



dRICH elettronica e dintorni (DAQ, cooling, ...)




P. Antonioli – INFN Bologna

Perche' questo talk?



- Talk in parte basato su talk dato a DAQ WG (richiedevano dRICH plans e li' avevo riportato varie idee senza una chiara discussione interna)
- Recap per “allineare informazioni”
 - dRICH baseline (as per ATHENA & ECCE proposals)
 - dRICH DAQ status (@ATHENA proposal time) + the ASIC
 - radiation, SiPM cooling and annealing: why DAQ matters?
- No recap su risultati esistenti su SiPM ben documentati in EIC_NET o conferenze:
 - Luigi @ [EIC_NET Giornata nazionali 2022](#)
 - Roberto @ [dRICH meeting 8 August](#)
 - Roberto @ RICH2022 [A SiPM-based optical readout system for the EIC dual-radiator RICH](#)
- **Primissimo** step verso divisione task, responsabilita' e organizzazione next steps
- Tentativo di “districare” polpaccione “elettronica sensori cooling spazi geometria” su cui abbiamo/avremo pressione da Progetto e pre-TDR
- Planning eventuali incontri esterni (aziende, altri gruppi), ma soprattutto **interni**

Radiation damage & annealing



Radiation damage model (HPK S13360-3050 @ Vover = 3 V)

- **reasonable assumptions**

- radiation damage is additive
- does not know and care of the past damage
- annealing heals up to a certain fraction of damage, not more than that



- **numbers**

- DCR when new = 1.5 kHz
- DCR increase with radiation damage = 350 kHz / 10^9 neq
- DCR increase with online annealing = 35 kHz / 10^9 neq
- DCR residual after oven annealing = 3%

- **how it works?**

- start with DCR as new → NEW
- add DCR with increasing radiation → NEW + NIEL1
- heal with annealing → NEW + x NIEL1
- add DCR with increasing radiation → NEW + x NIEL1 + NIEL2
- heal with annealing → NEW + x (NIEL1 + NIEL2)

Roberto / 8 Aug





online annealing every 2×10^8 neq (500 times)

Message:

Max “tolerable” rate per channel determines frequency of annealing

Online annealing



explore solutions for in-situ annealing

- total fluence of $10^9 n_{eq}$
 - delivered in 5 chunks
 - each of $2 \cdot 10^8 n_{eq}$
- interleave by annealing
 - forward bias, $\sim 1 \text{ W}$ / sensor
 - $T = 175^\circ\text{C}$, thermal camera
 - 30 minutes
- preliminary tests
 - Hamamatsu S13360-3050



M. Cordelli^a E. Diociaiuti^a A. Ferrari^c S. Miscetti^a S. Müller^c G. Pezzullo^b I. Sarra^{a,1}


^aLaboratori Nazionali dell'INFN, via Enrico Fermi 54, 00044 Frascati, Italy

^bDepartment of Physics, Yale University, 56 Hillhouse, New Haven, CT-06511, USA

^cHelmholtz-Zentrum Dresden-Rossendorf (HZDR), Bautzner Landstraße 400, 01328 Dresden, Germany

<https://arxiv.org/pdf/1804.09792.pdf>

Online annealing



Message:

- Primordial study in 2022 but promising. It will continue in 2023
- We heat a sensor giving 1 W with currents $O(1 \text{ A})$
- if validated impact on HV design
- Homework needed before pre-TDR

dRICH baseline (ECCE & ATHENA)



- SiPM readout
- $O(3 \text{ m}^2)$ 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 select 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:

