The dRICH data stream
(and a little bit of history about the "interaction tagger")

P. Antonioli (INFN-Bologna)

ePIC Collaboration Meeting
Lehigh University 22-27 July 2024

Workfest == make people **interact**

- Why we discuss this? A little bit of history
- An interaction tagger or a dRICH interaction tagger or a dRICH-interaction tagger or a data reduction algorithm?
- The dRICH data stream
- Data reduction approaches

Workfest == make people **interact**

- Why we discuss this? A little bit of history
- Interaction tagger(s), physics and dRICH data reduction
- The dRICH data stream modeling
- Data reduction approaches

- mood

This is a talk to **trigger discussion** and **set context** for following talks

- mood

This is a talk to **tag** key topics (to **streamline** discussionn ;-))

A long time ago (just two years and half ago!)

dRICH estimates of throughput (ATEHENA)

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

From ATHENA proposal (from dRICH section)

The sensor choice for a dRICH is quite challenging. After careful consideration, the only viable choice is the well known SiPM technology. These devices are ideal in terms of quantum efficiency, sensitive wavelength range, and single-photon signal size. The difficult aspect regards dark currents that arise as the devices receive a radiation dose. It has been demonstrated by studies at INFN that this damage can be repaired by thermal annealing, while R&D towards in-situ annealing are ongoing. The ATHENA design thereby features SiPM-based photon sensors in both the dRICH and the backward proximity-focussing RICH. Detailed calculations of the worst-case dark current rates have been used as the basis of estimates for the DAQ needs of the dRICH. Figure 2.8 shows the layout and performance of the ATHENA dRICH. As described previously, the optimization of the dRICH location (panel a) and optics (panel b) have already achieved the performance requirements in the Yellow Report.

(from DAQ section)

The biggest challenge to the goal of fully reading out the ATHENA detector with no deadtime will be the dark currents from the SiPM readout, expected to gradually increase with accumulated radiation dose. The current estimates assume an average rate of 300 kHz/sensor over the full detector is reached before undertaking an annealing cycle that will partially restore initial conditions of 3 kHz/sensor. This dark current is indistinguishable from signals from single photoelectrons. We hope to reduce this by a factor of 3–5 in the FEBs using sample cuts relative to the bunch crossing time. Further reduction can be obtained by a software trigger applied in the DAQ computers. Requiring a collision to be present will provide a data reduction by a factor of at least 200 allowing the ATHENA DAQ to write all collision data to tape. Another option for data reduction is by machine learning techniques implemented in the FPGAs of the FELIX boards; dedicated development and feedback from initial data are needed.

The data stream estimates @ ATHENA time still holds: **O(1-2 Tbps)@300 kHz/sensor**

The three key ingredients listed **"software trigger"**, **"require collision"** and "**ML techniques**" are on the table today

And further reading the specific bits..

This is now implemented on ALCOR64 shutter (gating)

from signals from single photoelectrons. We hope to reduce this by a factor of 3–5 in the FEBs using sample cuts relative to the bunch crossing time. Further reduction can be obtained by a software trigger applied in the DAQ computers. Requiring a collision to be present will provide a data reduction by a factor of at least 200 allowing the ATHENA DAQ to write all collision data to tape. Another option for data reduction is by machine learning techniques implemented in the FPGAs of the FELIX boards; dedicated development and feedback from initial data are needed.

Software trigger means making a selection "after" data are on computers (beyond DAM)

Requiring a collision means look to physics → Elke's talk + "detectors" talks

ML on FPGA for RICH (NA62 expertise) → Alessandro/Luca talk

First (years) data will be critical: we don't need to reduce initial data

do you remember the DPAP?

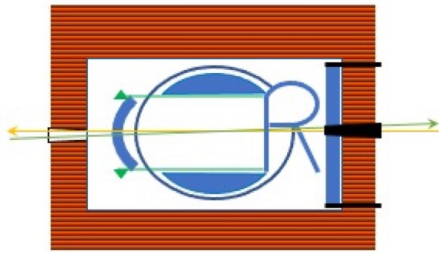

From **DPAP final report** (March 2022, two years and 4 months ago)!

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.

1

2

3

*Three steps recommended: we have now to work out **the last one***

The DAQ & electronics PDR (June 2024)

Report of the

Incremental Preliminary Design and Safety Review of the EIC Detector DAQ and Electronics

Responses to Questions

Charge Question 1:

Are the technical performance requirements appropriately defined and complete for this stage of the project?

