as "DAM" (Data Aggregator Module) [https://atlas-project-felix.web.cern.ch/atlas-project-felix/]

✓ ATHENA proposal considers as baseline ASIC: ALCOR from INFN-TO
✓ ECCE proposal considers as baseline ASIC: MAROC3 from Omega

TDC vs ADC approach

More info on these two ASIC:
Presentations/Datasheets available here:
ALCOR: https://agenda.infn.it/event/28762/contributions/146408/attachments/87284/116576/20211121_ALCOR4EIC.pdf
MAROC: https://www.weeroc.com/products/photomultipliers-read-out/maroc-3a

ATHENA estimates of throughput

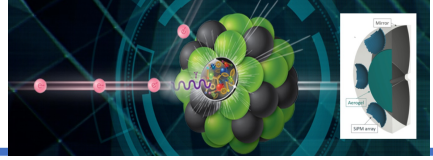

Table 2.5: Maximum data volume by detector.

Detector	Channels	DAQ Input (Gbps)	DAQ Output (Gbps)
B0 Si	400M	<1	<1
B0 AC-LGAD	500k	<1	<1
RP+OMD+ZDC	700k	<1	<1
FB Cal	4k	80	1
ECal	34k	5	5
HCal	39k	5.5	5.5
Imaging bECal	619M	4	4
Si Tracking	60B	5	5
Micromegas Tracking	66k	2.6	.6
GEM Tracking	28k	2.4	.5
uRWELL Tracking	50k	2.4	.5
dRICH	300k	1830	14
pfRICH	225k	1380	12
DIRC	100k	11	11
TOF	332k	3	.8
Total		3334	62.9

This was computed assuming an average 300 kHz DCR per pixel before moving to annealing cycles and a factor 3 reduction due to timing selection (it might be 5)

ALCOR: 64 bit per hit (TOT)

ALCOR: A Low Power Chip for Optical sensor Readout

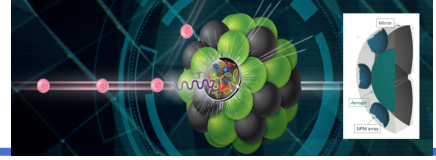

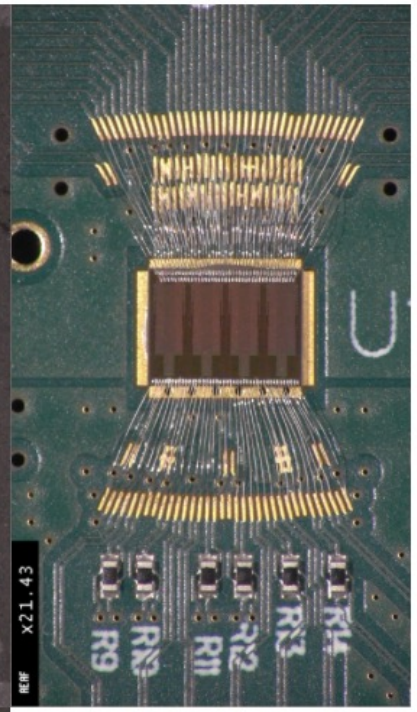

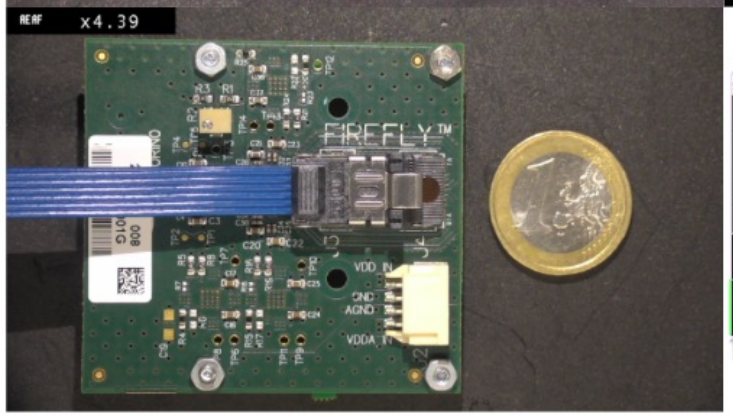

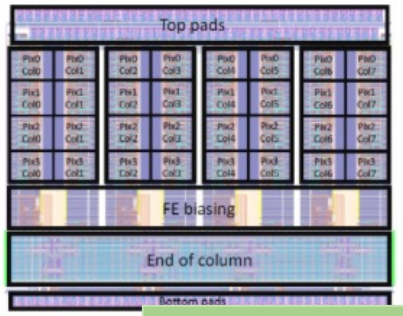

Developped originally for DarkSide (cryo)
expected several "branches" including for EIC

