Summary of measurements on FEAST2 modules to understand the failures observed in the CMS pixel system

F. Faccio, S. Michelis, G. Ripamonti
CERN EP-ESE-ME

October 15, 2018

1200 DCDC converters modules and based on the FEST2.1 ASIC have been installed in the CMS pixel detector system to power the front-end modules during the 2017 run. After a few months of smooth operation, some converters started to fail once the luminosity of the accelerator was increased, at the beginning of October. The failure pattern was very regular in the fact that converters were operating correctly until, in order to reset a control circuit, their output was disabled by the control system, after which the converters could not be started anymore. This has been described at different CMS and ATLAS weeks in the first 6 months of 2018, as well as at the common ACES workshop (https://indico.cern.ch/event/681247/contributions/2929069/attachments/1640160/2618524/Talk_DCDCs_ACES.pdf).

The present document summarises the main measurements performed on DCDC modules to uncover the origin of these failures, and find a way to mitigate their impact in the CMS pixel detector system. This intense study took place in the first 8 months of 2018, starting as soon as failed modules were retrieved from the CMS detector during the winter shutdown in late December 2017. It is important to understand that the FEAST2.1 ASIC had been extensively tested during the development and qualification phase, and no similar failures had ever been observed. This indicated that some of the specific conditions of use in the CMS pixel system were very relevant in triggering the failure. Unfortunately this environment is very complex, actually unique, in a good number of its characteristics, and it was therefore not straightforward to test individual modules in a similar system to observe the failure mechanism in more details. A list of the specific characteristics of the modules working in the CMS pixel system and of their environment includes:

- a module that is different than the one used for the tests done during the development and qualification of the FEAST2.1 ASIC (the CMS module was developed by RWTH in Aachen and, although similar to the FEASTMP module used at CERN, uses different components and a different board layout)
- a radiation background much larger than the one foresee for all other uses of the FEAST2.1 converter modules
- an operational temperature close to -20C
- a strong magnetic field
- an electromagnetic environment that is not well known, and that might be influenced by the proximity of the pulsed and intense LHC beam
- an electronics system around the module that is unique and whose properties (grounding, shielding) might influence the operation of the converters.

In order to observe if the combination of a few of these conditions could trigger the same failure, an irradiation run at the CERN IRRAD facility was planned at the earlier available date, in May. During this test, samples of both CMS/Aachen and FEASTMP modules could be exposed to an intense radiation field at -20C and in proximity of an intense and pulsed beam. However, while preparing a complex setup for this test run, a number of other measurements were taken while the converter was stressed in different conditions in the hope of reproducing the same damage mechanism. Tests in a 3T magnetic field in CERN Prévessin did not reveal any issue, and are therefore not discussed further here. Other stresses, instead, provoked damages that were similar - but never identical - to those that were observed in CMS.

After a first introduction of the symptoms of the CMS failures in Section 1, results from these other stress-tests will be reported. These have gradually helped to identify the damaged node, as well as the components that were damaged and the circuit's behaviour after the damage. With this knowledge, and when eventually the full failure mechanism was reproduced in IRRAD, it was possible to reach a clear understanding of the problem and propose solutions or at least applicable mitigation patches.

For those who are not familiar with the FEAST2.1 ASIC, detailed information can be found in the data sheet published in the public web page of the CERN DCDC project http://project-dcdc.web.cern.ch/project-DCDC/Default.html. The typical configuration of the ASIC in the application is illustrated in Fig.1. A node that will be often cited in this document is identified as V33Dr in this schematic: an external 200nF capacitance is required at that node, since it is internally connected to the the output of a linear regulator that regulate 3.3V to supply the buffers driving the power transistors (actually the value of 220nF is used in the final modules).

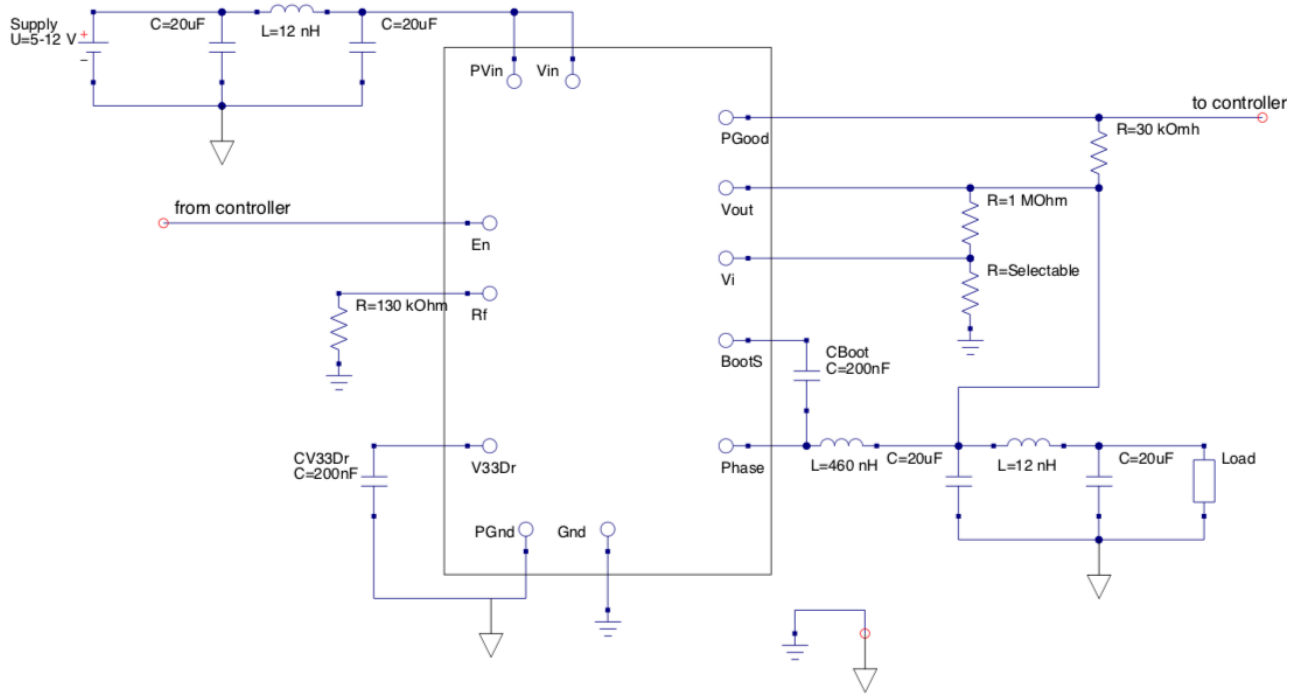

Figure 1: Typical configuration of the FEAST2 DCDC converter ASIC in the application.

1. The symptoms observed in modules from CMS

At the end of December 2017 we could measure the first modules removed from the CMS pixel detector. They were categorised and systematic measurements were done on a representative number of samples in each category. Table 1 summarises the categories and their main characteristics. It should be pointed out that no significant difference was observed between converters powering digital or analog lines (hence providing different output voltages), neither as a function of the load (unloaded modules were also failing). No correlation with the position was found either. The only real correlation was that the failures started when the luminosity of the run increased (larger radiation background, as well as change in the beam characteristics, in particular the emittance).

The “revival” procedure, after which broken samples were back to correct functionality, consisted in applying the 3.3V from an external supply with almost unlimited current (the basic OCP of the supply is several Amps). In one case we limited the current with a 100 Ohm resistor in series, and could observe an increase of the current from the external supply up to 28mA, after which the sample was revived. This suggests that when an excessive current is forced to ground from V33Dr for a sufficiently long time a parasitic current path is broken - and this same path is responsible for the malfunction of the converter. Once the path is removed, the on-chip linear regulator can again provide 3.3V to the power circuitry and the converter can resume operation. However, signs of damage remain in the form of a wrong UVLO_regs threshold and an excessive input current.

From these symptoms and an analysis of the schematics we argued that the damage occurs in the clamp transistors of the on-chip V33Dr regulator. Its schematic is shown in Fig.2. During normal operation and when the input voltage is above the UVLO_regs threshold, the two clamp transistors as well as the PMOS whose drain is connected to the Gate_PMOS node (top right green MOSFET) are open, and the feedback loop of the regulator tunes the Gate_PMOS node so

that enough current is provided by the PMOS pass transistor to the output. This situation changes if a current path opens along an hypothetically damaged clamp transistor. In Fig.2 we predict two possible parasitic paths from the drain of a clamp transistor. The black path sinks current to the gate, to the UVLO_regs node. This node is normally pulled down by a 5uA current, which is small and can be overwhelmed by the current flowing in the parasitic path. In that case, the UVLO_regs voltage rises and the Gate_PMOS node is pulled up - effectively turning off the PMOS pass transistor. In this case the V33Dr node would be stuck to a voltage well below 3.3V, in agreement with the behaviour observed in "broken" samples (note that the same thing can happen on the other clamp transistor, but a larger parasitic path is needed because the "Not_Vdd_OK" node is pulled down by an inverter).

Category	Thresholds	V33Dr	Comments
Broken	UVLO_regs < 3.2V most often	Stuck to 0.9-1.6V, with slight change above 5V (3rd threshold) most often	"Revival" possible. A few modules failing at the same time during T cycles had the same signatures.
"Revived"	UVLO_regs: No change UVLO_En: around 5V (larger than normal)	Back to about 3.3V most often (one case 2.4V)	Revived by the application of a large current to V33Dr.
High Current	UVLOs normally regular	Regular	Current below Enable 7mA to 23mA (normal is about 3mA)

Table1: Categorisation and main characteristics of the converters extracted from the CMS pixel detector in December 2017. Broken samples failed to provide the correct output voltage in the experiment, but could be "revived" during the measurements by the application of the 3.3V at the node V33Dr - they actually re-started correct functionality. UVLO stands for Under-Voltage Lock-Out, a function that prevents the converter from operating when the input voltage is below a pre-defined threshold. The circuit has two such levels: UVLO_regs is the threshold above which the on-chip regulators turn on and provide 3.3V to the analog, digital and power circuitry, whiled above the higher UVLO_En threshold the converter is allowed to start operation if enabled by the appropriate control pin.

The high current condition observed in other samples, instead, can be explained by the opening of a current path from V33Dr to ground. This path does not interfere with the feedback mechanism of the regulator, which can still operate correctly. To verify the possibility of this type of damage in transistors in this technology, we applied high voltage spikes to the drain of individual MOSFETs available in separate test chips. We indeed observed a drain-current path typically followed by the appearance of a gate current when further stress was applied.