Yes. The teams are working towards well-defined specifications of the front-end ASICs for processing the detector signals. The requirements for ReadOut/DAQ/controls are less mature but acceptable for this stage of the project. The remaining steps to bring these to maturity appear to be well defined.

Comments

- We applaud the adaptation of the readout architectures to match detector environment specs, in particular for the dRICH where steps to mitigate the high DCR will be implemented in the ASIC together with other handles such as triggering within the DAM.
- We appreciate these steps towards robustness. However, it would be reassuring to see simulations or measurements to demonstrate the effectiveness of the shuttering mechanism for dRICH.
- The RO/DAQ architecture is evolving in the right direction, but some specifications are yet to be fully defined and this should now move ahead at speed to allow implementation (RDOs in particular) to start.
- We feel that a dedicated overview and discussion on the slow-controls would have been helpful for the review, in particular from the point-of-view of hardware/firmware.

(a little bit of celebration...)

dRICH DAQ (from RDO to DAQ servers)

Full dRICH data modeling understood: throughput will be really high only after several years of operations
- see ageing modeling and annealing procedures for SiPM DCR (backup and in pre-brief)
- see detailed data modeling throughput (backup)
- see J. Landgraf presentation

From electronics & DAQ PDR

dRICH DAQ design foresees two levels of DAMs:

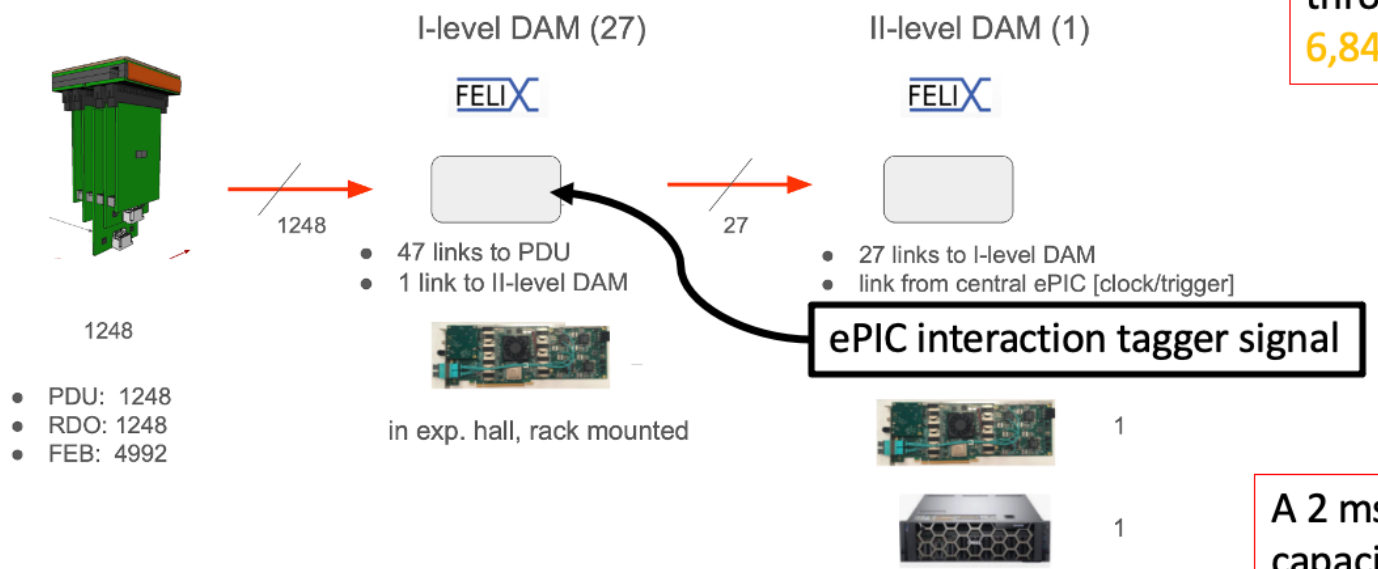

Feasibility studies for an interaction tagger are on-going providing a factor 100-200 reduction reducing throughput out of DAM (PCI memory) to:
6,84 Gbps - 70 Mbps