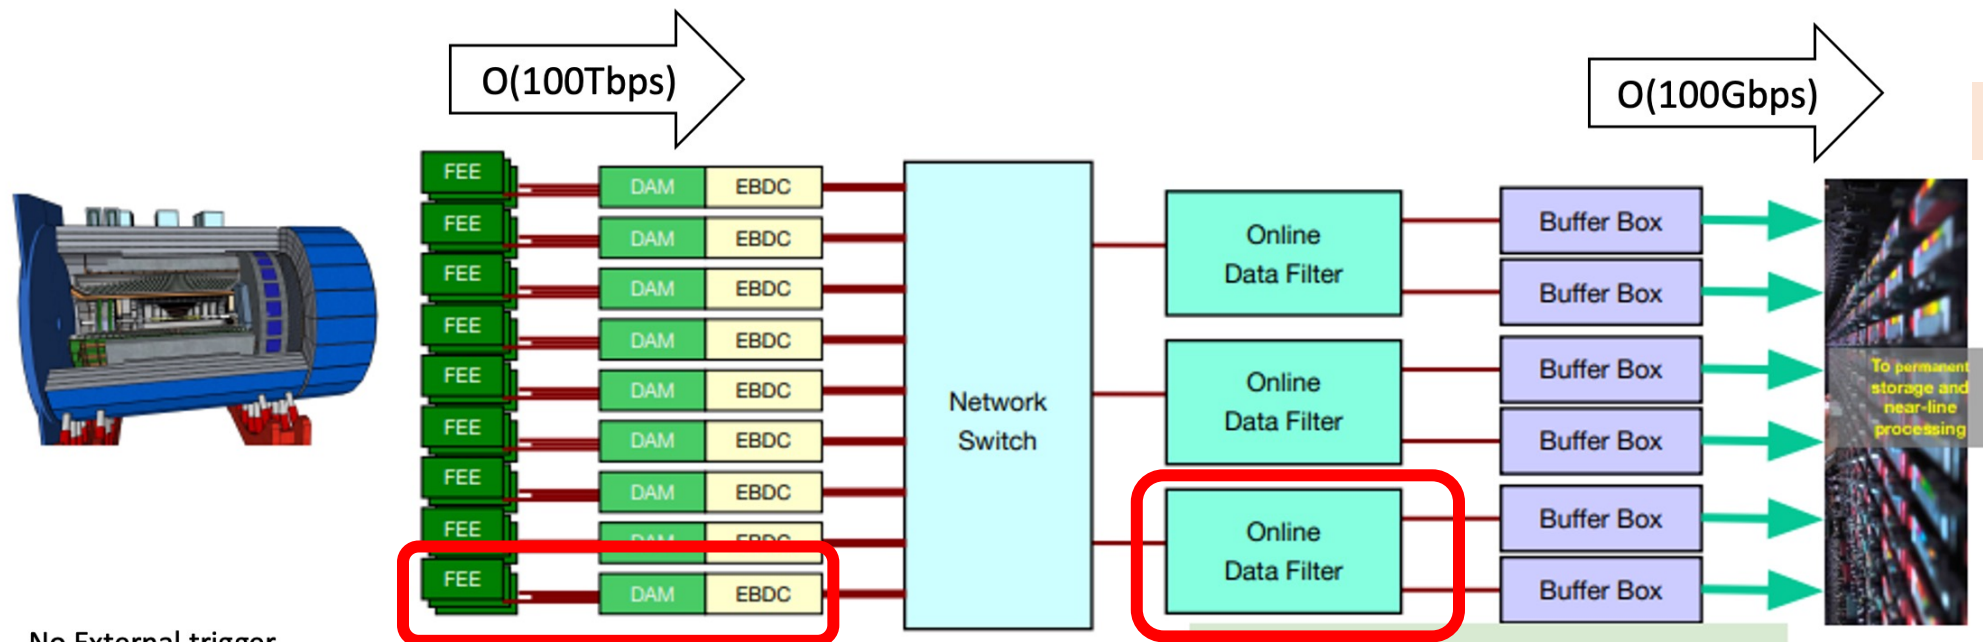
ALCOR: https://agenda.infn.it/event/28762/contributions/146408/attachments/87284/116576/20211121_ALCOR4EIC.pdf

MAROC: <https://www.weeroc.com/products/photomultipliers-read-out/maroc-3a>

EPIC Streaming DAQ



J. Lajoie



- No External trigger
- All collision data digitized but aggressively zero suppressed at FEE
- Low / zero deadtime
- Event selection can be based upon full data from all detectors (in real time, or later)
- Collision data flow is independent and unidirectional-> no global latency requirements
- Avoiding hardware trigger avoids complex custom hardware and firmware
- The “Front End Processing”, programmable hardware between the FEEs and the DAQ computers, is deemphasized relative to the yellow report, but should not be precluded.
- Data volume is reduced as much as possible at each stage

FEE = Front End Electronics
DAM = Data Aggregation Module
EBDC = Event Buffer / Data Compressor

Stored data volumes
 $\sim O(100\text{Pb})$ per run

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


ASSUMPTIONS

This was computed assuming an average 270 kHz DCR per pixel MAX before moving to annealing cycles given limitations on ALCOR and DAQ bandwidth

We considered already a factor 3 reduction due to timing selection (it might be 5 or 10...) [at the time of proposal it was ambiguous if it was at ALCOR or FPGA level]

Throughput assumed 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

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(4 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

