UNIVERSITÄT BONN

ATLAS ITK PIXEL

MODULE POWER ANALYSIS

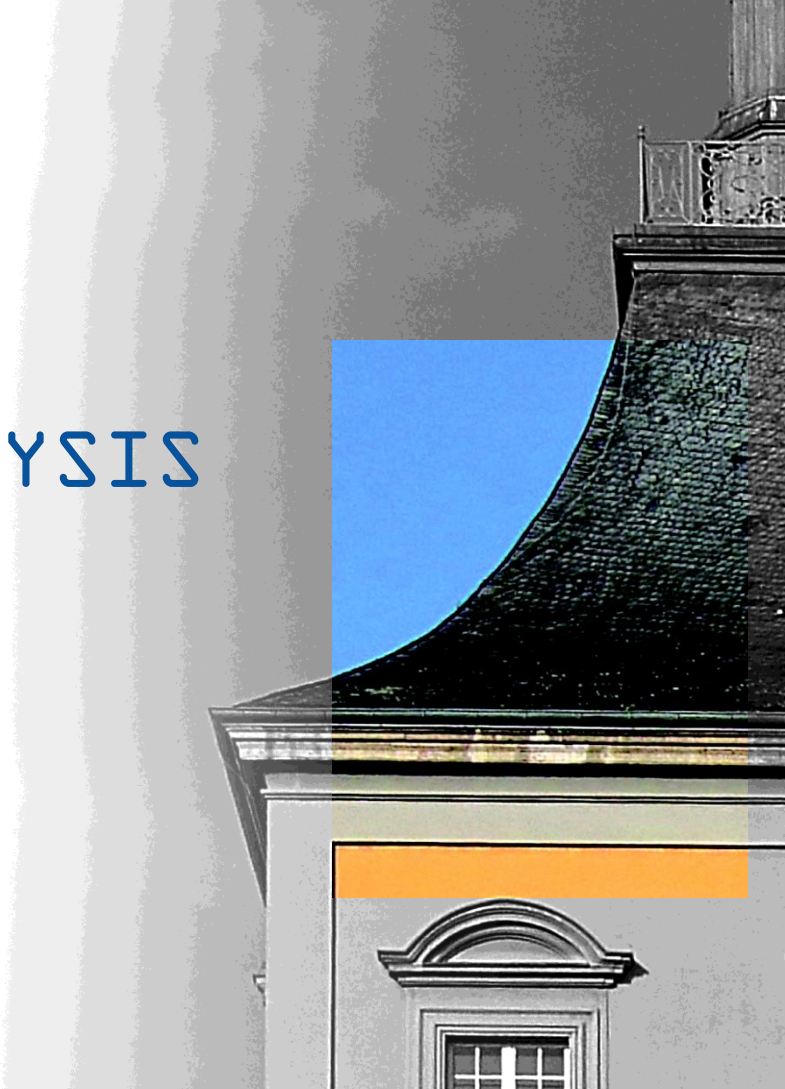

Matthias Hamer

RD53B Testing Meeting, 19 May 2021

UNIVERSITÄT BONN

INTRO

- update on the SLDO configuration space scans that were done before
  - immediate motivation for ATLAS to do this now was to get a reasonable and conservative power estimate for the thermal FEA
  - in the not so far future we need to converge on baseline choices for our modules (FDR is approaching, pre-production to start after that)
- for previous analyses we have considered the ‘LSD plots’
  - binary plots, either ‘allowing’ a set of parameters, or not
  - gave us a good idea of the parameter space we can expect to use
  - do not give us the full story, in particular for our power estimates
- today presenting an updated version of this study with somewhat more quantitative output

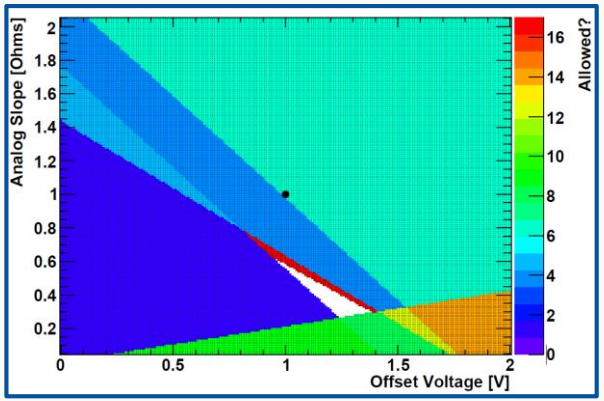

ITK PIXELS IN A NUTSHELL

- ITk Pixel Detector
  - 5 layers of pixel modules
  - either quad chip modules or triplets, i.e. 3 or 4 FE chips powered in parallel on a serial powering unit
  - between 3 and 14 such serial powering units will be powered in series; modularity largely follows from mechanical and electrical constraints with almost no degrees of freedom left
  - 8.372 SP Units in 912 Serial Powering Chains, on 360 local supports, i.e. average chain length is 9 and 2.5 SP chains per local support
  - 3 basic flavours of power modules:
    - triplets with 3D sensors in L0 → special HV distribution
    - quads with 100 um thick planar sensors in L1
    - quads with 150 um thick planar sensors in L2-L4

will share the same hybrid flex

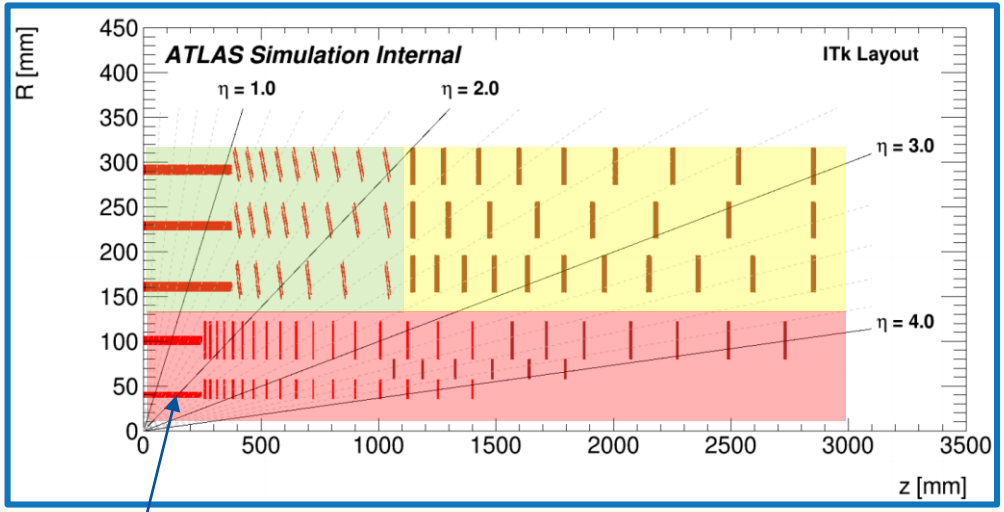

Deprecated Picture:

L0 Radius has been reduced by a few mm from 39 mm to 34 mm (barrel) and from 37 mm to 33 mm (inner radius for endcaps)

Outer Barrel:
4472 Quad Chip Modules
448 Serial Powering Chains
~7.2 m2 of active silicon

Outer Endcaps:
2344 Quad Chip Modules
224 Serial Powering Chains
~3.75 m2 of active silicon

Inner System:
1160 Quad Chip Modules + **396** Triplets
240 Serial Powering Chains
~2.5 m2 of active silicon

UNIVERSITÄT BONN

USED NUMBERS

- precise requirements depend on where in the detector a module is located
  - for previous iterations, used a detailed current consumption estimate for each layer
  - for this study have reduced the number of different sets to a total of 3 (instead of 15)
  - based on measurements provided by Timon:
    - L0: 750 mA analog, 770 mA digital
    - L1: 648 mA analog, 700 mA digital
    - L2: 540 mA analog, 650 mA digital
  - assuming these are the required currents, scanned the parameter space for Vofs and Slopes and determined the module yield for different overhead currents
- this is still a very parametric model
  - not all modules will be the same – there will be variations from chip to chip and from module to module
  - there is usually more than one source for differences → at the moment, the model I'm using is merging all of them in one big mess; a more detailed analysis is on my list, but I'm not sure when I get to that.