A 2 ms latency requires less than 1 MB buffering capacity in each DAM

dRICH data stream modeling: sensors + ALCOR

ALCOR parameters INPUT		Notes
Front end limit [kHz]	4000	
ALCOR Clock [MHz]	394,08	It will be 394.08 MHz or 295.55 MHz
Channels/serializer	8	
Bits per hit	64	2 32-bit words per hit (also TOT)
Bits per hit encoding 8/10	80	
Serializer band limit [Mb/s]	788,16	
Theoretical Serializer limit/ channel [kHz]	1231,5	this would be with 0 control words
Serializer limit single ch [kHz]	800	this is expected to improve with ALCOR v3
Number of serializer per chip	8	
Channel/chip	64	
Shutter width (ns)	2	(if you put 10 ns == no shutter)

relevant for data reduction. We cut a factor 5 here.

dRICH data stream modeling: DAQ parameters

dRICH DAQ parameters	
RDO boards	1248
ALCOR64 x RDO	4
dRICH channels (total)	319488
Number of DAM L1	27
Input link in DAM L1	47
Output links in DAM L1	1
Number of DAM L2	1
Input link to DAM L2	27
Link bandwidth [Gb/s] (assumes VTRX+)	10
Interaction tagger reduction factor	50
Interaction tagger latency [s]	2,00E-03
EIC parameters	
EIC Clock [MHz]	98,522
Orbit efficiency (takes into account gap)	0,92

with final design we might end up with something less

this is "baseline" / minimal configuration

dRICH data stream modeling: reduction factor

dRICH DAQ parameters	
RDO boards	1248
ALCOR64 x RDO	4
dRICH channels (total)	319488
Number of DAM L1	27
Input link in DAM L1	47
Output links in DAM L1	1
Number of DAM L2	1
Input link to DAM L2	27
Link bandwidth [Gb/s] (assumes VTRX+)	10
Interaction tagger reduction factor	50
Interaction tagger latency [s]	
EIC parameters	
EIC Clock [MHz]	
Orbit efficiency (takes into account gap)	

Foglio1

with final design we might end up with something less

this is "baseline" / minimal configuration

1 no data reduction

5

10

50

we might target something in this range

200 physics: BC rate / IR interaction rate = (100 MHz)/(500 kHz) = 200

However, only a fraction of the phsyics (IR = $\mathcal{L}^* \sigma$) gives a charged particle above threshold on dRICH → dIT (Marco's talk)

dRICH data stream modeling: DAQ parameters

dRICH DAQ parameters	
RDO boards	1248
ALCOR64 x RDO	4
dRICH channels (total)	319488
Number of DAM L1	27
Input link in DAM L1	47
Output links in DAM L1	1
Number of DAM L2	1
Input link to DAM L2	27
Link bandwidth [Gb/s] (assumes VTRX+)	10
Interaction tagger reduction factor	50
Interaction tagger latency [s]	2,00E-03
EIC parameters	
EIC Clock [MHz]	98,522
Orbit efficiency (takes into account gap)	0,92

with final design we might end up with something less

this is "baseline" / minimal configuration

Here we take (not full) advantage of EIC beam structure

About EIC beam structure and data reduction

Beam structure and bunch crossing timing

- Beam structure repeats every ~12.7886 μs
  - → Revolution frequency: ~78.195 kHz
- There are 1260 clock ticks in each revolution
  - → Clock period: ~10.14968 ns
    - Frequency: ~98.52525 MHz
  - → 1160 "bunch" ticks *approximately* marking *potentially* active intervals with particles
    - *Approximately*: electron beam structure is much more stable compared to hadron beam structure
    - *Potentially*: there can be 290 or 1160 active bunches with particles
  - → 100 "gap" ticks without particles