A-1	B-1	C-1	D-1	E-1	F-1	G-1	H-1
A-2	B-2	C-2	D-2	E-2	F-2	G-2	H-2
A-3	B-3	C-3	D-3	E-3	F-3	G-3	H-3
A-4	B-4	C-4	D-4	E-4	F-4	G-4	H-4
A-5	B-5	C-5	D-5	E-5	F-5	G-5	H-5
A-6	B-6	C-6	D-6	E-6	F-6	G-6	H-6
A-7	B-7	C-7	D-7	E-7	F-7	G-7	H-7
A-8	B-8	C-8	D-8	E-8	F-8	G-8	H-8


dRICH tile



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


dRICH tile 5.16 x 5.16 cm²

dRICH FEB (Front-End Board)



current thinking: 1 FEB serving 1 dRICH tiles


ASSUMES ALCOR-64




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 cm²**
- ALCOR packaging (BGA)





proto-readout-tile
(Peltier cell?)
cooling
front-end ASIC




flex PCB



2023: first proto-tiles




Progresses on flex-PCB + next steps





Flex “si piega bene”

Curve I-V a temperature ambiente con/senza flex indistinguibili



readoutBox (for test beams + tests)



Message:


- Essential to study/test compact solution
- Essential to study (partially) cooling solution
- In 2023 we will not have ALCOR-64 packaged, additional challenge!

2023

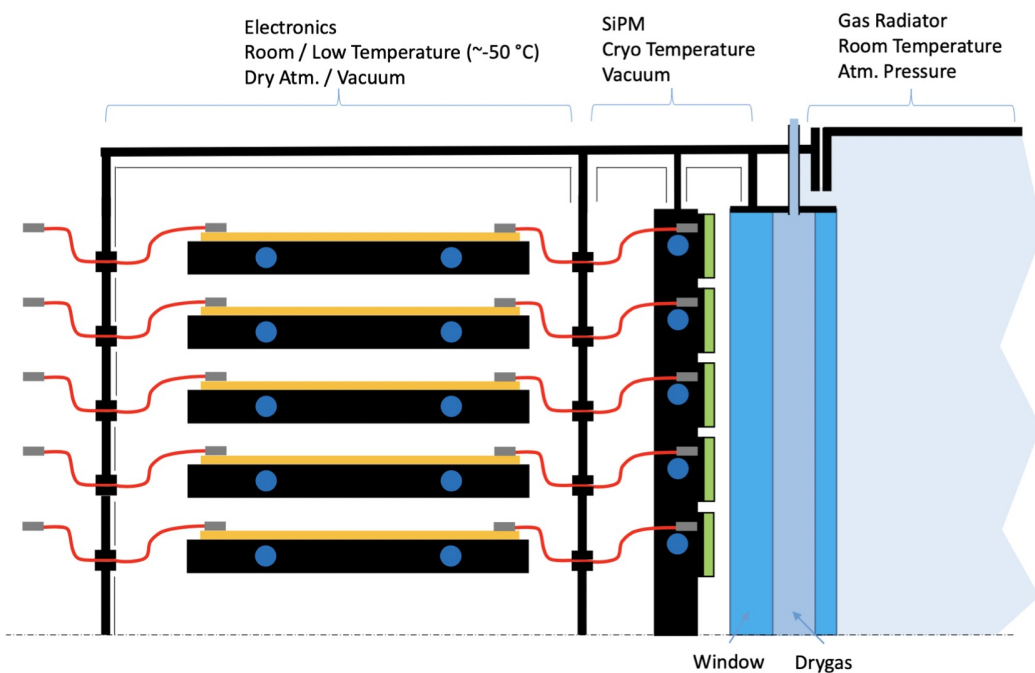
dRICH protoTile
(256 or 1024 channels)
with improved cooling

eRD102
eRD110

We are not alone...




Global cooling strategy: cryostat



Interessante [talk](#) di Roberta Cardinale (LHCb-GE) su cooling per RICH2 con SiPM a RICH2022