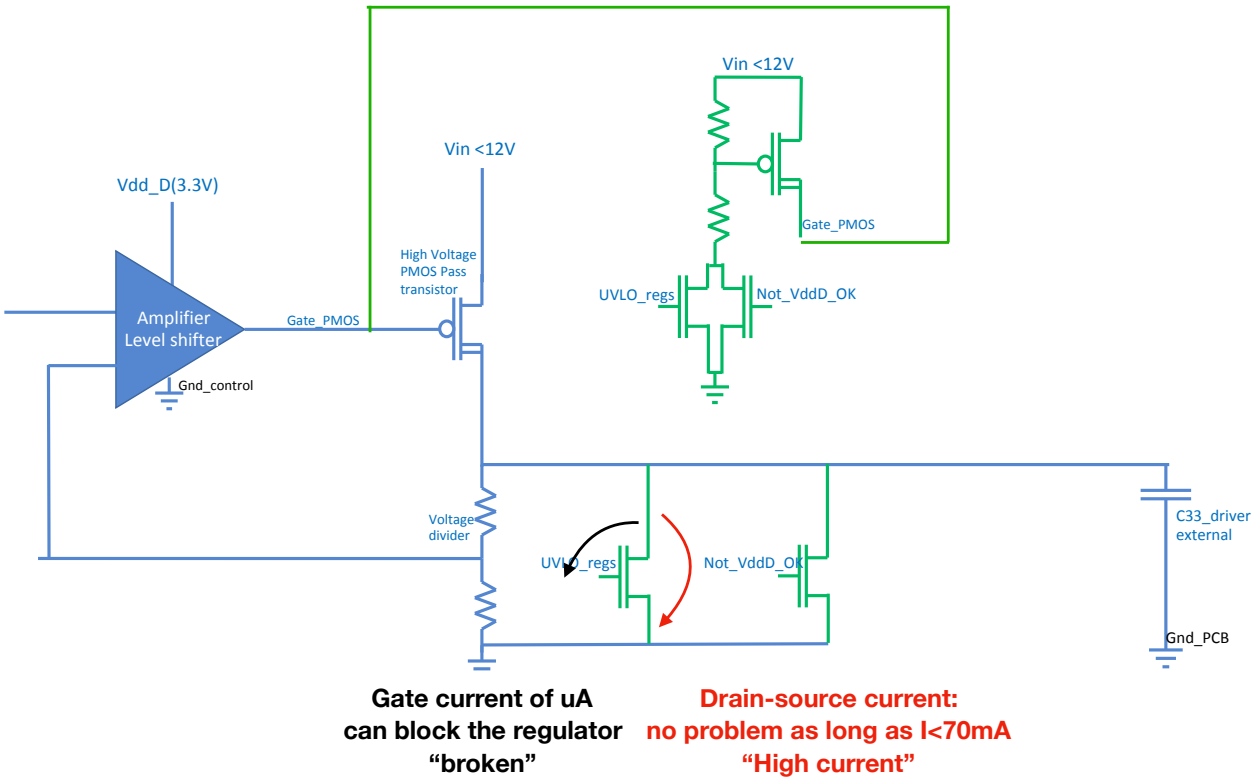

Figure 2: Simplified schematic of the on-chip regulator supplying the drivers of the power transistors. The C33_driver capacitance is the 200nF element connected to node V33Dr in Fig.1. Two possible parasitic current paths, opened by a damage mechanism, can be responsible for the symptoms observed in samples retrieved from the CMS pixel detector.

2. Is there a problem with the external C33_driver 220nF capacitor?

The first attempt at damaging the samples was done by acting on the 220nF SMD capacitor connected to the V33Dr node. In the CMS modules this is a tiny 0201 SMD element, and it seemed obvious to start from this component because it is directly connected to the node where the damage appears. When the capacitor is fully removed, the V33Dr node voltage is actually widely moving at every clock cycle, reaching peaks at 8V. This happens because of charge sharing with the bootstrap capacitance. However the converter manages to operate correctly and continues to be functional even after long periods of time. Intermittently connecting and disconnecting the capacitance produces a first burst of peaks that are larger in voltage, and in that case it was possible to damage the converter. The produced damage was similar (although not fully identical) to the one observed in samples from CMS. The capacitance, as well as its soldering to the PCB, were examined in different manners without revealing any obvious problem. So the hypothesis where a defect in the capacitance, or its assembly on the PCB, is responsible for the damage was unconvincing.

3. ESD injection and RF tests

Another attempt at reproducing the damage in fresh samples was via Injection of electrostatic discharges (typically we used sample modules discarded from our production lots). This was done using an ESD gun kindly provided by F. Szoncso (HSE-DI) that could produce 1.2kV voltage peaks for the discharge, either manually (individual spike) or in bursts at 5kHz. An image of the tip of the gun during the test of a FEASTMP module is shown in Fig. 3. This stress managed in some conditions to provoke damages similar to those observed in samples from CMS.

The following is a list of the main characteristics of the damage:

1. To produce the damage, it was necessary to avoid the contact of the tip of the ESD gun with the target: the discharge is faster and more severe if a spark develops and only in this condition a damage could be observed
2. The damage could be produced for discharges at the Rf, Enable and PowerGood pins. Probably all signal pins are sensitive, however we concentrated our tests on the Enable and PowerGood only, since these are connected to long lines in the application. Power pins were not sensitive, however
3. Modules unpowered or with Vin below the UVLO_regs threshold were never damaged irrespective of the number of discharges and the location of the discharge
4. Modules installed on the standard test bench (as in Fig.3) required more discharges to be damaged with respect to modules installed on a small board like those used for irradiation, where the enable is a resistive bridge from Vin:
    a. When installed on the standard test bench, modules would be damaged in High Current mode after 2-20 single discharges. With more discharges some modules could fail in the broken state
    b. When modules were installed on the small irradiation board, damage was observed during the first few discharges, and the damage could be either the High Current or the Broken mode
5. High-current modules had a different signature from those in CMS: their excess current was the same when the DCDC was enable or disabled (this was not the case in the CMS and IRRAD modules).

Some waveforms taken with the oscilloscope on the Enable and V33Dr nodes during the discharge (on the enable) are illustrated in Fig.4.

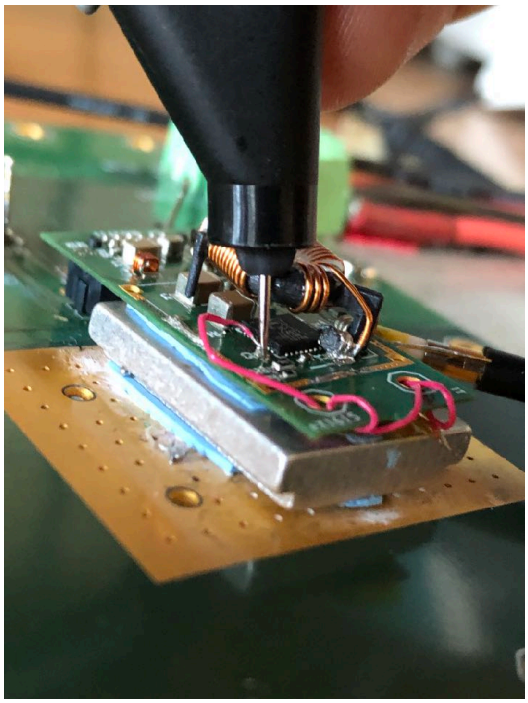

Figure 3: Tip of the ESD gun during injection tests on a FEASTMP module.

Other tests were performed by injecting disturbances in different ways on modules connected to mockups of the CMS system (at CERN building 186, and in a shielded cage in the ESD lab). These revealed that the chip is rather robust to electromagnetic noise injection, but when the power of the disturbance is very large, damage is eventually produced. During the tests in building 186, input and output cables as well as control lines were surrounded by a capacitive device and disturbances were injected via capacitive coupling (up to 5kV signals). Damage was never induced. On the contrary, some (2) samples were damaged in the shielded cage when fast transient pulses were produced by a powerful antenna. The damage signature, however, was different than the one observed in samples damaged in CMS. All these tests were done with the assistance of F. Szoncso (HSE-DI) and D. Valuch (BE-RF-FB).

The hypothesis of noise pick-up from the long communication lines (Enable or PowerGood) was only discarded by the tests performed on a large number of converters in the IRRAD irradiation facility that will be described later.

Figure 4: Screenshots of the oscilloscope used to observe the Enable and V33Dr nodes during ESD discharge tests. The yellow trace (with red dots) is the enable, the purple is V33Dr. The negative peak of the enable is 114V in the top and 122 Vin the bottom image. V33Dr can raise to almost 9V (bottom image), making a damage on that node plausible.

4. Different damage in samples from CMS and from ESD discharge tests

If the damage is indeed in the clamp NMOS transistors (Fig.2), it should then be possible to observe a significant modification in the static characteristics of the transistor. Given the configuration of the circuit, it is possible with a simple measurement to observe the output characteristics ($I_{ds}=f(V_{ds})$) of the clamp NMOS, and compare it in fresh and damaged samples. When the input voltage is below the UVLOS_regs threshold (about 4.2V), the left NMOS clamp in Fig.2 is turned on and keeps the V33Dr node shorted to ground. At the same time the right clamp transistor is still off, since the inverter that drives its gate is not powered yet (it is powered by the internal VddD bus, which is still off in these conditions). Sweeping the V33Dr voltage with an external generator and measuring the resulting current, it is possible to draw the curve for the output characteristics of the left NMOS clamp. In fresh samples, this produces a curve that is comparable to the result obtained with transistor-level SPICE simulations, where the maximum current at a V_{ds} of 3.3V is about 1.7mA.

For damaged samples, this measurement is performed only on High Current circuits because in the case of the Broken samples the application of a voltage on V33Dr induces a large current that breaks the parasitic path and leads to the functionality recovery.

The comparison of the results from samples damaged in the CMS experiment and those damaged in ESD discharge test leads to the conclusion that there is a notable difference between the two cases. Fig.5 reports the result for a typical circuit that exhibited the High Current behaviour after an ESD discharge (spark close to the enable line). The pre-ESD curve is the expected $I_{ds}=f(V_{ds})$ for a fresh sample, with a peak current of about 1.7mA. After the ESD discharge that has introduced the High Current behaviour, the blue curve shows instead an almost linear (Ohmic) behaviour suggesting that a resistive path from V33Dr to ground has been opened. The equivalent measurement for a High Current sample retrieved from CMS in Fig.6 shows instead that the linear behaviour starts at about 0.7V - below this voltage the current superposes well to the output characteristics of a fresh NMOS. In the two figures it is also possible to see how the $I_{ds}=f(V_{ds})$ curve evolves with successive “revival” cycles. This operation is typically done by forcing an external voltage to the V33Dr node. In High Current samples, this voltage needs to be superior to 3.3V, and the operation does not systematically work (some samples are not recovering). However, when samples recover it is possible to notice a remarkable difference between devices damaged with ESD or in CMS. This is summarised in the table 2 below, that lists the sequence of operations.

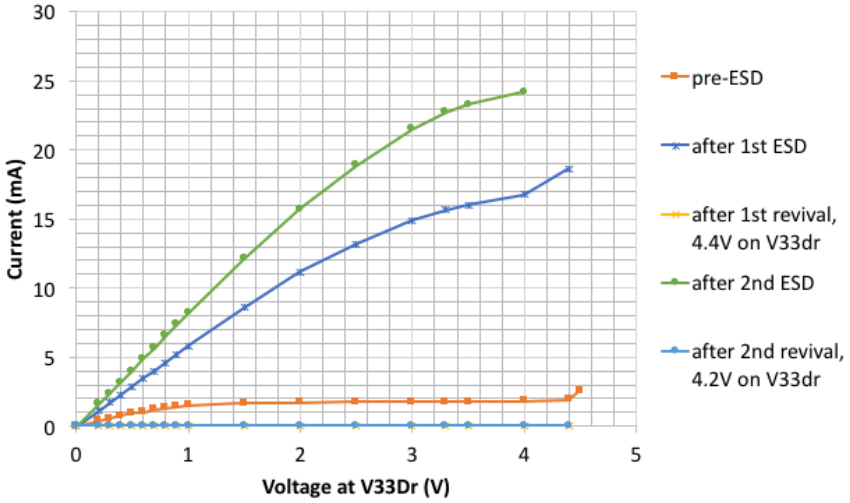

Figure 5: Result of the measurement of the output characteristics of the clamp NMOS transistors in a sample damaged in the High Current mode with ESD discharges. The yellow curve, (after 1st revival) is perfectly superposed to the orange (pre-ESD).