ANALYSIS STEPS

- for each set of requirements (L0, L1, L2), scanned the Vofs – Slope Analog plane
  - the ratio of **nominal** slopes for analog and digital parts was fixed according to the estimated current consumption:
    - L0: 0.974 (= RextD/RextA nominal)
    - L1: 0.926 (= RextD/RextA nominal)
    - L2: 0.831 (= RextD/RextA nominal)
  - for each bin in the Vofs-Slope-Analog plane, 10.000 pseudo modules were generated
    - Vofs according to a Gaussian with a width of 1%
    - Slopes with a width of 5%

    } parametrization of different contributions

    chip-to-chip process variation, tolerance of SMD resistors, differences in trace resistances on module flex for Vin and GND, . . .
    - the mean values of all distributions are the nominal values

I assume we cut on slopes during wafer probing → 5% is the _max_ deviation from the nominal value for each individual slope (→ expect the dominant contribution to the total width is chip-to-chip variation, not flex design)
there is no cut on the Vofs, 1% is the 1-sigma width, actual deviation can be larger in the MC

→ newest measurements from wafer probing indicate that the width of the slope distributions is overestimated and the width of the Vofs distributions is underestimated if we assume no offset-sharing (not sure about the case of offset sharing)

→ will improve the model here to consider different sources for the variations

ANALYSIS OF PARAMETER SPACE

- a module is considered to pass our QC, based on a few criteria
  - module voltage in normal operation with a certain global overhead
  - module voltage with an open FE on the module
    - always the same FE, not necessarily conservative for that reason
    - assuming the offset voltage does not drop, overly conservative for that reason
  - local overhead current for each regulator (not clear to me what a reasonable boundary is for this cut)
- output in this case are yield plots
- based on those yield plots, we defined a couple of candidate configurations and ran a more detailed MC with 1.000.000 modules each to get a few distributions of interest
- will go through yield plots first, then show a few details for some of the **candidate scenarios**

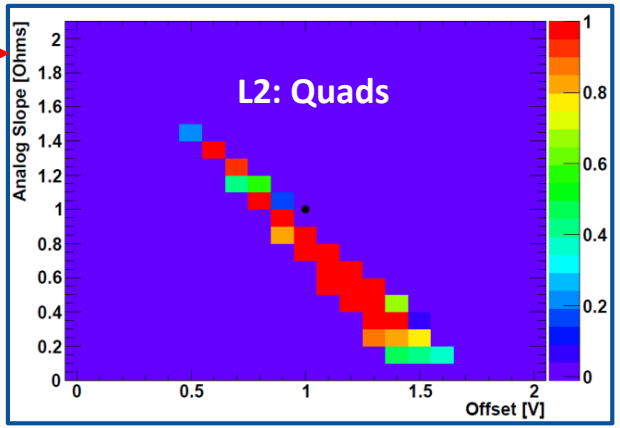

CLOSER LOOK AT YIELDS

reason to reject a module:
0: Vin > 1.8 V
1: Vin < 1.45 V
2: local headroom < 5%
3: Vin > 1.8 V with an open FE
multiple entries possible