INFN Istituto Nazionale di Fisica Nucleare SEZIONE DI TORINO

R. Kugathasan et al, https://doi.org/10.22323/1.370.0011

- 32-pixel matrix mixed-signal ASIC in 110 nm CMOS technology
- the chip performs
  - ✓ signal amplification
  - ✓ conditioning and event digitisation
- each pixel features
  - ✓ dual-polarity front-end amplifier
    - ➢ low input impedance
    - ➢ 4 programmable gain settings
  - ✓ 2 leading-edge discriminators
  - ✓ 4 TDCs based on analogue interpolation (50 ps LSB)
- single-photon time-tagging mode also with Time-Over-Threshold
- fully digital output: 4 LVDS data links

- signs of ALCOR saturation (bandwidth) seen at O(1 MHz)
- plans for 64 ch version, optimized AC-coupling with selected sensors, packaging, increase bandwidth?

→ eRD109

How would fit ALCOR in dRICH readout?

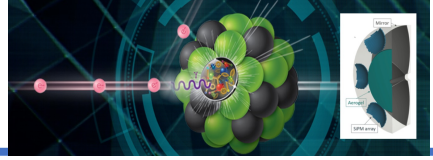

caveat on "current scheme": it was used for proposal/costing not necessarily the final one.
A lot of work ahead, but useful to focus on needed R&D including DAQ requirements