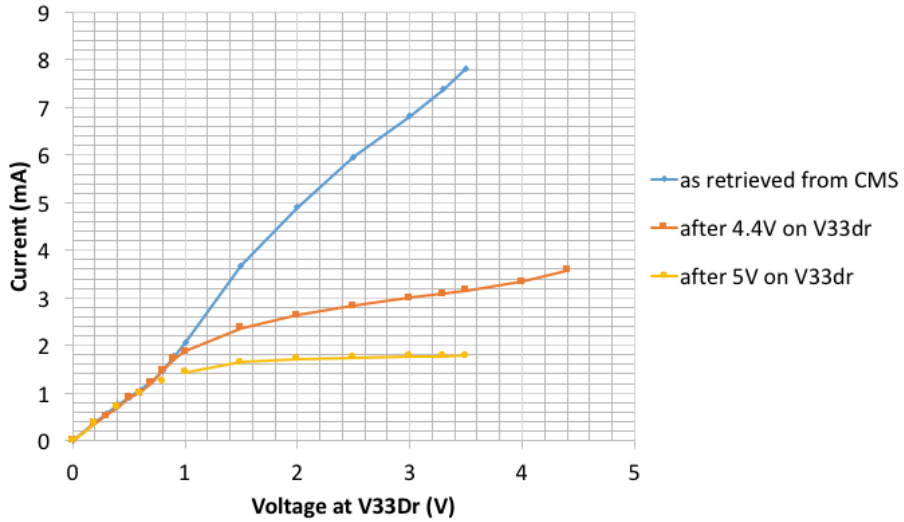

Figure 6: Result of the measurement of the output characteristics of the clamp NMOS transistors in a sample damaged in the High Current mode in CMS.

	High-current from ESD discharge	High-current from CMS
1. Measurement of IdVd of damaged samples	IdVd shows a resistive behaviour.	IdVd follows the typical pre-rad curve (transistor's output conductance) until about 0.7V, then it becomes resistive.
2: Revival by applying a large voltage to V33Dr	IdVd is flat, there is no current flowing.	IdVd comes back to the pre-rad curve (transistor's output conductance).
3: Second ESD gun discharge	IdVd shows again a resistive behaviour.	
4: revival by applying a large voltage to V33Dr	IdVd is flat, there is no current flowing.	
5: More ESD gun discharges	Regardless the number of discharges, the sample is not damaged anymore.	

Table 2: Result of the measurement of the $I_{ds}=f(V_{ds})$ characteristic of the NMOS clamp transistors for samples damaged with ESD discharges or in the CMS experiment (High Current samples only).

In the samples damaged with the ESD discharge, the result clearly indicates that damage occurs in the clamp NMOS transistors. After the first discharge, the left clamp in Fig.2 is broken and presents an ohmic path to ground. A successful "revival" operation breaks this path, and as a result no current can be measured anymore in the $I_{ds}=f(V_{ds})$ curve (the other transistor is turned off in the measurement's conditions). If a second ESD discharges leads to a High Current state again, then we can observe once more a linear current behaviour indicating that the second clamp (right transistor) has been damaged. A second successful "revival" breaks this second path - but now there is no clamp transistor anymore, hence the sample can not be damaged by other discharges. In the samples damaged in CMS, instead, it looks like the additional "high" current flows in a diode - below about 0.7V the curve follows the pre-damage characteristics. Successive "revivals" gradually increase the resistance of the additional path (orange line in Fig.6), then fully break it (yellow line) and the output characteristics of the NMOS clamp is again as in fresh samples - so this device has never been damaged. As stated above, this indicates that the High Current damage mechanism is different in the two cases.

There is yet another distinctive feature differentiating the High Current samples from the ESD discharge and from CMS. In the former case, the excessive current is independent on the functionality of the converter: if the current increases by 7mA in the disabled state, it also increases by 7mA when the converter is enabled. In the samples from CMS, instead, the increase in the enabled state is considerably smaller (often limited to only 1-2mA). This characteristic will be reproduced only in samples damaged during irradiation experiments, as it will be pointed out later.

5. Where is the additional current flowing in High Current samples?

The measurements in section 4 have ascertained that the clamp is still intact in High Current samples from CMS, where the current increases in the output characteristics after a voltage corresponding to the typical turn-on of a diode (0.7V). To understand where this current is actually flowing we have performed a Failure Analysis experiment at a specialised company (Maser, the Netherlands). With the techniques used during the analysis, it was possible to gather some evidence that the current flows in the bootstrap diode. This is a large element used in the bootstrap circuitry that allows to correctly drive the gate of the High-Side power transistor.

Fig.7 compares an emission scan image from a fresh (left) and a High Current sample from CMS (right). The emission signal is much larger in the bootstrap diode, the large slot on the top-right of the images, for the CMS converter. The images are taken with the converters disabled, where the

High Current (HC) sample has a considerably larger consumption: 14mA versus 3.58 in the fresh device.

Unfortunately the techniques used for the analysis do not allow for resolving very small currents, so it was not possible to see where the current flows after the bootstrap diode: either directly to ground somewhere or via a distributed path. However, a current path via the bootstrap diode could explain why the current increase is larger when the converter is disable than when it is enabled. This current might depend on the state of logic gates after the diode, steadily flowing when the converter is disabled and flowing only during a part of the clock cycle otherwise.

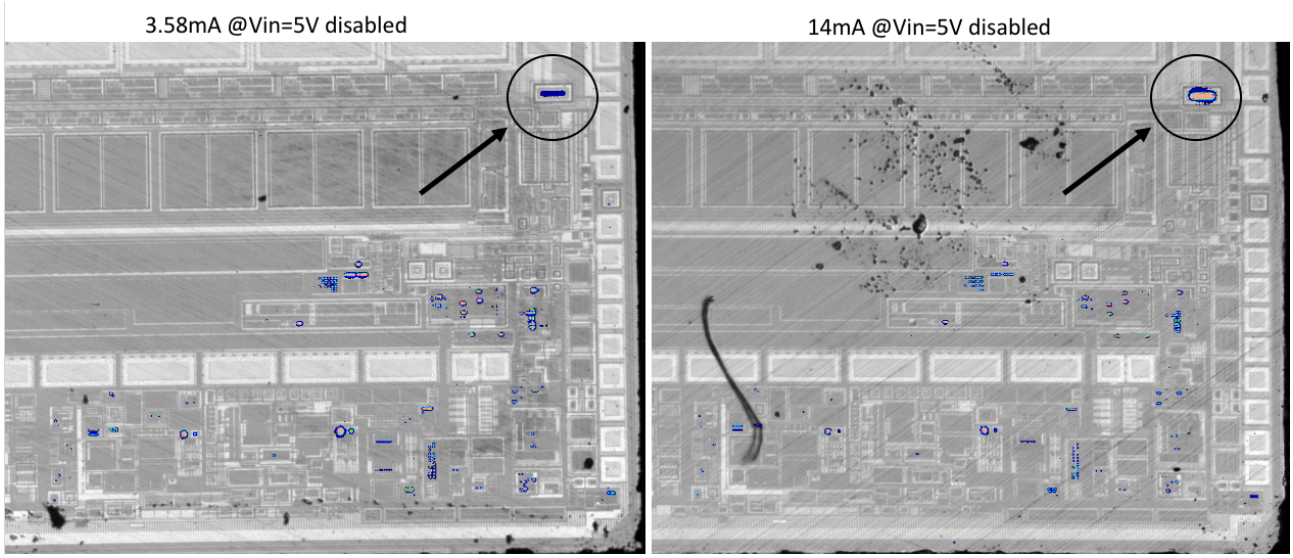

Figure 7: Emission scan images for a fresh (left) and a High Current sample from CMS (right). The coloured areas indicate locations with relatively large current flow. To the top right, the bootstrap diode drives much more current in the damaged sample: the orange colour stands for a larger current density.

6. … and where is the additional current flowing in Broken samples?

The same Failure Analysis was also carried on Broken samples at Maser. In these samples the current appears to be flowing in the clamp NMOS transistors as well as in an inverter that is connected to them. An example image is shown in Fig.8, where the red dot represents a current flowing in one of the clamp NMOS transistors.

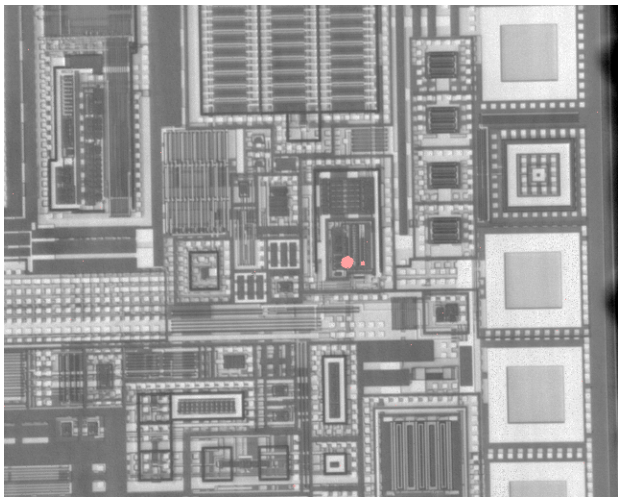

Figure 8: Emission scan image for a Broken sample from CMS: the red dot indicates that a relatively large current flows in one of the two clamp transistors.

It is extremely interesting to look at the emission scan images taken during the “revival” operation of a Broken sample. In this controlled experiment the voltage on V33Dr is limited by a resistance inserted in series to the voltage supply, and in this way it is possible to measure the current flowing from the supply to the regulator (via the parasitic path) in time. At the same time, the emission scan is performed and the sequence of the images produced is shown in Fig.9. From top left to bottom right the voltage on V33Dr is increased; the value of the voltage and the resulting current are both reported on each image. As the voltage and current increase, a large spot is observed in the images in a location corresponding to the clamp transistor. When the current reaches about 18-20mA, and after some time, the current suddenly drops to only a few uA and the current spot disappears from the image. At the same time, the sample came back to full functionality (“revival”). This clearly shows that a current path has been removed by the large current forced across V33Dr and via the clamp transistor.

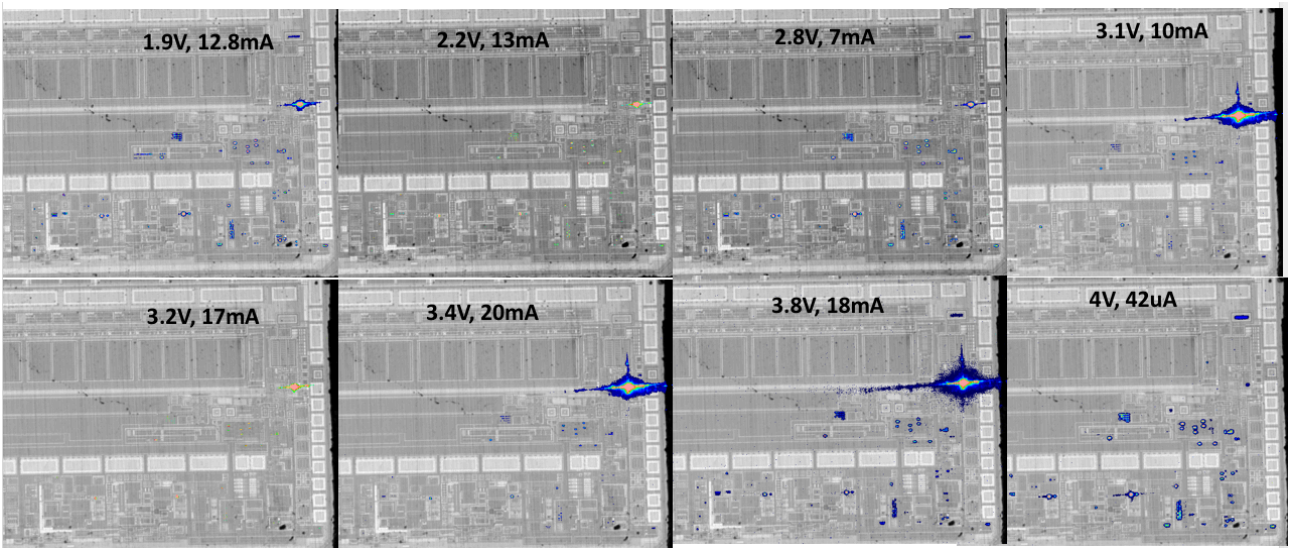

Figure 9: Sequence of emission scan images taken while increasing the voltage applied to V33Dr on a Broken sample (“revival” procedure, in this case with a resistance in series to limit the current). After a current of about 20mA is reached for some time, the parasitic current path is broken and the converter is back to functionality.