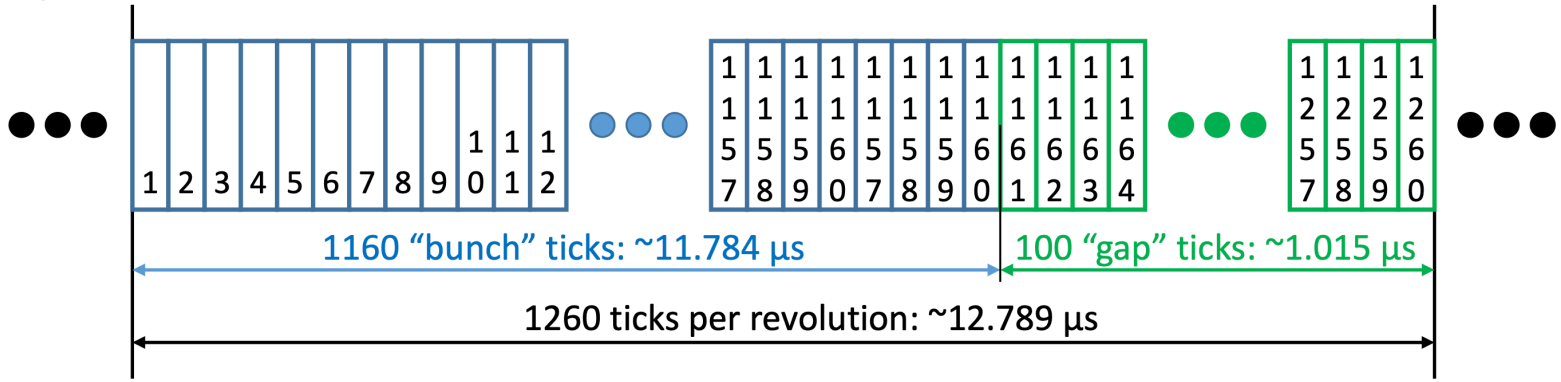

- Assume a stable clock can be derived from Machine with this frequency despite of bunch time variations
  - → This clock or its (sub)multiples is distributed to frontend electronics as "System clock"
    - With bunch-phase recovery mechanism
    - Used for bunch level synchronization, coarse timestamp bookkeeping, serial communication and possibly timing measurements.

Sep 8, 22

This could add a factor 4 reduction, not implemented in the model so far

100/1260=7.9% of BC are empty: implemented in the model

dRICH data stream: the DCR/sensor parameter

Interaction tagger reduction factor	50		Channel/chip	64
Interaction tagger latency [s]	2,00E-03		**Shutter width (ns)**	2
EIC parameters				
EIC Clock [MHz]	98,522			
Orbit efficiency (takes into account gap)	0,92			

dRICH data stream analysis		**Limit**	**Comments**
Sensor rate per channel [kHz]	20,00	4.000,00	
Rate post-shutter [kHz]	3,68	800,00	
Throughput to serializer [Mb/s]	2,30	788,16	
Throughput from ALCOR64 [Mb/s]	18,40		limit FPGA dependent: - check with RDO
Throughput from RDO [Gb/s]	0,07	10,00	based on VTRX+
Input at each DAM I [Gbps]	3,38	470,00	
Buffering capacity at DAM I [MB]	0,86		to be checked but seems manageable
Throughput from DAM I to DAM II [Gbps]	0,07	10,00	this might be higher (from FELIX to FELIX)
Output to each DAM II [Gbps]	1,82	270,00	

Aggregated dRICH data throughput		**Comments**
Total input at DAM I [Gb/s]	91,21	This is only "inside" DAM, not to be transferred on PCI
Total input at DAM II [Gb/s]	1,82	This is based on aggregation above + reduction factor of the interaction tagger
Total output from DAM II [Gb/s]	1,82	Further reduction possible to be investigated (FPGA level?)

At given DCR/sensor, shutter width and DAM data reduction factor, you know what happens...

Enjoy: https://docs.google.com/spreadsheets/d/1P3qoogFWuicXDgojwvhaFL2EnwQ7BEmGlITg1fwDUkE/edit?usp=sharing

dRICH data stream at EIC day-0

Interaction tagger reduction factor	1
Interaction tagger latency [s]	2,00E-03
EIC parameters	
EIC Clock [MHz]	98,522
Orbit efficiency (takes into account gap)	0,92

Channel/chip	64
Shutter width (ns)	10

We first learn how to calibrate the shutter

dRICH data stream analysis		**Limit**	**Comments**
Sensor rate per channel [kHz]	3,00	4.000,00	
Rate post-shutter [kHz]	2,76	800,00	
Throughput to serializer [Mb/s]	1,73	788,16	
Throughput from ALCOR64 [Mb/s]	13,80		limit FPGA dependent: - check with RDO
Throughput from RDO [Gb/s]	0,05	10,00	based on VTRX+
Input at each DAM I [Gbps]	2,53	470,00	
Buffering capacity at DAM I [MB]	0,65		to be checked but seems manageable
Throughput from DAM I to DAM II [Gbps]	2,53	10,00	this might be higher (from FELIX to FELIX)
Output to each DAM II [Gbps]	68,41	270,00	

Aggregated dRICH data throughput		**Comments**
Total input at DAM I [Gb/s]	68,41	This is only "inside" DAM, not to be transferred on PCI
Total input at DAM II [Gb/s]	68,41	This is based on aggregation above + reduction factor of the interaction tagger
Total output from DAM II [Gb/s]	68,41	Further reduction possible to be investigated (FPGA level?)

dRICH data stream at EIC day-? (few months..)