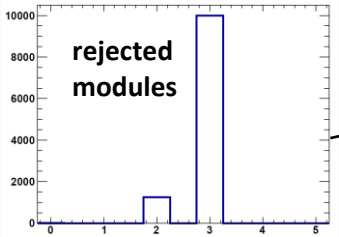

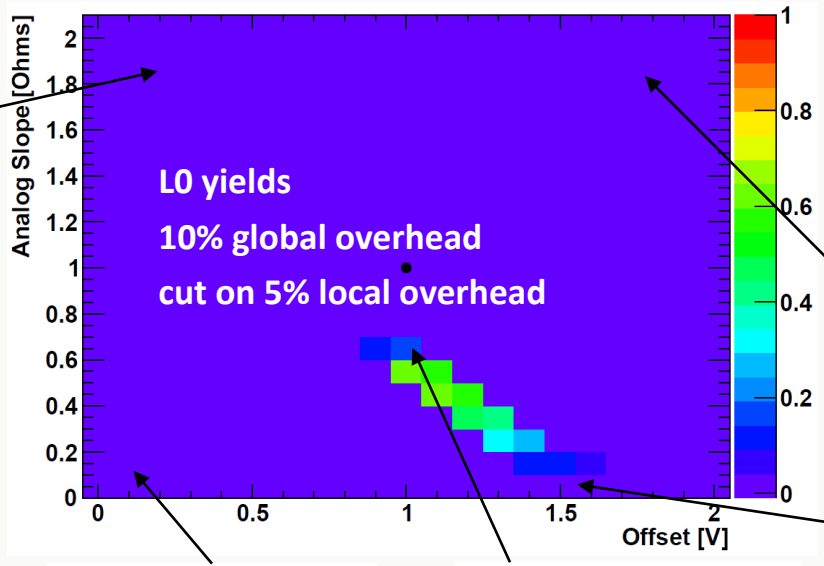

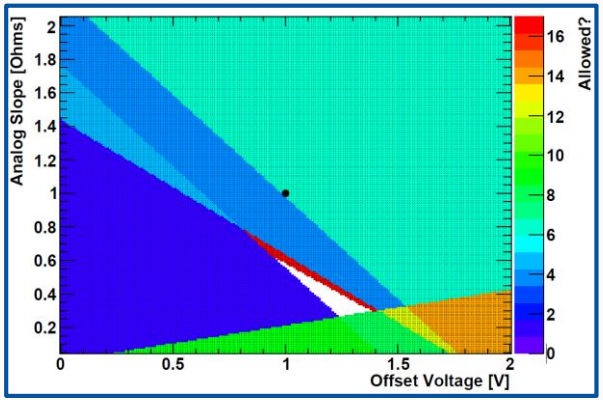

yield plots are MC version of the LSD plots

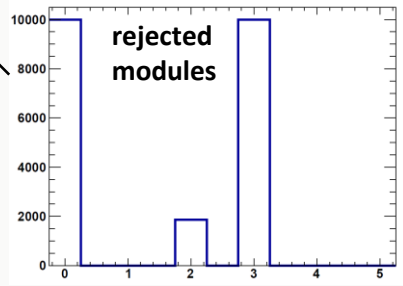

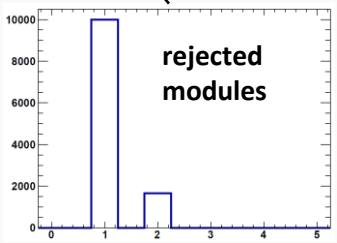

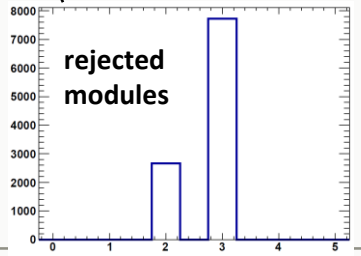

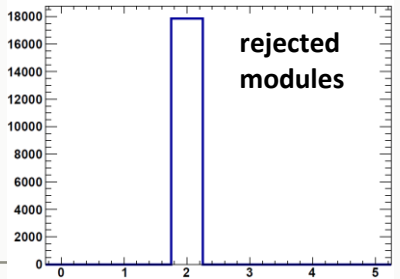

EXPECTED MODULE YIELDS - 20% HEADROOM

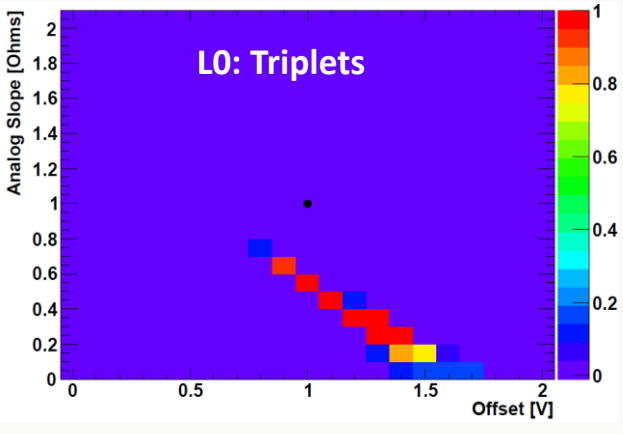

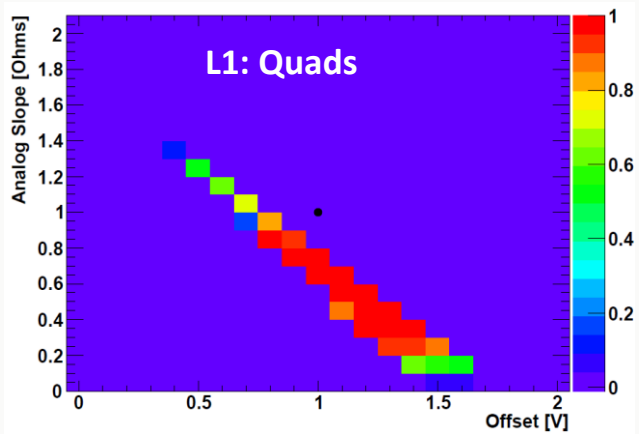

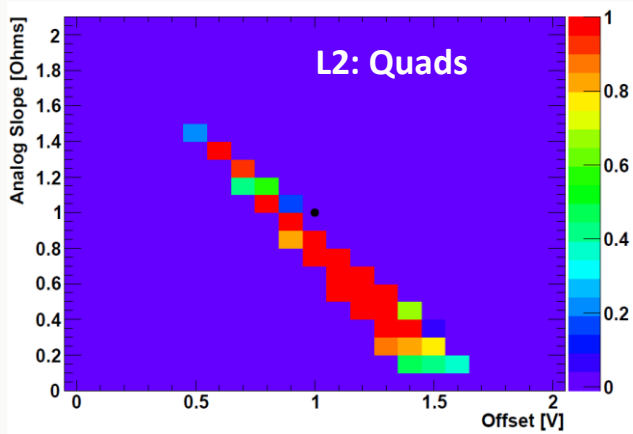

1. ~~Vmod < 1.8 V for 10% global current overhead~~
2. **Vmod > 1.45 V for 20% global current overhead**
3. **overhead > 5% for every single regulator for 20% global current overhead**
4. ~~Vmod < 1.8 V for 10% global current overhead and one FE open (not necessarily conservative)~~
5. Vmod < 1.8 V for 20% global current overhead
6. Vmod < 1.8 V for 20% global current overhead and one FE open (not necessarily conservative)

EXPECTED MODULE YIELDS - 10% HEADROOM

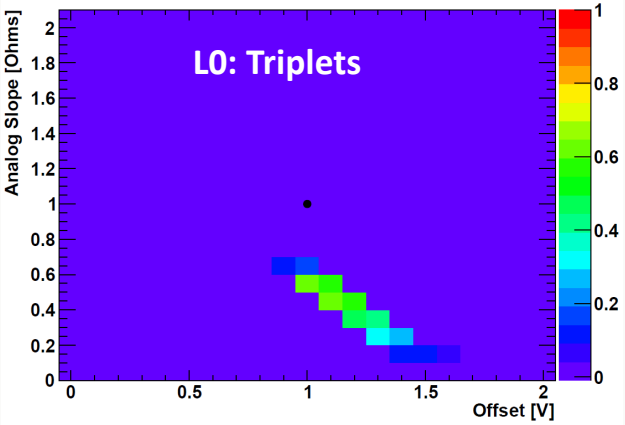

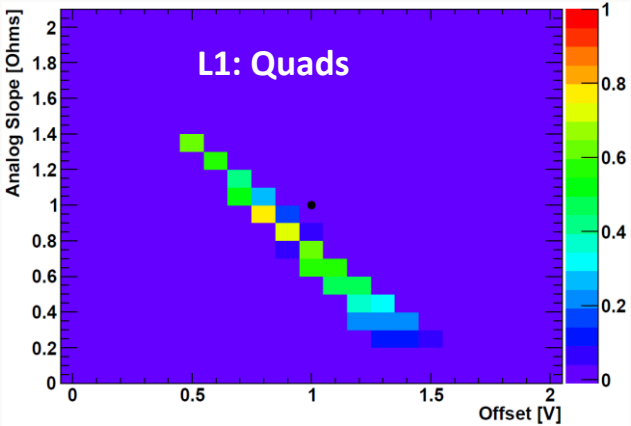

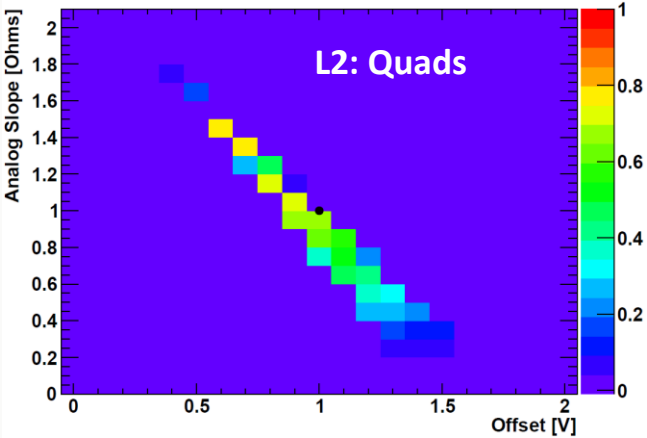

1. Vmod < 1.8 V for 10% global current overhead
2. Vmod > 1.45 V for 10% global current overhead
3. overhead > 5% for every single regulator for 10% global current overhead
4. Vmod < 1.8 V for 10% global current overhead and one FE open (not necessarily conservative)
5. ~~Vmod < 1.8 V for 20% global current overhead~~
6. ~~Vmod < 1.8 V for 20% global current overhead and one FE open (not necessarily conservative)~~