This evidence confirms the initial suspicion that the damage of one of the clamp transistors is responsible for the failure of the converters in CMS: when the transistor is damaged, the voltage at either the UVLO_regs or Vdd_OK in Fig.2 rises, which prevents the correct functionality of the PMOS pass transistor. As a consequence, the V33Dr regulator does not work anymore and there is not enough voltage for the power train transistors to turn-on.

7. A first irradiation run at the IRRAD facility brings a breakthrough

Up until this point (early May 2018) it was impossible to perfectly reproduce the failure observed in the CMS system. This situation changed with the first test in the IRRAD facility, where a large number of converters was exposed to a mixed radiation background at cryogenic temperature. One of the reasons to perform this test in IRRAD is the availability of a cold box where samples could be kept at -25°C during the irradiation, as well as the mixed field present during the test. These are two characteristics that reproduce the conditions of the samples in the CMS pixel detector. To produce a mixed radiation field, a 1cm Cu target was added on the path of the beam in zone 2 of IRRAD, and the samples in the cold box were exposed in the field dominated by the products from the interaction of the proton beam with the copper.

Due to the design of the test boards, that were developed in view of a further test on the proton beam line of PSI, and to the geometry of the radiation field, there was a large difference in the field intensity where the different samples were exposed. Fig.10 illustrates the position of the 4 motherboards, each containing 4 FEAST modules, in the cold box itself positioned in the simulated radiation field in IRRAD. Each motherboard contained 8 modules: 4 modules below the black line indicating the mid-point of the boards, and 4 considerably above it.

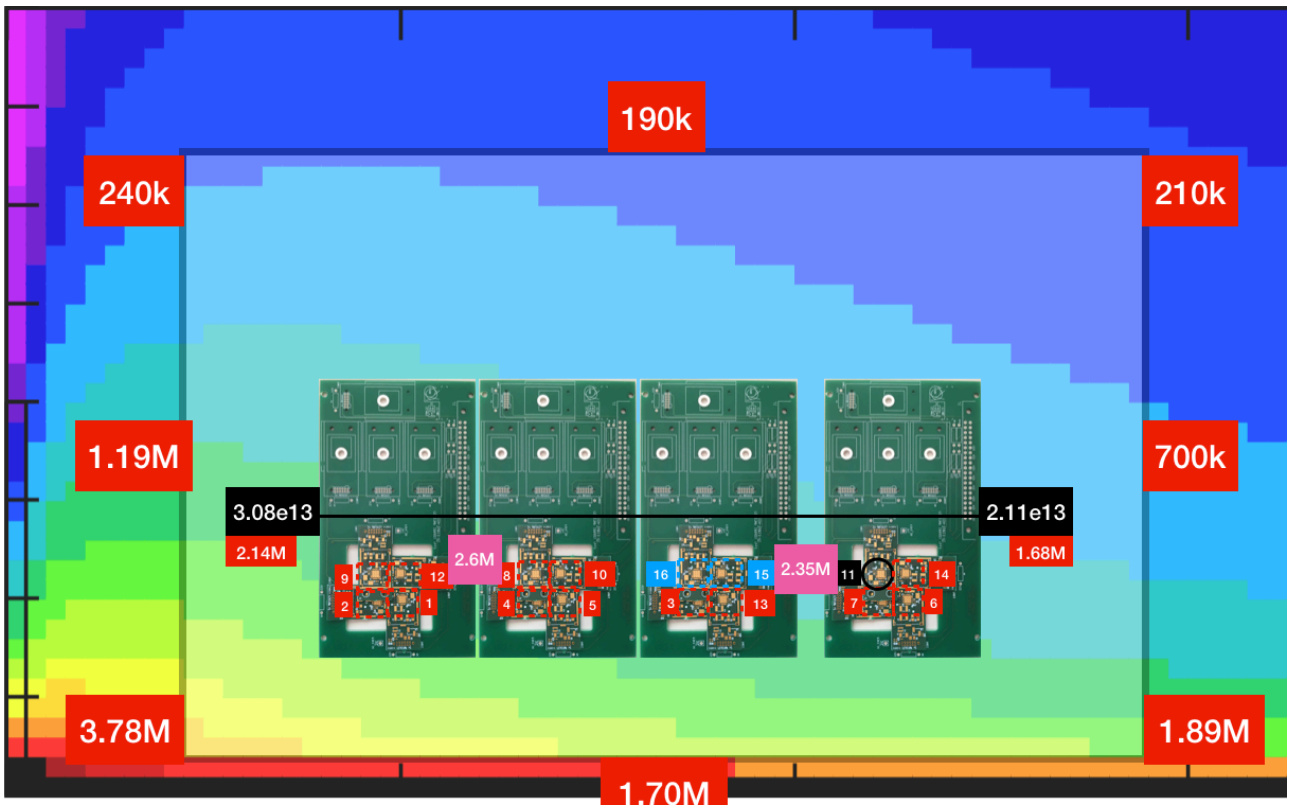

Figure 10: Position of the 4 motherboards with 8 modules each in the simulated radiation background created in the IRRAD facility by the insertion of a 1cm Cu target along the proton beam. Numbers in red/black/blue squares close to each module represent the chronological failure order. The other numbers in larger red/purple squares report the dose in rad registered by passive dosimeters in these positions, while in the black squares the 1MeV n-equivalent fluence from active dosimeters (diodes). All dosimeter readings refer to the integrated doses/fluences for the full exposure of about 20 days.

The test setup was designed to keep the converters running most of the time, while periodically performing an enable-disable operation and an I-V curve to record the currents below UVLO thresholds - the system to discover samples damaged in the High Current state. Software was developed to control the instruments (in Labview) and to observe the functionality of the converters from the produced files (in python). We exposed a total of 16 FEASTMP modules and 16 CMS modules, all with only a 100Ohm load, at a constant temperature of -25°C. During the test a cycle of about 2 hours was repeated continuously. Within each cycle each converter was disabled-enabled twice (once every hour) and the I-V was measured once at the beginning of the cycle. During the rest of the cycle, about 97% of it, the converters were running with an input voltage set at 12V from the power supply (but the effective voltage at the input of the module dropped to about 11V during operation due to the long cables).

Some damages started to be observed after about 10 days, from which point there was a continuous rate of damaged converters - almost all of the High Current type (only one sample was actually Broken). We found a very good correlation with the geographical position of the converter, which means that the radiation background was instrumental to provoke the damage. The other very important result was that all failures happened during the disable-enable cycle: no converter seemed to have problems while steadily running at 11V, although this condition was kept for 97% of the time. The map of the damages is shown in Fig.11, where the number indicates the chronological order of the damage detection. In red are the high-current devices, in black the broken sample, and in blue the samples that went to high-current after the end of the irradiation while the samples were still kept at -25°C. These 2 blue damages happened in the 2-4 days following the end of the exposure (the Cu target was removed). No significant difference was found between modules FEASTMP and from CMS.

Another important information from the test is that, similar to damaged (High Current) samples from CMS, the excess current in the samples damaged in IRRAD is different when the converter is

enabled or disabled. As we have seen in sections 4 and 5 above, this is an indication of the fact that the clamp transistors are still intact - which was not the case in high-current devices from ESD stresses. All other measured characteristics in the damaged samples from IRRAD matched those observed in similarly damaged modules retrieved from the CMS detector.

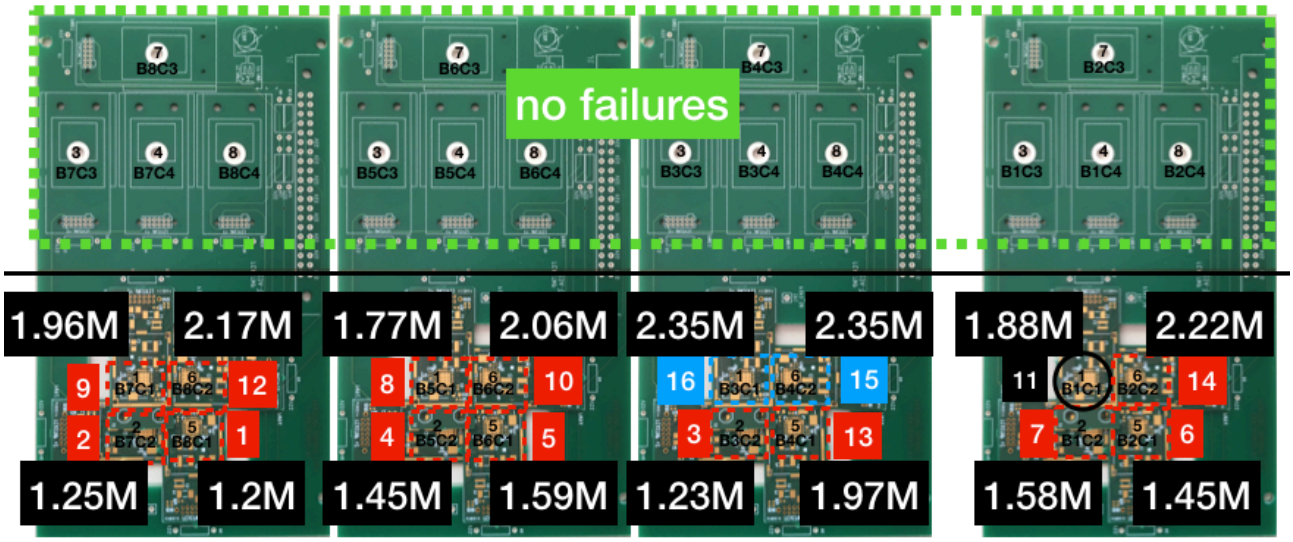

Figure 11: Detail of the modules damaged during the first irradiation run in IRRAD. All samples highlighted in red went to the High Current state at a dose whose best estimate is reported (in Mrad) in the black squares. The sample circled in black was the only one actually failing (Broken state), while those highlighted in blue went to High Current shortly after the end of the irradiation, while running at the same -25°C temperature. None of the 16 samples above the mid-line of the boards and situated farther from the beam (smaller TID) had any observable damage. Note that these 16 samples are not shown, but are represented by the 16 numbered white holes on the top of the boards.

This irradiation run represents a breakthrough because for the first time the same damage observed in CMS samples could be reproduced outside of the specific environment, where environment refers to the combination of radiation, temperature, electromagnetic background, functionality cycle, and electrical conditions (grounding, shielding, etc.). Because most of this environment differs in the IRRAD test and in the CMS pixel system, the results powerfully indicate that the damage requires “only” the combination of 2 elements:

1. a “sufficient” radiation background (dose or fluence)
2. a disable-enable sequence has to be applied.

However it is not yet clear if TID or fluence (Non-Ionising Energy Loss, NIEL) is to be blamed, but this can be easily found by separately exposing modules to either TID or non ionising particles.