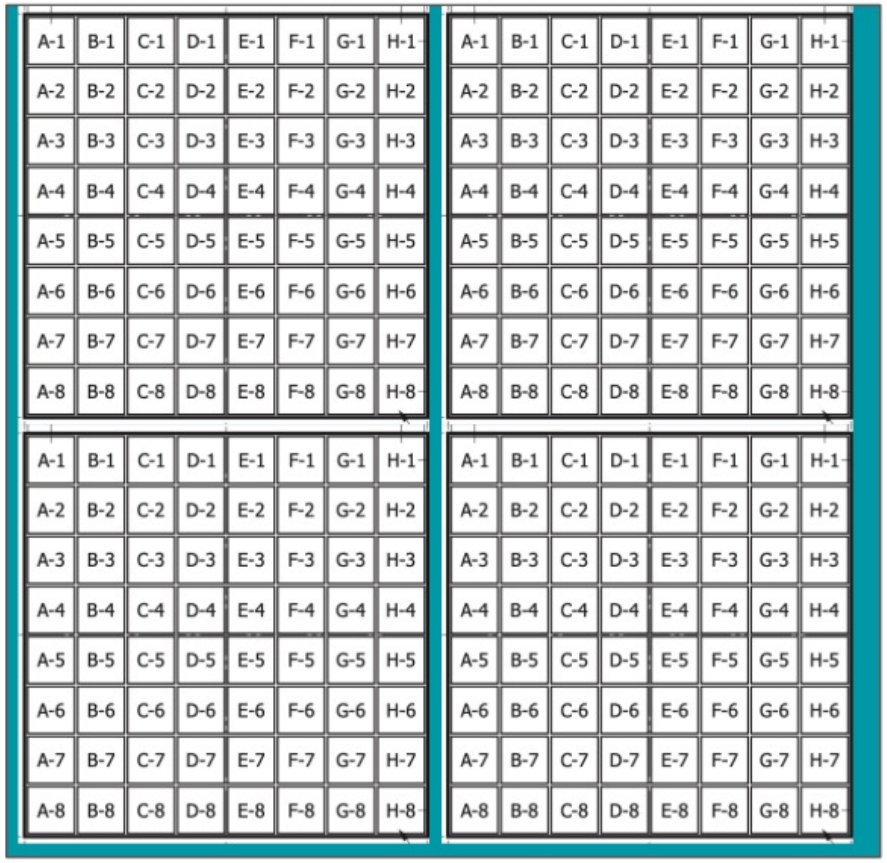

dRICH tile

dRICH tile 5.6 x 5.6 cm^2

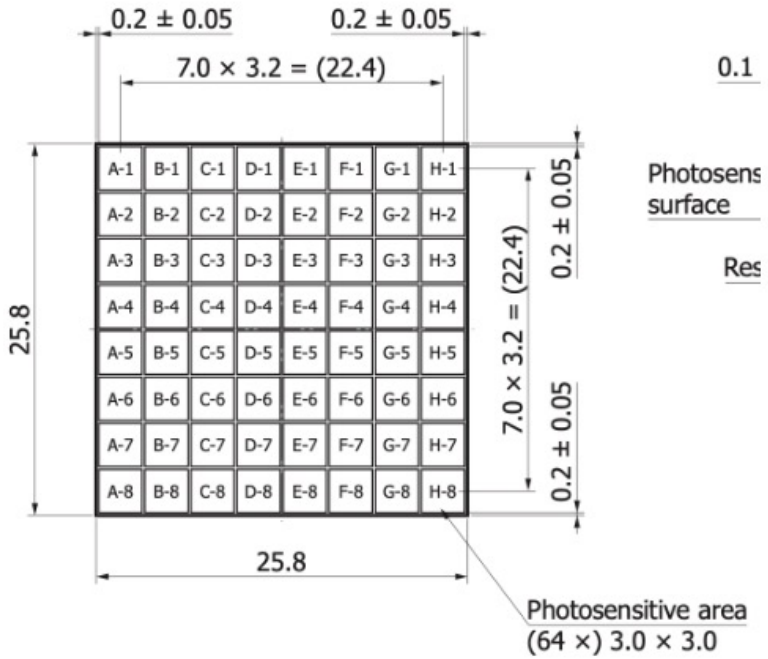

Based on "popular" 8x8 HPK matrix
S13361-3050NE-08

dRICH FEB (Front-End Board)

current thinking: 1 FEB serving 1 dRICH tiles

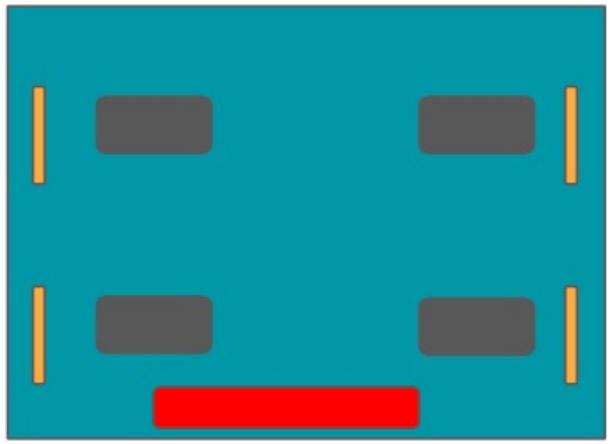

SiPM bus
readout bus + LV bus

- dRICH FEB with external LV connector + routing to dRICH tile
- readout bus to FPGA (dRICH ROB)
- potential area for dRICH FEB: **5x5 cm^2**
- ALCOR packaging (BGA)

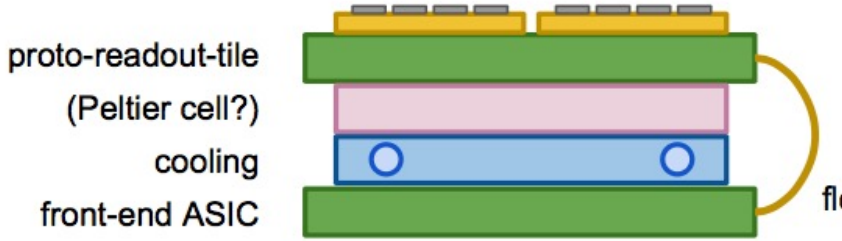

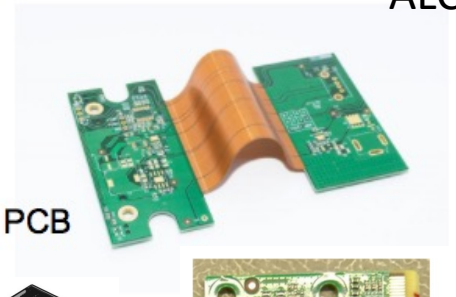

2023: first proto-tiles

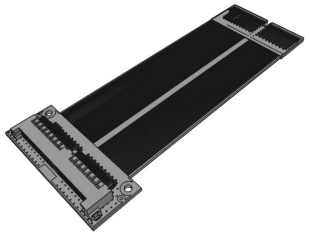

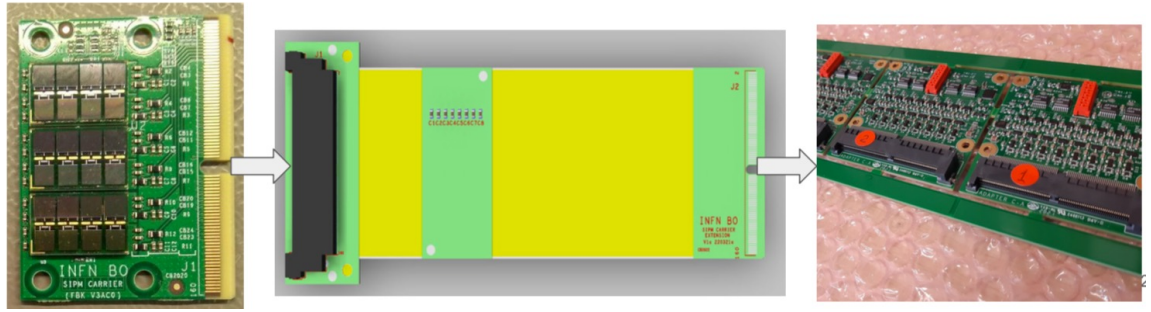

dRICH ROB (read-out board)