EXPECTED MODULE YIELDS

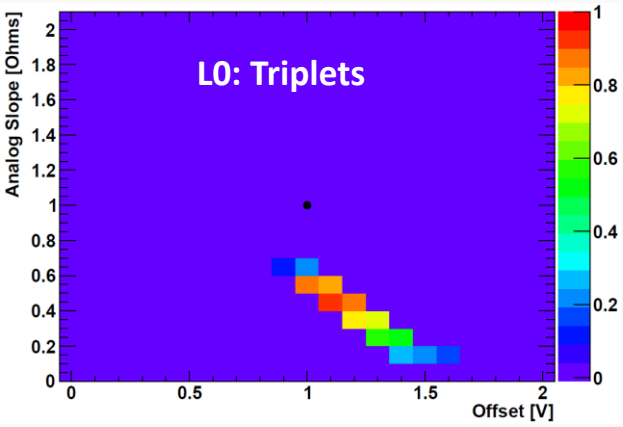

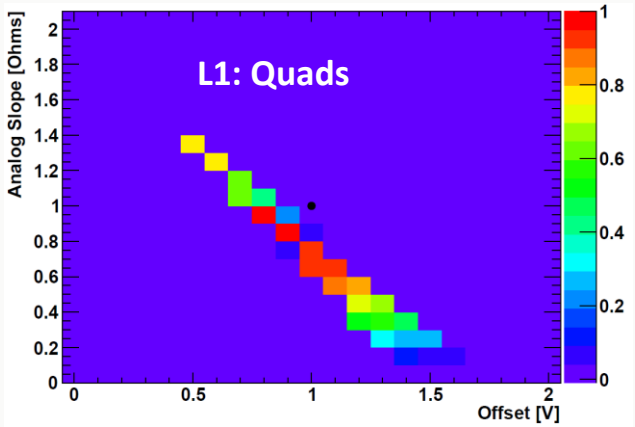

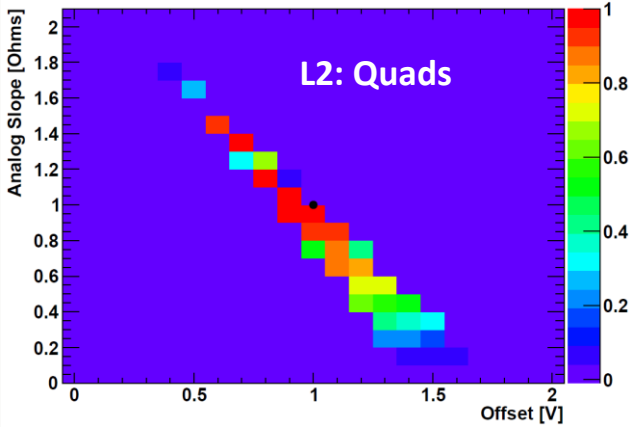

1. Vmod < 1.8 V for 10% global current overhead
2. Vmod > 1.45 V for 10% global current overhead
3. **overhead > 3% for every single regulator for 10% global current overhead**
4. Vmod < 1.8 V for 10% global current overhead and one FE open (not necessarily conservative)
5. ~~Vmod < 1.8 V for 20% global current overhead~~
6. ~~Vmod < 1.8 V for 20% global current overhead and one FE open (not necessarily conservative)~~

YIELD PLOTS FOR REDUCED DISPERSIONS

- reduced widths for Vofs and Slope distributions by a factor of 3
  - cutting on 3 sigma in the slope distributions in this case (i.e. accepting a 3-sigma deviation instead of a 1-sigma deviation in this case)

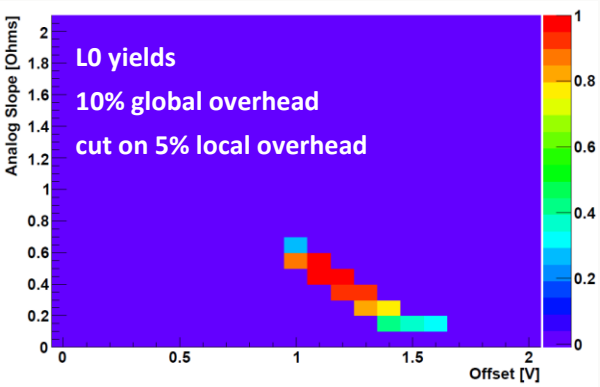

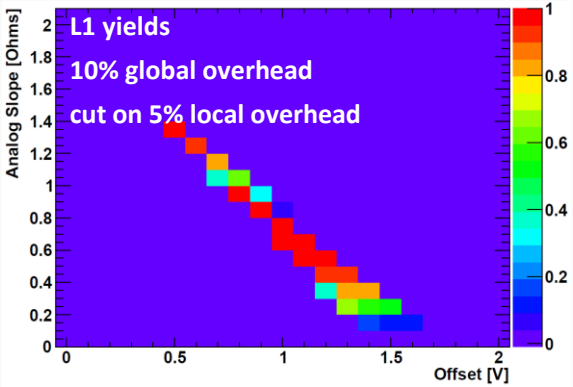

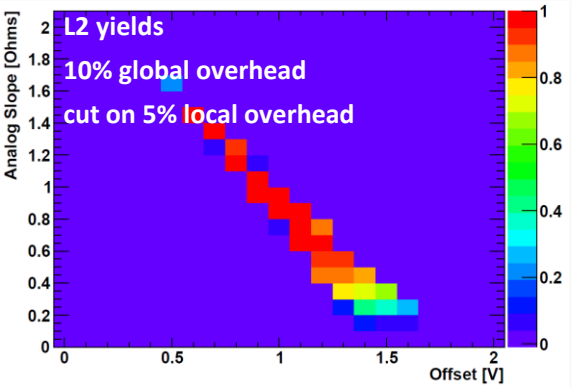

UNIVERSITÄT BONN

ITK CANDIDATE SCENARIOS

- before going into this, it should be understood that the results of the above scans are still to be considered somewhat random, due to the assumtions made on slide 5
- nevertheless, defined a few candidate scenarios, one of which we used to run a thermal FEA for our modules on the local supports
  - SC1: 20% global overhead current, high offset voltages, low slopes
  - SC2: 20% global overhead current, moderate offset voltages, moderate slopes
  - SC3: 10% global overhead current, moderate to high offset voltages, moderate to high slopes
  - Low Power Mode for Scenario SC2

PROPOSAL FOR NOMINAL OFFSETS AND SLOPES - 20% SCENARIO 1

- lower offsets and higher slopes are favoured as long as a the actual offset voltages differ from chip to chip (and the spread in the actual voltages is truly relative...)
- for the power analysis, I'll assume the following settings:
  - L0: Vofs = 1.30 V, Analog Slope = 0.35 Ohms
  - L1: Vofs = 1.25 V, Analog Slope = 0.45 Ohms
  - L2: Vofs = 1.25 V, Analog Slope = 0.55 Ohms