8. X-ray irradiation tests eventually unveil the mystery

Given the easy availability of an X-ray source at CERN, it seemed natural to start by exposing samples to TID alone. This had of course been done before. During the development of the FEAST ASIC family, a large number of samples had been exposed to TID levels of up to 700Mrad. In that case, however, samples were not disabled and enabled often during the test. In November 2017, when failing modules were reported from CMS, we tried to reproduce the failure by exposing 2 samples in the X-ray machine at -30°C. This test was run up to a TID of 1.5Mrad and this time the converters were periodically disabled and enabled regularly. Since at that time the existence of the High Current damage was ignored, the test was tailored to reveal functional failure only and not to measure the current consumption below the UVLO_regs threshold. However, at this point we recovered the two samples exposed in November and discovered that one of them presented the High Current damage, which was already a promising indication for a TID- induced problem.

A new irradiation test was hence started, where the DUT was disabled-enabled every 100krad and the I-V curve was measured every 500krad to unveil a possible High Current type of damage. It turned out that this was a perfect way to damage the samples: all exposed modules were effectively systematically damaged. Additionally the damage always occurred during the disable-enable sequence, a characteristic identical to the typical signature in both CMS and in the IRRAD test. The list of samples damaged during this X-ray study included 3 High Current and 1 Broken module, all showing the first evidence of the damage at a TID of about 0.9-1Mrad. Two other modules were irradiated with a different sequence (no disable-enable at large input voltage) and did not fail until the sequence was changed to the usual one - at which point they both went to the High Current state.

Having demonstrated that the combination of TID and a disable-enable sequence was the culprit, it was natural to wonder if the same threat also affected the bPOL12V ASIC that is being developed in view of tracker application at the HL-LHC. This circuit largely uses the same architecture and detailed implementation in a different CMOS technology (but very similar to the one used for the FEAST family). The damage also appeared in bPOL12V samples, but with several important differences:

- the failure appears at larger TID levels in bPOL12V (2.2-2.7Mrad, versus 0.9-1Mrad in FEAST2.1)
- although the statistics is low, it seems that the bPOL12V is not sensitive to the High Current damage mode: all 3 damaged samples directly failed to provide output voltage (Broken state)
- contrary to the FEAST2.1 samples, it was not possible to “revive” the broken bPOL12V with the application of an external V33Dr voltage.

To understand the mechanism leading to the damage, we observed the voltage of the V33Dr node with the probe of an oscilloscope while the sample was exposed. Fig.12 shows the result for a bPOL12V sample where successive voltage curves taken at increasing levels of TID are superposed (results were even more spectacular for the FEAST2.1, where the voltage peak reached 8.5V already before a TID of 1Mrad). These curves all correspond to the time variation of the V33Dr voltage immediately after a disable command is sent to the module. As soon as the sample is disabled, the voltage rises and reaches a TID-dependent peak increasing with the dose. The duration of this pulse also increases with TID. This voltage clearly exceeds the 3.3V rating of the core transistors in the used technology, so it is not surprising that a device connected to the V33Dr node breaks (such as the clamp NMOS transistors).

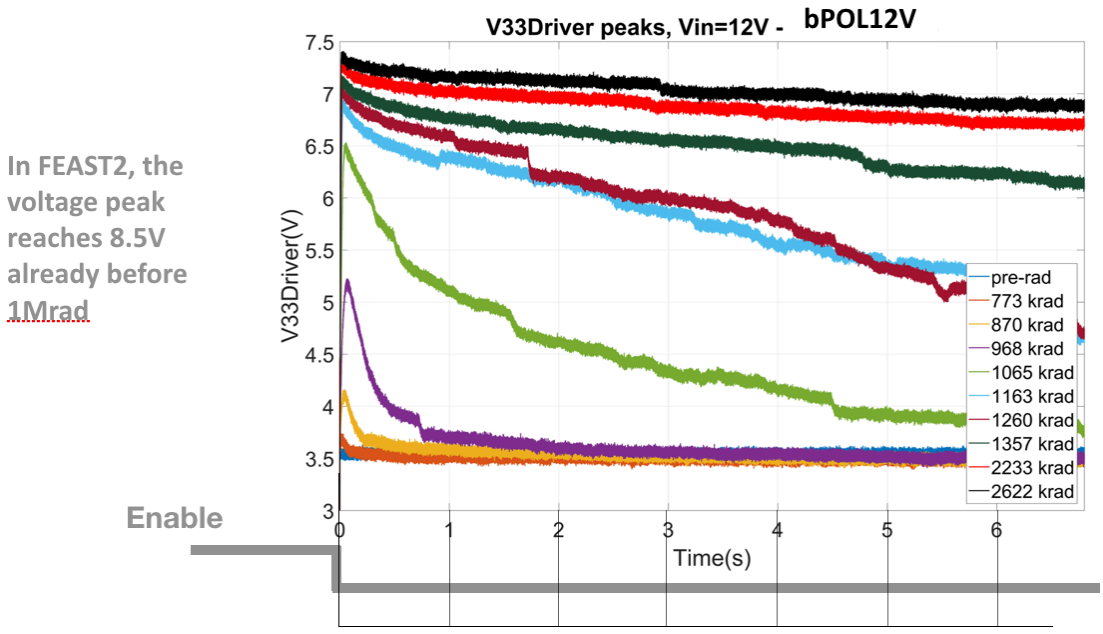

Figure 12: Voltage on the V33Dr node observed with an oscilloscope as soon as the sample of bPOL12V is disabled at increasing levels of TID. Results were even more dramatic for FEAST2.1 samples (larger voltage peaks at lower TID levels). All these measurements were done with a large dose rate of about 4.5Mrad/hour.

Very interestingly the TID level when this excess voltage appears on the V33Dr node coincides with the one where we know leakage currents reach their peak value in fast irradiations at room temperature. Looking for a “leaking” path in the circuit rapidly led to the discovery of the culprit: an LNDMOS transistor in the voltage shifter of the V33Dr regulator. This is shown in Fig.13. With respect to the representation in Fig.2, this schematic also shows the voltage shifter at the output of the amplifier in the feedback. Precisely in this shifter, the LNDMOS transistor circled in red is known to be sensitive to TID-induced leakage current. When the regulator is providing current to its output (to the buffers of the power transistors during operation) the LNDMOS is turned on and the leakage is adding up to the main current. This mainly alters the Vgs necessary for that transistor. However, when the converter is disabled the pass PMOS transistor should only provide the 5uA needed by the 2 resistors in the voltage divider (600kOhm overall). If the leakage in the LNDMOS is larger than some tens of nA, the mirrored current in the pas PMOS exceeds the 5uA and the current can only flow in the C33_driver external capacitance, where it is integrated as a voltage. This is illustrated in Fig.14, where the waveform of the V33Dr (blue) and of the bootstrap (red) nodes during the rise of the voltage is shown. At first only the V33Dr capacitance gets the current, and the rising slope is larger. When the bootstrap diode can flow current, the 2 capacitors are in parallel and share the charge, then they both charge up. From the slope of the curve, we estimate a current of 2-300uA from the pass PMOS transistor. This is a reasonable value: it requires several 100nA leakage from the LNDMOS.

It should be noted that 2 core ELT transistors are in series to the leaking LNDMOS. This type of structure is used purposely to interrupt the possible leakage path of the LNDMOS, but in this case it is useless because the 2 transistors are closed above the UVLO thresholds - so they are turned on during disable.

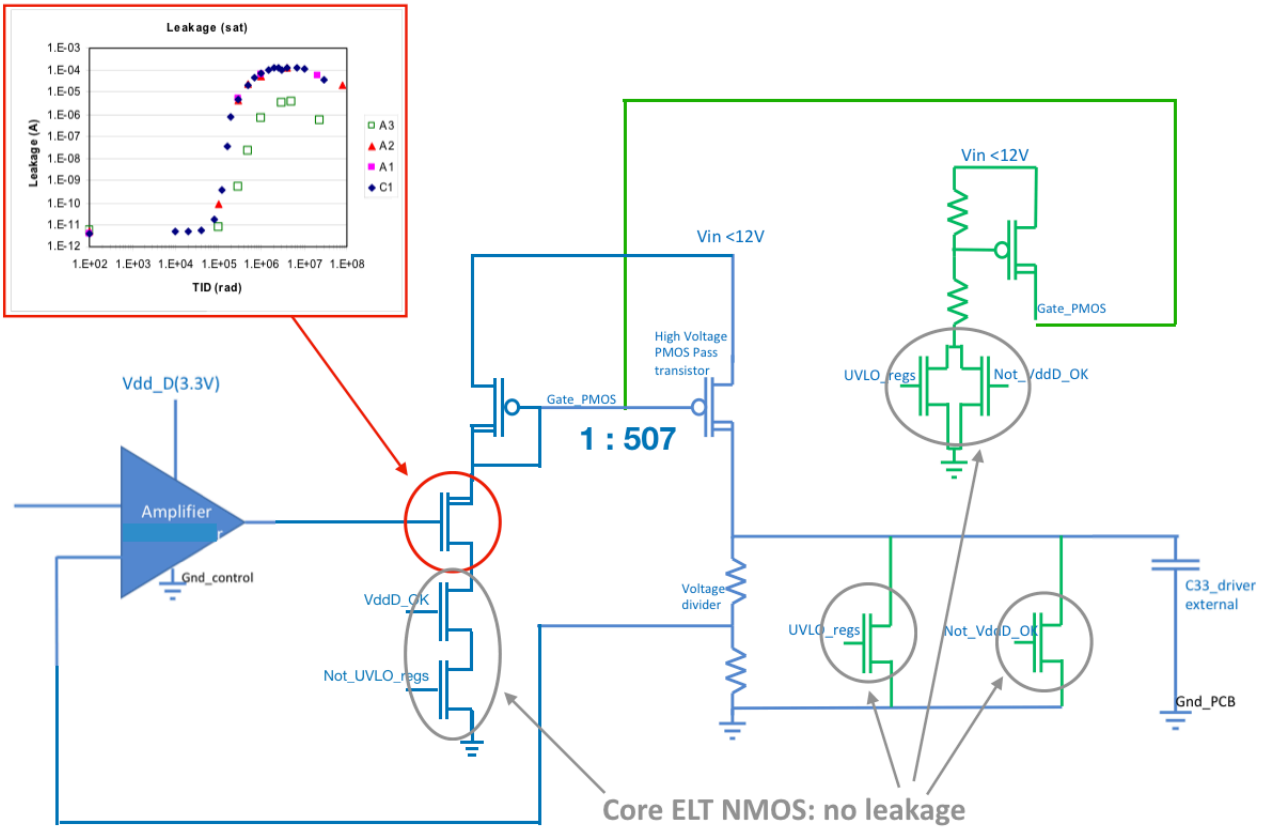

Figure 13: A more complete schematic representation of the regulator for the power drivers (V33Dr) that includes the voltage shifter as output stage of the amplifier. The LNDMOS encircled in red is subject to radiation-induced leakage current (shown in the red squared image) that can not be prevented by layout: changing the layout to Enclosed Layout Transistor (ELT) is not compatible with the high voltage rating of the LNDMOS. The core NMOS are instead all ELTs and are not subject to leakage currents.