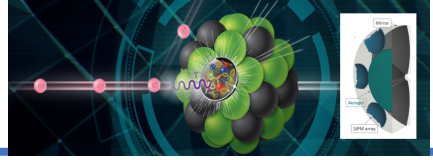

- based on readout/throughput considerations 4 dRICH FEB (1024 ch) should be read-out by 1 dRICH ROB (4096 channels)
- ROB acts as concentrator + data reduction (BC timing) (factor 3-5: EIC 1 BC every 9.6 ns: just get a fraction (like 2 ns of window of interest or possibly less: potential spread is 150 ps but bunch length 0.3-0.4 ns!)
- ROB potential area 10x10 cm^2

current scheme assume ATLAS
FELIX board on PC
+ LpGBT link on ROB @ 10.24 Gbps

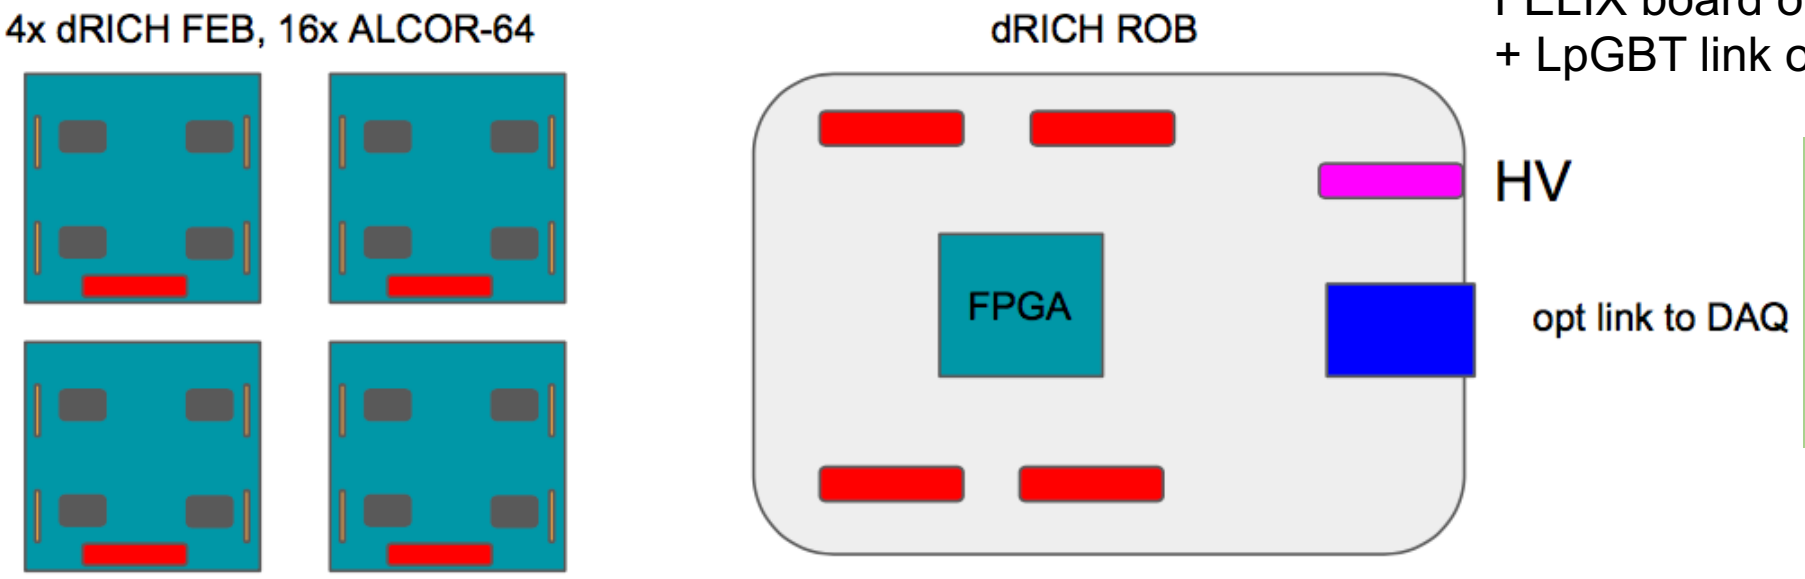

- possibly in-FPGA, not clear
- we need lpGBT ASIC)
- 10 Gbps seems big limitation for an experiment for next decade