Interaction tagger reduction factor	1		Channel/chip	64
Interaction tagger latency [s]	2,00E-03		**Shutter width (ns)**	2
EIC parameters				
EIC Clock [MHz]	98,522			
Orbit efficiency (takes into account gap)	0,92			

We then train our tagger/data reduction algorithm

dRICH data stream analysis		**Limit**	**Comments**
Sensor rate per channel [kHz]	3,00	4.000,00	
Rate post-shutter [kHz]	0,55	800,00	
Throughput to serializer [Mb/s]	0,35	788,16	
Throughput from ALCOR64 [Mb/s]	2,76		limit FPGA dependent: - check with RDO
Throughput from RDO [Gb/s]	0,01	10,00	based on VTRX+
Input at each DAM I [Gbps]	0,51	470,00	
Buffering capacity at DAM I [MB]	0,13		to be checked but seems manageable
Throughput from DAM I to DAM II [Gbps]	0,51	10,00	this might be higher (from FELIX to FELIX)
Output to each DAM II [Gbps]	13,68	270,00	

Aggregated dRICH data throughput		**Comments**
Total input at DAM I [Gb/s]	13,68	This is only "inside" DAM, not to be transferred on PCI
Total input at DAM II [Gb/s]	13,68	This is based on aggregation above + reduction factor of the interaction tagger
Total output from DAM II [Gb/s]	13,68	Further reduction possible to be investigated (FPGA level?)

dRICH data stream at maximum DCR damage

Interaction tagger reduction factor	1		Channel/chip	64	
Interaction tagger latency [s]	2,00E-03		**Shutter width (ns)**	2	(i
EIC parameters					
EIC Clock [MHz]	98,522				
Orbit efficiency (takes into account gap)	0,92				

dRICH data stream analysis		**Limit**	**Comments**
Sensor rate per channel [kHz]	300,00	4.000,00	
Rate post-shutter [kHz]	55,20	800,00	
Throughput to serializer [Mb/s]	34,50	788,16	
Throughput from ALCOR64 [Mb/s]	276,00		limit FPGA dependent: - check with RDO
Throughput from RDO [Gb/s]	1,08	10,00	based on VTRX+
Input at each DAM I [Gbps]	50,67	470,00	
Buffering capacity at DAM I [MB]	12,97		to be checked but seems manageable
Throughput from DAM I to DAM II [Gbps]	50,67	10,00	this might be higher (from FELIX to FELIX)
Output to each DAM II [Gbps]	1.368,14	270,00	

Aggregated dRICH data throughput		**Comments**
Total input at DAM I [Gb/s]	1.368,14	This is only "inside" DAM, not to be transferred on PCI
Total input at DAM II [Gb/s]	1.368,14	This is based on aggregation above + reduction factor of the interaction tagger
Total output from DAM II [Gb/s]	1.368,14	Further reduction possible to be investigated (FPGA level?)

Per se at maximum damage we "just" have 50 Gbps inside DAM-1 manageable to go to PCIe?

dRICH data stream getting only "true" interactions

Interaction tagger reduction factor	200		Channel/chip	64	
Interaction tagger latency [s]	2,00E-03		**Shutter width (ns)**	2	(i
EIC parameters					
EIC Clock [MHz]	98,522				
Orbit efficiency (takes into account gap)	0,92				

dRICH data stream analysis		**Limit**	**Comments**
Sensor rate per channel [kHz]	300,00	4.000,00	
Rate post-shutter [kHz]	55,20	800,00	
Throughput to serializer [Mb/s]	34,50	788,16	
Throughput from ALCOR64 [Mb/s]	276,00		limit FPGA dependent: - check with RDO
Throughput from RDO [Gb/s]	1,08	10,00	based on VTRX+
Input at each DAM I [Gbps]	50,67	470,00	
Buffering capacity at DAM I [MB]	12,97		to be checked but seems manageable
Throughput from DAM I to DAM II [Gbps]	0,25	10,00	this might be higher (from FELIX to FELIX)
Output to each DAM II [Gbps]	6,84	270,00	