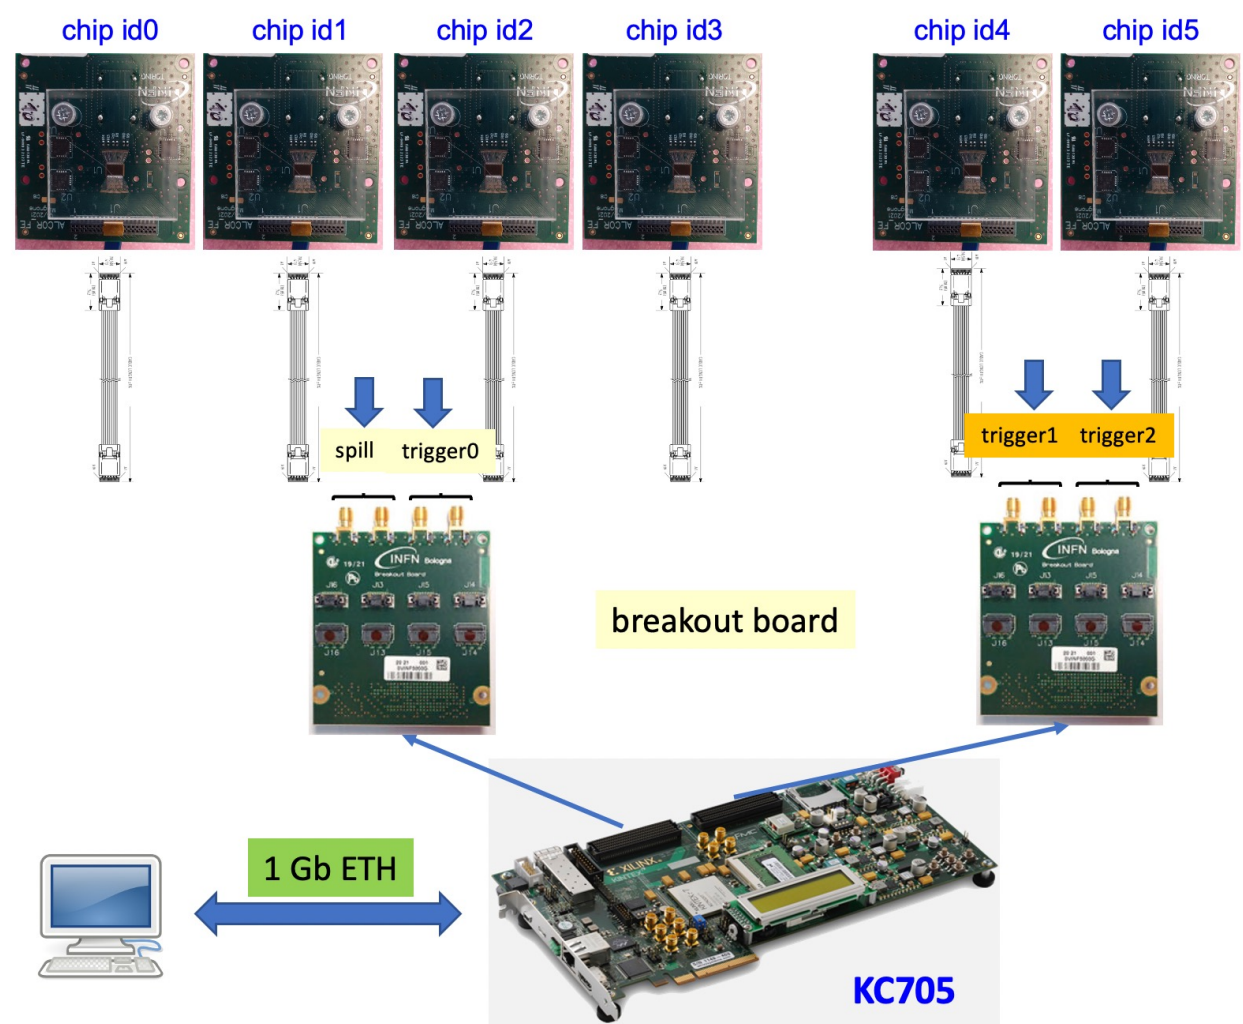
ALCOR readout (generic info + “current setup”)



Alcor DAQ hardware

12 LVDS pairs per chip

- clock_in
- data (4)
- SPI (4)
- clock_out
- reset
- testpulse




Davide (ARCADIA based)

(mandatory)
(optional)

ALCOR bus (LVDS pair)

Signal	LVDS pairs	I/O
CLKIN	1	I
CLKOUT	1	O
DBUS – COLUMN 0	1	O
DBUS – COLUMN 1	1	O
DBUS – COLUMN 2	1	O
DBUS – COLUMN 3	1	O
TEST PULSE / SHUTTER	1	I
DOIT	1	O
DIN	1	I
SPICLK	1	I
SPIENA	1	I