										One Chip 20% headroom								Average Power Density [W/cm2]							
	Required Current [A]			Required Current [A]			Proposed Configuration			Nominal				1 FE Open				Nominal		No Config		One FE Open		One FE Open - No Config	
	Analog	Digital	Total	Matrix	Periphery	Total	Vofs [V]	SlopeA [Ω]	SlopeD [Ω]	Current [A]	Voltage [V]	Power [W]	P Density [W/cm2]	Current [A]	Voltage [V]	Power [W]	P Density [W/cm2]	Matrix	Periphery	Matrix	Periphery	Matrix	Periphery	Matrix	Periphery
L0	0.75	0.77	1.52	1.14	0.38	1.52	1.30	0.35	0.34	1.824	1.615	2.946	0.696	2.736	1.773	4.850	1.146	0.354	4.361	0.000	8.142	0.354	9.623	0.000	13.404
L1	0.648	0.7	1.348	0.998	0.35	1.348	1.25	0.45	0.42	1.618	1.600	2.588	0.612	2.157	1.717	3.702	0.875	0.310	3.843	0.000	7.153	0.310	6.923	0.000	10.233
L2	0.54	0.65	1.19	0.85	0.34	1.19	1.25	0.55	0.46	1.428	1.606	2.294	0.542	1.904	1.725	3.285	0.776	0.264	3.521	0.000	6.340	0.264	6.260	0.000	9.079

20% SCENARIO - L0, 1.000.000 MODULES

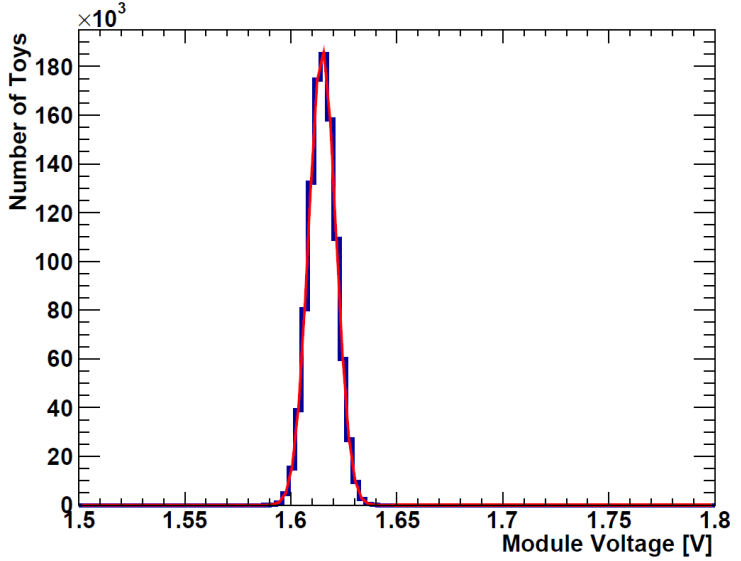

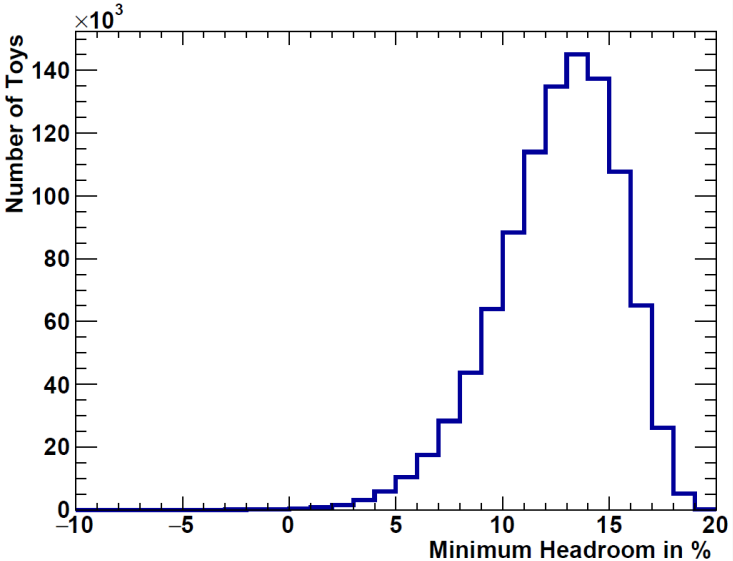

- Module Power: 8.84 +- 0.04 W (1 sigma)
- Total L0 Power: 3500 +- 0.72 W (1 sigma)

20% SCENARIO - L1

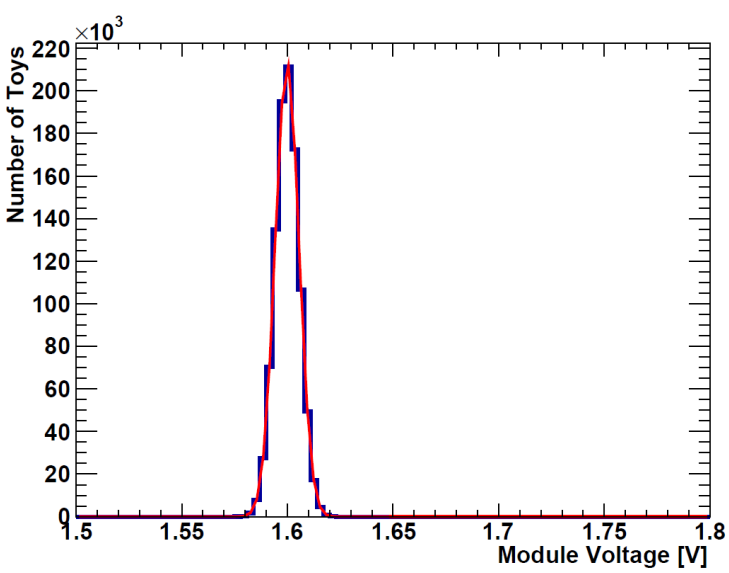

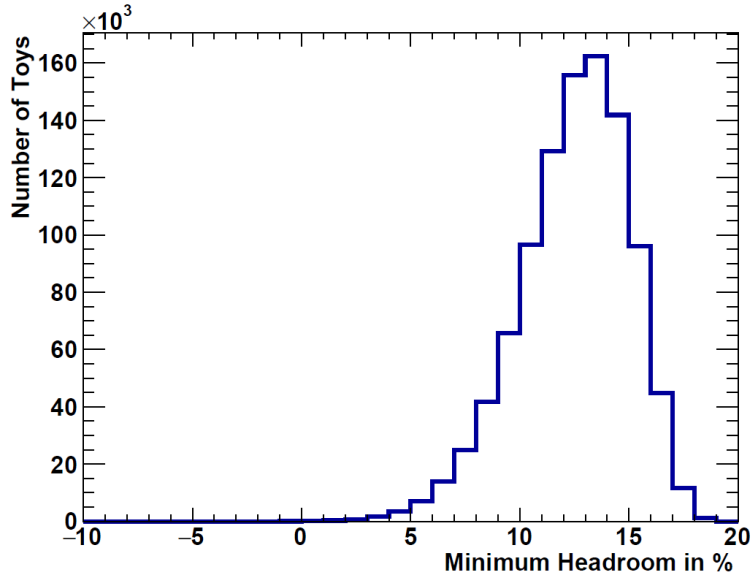

- Module Power: 10.35 +- 0.04 W (1 sigma)
- Total L1 Power: 12000 +- 1.29 W (1 sigma)