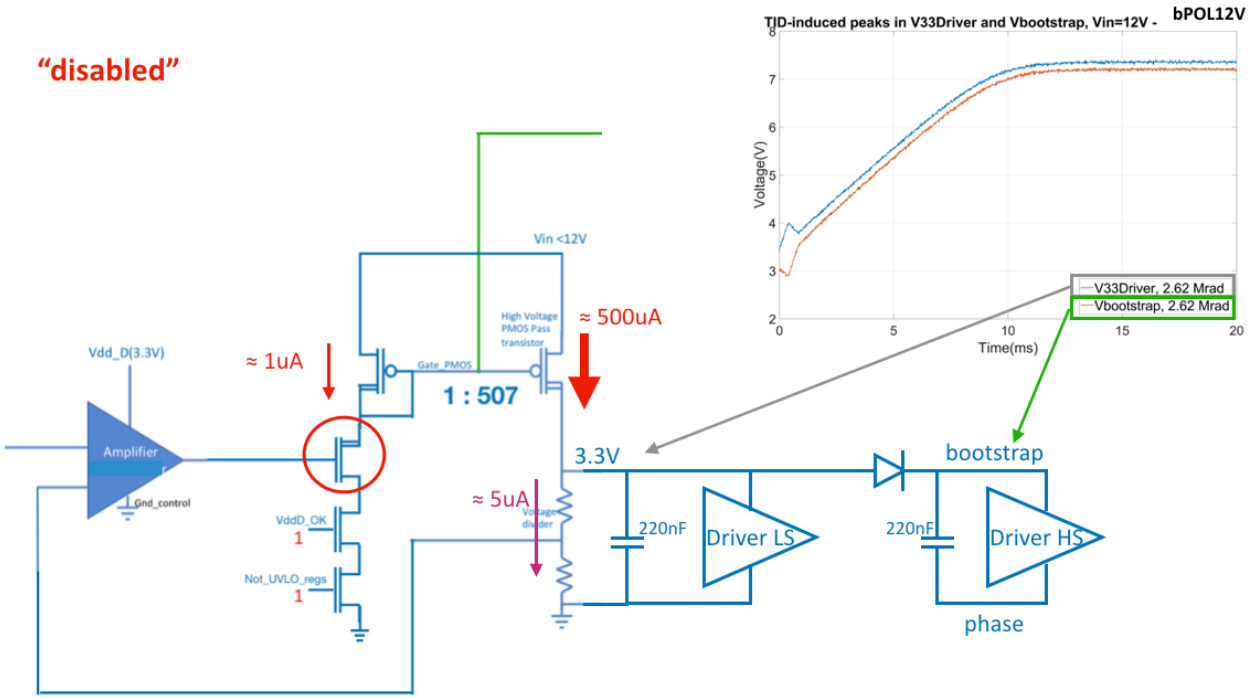

Figure 14: When the converter is disabled, the leakage current of the LNDMOS (hypothetically 1uA here) is mirrored with a ratio 500 to the V33Dr node where it is integrated on the 220nF capacitance. The voltage of the V33Dr node rises linearly until the bootstrap diode is forward biased, at which point the 2 220nF capacitances are in parallel and the slope of the voltage rise is halved. The voltage rise is shown in the top right corner for both the V33Dr (blue) and Vbootstrap (red) nodes.

In the impossibility of cutting the leakage path, there are 2 provisions that can be implemented to prevent the damage: lowering the voltage peak and never turn the V33Dr regulator off. For the first, it is sufficient to lower the input voltage before disabling the converter (to 5V, for instance). The peak voltage at V33Dr can not reach the Vin at the disable time, because the pass PMOS transistor gets squeezed when its drain voltage approaches Vin.
The second provision can be implemented by the addition of an adequate resistance in parallel to the C33_Dr capacitance. In our experiments at the X-ray machine a 3kOhm resistance was adequate even at the very large dose rate of 4Mrad/hour and at room temperature. This requires a current of about 1.1mA, and no peak was observed in this condition on the V33Dr node during exposure. We tried a 15kOhm resistor in the same conditions, and peaks were instead still observed.

The maximum voltage reached at V33Dr during the peak has been measured in a series of experiments where the dose rate and temperature were varied. In these experiments the aim was not to damage the converter, so they were run keeping an input voltage of 8V, which limits the maximum peak of the V33Dr transient at disable. During the test, up to a TID of about 4Mrad, the irradiation has been interrupted every 100krad to perform a 7 s long disable while an oscilloscope monitored the V33Dr node. The maximum amplitude of the peak has been registered, and the evolution of the peak with TID is reported in Fig.15. The dose rate was varied between about 570 and 72 krad/hour, and the temperature was set to either 25 or -30°C.

Samples exposed at -30°C (dashed lines) show a sharp increase in the V33Dr voltage peak at doses between 500 and 800 krad at both of the two used dose rates (70 and 500 krad/hour). The comparison of these results with the equivalent ones obtained at 25°C evidences that the lower temperature is a worst case condition, as expected from the normal behaviour of source-drain leakage current. When using the largest rate at 25°C, the voltage peak increase is comparable; however when the rate is decreased the threshold TID for the peak to appear shifts to larger

doses (about 1Mrad), the increase with TID is more gentle, and the samples are not damaged (the abrupt V33Dr voltage decrease in the blue lines happens when the High Current damage occurs). The figure also shows that, at lower dose rates, the voltage peak decreases gently when the TID is larger than 2-3Mrad. This is also in agreement with the observation that the source-drain leakage in single LNDMOS transistors reaches a peak around 1-2Mrad and then decreases with the accumulation of further dose.

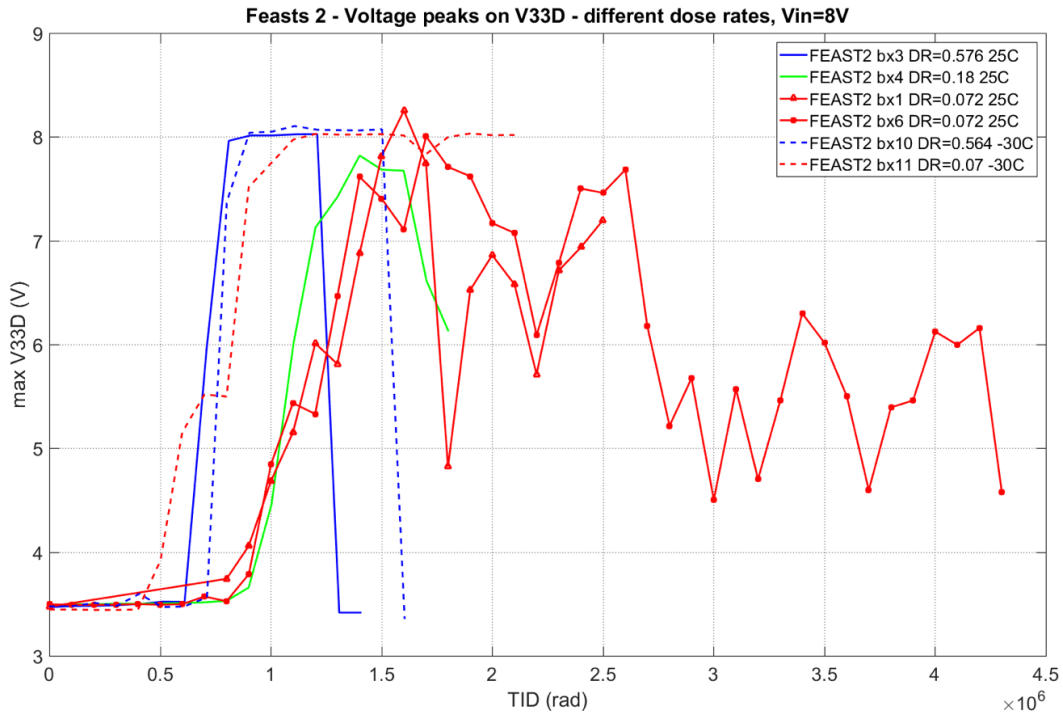

Figure 15: Voltage peaks on V33Dr recorded on FEAST2.1 samples exposed in different conditions of temperature and dose rate. In the legend, the Dose Rate (DR) is indicated in Mrad/hour. Even the lowest rate in the test, 70krad/hour, still leads to the accumulation of 1Mrad in less than a day - so it should still be considered as a very high dose compared to those expected in the LHC experiments.

All other characteristics of damaged samples from these X-ray tests are similar to those observed in both the samples from CMS and from IRRAD. Also in the X-ray case the current increase after the High Current damage is larger when the circuit is disabled than when it is enabled (6mA and 2mA respectively in a sample case). Moreover, High Current samples can be “revived” by the application of an external voltage to the V33Dr node (we did not have Broken samples available to verify that this was also the case).

9. Confirmation of the damage model during a second irradiation test at IRRAD

A second irradiation run at the IRRAD facility, where the facility was configured in the same conditions, was organised with several purposes:

- measure with some precision the TID level at which failure occurs in order to work out a safe area of operation for the FEAST2.1 converters (the dosimetry was very poor in the first run, where converters were spread widely in different locations of the field and an insufficient number of dosimeters had been placed). To this purpose converters were positioned as close as possible to each other, to be exposed to very comparable levels of radiation, and many more dosimeters were positioned
- verify that possible noise pickup on the enable or PowerGood lines was not responsible for the damage. To this purpose, the PowerGood was not routed at the output of the board, and an RC filter advised by the CERN RF specialist D. Valuch (BE-RF-FB) was added on the enable line very close to the connector to the module. This filter was alternately used for half of the modules only, the other half remaining unprotected as for the first run

- confirm that the addition of a 3kOhm resistor on the V33Dr node was a sufficient provision to protect the modules
- since a modified version of the ASIC, FEAST2.2, used a 15kOhm resistor on the node, some samples with this value of protection resistor were used to observe if it is sufficient to prevent the damage in the IRRAD irradiation conditions
- confirm that the other suggested provision, namely the interruption of the supply voltage Vin rather than the disable, was efficient in preventing the damage
- observe the difference in the damage between samples kept at -25°C in the cold box, and samples at room temperature
- observe if samples kept at a constant input voltage of 8V at room temperature were subject to the damage, and at which level of TID (these are the conditions of converters in some applications where the expected level of TID is around 500krad). Because of the voltage drop along the cables, the actual voltage on the module was close to 7.2V when the converters were running, rising to 8V when during the disable.

During this run the setup was duplicated and a total of 64 modules were exposed, half inside the cold box and half outside. Both Aachen and FEASTMP modules were used, and again no significant difference was found between the two. The motherboards used for the test were also redesigned in order to have a much more uniform dose distribution on all converters.

Only “high-current” damages were observed during the run, and exclusively on modules that were not protected by either an additional resistor (3 or 15kOhm) on the V33Dr or by the power-down sequence. Inside the cold box the damaged modules were all on the same board - but this was the only board containing unprotected FEAST2.1 samples. We can hence conclude that, as for the run1, if the TID is sufficient all the converters are eventually damaged. Fig.16 illustrates the 4 boards under test at -25°C and highlights in red the damaged modules. The third board from the left, labelled “Power Down”, also had standard modules but here the supply was turned off to interrupt the modules’ output voltage (no disable was sent). In this way, the converters are disabled by the UVLO_enable threshold which is below 5V - so this corresponds to the protection provision where the converters are disabled at lower voltage.