- This choice (see later) for throughput modelling keeps bandwidth on opt link to DAQ < 10 Gbps (current limitation)
- On each FEB-ROB bus expected throughput at 4 Gbps (at maximum damage from rad before annealing) if no veto on ALCOR is possible
- On each opt-link (after data reduction via timing): 5.9 Gbps (to be studied a further data reduction, if possible, via coincidences)

Throughput model

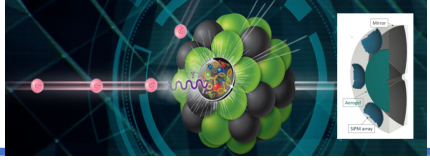

Key inputs:

- DCR: 300 Hz/mm^2 (Hama @ -40C)
- DCR sensor (9 mm^2): 2.7 KHz
- average max tolerable DCR sensor rate: **270 KHz** (reached with some 10^9 MeV-neq… then you need to anneal!)
  ---> this gives already 1.8 Tbps from dRICH to DAQ (with data suppression - gated BC - by a factor 3)
- 317440 sensors

We aim for ALCOR hit rate (with TOT) up to 1 MHz
per single channel appears as a minimum requirement.

- (average) ALCOR channel: 270 kHz → 1 MHz
- 1 hit (including ToT)=64 bits
- 64 channels
- (with 1 MHz) → 2.0 Gbps/chip (current maximum)

Assuming average 270 kHz and TOT:

4.4 Gbps/FEB
18.6 Gbps/ROB ---> data reduction (BC) by a factor 3-5
(**veto signal on ALCOR exists but to be tested at high rate this could decrease FEB rate**)

→ 310 links to DAQ

Excerpts from DPAP report (and my views)

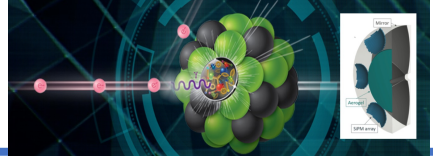

Common Risks and Challenges

The implementation of a streaming readout comes with a certain number of challenges and uncertainties. Within such an architecture, bad/noisy detector channels or unexpected backgrounds will lead to a significant increase in data rate that could potentially exceed the data transfer capability between the FEEs and readout computers. Both ATHENA and ECCE have pointed out the particular importance of having the ability to monitor for bad/noisy channels or modules and reset/disable them, of excellent control of FEE level detector calibration related to zero suppression, and of having the ability to implement common mode noise removal. Risks are further mitigated by both ATHENA and ECCE by optimizing the scale of data aggregation for each subdetector system in such a way as to ensure a reasonable throughput safety margin between the FEE and the readout computer farm, and by the possible implementation of software event filtering and/or software triggers. Further risk mitigation strategies specific to individual proposals include, for ECCE, a detector design relying as much as possible on low-noise detectors (e.g. avoiding SiPM for mRICH), and for ATHENA, the possibility to retain the capability for operating with a hardware trigger via the timing system.

This MUST be implemented. It makes no sense with BC every 10 ns, read all the 10 ns

the DAQ must anticipate the larger DCR that will come

For all three proposals, one of the challenges to a streaming readout is expected to come from the dark currents from the SiPM that will gradually increase with accumulated radiation dose. Both ATHENA and ECCE explicitly plan to have an additional throughput safety margin between the FEE and readout computer farm to account for this effect. We note that the ATHENA proposal foresees a number of readout channels from SiPM that is a factor of about 3 to 5 times larger than that of ECCE and CORE proposals, likely requiring the need for additional mitigation strategies to maintain the ability to operate a streaming readout as function of time. Additional mitigation strategies put forward and studied by ATHENA include the undertaking of an annealing cycle of the SiPM to partially restore initial dark current conditions, the implementation of timing selection cuts in the FEE, and the possible further downstream data reduction based on algorithms implemented in the FPGAs of the FELIX-like boards and/or in software. We note that the in-situ thermal annealing of SiPM, proposed for the RICH detectors present in all proposals, still requires further R&D work for its successful implementation.

Please forget we can provide all this R&D by CD 2/3a
The Oct 2023 baseline will indicate we have R&D in place to identify the solutions by 2025, but...