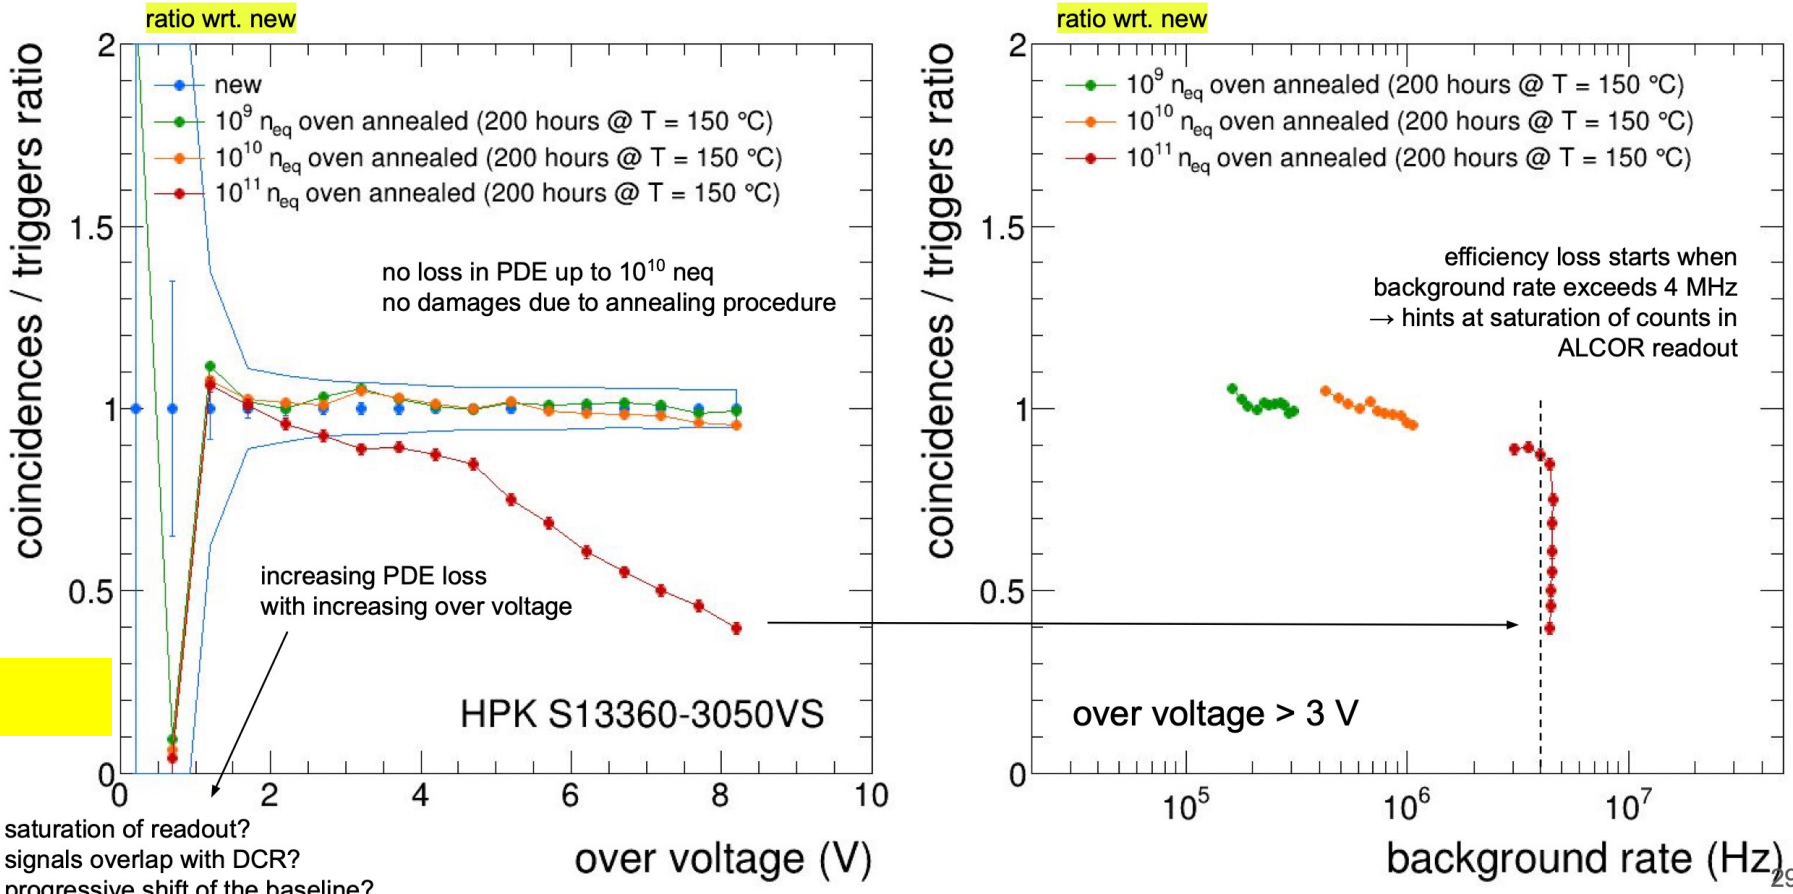
ALCOR saturation?