Aggregated dRICH data throughput		**Comments**
Total input at DAM I [Gb/s]	1.368,14	This is only "inside" DAM, not to be transferred on PCI
Total input at DAM II [Gb/s]	6,84	This is based on aggregation above + reduction factor of the interaction tagger
Total output from DAM II [Gb/s]	6,84	Further reduction possible to be investigated (FPGA level?)

- This is reduction case "by physics"
- We can of course always allow some background events to pass
- A full ep 500 kHZ tagger needs – necessarily – to tag also events where we don't have charged tracks in dRICH (low-Q^2 etc.)

How we presented things to DAQ/Elec. PDR

data throughput from sensor to RDO opt. link

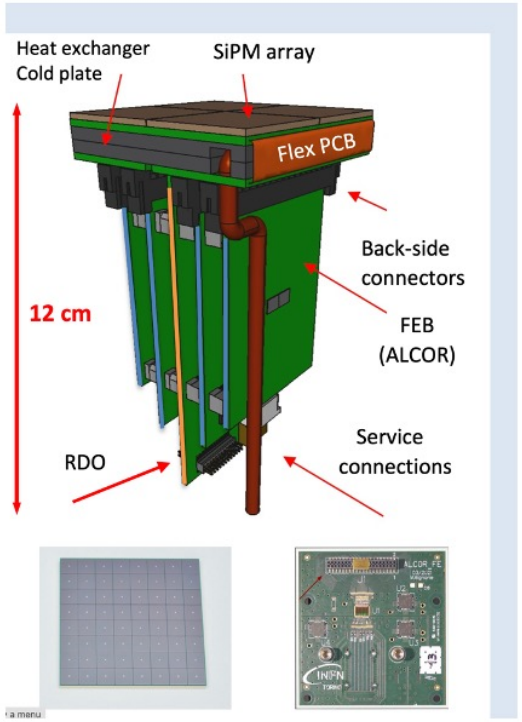

SiPM sensor

x 64

SiPM matrix

ALCOR64 (ALCORV3)

x 4

ALCOR SHUTTER
REDUCTION (**5**)
INCLUDED HERE

3x3 mm^2

300 kHz
3 kHz

34.5 Mbps each lane
0.35 Mbps each lane

x 8 differential lanes

276 Mbps each ALCOR64
2.76 Mbps each ALCOR64

At maximum rad. damage
At EIC start

AMD ARTIX

x 1

1.08 Gbps RDO output through VTRX+
10 Mbps output through VTRX+

opt.transceiver/FPGA requirement: it has to be able to implement optical link/Multigigabit transceiver **up to 10 Gbps**

- see F. Cossio presentation about the shutter
- detailed modeling of data throughput (down to DAM) available on backup

Global dRICH throughput out of RDO:
1.4 Tbps (14 Gbps)

So what? Many possibilities!

- check if some tracks is entering dRICH → "dIT"
- combine different detec

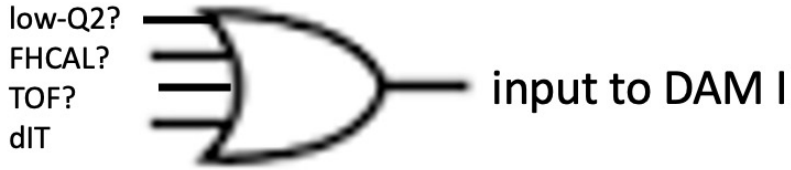

- data Reduction Algorithms: the need to aggregate:

In a DAM-I we have 47 links input (approx 20% of a dRICH sector)
→ approx 1.1 hit (signal) in 2 ns (in 1 BC)
→ approx 0.7 hit (DCR) in 2 ns (in 1 BC)

(you can even be artificially intelligent but you need to aggregate data before taking decisions, you might may be apply a Hough Transfrom but not the analysis of the pattern (that needs aggregated data))

- FPGA or GPU?