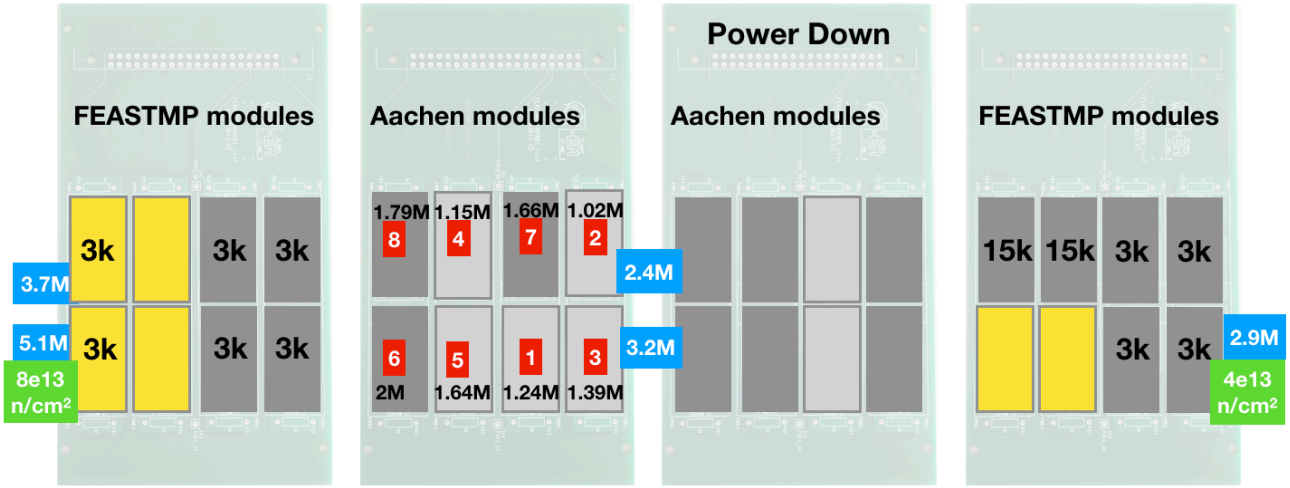

Figure 16: Summary of the results from the 32 modules exposed during the second IRRAD run at -25°C, inside the cold box. FEAST2.1 samples are in grey, while bPOL12V samples are in yellow. Light grey modules were irradiated already in the first IRRAD run, but were undamaged during that run (they were in the farer section of the motherboards of Fig.11). Failing samples are marked in red, and the number is their chronological order of failure during the test. Close to each failure, the number in black reports the best estimate for the dose in Mrad at the time of failure. Samples protected by a resistor are distinguished by the text specifying the resistor value in Ohm (either 3 or 15kOhm). The third module from the left was protected by the a different provision: here the disable command was not sent but the supply to the module was turned off instead. Finally, in blue (TID in Rad) and green (fluence in 1MeV n-equivalent) boxes we report the best estimate of the radiation field from the available dosimeters in different locations.

Outside the cold box, again only the unprotected modules powered at 12V were damaged, but not all of them. This is shown in Fig.17, where again the damaged samples are highlighted in red. Some samples in the first board to the right, as well as the not-highlighted modules in the second, were not protected but they were not damaged either - at least not at the dose reached during the test. None of the samples in the last board to the right, unprotected and using the same disable sequence as the others but powered at 8V (7.2V when the modules were enabled), was damaged.

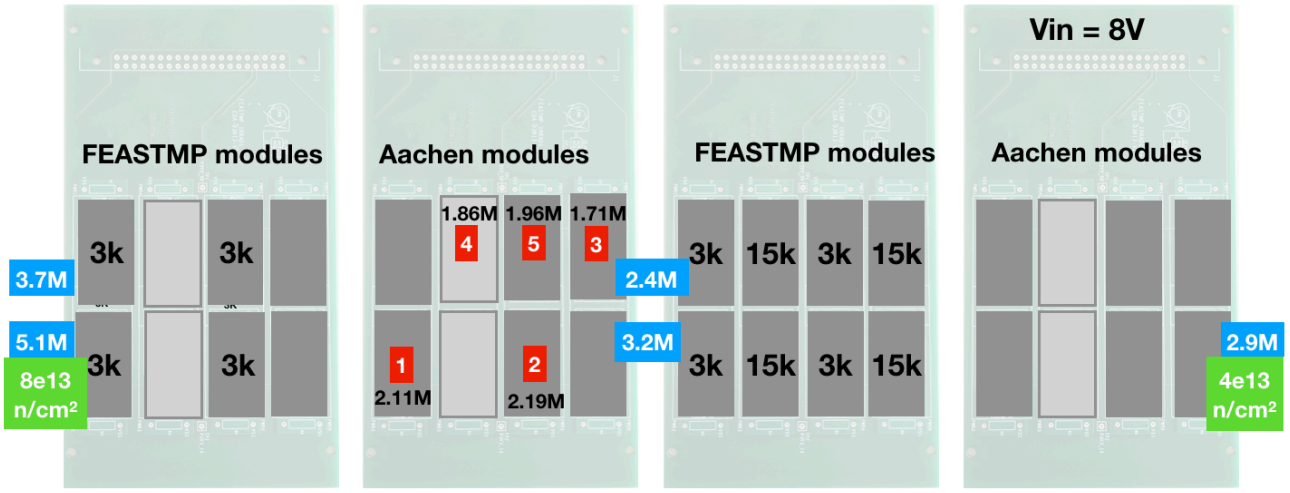

Figure 17: Summary of the results from the 32 modules exposed during the second IRRAD run at room temperature (outside the cold box). All notations are the same as in Fig.16.

Although a much denser network of dosimeters was installed, dosimetry in this complex radiation environment is not easy and therefore the TID levels are affected by a relatively large error. The passive dosimeters used for the two runs in IRRAD were of two technologies: alanine PADs and RPL. They were measured after irradiation, and the measured change in their properties compared with a calibration table where the same devices were exposed in a ^{60}Co source. Therefore the resulting TID levels indicate the equivalent dose of photons in air needed for the dosimeters to show a similar change in properties.

In order to define a safe area of operation for the FEAST2.1 converters, it is important to summarise all results obtained in the two irradiation runs. This is done in Fig. 18, that reports the best estimate for the TID to failure for all modules. The image has to be taken with caution, since the dose levels are given by a combination of passive dosimeter results and extrapolation to different locations and times during the experiment. Also, the size of each data point is proportional to the uncertainty on the dose added by the physical distance of the device from a passive dosimeter - and only to that. The closest the module to a dosimeter, the smaller the uncertainty added by this parameter to the overall large error on the TID estimates. Some general conclusions can be drawn from these results:

- damage only appears in all cases at estimated TID levels above 1MRad
- samples at room temperature start to show damage at a dose larger than those kept at -25°C
- samples already exposed in run 1 are damaged at lower doses when irradiated again at -25°C, but this tendency is not observed at room temperature.

This test allowed the following conclusions, that are discussed in more details in specific report dedicated to the IRRAD irradiation tests:

1. The possible contribution from noise pickup along the enable line is excluded: there was absolutely no influence of the RC filter on the enable line on the damage
2. FEAST2.1 chips can be efficiently protected by the addition of a 3 kOhm resistor on the V33Dr node. This is in good agreement with the X-ray tests where the peaks were eliminated by the use of this resistor. Even a 15 kOhm resistor seems to be sufficient - at least in the conditions of this second IRRAD experiment
3. In the impossibility of adding a resistor, the second recommended provision to protect the chips, lowering the voltage at which the disable command is sent, is efficient. This can also be

done by turning off the supply to the FEAST module while keeping the enable signal high, although it might lead to oscillations in systems where long cables are used to bring the power to the module

4. In order for the FEAST2.1 converters to fail, a TID of at least 1Mrad had to be accumulated for modules at -25°C during a time frame of 20 days in IRRAD, which is a dose rate far superior to the one expected outside the LHC trackers. This dose to failure is even larger at room temperature, where no failure was observed below 1.7Mrad. Even taking into account uncertainties in the dosimetry, it appears safe to use FEAST2.1 modules in environments where a TID below 500krad is accumulated in years while the modules are air or water cooled.

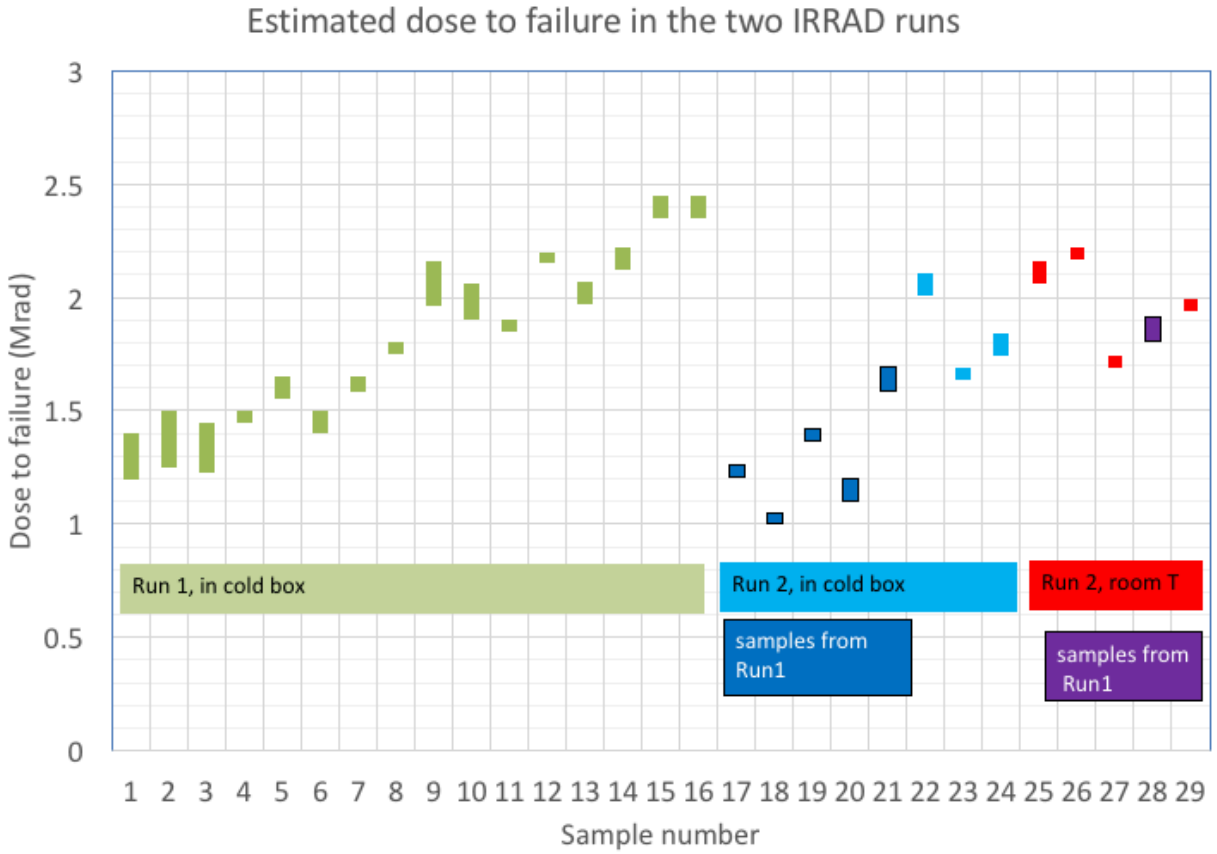

Figure 18: Estimated dose to failure for all samples damaged during the two IRRAD runs. Estimates are based on all the available dosimetry results, and the size of each data point only represents additional uncertainty proportional to the physical distance of the module to the nearest dosimeter. Green data points refer to the first irradiation run, blue (for modules in the cold box) and red (for modules at room temperature) refer to the second irradiation run. Samples exposed in both runs are represented with darker colours.

10. FEAST2.2 is a considerably more robust version of the ASIC

A modified version of the FEAST2.1 chip was designed at a time when the origin of the CMS failures was completely obscure, and the production of 8 wafers of this new design was started. The modifications can be summarised in the 2 following points:

- the clamp transistors at the V33Dr node were removed, and replaced by a 15 kOhm resistor.
- a dedicated full ESD protection structure was added at the V33Dr pad.

No sample was available yet at the time of starting the second IRRAD run, but prototypes were evaluated using the ESD gun and with X-ray irradiations. The new design proved to be much more robust agains ESD discharges, as it was almost impossible to damage it with sparks from the ESD gun. Only after more than 10 minutes of injection on all pads with bursts at 5 kHz, a sample was finally permanently damaged - but the signature of the damage was different than what observed in FEAST2.1, where damage was rather easily introduced by single spikes on some pins (enable, power good).