- A. Analogic stage: 5 MHz/channel ? {is it a hard limit?}
- B. Serializer (1 over 8 channels): 16 MHz --> 8 MHz (TOT) → 1 MHz [→ this would limit at **250 kHz** or **500 kHz**]
[“togliendo le parole di status si arriva senza problemi a 800 kHz”]
- A. What we see in data (1 over 8 active) seems to point at 4 MHz, it might be we reached limit (A)?




Light response after irradiation and annealing



Roberto @RICH2022



saturation of readout?
signals overlap with DCR?
progressive shift of the baseline?

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 in ATHENA proposal: EIC 1 BC every 9.6 ns: just get a fraction)
- ROB potential area 10x10 cm²

4x dRICH FEB, 16x ALCOR-64




proposal scheme assumed ATLAS
FELIX board on PC
+ IpGBT link on ROB @ 10.24 Gbps

- possibly in-FPGA, not clear we need IpGBT 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 (I)




- ALCOR clock 320 MHz ma dati serializzati a frequenza DDR a gruppi di 8 canali → le parole di 32 bit sono encoded a 40 bit
- Questo porta a un limite massimo di 640 Mb/s che corrisponde a un massimo rate su 8 canali di 16 MHz → quindi siamo ora limitati a 2 MHz su singolo canale (averaged su 8)

Starting to play with numbers and known limitations:

- Below no TOT
- Note that 10 Gbps limitation might be overcome (20 Gbps seems reachable even if rad. Tol. might be problem) to be investigated
- Exploit timing reduction

Singolo canale	serializer 8 canali ALCOR column			FPGA (16 ALCOR-64). Timing reduction		Opt Link
	16 MHz	8 ch	→ ALCOR-64			
5 MHz	2 MHz	64 MB/s.	→ 512 MB/s	8 GB/s → 64 Gbps	1	10 Gbps
300 kHz	300 kHz	9,6 MB/s.	→ 76,8 MB/s	1,2 GB/s → 9,8 Gbps	1	10 Gbps
300 kHz	300 kHz	9,6 MB/s.	→ 76,8 MB/s	1,2 GB/s → 9,8 Gbps	3	3,3 Gbps

Throughput (II)



ALCOR test pulse (in inverse polarity) can act as “inhibit” of the digitalization. This could help greatly to reduce data throughput
Note [EIC beam bunch timing](#) presentation by Todd (Sep. 2022)

Short summary:

EIC Orbit	12.78 usec
Bunch spacing (Nb 290)	40.599 ns
Bunch spacing (Nb 1160)	10.150 ns
Gap length	1.01 usec
Collision spread (bunch length)	23-30 ps (ESR) 250-200 (ps) HSR


Notes:

- we will need to implement disable during gap region which is sizeable (8% data reduction “for free”)
- Bunch crossing every 10 ns! With a bunch length of 250 ps if we could select just 1 ns data reduction by a factor 10 before serializing stage in ALCOR

Critical measurement to be done in Turin:


- How react ALCOR to such “short” shutter cycles?
- What happens to the TOT measurement?
- What happens if trailing edge is ON and leading edge was OFF



Note LHCb is thinking something very similar with [FastRich](#)



If it works then requirement on precise clock alignment / phase shift inside FPGA sending shutter to ALCORs

Throughput (III)





		serializer 8 canali ALCOR column				
Singolo canale 5 MHz	16 MHz		8 ch	→ ALCOR-64	FPGA (16 ALCOR-64). Timing reduction	Opt Link
	2 MHz		64 MB/s.	→ 512 MB/s	8 GB/s → 64 Gbps	1  10 Gbps
	300 kHz		9,6 MB/s.	→ 76,8 MB/s	1,2 GB/s → 9,8 Gbps	1 10 Gbps
	300 kHz		9,6 MB/s.	→ 76,8 MB/s	1,2 GB/s → 9,8 Gbps	3 3,3 Gbps
5 MHz	500 kHz		16 MB/s.	→ 64 MB/s	1,0 GB/s → 8 Gbps	1/0.92 7,4 Gbps

Hardware shutter reduces by factor 10


Orbit gap

The implementation of the shutter might reduce annealing cycles by a factor 10
Question for ASIC experts: will it really work this way?
Question for ALL: we will might have in 1 ns a 5 MHz DCR but... we will then still see the rings?
→ simulation with flat noise is a must!!

We need to add the DCR and check!