20% SCENARIO - L2

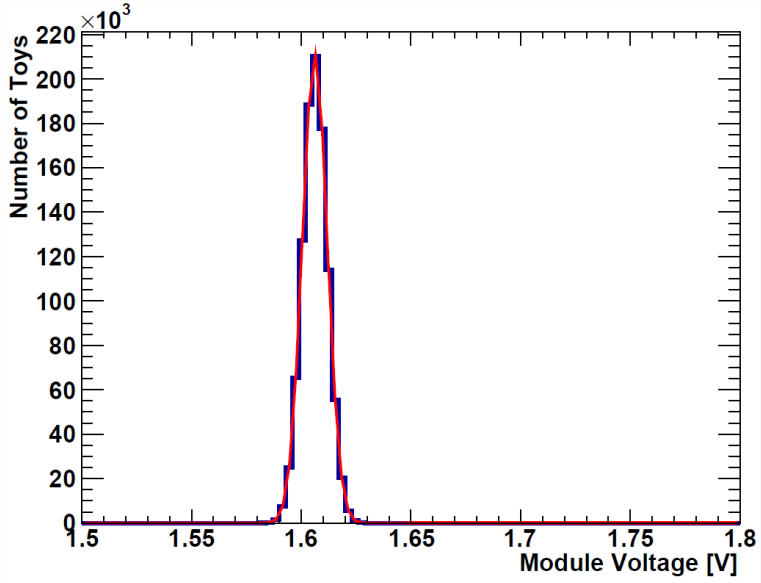

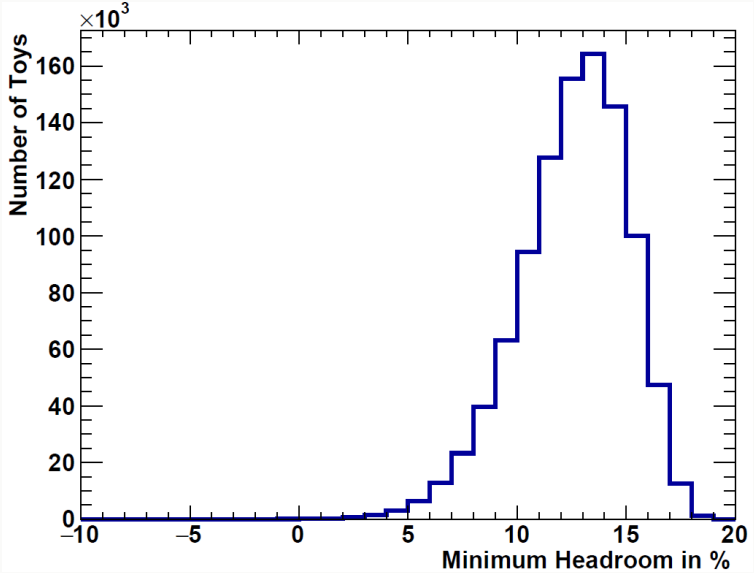

- Module Power: 9.18 +- 0.03 W (1 sigma)
- Total OS Power: 62589 +- 2.8 W (1 sigma – assuming all layers running with L2 settings)

PROPOSAL FOR NOMINAL OFFSETS AND SLOPES - 20% SCENARIO 2

- lower offsets and higher slopes are favoured as long as a the actual offset voltages differ from chip to chip (and the spread in the actual voltages is truly relative...)
- for the power analysis, I'll assume the following settings:
  - L0: Vofs = 1.10 V, Analog Slope = 0.55 Ohms
  - L1: Vofs = 1.00 V, Analog Slope = 0.77 Ohms
  - L2: Vofs = 1.00 V, Analog Slope = 0.93 Ohms

										One Chip								Average Power Density [W/cm2]							
										20% headroom															
	Required Current [A]			Required Current [A]			Proposed Configuration			Nominal				1 FE Open				Nominal		No Config		One FE Open		One FE Open - No Config	
	Analog	Digital	Total	Matrix	Periphery	Total	Vofs [V]	SlopeA [Ω]	SlopeD [Ω]	Current [A]	Voltage [V]	Power [W]	P Density [W/cm2]	Current [A]	Voltage [V]	Power [W]	P Density [W/cm2]	Matrix	Periphery	Matrix	Periphery	Matrix	Periphery	Matrix	Periphery
L0	0.75	0.77	1.52	1.14	0.38	1.52	1.10	0.55	0.536	1.824	1.595	2.909	0.688	2.736	1.843	5.041	1.191	0.354	4.260	0.000	8.041	0.354	10.152	0.000	13.933
L1	0.648	0.7	1.348	0.998	0.35	1.348	1.00	0.77	0.713	1.618	1.599	2.586	0.611	2.157	1.798	3.879	0.917	0.310	3.838	0.000	7.148	0.310	7.410	0.000	10.720
L2	0.54	0.65	1.19	0.85	0.34	1.19	1.00	0.93	0.773	1.428	1.603	2.289	0.541	1.904	1.804	3.434	0.812	0.264	3.506	0.000	6.326	0.264	6.672	0.000	9.491
																		Case 2		Case 3		Case 4		Case 5	

20% SCENARIO - L0, 1.000.000 MODULES

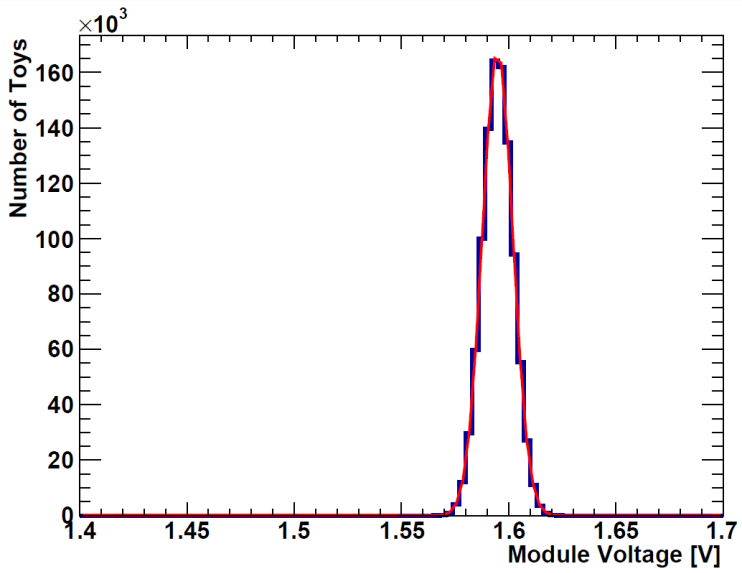

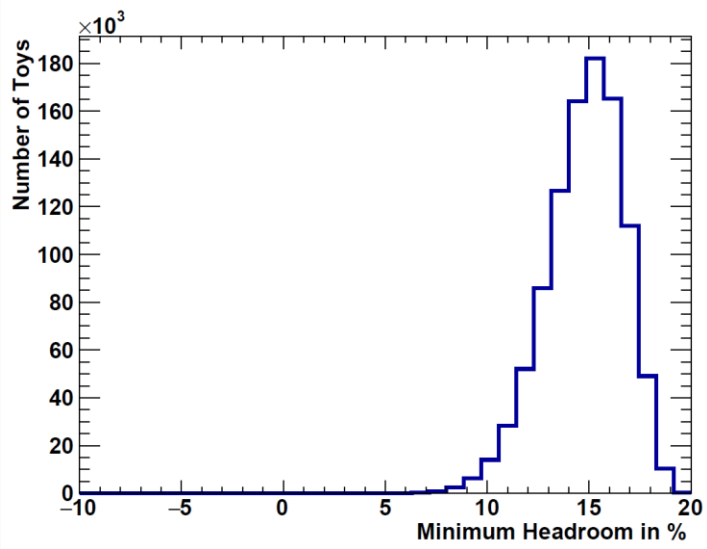

- Module Power: 8.72 +- 0.04 W (1 sigma)
- Total L0 Power: 3456 +- 0.76 W (1 sigma)