Robust R&D in place to study mitigation strategies

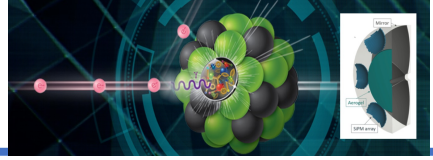

Recent presentations:

P. A, eRD110 presentation at AGS/RHIC users meeting:
https://indico.bnl.gov/event/15479/contributions/64010/attachments/41447/69438/eRD110-20220608.pdf

R. Preghenella, at NDIP 2022 Troyes conference: https://www.ndip.fr/presentations

- SiPM "baseline" selection
- Annealing "performance"
- Annealing protocol "from oven to in-situ"

"DAQ comments":
- DCR after 10^9 already @ O(1 MHz)
- "frequent" annealing (inverse current?) might be key

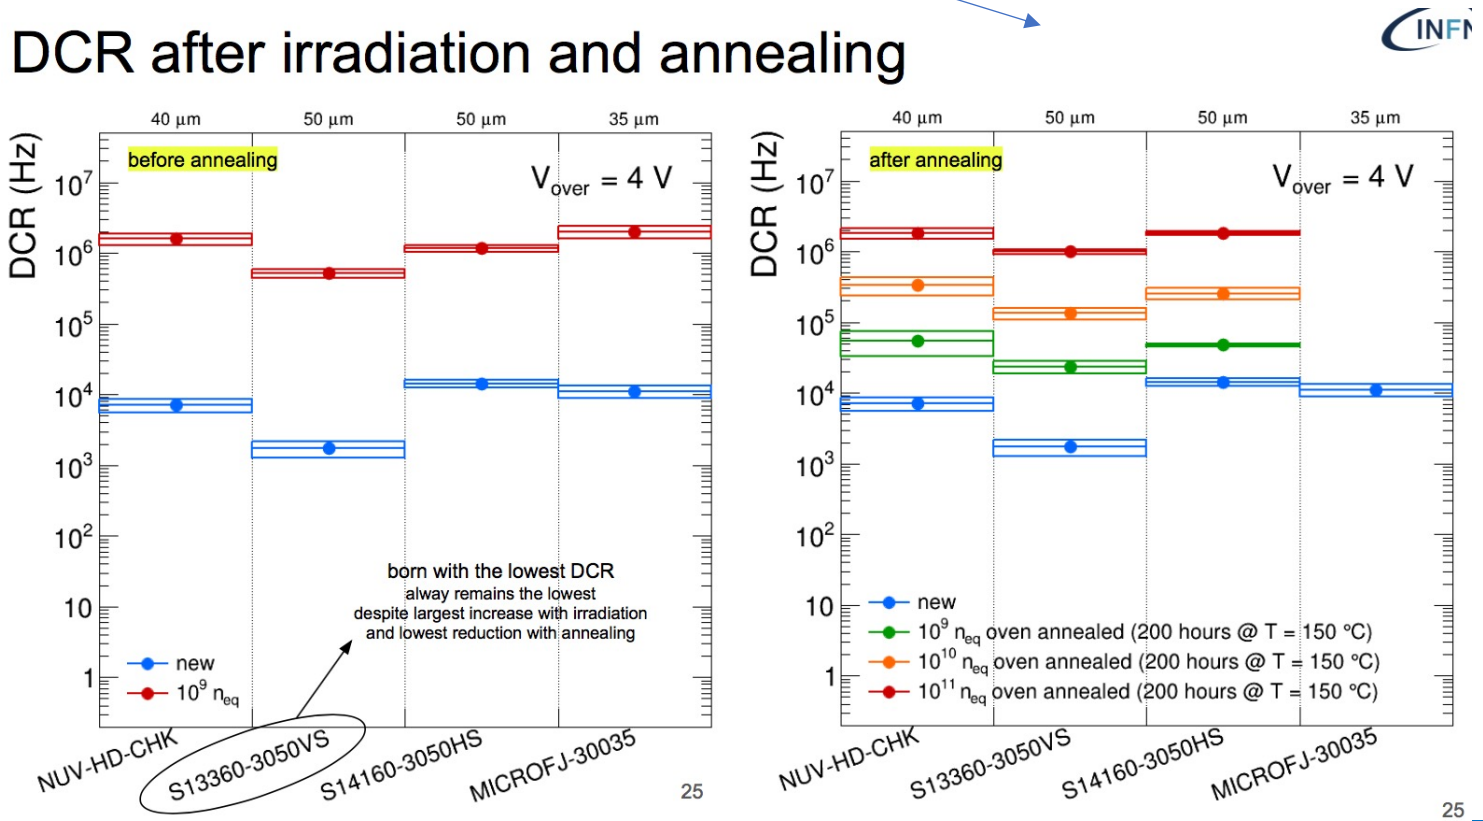

What next for mitigation strategies?