X-ray irradiation confirmed also that the removal of the clamp transistors and the introduction of the resistor load on the V33Dr node is a very efficient way of protecting the chip. As for the FEAST2.1 chips in section 8 above, samples of the 2.2 version have been exposed to X-rays at different dose rate and temperature, setting an input voltage of 8V and observing the V33Dr voltage with an oscilloscope when the chips were disabled (every 100 krad, just after having stopped the irradiation). The evolution of the voltage peak is shown in Fig.19. Only at the highest dose rate of the X-ray machine, 9 Mrad/hour, the peaks reached the input voltage of 8V. When this irradiation was run at -30°C, the converter went in the High Current state. For the more reasonable dose rates of 500 to 70 krad/hour, and at both -30 and 25°C, the peak voltage is limited well below the input voltage. The available data still do not demonstrate with certainty that samples exposed at -30°C at particularly high dose rates (500 krad/hour or more) will not be damaged - although in this case the damage should always be the High Current state, since the clamp transistors have been removed. However, when looking at both the cold and room temperature data, it clearly appears that a voltage peak is only observable for samples exposed to a dose rate in excess of 180 krad/hour, a dose rate about 2 orders of magnitude superior to the one the converter modules are exposed to in the CMS pixel system.

As an additional remark, it is interesting to point out that, despite the absence of the clamp transistors, samples have been damaged in the High Current mode. This confirms that the current for this type of damage does not flow in the clamp transistors, as demonstrated above in sections 3 and 4 where the images from the emission scan pointed at a current path via the bootstrap diode.

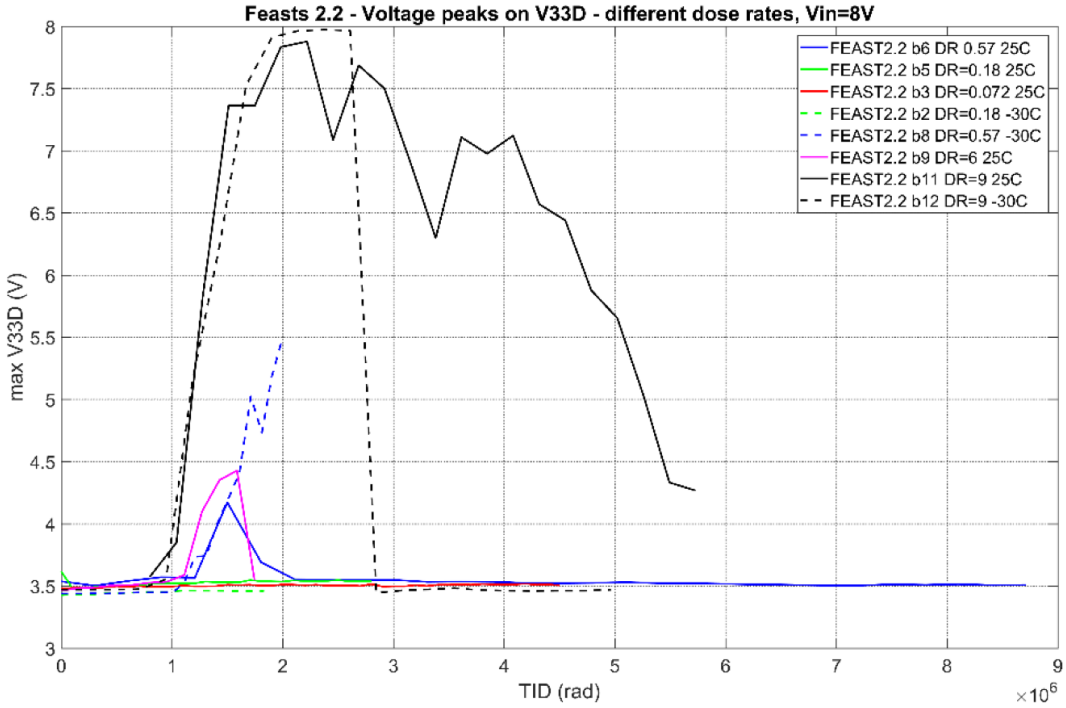

Figure 19: Voltage peaks measured with an oscilloscope on the V33Dr node of FEAST2.2 samples exposed to X-rays at different dose rate and temperature. The Dose Rate (DR) is expressed in Mrad/hour. At rates of 180krad/hour there is no evidence of voltage increase at both 25 and -30°C.

11. Conclusion

This 8-months long study revealed a number of ways to damage a DCDC module built around the FEAST2 and 2.1 converter ASIC. Table 3 is an attempt to summarise them all.

Aggression	Details	Resulting damage
Excessive Vin	The Vin needed to produce damage is above 17V (samples damaged only from 17.5V, often only above 18V). For reliability evaluations, samples have been running for months at 13, 13.5 and 14V without any failure observed.	Full damage, nothing works anymore.
ESD	Damage is only produced with the ESD gun (1.2kV) if: - the converter is powered with an input voltage above the UVLO_regs threshold (enabled or disabled) - the discharge takes place on either the “enable” or the “PowerGood” lines - the tip of the gun is not touching the module, and a spark is created (this condition is worse, the discharge is more intense and more rapid)	High current or V33Dr stuck (both damaging the V33Dr regulator)
Missing or intermittent capacitor	Removal of the V33Dr capacitance, or intermittent contact, have the capacity to create over-voltage transients on the node.	High current or V33Dr stuck
Over-V on the enable	Over-Voltage spikes on the enable line from old/defective power supplies. Observed in some measurement setups, with old supplies that were randomly outputting spikes of > 12V when turned on.	No voltage present on the internal regulators
RF coupling	Only when the aggression is produced with an antenna in the form of very fast transients (sine waves at a single frequency are ineffective) and at very large power.	1 sample: VregHV is present, but not the bandgap
TID	Damage is produced by instantaneous over-voltage peaks at the V33Dr node when the converter is exposed to a TID level close to the one producing a leakage current peak (1-2 Mrad) and disabled while the input voltage is above about 8V. This is the mechanism that has created the CMS problems.	High current or V33Dr stuck

Table 3: Summary of the different aggressions that proved efficient in damaging in any way the modules built around the FEAST2.1 ASIC.

The failures reported in the CMS pixel detector system in late 2017 have been traced to the effect of Total Ionising Dose (TID), and in particular to the radiation-induced source-drain leakage current in the high-voltage (12V) LNDMOS transistors. These transistors can not be laid out as enclosed layout transistors (ELTs): early measurements on dedicated test structures at the time of qualifying the technology have shown that when the original layout is modified the devices loose their high-voltage capability after proton irradiation. Because of TID, positive charge accumulates in the lateral Shallow Trench Isolation (STI) oxide and eventually open a source-drain leakage current path: a current flows even when the transistor is turned off. In dose rate conditions of practical experiments and in the special case of the pixel detector, this happens at a TID of about 1 Mrad (and tend to decrease when further dose is accumulated).

The linear regulator providing the 3.3V for the power circuitry in FEAST2.x (all versions) has an unprotected high-voltage NMOS that is subject to leakage. When FEAST2 is disabled, this small current is mirrored with a ratio 500 and flows in the V33Dr node. The integrated charge on the V33Dr capacitance (external 220nF) originates voltage peaks every time the converter is disabled, and these peaks can reach the input voltage Vin - a voltage clearly exceeding the 3.3V rating of all circuitry attached to the node. Eventually, some component fails and a damage is observed. Depending on the component that fails first, the FEAST2 sample can either go in the “High Current” state, or completely fail to provide voltage (“Broken” state, associated to the lack of appropriate voltage on the V33Dr node).

This model for the failure has been confirmed by all observations reported in this document, and failures on fresh components have been consistently reproduced in controlled conditions in both the X-ray and the IRRAD facilities. Results obtained by the CMS pixel collaboration during tests at the Castor table are also in agreement.

Different strategies are proposed to protect the FEAST2.1 converter modules from this threat, when needed. The need might arise only for applications where FEAST2.1 samples are exposed to a TID in excess of 500krad. This level is chosen with a large safety factor, since no failure has been detected below 1Mrad in tests where this dose was accumulated in less than 20 days on modules kept at -25ºC, and below 1.7Mrad for modules at room temperature. The best strategy is to protect the V33Dr node by adding a 3 kOhm resistor on the module. If this is not possible, then alternative strategies rest on the avoidance of disable cycles at large voltages. This can be achieved by:

- lowering the input voltage to the module as much as possible below 8V before sending a disable command, and enabling the converter before rising the voltage back to the operational level
- if the above can not be done, then completely turning off the voltage supply to the module rather than disabling it achieves the same result. This however can potentially provoke problems because when the voltage is close to the UVLO threshold the converter will sequentially turn on and off, with pulsed currents in the whole distribution system to the module.
- lowering the input voltage of the module during regular operation below 8V. This has to be seen as a mitigation technique only, because there is no data with statistical significance certifying that the converters are safe when the voltage peak on V33Dr is below 8V. Available data indicates however that the failure rate should be considerably smaller at least.

Finally, a new version of the FEAST2 ASIC has been designed and measured. This version, FEAST2.2, eliminates by design the clamp transistors on the V33Dr node, hence removing a vulnerable component and the feedback that caused the catastrophic failure in CMS. In principle, this new version could only be failing in the “High Current” state, which does not compromise the correct functionality of the circuit (in fact, no Broken converter could be produced during X-ray exposures, but the number of samples tested is not statistically significant for the large population expected in the application). The addition of a 15 kOhm resistor on the V33Dr node in this new design, as well as a dedicated main ESD protection structure on the V33Dr pad, have been verified to considerably improve the resilience of the ASIC to both ESD events and to the TID-induced voltage peaks on V33Dr. Modules with FEAST2.2 chips have been irradiated with X-rays, and the voltage peak only appears at dose rates above 180 krad/hour at both room T and -30ºC.

As a last note, we point out that the same damage mechanism has been observed during X-ray irradiation of samples of the bPOL12V ASIC converter. This circuit is a new version of the DCDC ASIC that is being developed in view of application in HL-LHC trackers and that features a considerably larger tolerance to displacement damage. Since this circuit is still in the final prototyping phase, the required modifications to eliminate the sensitivity discovered in FEAST2.1 will be implemented in future design releases.

Acknowledgements

All the results discussed above, and many more that have been excluded from this report for the sake of clarity, were made possible by the work of a large number of individuals, to whom the authors are deeply grateful. We would like to thank the members of the CMS pixel collaboration who contributed with commitment and efficiency by providing samples, preparing them for the Failure Analysis, and helped with the measurements of the modules failing in the experiment. A special acknowledgement goes here to N.Bacchetta (EP-UCM), who coordinated the whole effort. The IRRAD setup, as well as a long-term stress system that was used to exclude voltage stress as the cause of the failures, were built with the help of colleagues from EP-ESE and from CMS: D.Porret (EP-ESE-ME) and S.Cuadrado Calzada (EP-CMO) took care of the most of the hardware, A.Karneyeu (EP-UCM) wrote the Labview program for data acquisition in IRRAD, T.Prousalidi (EP-CMX-DA) wrote the Python code for data analysis in IRRAD, M.Hansen (EP-ESE-BE) participated in the design of the IRRAD test system and supported strongly this work, G.Borghello (EP-ESE-

ME) wrote the Labview program for the stress test system, M.Smith (EP-ESE-ME) surveyed the whole second IRRAD run. We are also deeply grateful to our CERN colleagues F. Szoncso (HSE-DI) and D. Valuch (BE-RF-FB), who drove the study of possible sensitivities to noise injection in the complex CMS environment. A warm “thank you” goes to F.Ravotti and G.Pezzullo (both with EP-DT-DD) for all their proactive support with the IRRAD facility. Finally, while the investigation described in this report was taking place, other colleagues in the CMS pixel collaboration were hardly and efficiently working to ensure that data taking was possible in 2017, and that the detector was ready for the 2018 run. Their effort minimised the consequences of this radiation-induced effect on data taking.