20% SCENARIO - L1

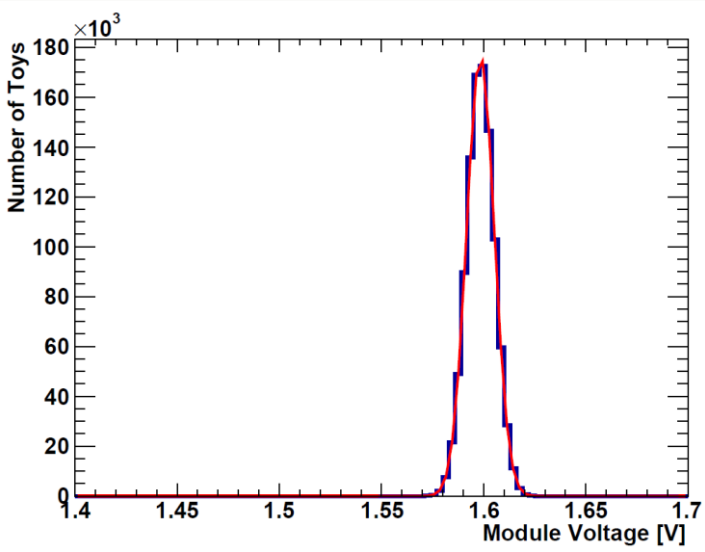

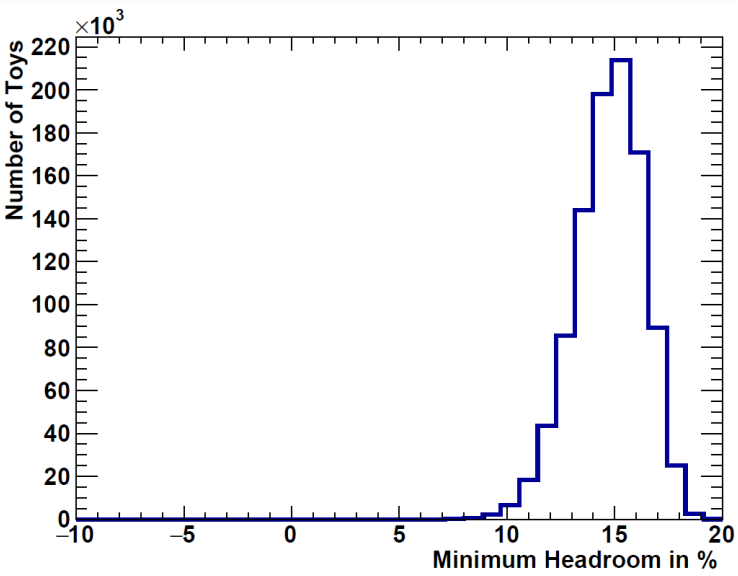

- Module Power: 10.35 +- 0.04 W (1 sigma)
- Total L1 Power: 12005 +- 1.5 W (1 sigma)

20% SCENARIO - L2

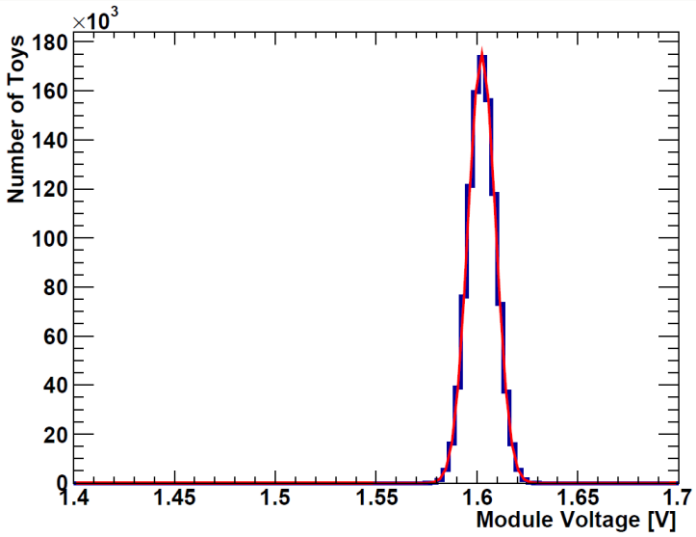

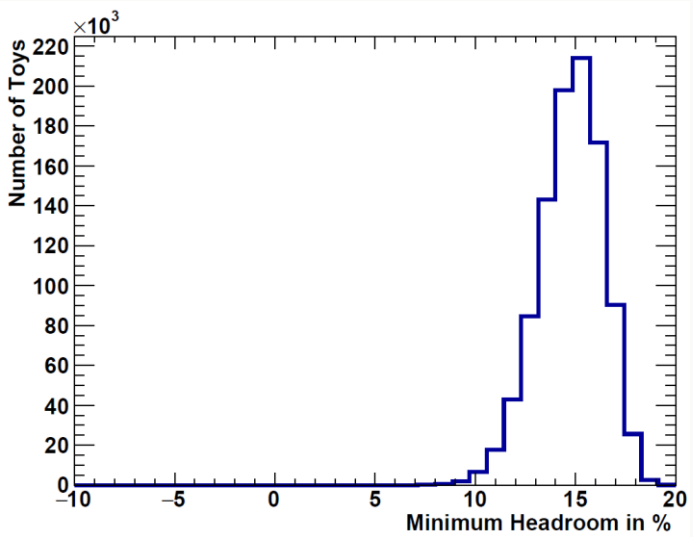

- Module Power: 9.16 +- 0.04 W (1 sigma)
- Total OS Power: 62434 +- 3.1 W (1 sigma – assuming all layers running with L2 settings)

PROPOSAL FOR NOMINAL OFFSETS AND SLOPES - 10% SCENARIO

- lower offsets and higher slopes are favoured as long as a the actual offset voltages differ from chip to chip (and the spread in the actual voltages is truly relative...)
- for the power analysis, I would assume the following settings:
  - L0: Vofs = 1.1 V, Analog Slope = 0.45 Ohms
  - L1: Vofs = 0.9 V, Analog Slope = 0.85 Ohms
  - L2: Vofs = 1.0 V, Analog Slope = 0.95 Ohms

One Chip

10% headroom

	Required Current [A]			Required Current [A]			Proposed Configuration			Nominal				1 FE Open				Average Power Density [W/cm2]					
	Analog	Digital	Total	Matrix	Periphery	Total	Vofs [V]	SlopeA [Ω]	SlopeD [Ω]	Current [A]	Voltage [V]	Power [W]	P Density [W/cm2]	Current [A]	Voltage [V]	Power [W]	P Density [W/cm2]	Nominal		No Config		One FE Open		One FE Open - No Config	
																		Matrix	Periphery	Matrix	Periphery	Matrix	Periphery	Matrix	Periphery
L0	0.75	0.77	1.52	1.14	0.38	1.52	1.10	0.45	0.44	1.672	1.471	2.460	0.581	2.508	1.657	4.155	0.982	0.354	3.018	0.000	6.799	0.354	7.704	0.000	11.485
L1	0.648	0.7	1.348	0.998	0.35	1.348	0.90	0.85	0.79	1.483	1.506	2.233	0.528	1.977	1.708	3.377	0.798	0.310	2.862	0.000	6.172	0.310	6.022	0.000	9.333
L2	0.54	0.65	1.19	0.85	0.34	1.19	1.00	0.95	0.79	1.309	1.564	2.048	0.484	1.745	1.752	3.059	0.723	0.264	2.840	0.000	5.660	0.264	5.634	0.000	8.454