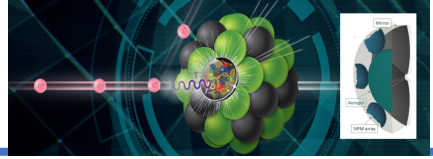

- qualify ALCOR with respect to "time selection" we could cut throughput at FEB-ROB level by factor 5
- I didn't cover "cooling" but 10 degrees down is factor 2/2.5 decrease in DCR. LHCb SciFi has SiPM cooled at -50 C. Our counts are for the moment at -30 C
- is TOT really needed (factor 2 in throughput)
- explorative work on FPGA / AI-based selection starting next year at INFN-GE/RM2 ("data reduction at DAM level")
- we need to keep the glory old trigger option open: LGAD-TOF (or a MPGD layer after dRICH?) or EMCAL could provide "quick loose trigger?" if we select just the BC when we have an interaction this is factor 200...

A side comment on radiation

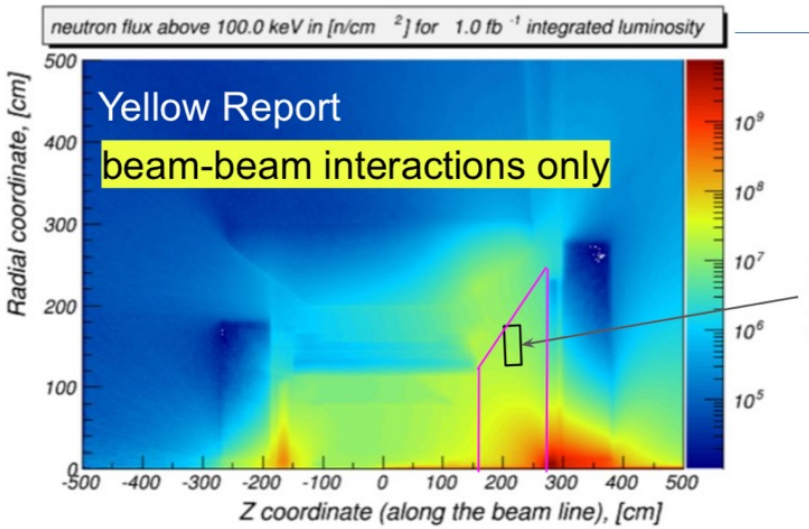

having as target 100 fb^{-1} (several years at maximum luminosity) this brings
10^{11} n/cm^2 1 MeV-neq as "maximum"

- 10 fb^{-1} in 30 weeks of operations at 10^{34} s^{-1}cm^2
- 100 fb^{-1} in 10 years → 1.5 10^9 n/cm^2

potential location of sensors in ATHENA design. To be revised in ECCE (180<z<280) but order of magnitude will not change.
≈ 1.5 10^7 n/cm^2 (100 keV ≈1 MeV-eq) every 1 fb^{-1}

Foreseen radiation levels@ECCE seem somehow lower than YR/ATHENA (or we were hyperpessimistic)

Still discussing with relevant people. But if we expect 10^7 100-kev n/cm^2 at 100 fb-1 (10 years), life will be simpler....