data Reduction approaches: baseline

Blue sky exercise

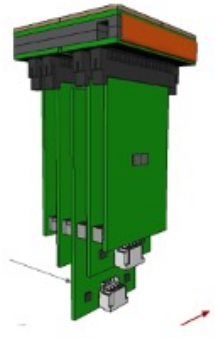

1248

- PDU: 1248
- RDO: 1248
- FEB: 4992

I-level DAM (27)

FELIX

1248

- 47 links to PDU
- 1 link to II-level DAM

27

in exp. hall, rack mounted

II-level DAM (1)

FELIX

27

- 27 links to I-level DAM
- link from central ePIC [clock/trigger]

ePIC interaction tagger signal

1

1

In **Alessandro's talk** different configuration explored with slight increase of DAM Variants but **the approach is FPGA based**

Data reduction model: the GPU approach

Blue sky exercise

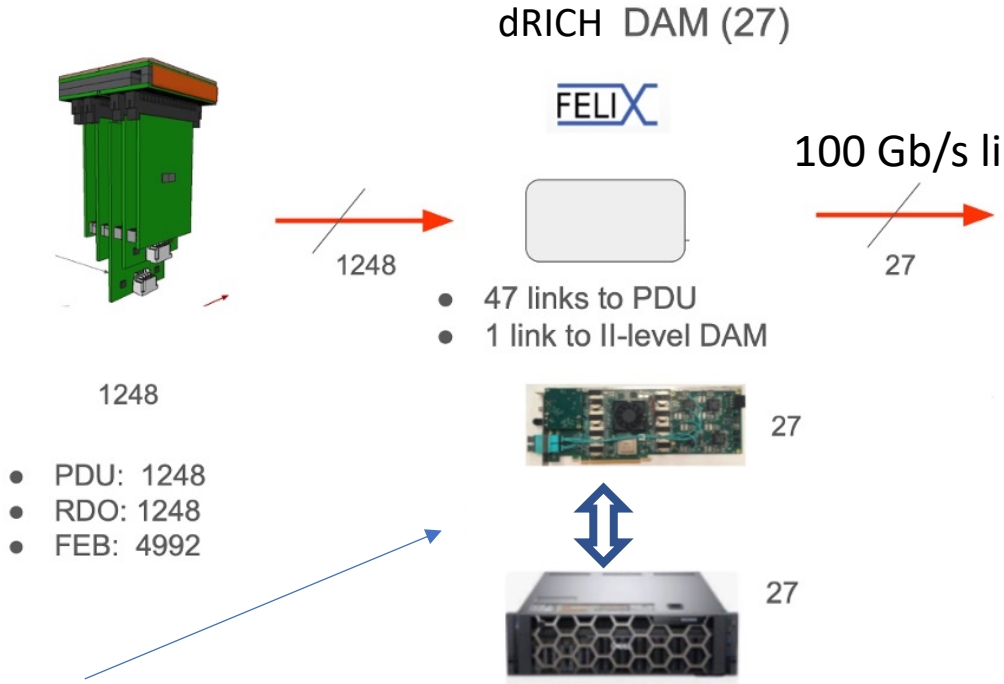

High end switch

"BC dRICH fragments" sent to farm with GPU with round-robin algorithm

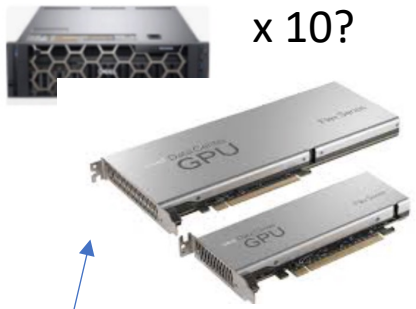

(just an example from ARISTA)

the bottleneck here can/would be
the maximum transfer/rate between FLX-155 trhough PCIe?
(PCIe 5.0 x16 Bandwidth 64 GB/s *per se*)
but 100 Gbps (on eth Eth) is ultimate botttleneck

Information from "taggers" should go here (I guess)
(via data stream?)