→ distributions just more of the same – won't go through them here

LOW POWER MODE ANALYSIS FOR SCNEARIO 2

- assuming a *required* current in the LP mode of

  I_analog = 0.15 A

  I_digital = 0.30 A

→ scanned headroom-VofsExtra parameter space to identify LP settings (required 1% leftover local overhead): below are the yield plots in LP mode

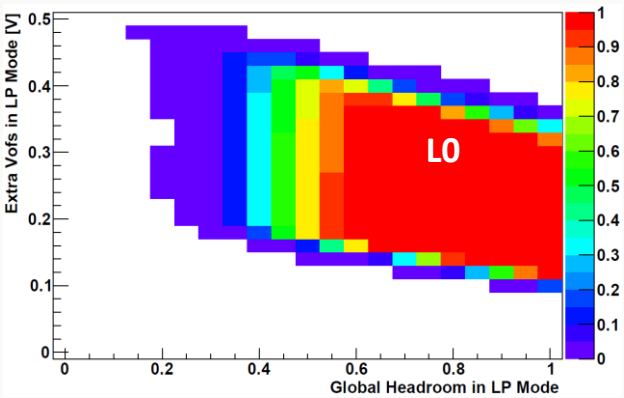

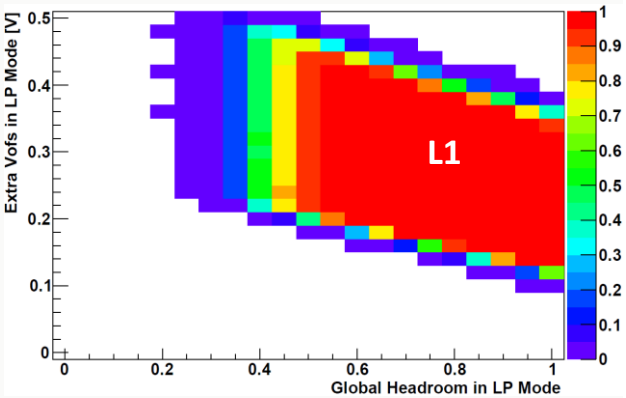

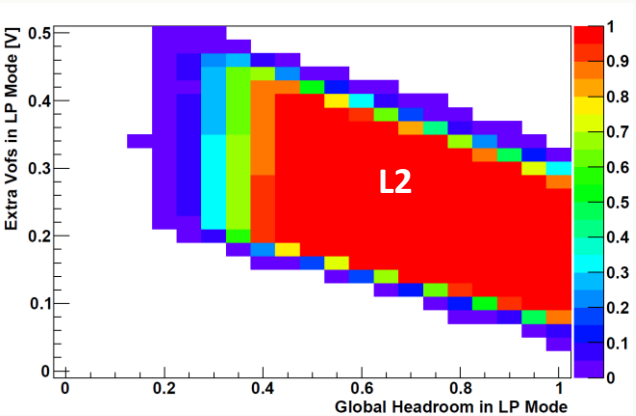

LOW POWER MODE ANALYSIS FOR SCNEARIO 2

- identified minimum current headroom to get at least 99% of all modules working in LP mode
- set Vofs as close to 1.32 V as necessary to get an input voltage as close to 1.6 V as possible
  → conservative estimate, if for LP mode testing a voltage < 1.6 V is okay, power could go down
  → Michael K. is checking if these offset voltage (or even higher ones) are okay (drop-out voltage in Vofs-generation circuit, overvoltages in Vofs injection?)

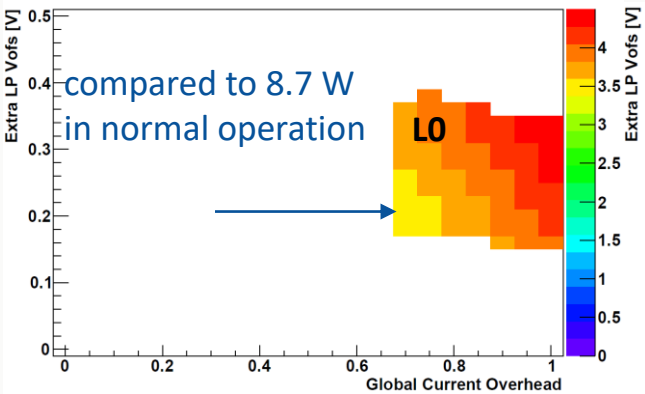

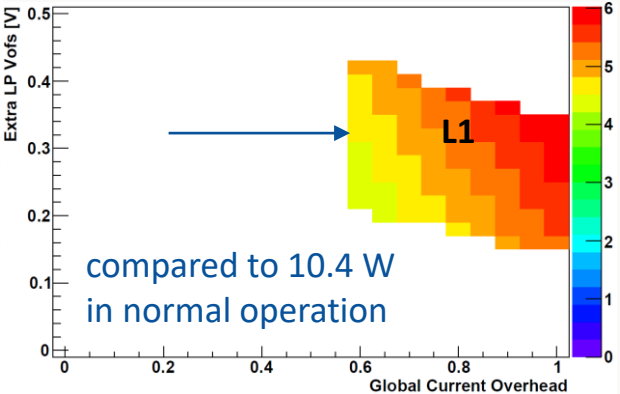

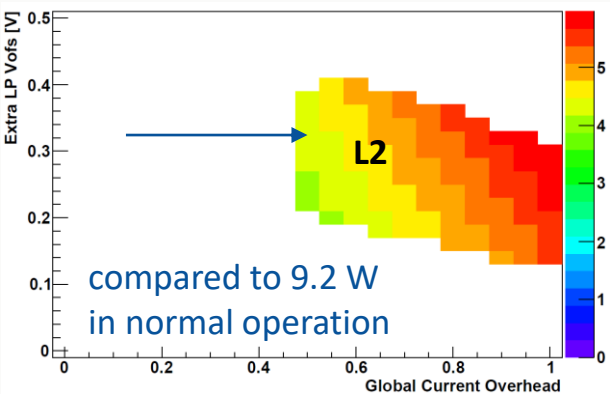

UNIVERSITÄT BONN

SUMMARY

- moved on from LSD plots to more quantitative analysis
- defined benchmark scenarios for ATLAS ITk Pixels (still need to be signed off)
  - tend to be a tad more conservative than what we may be able to achieve
- will have to update the analysis based on new numbers for the distributions of offset and slopes (waferprobing → Mark working on that)
  - slope dispersion
  - offset dispersion after tuning (Vofs vs VRext)
- analysis of high offset voltages ongoing (Michael)
- analysis of low power mode → ~40% of nominal power (allows us to check if LP mode is actually useful)
- will run final cross-check for actually available baseline values
  - need additional measurements before we can fix those values