Miscellanea on ALCOR and other ASICs



(Giulio) Al momento il chip che abbiamo testato non sappiamo bene fino a che rates riesce ad arrivare per i problemi della logica del tdc.

In simulazione comunque vedo che il readout non perde eventi fino a 500 kHz per pixel (in leading edge)– Oltre le FIFO del fine colonna vanno full con le parole di status, perchè in quel momento c'è un picco di trasmissione dati. Togliendo le parole di status si arriva senza problemi fino a 800 kHz. In una prossima versione posso aumentare le dimensioni delle ram e ottimizzare un po' i tempi morti del readout. Ad esempio le parole di status possono essere tolte dai dati e i singoli registri interrogati tramite spi... ci si può pensare.

A livello di singolo pixel credo che Fabio (che aggiungo in copia) avesse fatto delle simulazioni con rates anche superiori, bisogna solo vedere quale sarà il limite del readout di colonna. mi pare che ora riesca ad arrivare a 10 MHz,

Other ASICs for comparison:

- LHCb is now developing FastRICH (derived from FastIC but with digitization): **16 ch/25 ps LSB/shutter** at 600 ps
Plans to have it deployed for RUN4
- Weeroc is realising RADIROC these days as FEA
- CAEN to commercialize RADIROC+picoTDC
- Weeroc TRIROC very slow (100 kHz)

ALCOR is in competitive position! Key points:

- Shutter
- Throughput
- Packaging and ALCOR-64 optimization




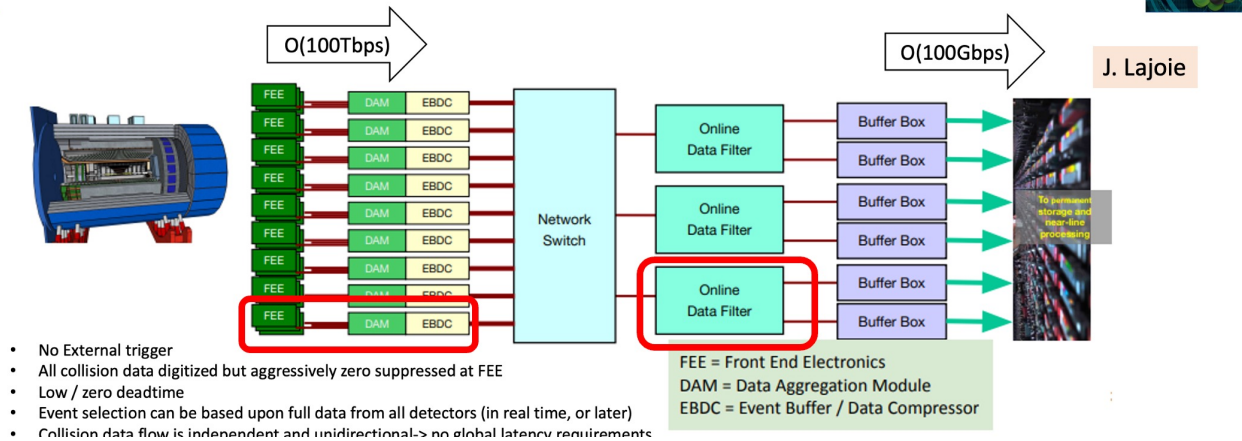
Detector Read-Out	SiPM, SiPM array
Number of Channels	64
Signal Polarity	Positive
Sensitivity	Trigger on 1/3 of photo-electron
Timing Resolution	Better than 35 ps FWHM on single photo-electron
Dynamic Range	Up to 2000 photo-electrons @ 10 ⁶ SiPM gain
Packaging & Dimension	BGA 20x20 mm ²
Power Consumption	310 mW – Supply voltage: 1.2 V
Inputs	64 analogue inputs with independent SiPM HV adjustments
Outputs	2 direct outputs per channel, selectable channel-per-channel, either : <ul style="list-style-type: none">• 1 LVDS triggers• 2 TTL triggers• 1 TTL triggers and 1 analog outputs• 2 analog outputs 2 multiplexed analogue outputs 3 NOR64 trigger outputs
Internal Programmable Features (I2C)	64 HV adjustment for SiPM (64 x 8 bits), 3 trigger threshold tuning (10bits), channel-by-channel gain and shaping time adjustment (τ = 20 ns to 300 ns), individual trigger masking and cell powering.

In FELIX board (or equivalent) we aim to pump 2 Tbps

FPGA algorithm:

- to start to simulate this we need (first) on simulation generate clean data + DCR/white noise
- Segmentation is important (FPGA will currently see just 1/10 -1/20 of dRICH at current segmentation)


EPIC Streaming DAQ



Online data fitter

Not clear to me how much we can reduce “online”

A long (incomplete) list of TODO....