Last comment:

Ageing model

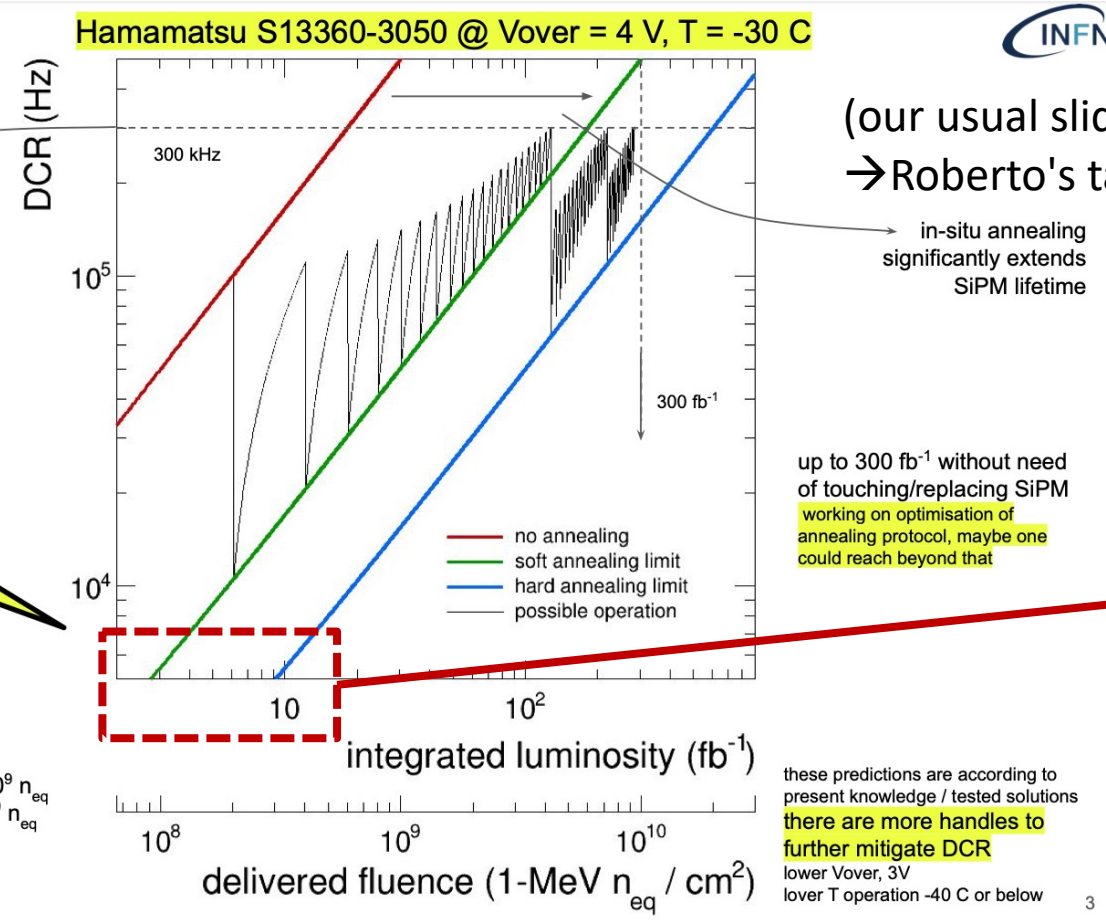

max acceptable DCR for Physics performance
~ 10 noise hits / sector within 500 ps

in-situ annealing significantly extends SiPM lifetime

up to 300 fb^{-1} without need of touching/replacing SiPM
working on optimisation of annealing protocol, maybe one could reach beyond that

the "possible operation" scenario shown here has 44 soft-annealing cycles and 3 hard-annealing cycles

model input from R&D measurements (up to 2022)
- DCR increase: 500 kHz/10^9 n_{eq}
- residual DCR (online annealing): 50 kHz/10^9 n_{eq}
- residual DCR (oven annealing): 15 kHz/10^9 n_{eq}

1-MeV neq fluence from background group
- 1.75 10^7 n_{eq} / fb^{-1}
- add an extra 2x safety factor

these predictions are according to present knowledge / tested solutions
there are more handles to further mitigate DCR
lower Vover, 3V
lover T operation -40 C or below

3

(our usual slide about ageing model
→Roberto's talk yetserday

it would be very useful having an estimate of radiation levels (from background group) for first years (**"early EIC Physics Program"**)! This would give a better understanding of time scale...

Conclusions

- model of dRICH "data stream" is defined and helpful to design back-end solutions ("beyond RDO")
- the so-called dRICH issue (high trhoughput) is understood and under control
- we have now to define the approach and the exact design of the dRICH data reduction solution embedded in streaming DAQ and computing