WE NEED TO AGREE A NUMBER BY CD 2/3a

ECCE radiation levels

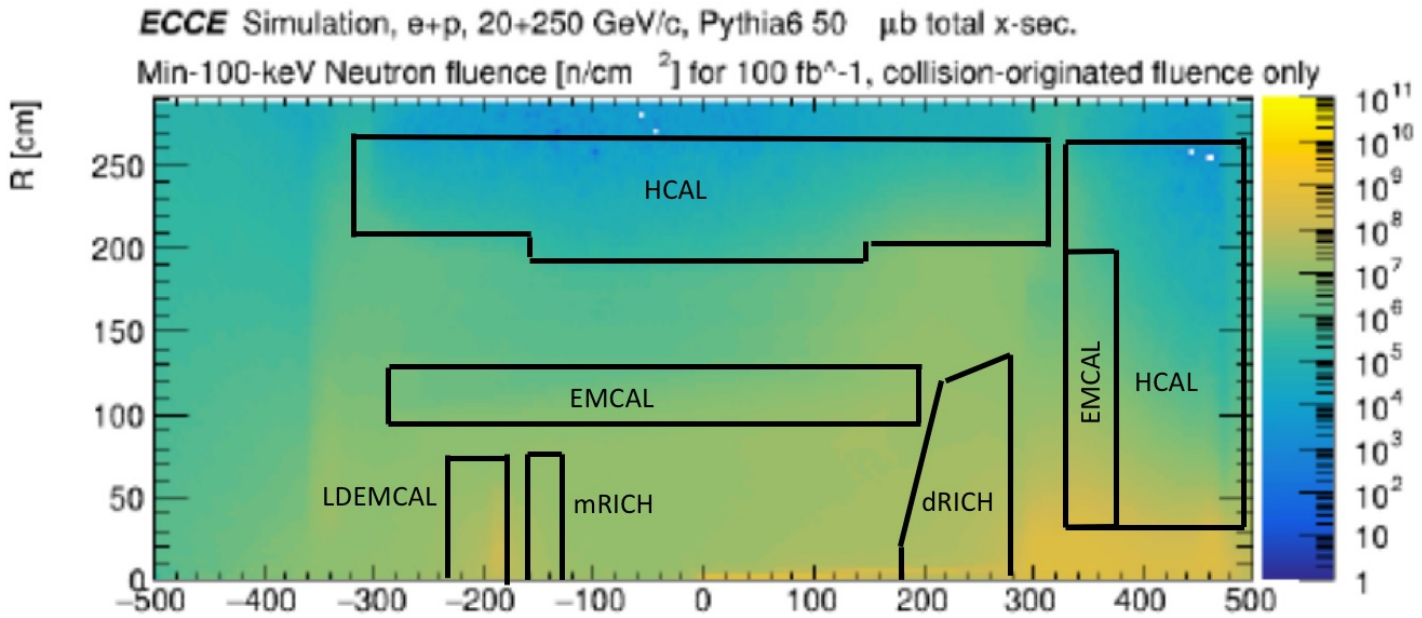

Sources:
ECCE radiation levels: https://indico.bnl.gov/event/14715/contributions/59782/attachments/39682/65822/SiPMs%20for%20EIC%202-4-2022.pdf
Detector positions: https://physdiv.jlab.org/EIC/Menagerie/docs/DetectorParameterTable.pdf

Blue sky exercise

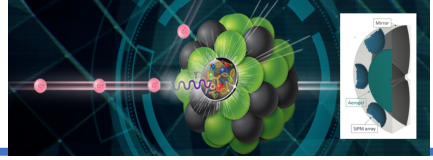

- a 10 Gbps link for 2030/2040 operations might be an underestimation
- is the custom CERNbased lpGBTX really needed?
- Why not look on COTS with 40 GbE now already available?
- Concerned about the syncronization? what about White Rabbit technology?
  https://white-rabbit.web.cern.ch

https://ohwr.org/project/wr-std/wikis/home

DISCLAIMER: this is very brainstorming mode
I need to study a lot before to "defend" all this, happy to get criticism!

Is FELIX limiting us? Do we need a "custom" DAM or a high end switch could manage dRICH (and Detector 1) needs?

Commercial routers/switches in 2022 offer already much more than FELIX...
https://www.arista.com/en/products/spine-routing

Such "aggregators" (possibly running custom firmware) could then route data to a SRO farm without any custom hardware.
Route all packets of a given BC to a machine, and so on...

Implementing 50 G interfaces could scale down dRICH ROB by a factor 5...

Arista 7500R3 Series

System Performance			
	7512R3 	7508R3 	7504R3 
Description	Delivers high density, low power non-blocking Ethernet switching system.	Delivers high density, low power non-blocking Ethernet switching system.	Optimum form-factor for many mid-size deployments.
Switching Capacity	230Tbps	153Tbps	76.8Tbps
Linecard Capacity	9.6 Tbps	9.6 Tbps	9.6 Tbps
10G Interfaces	2304	1536	768
25G Interfaces	2304	1536	768
40G Interfaces	432	288	144
50G Interfaces	2304	1536	768
100G Interfaces	432	288	144
400G Interfaces	288	192	96
Forwarding Rate	48 Bpps	32 Bpps	16 Bpps
Total Buffer	96GB	64GB	32GB
Rack Units	18	13	7
Airflow	Front to Rear	Front to Rear	Front to Rear

Not a summary

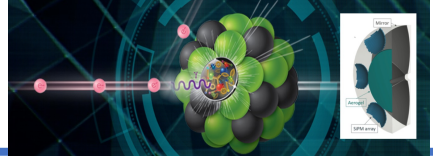

We still need a lot of data before to go for a convincing solution

DAQ design for dRICH will depend on:

- neutrons (and precise simulation of material budget...)
- SiPM radiation "performance"
- ASIC bandwidth performance on single channel
- annealing performance
- cooling performance
- data reduction/selection performance (when noise is properly simulated)
- choice of aggregator technology "up to next decade"