- Space: check space in simulation (Chandra/Roberto)
- Space: how much space we will need for the electronic sandwich?
 - Design at pre-TDR level + implement “dRICHtile”
- Cooling: how and how much to cool? [LHCb, ATLAS ITK, ...]
- ASIC shutter + readout limit + plan ALCOR-64
- FEB+ROB develop a design for pre-TDR – as much as possible anticipating it in reality for 2023 prototype
- DAQ maintain contact with EPIC WG + have “something” on data reduction (DAM/fitter)
- HV requirements / segmentation / distribution + requirements for online annealing
- dRICH thermal isolation between SiPM and gas → no temperature gradient inside!
- Trigger bunch crossing at 100 MHz – collisions at 200 kHz [keep it open and prepare something in pre-TDR?]

Investigate with industrial partners (CAEN (readout and HV + online annealing) / FBK (cooling for SiPM?) / other industries for Cooling (Huber, ...) + expertise (CERN...)

Note: for many things it is “early” but:

- important for groups to start to identify tasks that will closely follow → RN meetings with groups
- pre-TDR is horribly close (draft May 2023). We need to understand what is needed at that level
- some task to start with:
 - can be naturally discussed in sipm-electronica “group”/”mailing list” [weekly → Roberto martedì ore 9:00
 - others need to be part of dRICH general planning
 - having separate discussions on specific topic (“less meetings to work more...”)



Plans for Design Maturity


System	Estimated Design % Complete	Estimated Design % Complete	Estimated Date for Final Design Complete	Comments
	Now	CD-2 / 3A		
6.10.02 Detector R&D/Physics Design	0%	60%	06/30/2026	Project R&D just started
6.10.03 Tracking	10%	50%	12/31/2026	Need only late
6.10.04 PID	15%	50%	03/31/2026	hpDIRC well underway
6.10.05 EmCal	20%	85%	12/31/2024	eEMCAL far ahead
6.10.06 HCal	15%	70%	06/30/2025	Barrel Hcal reuse, rest delayed
6.10.07 Magnets	30%	100%	12/31/2023	LLP, completed 30% design
6.10.08 Electronics	10%	50%	03/31/2027	ASICs/electronics can come in late

What Does a Preliminary Design mean

Example: dRICH

- need to define the sub-detector technology to a level of detail that we can baseline cost, schedule and workforce and functional requirements needs
- what do we build: a CF-gas + Aerogel RICH or is the CF-gas replaced with a pressurized or cooled Argon
 - vessel design needs to be well advanced
- geometry of the subsystem and how it is integrated in the overall detector
- photon-sensor technology and # of readout – channels
- what is the front-end electronics, what ASIC will be used
- define mirror - system
- what needs to be cooled and how
- 3d-CAD of the detector with details how the detector will be assembled, drawings of the different components but not on fabrication quality
- design of gas system
- slow control and monitoring of hardware systems are needed, how do we realize it
- A worked-out concept (but no detailed plan) of assembly and service needs

There can still be some open questions (but not affecting costs and schedule in major way), further engineering design to be done, detailed drawings to be done, etc.



Miscellanea things from DAQ meeting presentation

Throughput model in ATHENA proposal



Key inputs:

- DCR: 300 Hz/mm² (Hama @ -40C)
- DCR sensor (9 mm²): 2.7 KHz
- average max tolerable DCR sensor rate: **270 KHz** (reached with some 10⁹ 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

A side comment on radiation



having as target 100 fb⁻¹ (several years at maximum luminosity) this brings
10¹¹ n/cm² 1 MeV-neq as "maximum"

- 10 fb⁻¹ in 30 weeks of operations at 10³⁴ s⁻¹cm²
- 100 fb⁻¹ in 10 years → 1.5 10⁹ n/cm²

potential location of sensors in ATHENA design. To be revised in ECCE (180<z<280) but order of magnitude will not change.
≈ 1.5 10⁷ n/cm² (100 keV ≈ 1 MeV-eq) every 1 fb⁻¹


ECCE radiation levels



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


Still discussing with relevant people. But if we expect 10⁷ 100-keV n/cm² at 100 fb⁻¹ (10 years), life will be simpler....

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



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>

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...

Blue sky exercise



- a 10 Gbps link for 2030/2040 operations might be an underestimation
- is the custom CERNbased IpGBTX really needed?
- Why not look on COTS with 40 GbE now already available?
- Concerned about the sincronization? 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

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...

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"