

Dear Lea, dear Danek, dear Katja,

first of all let me congratulate you on a very comprehensive, informative and well-written paper.

I have read the paper draft and you can find my comments to it below, most of them are small and I hope that they are useful.

I wish you all the best moving forward with the paper.

Should you have any questions or the need for clarifications on any of my comments, please do not hesitate to get back to me.

Best regards

Erik

Type A comments:

Abstract

Winter shutdown -> throughout the paper "year end technical stop" is used mostly and is what CERN also uses for these stops. Suggest to use it also here. → ok, changed in all places

L56: on each side -> on each end → ok

L60: have more than doubled -> suggest to add "compared to the design values" as it's tricky to compared 50 ns run 1 to run 2 → ok

L64: shutdown -> technical stop, see above → ok

L76++: several times you use Section in places where "Sec." is recommended by the CMS style guide, is this deliberate for this paragraph? → changed to Sec.

L87++: in principle adding the fourth layer was not contingent on moving the first layer further in, was it? The phrasing somehow seems to imply this. The phrasing "innermost layer closer to the interaction point" is inconsistent with the phrasing in the summary (L1659) where "an additional innermost barrel layer" is used. I suggest to see the 4th layer as the additional one consistent with the phrasing here. → removed "as a result here" and removed "innermost" in summary.

L97: could introduce "ROC" at this point → ok

L107: could introduce "EYETS" here → prefer to not use this acronym

L109: suggest to add "also" as "refurbishment also during the short periods of [...]" → ok

L110 & L114: we first say "connected to four service half cylinders" and just below "installed in a service half cylinder" one sounds like being connected to an outside object, the other like the cylinder is structural, maybe we can make this either consistent, or, if for the FPIX case we want to say sth different, not say "connected to" for both in L110. →

changed to "supplied by"

L122: maybe add: uses "point-of-load" DCDC converters → prefer to keep as is

L131: either "are placed at higher pseudorapidities" (drop 'regions') or "are placed in higher pseudorapidity regions" → ok

L131: "outside the sensitive tracking volume" → maybe "outside the tracking acceptance"? (e.g. they are still within the volume of the strip tracker which also belongs to the sensitive tracking volume) → ok

L135:

"from the beam line, the detector has to withstand" → as the radiation levels differ very much, more precisely could say "the modules in this region" (i.e. at 29 mm from the beam line) → ok

L147: at twice as large → at twice as high? → rephrased

L154: "at some distance" add: "from each other"? → ok

L156: flex printed circuit → could introduce "PCB" here? → but, the "B" is missing.

L158: "in order to facilitate" the verb "facilitate" sounds too strong for me as this suggests you could not manage otherwise. Is this implied here because of the larger number compared to phase-0? Otherwise one could say "ease module production" or "reduce the complications during module production" or sth along these lines → changed to "simplify"

L229: FZ-wafers in <100> orientation → with <100> orientation? → ok

L248/249: power supplies will be → suggest: "power supplies have been" as this is the state at the time of writing. Suggest to add "during LS2" at the end of the sentence → ok

L255: suggest here and further down to introduce a named paragraph (or sub-subsection) PSI46dig as this makes especially the start of the PROC600 part easier (less important here) → had this earlier, but removed due to reviewer comments.

L258: I find the phrasing "26x80 double columns" confusing. How about sth along the lines of "in 26 double columns of 2x80 pixels each" → ok

L299: here we have "time walk" as opposed to "time-walk" e.g. in L356, please choose one and be consistent. → This is correct. L299: noun, L356: adjective

L313: suggest to add named paragraph or sub-subsection "PROC600" here, see comment above → see above

L332/33 and L338/39/...: some parts of the description of the double column synchronization loss seem to be described twice in short succession → indeed. Deleted part of the second sentence.

L352: is passed to the ROCs → would suggest: is passed to each ROC in turn → ok

Figure 7: the plot is missing a label like "CMS", "CMS Pixel" or the like → this is not made with CMS data, so does not need a label. All plots with calibration data do not have labels.

L397: by the CMS trigger and control system -> by the CMS trigger control and distribution system (TCDS) → ok

L420:
an adjustable, up to 32 ns, delay -> an adjustable delay of up to 32 ns → ok

L422:
used in new modules after replacement -> for the reader it might be unclear whether this encompasses only on the new L1 or also on other modules that might be replaced, suggest to make this explicit. → ok added "in BPIX L1"

L435: I find the footnote mark on the percent unfortunate. One could e.g. split the two statements like "The [overall] yield was 93% for the ROC and 53% for the TMBs\footnote{[...]}" → ok

L496: used the fiducials -> used the fiducial markers → ok

L497: were wire-bonded, wire-bonds were encapsulated
-> not nice, please try to find a different phrasing or maybe add at least "and the wire-bonds → removed "after the modules were wire-bonded," since this is sort of obvious.

L498: suggest to inhibit line break between Sylgard and 186 → ok

L504 and L506: repetition of "in turn connected", suggest to change one of them → ok

L513: pixel unit cell -> suggest to use plural "cells" → ok

L540: Then the trimming -> suggest "Next the trimming" → rephrased

L557: due to THE larger area pixels → ok

Figure 12/13/14: the plot is missing a label like "CMS", "CMS Pixel" or the like → see answer above

L631: "the damage that occurred during the dicing of the ROC wafer" suggest to add sth like "described earlier" → ok

L641: detector mechanical structure -> plural "structures"? → ok

L647: suggest "cylindrical layer EACH formed by" → ok

L667/668: have the optohybrids already been described at this point? → changed to optoelectronics components (we do say that the data-transmission is optically earlier)

L668: "space for the module connections" I find this rather unclear as it is not described what connection is being made here → The details are described in Section 5. Not needed to add more details here.

Figure 18/24: I don't have a really good suggestion, but it seems to me that Figure 18 is almost redundant and Figure 24 is a lot more instructive. Already in this section we are referencing Figure 24 as things cannot be seen well from Figure 18. →yes, but one is a

picture of the real supply tube.

L673/74: similar comment, section D is not visible at all in Fig. 18 → yes, but we don't say so. Don't think it is worthwhile to add a figure. Can point to Fig. 24.

L679: "the inner shielding" should we somehow mention its function or this too obvious? → it is obvious.

L697: an irradiated detector at highest luminosities -> I assume this would be "instantaneous" lumi and suggest to add this. → ok

L791: "tubing. Tubing" -> please try to rephrase this → ok

L821: "A CCU supports two I2C channels" the CCU itself can support many more(16, see also type B comment on line 822), so maybe we can make it clearer that we are using two? → ok

L825: "a microTCA system" as you say "VME-based" just the line before, suggest to also use "microTCA-based" here. → ok

L825: "links providing data rates" could say "links providing for data rates" → ok

L832: "providing the clock and trigger signals" -> I find this a bit misleading as mostly they are just relaying these while the phrasing seems to suggest they do this stand-alone. → kept. Is also used here: <https://inspirehep.net/literature/1276617>

L845: suggest to introduce PCB earlier → not needed.

→

L846: "(plus one for redundancy)" it is unclear what redundancy they provide for, i.e. which part can fail? → yes, but since it is the same as in Phase-0, do not want to go into this discussion here.

L862: "can be adjusted in hardware" just to be clear: does this mean they are adjusted once and for all, i.e. no further adjustment possible? → yes.

L901++: should the concept of "analog/digital converter" be introduced somehow? This is used later, but sounds a bit colloquial to me, but at least we should introduce it → is done on L941

L948: here we say "status lines" while later we speak about the "power good line" or maybe other things still. Suggest to either choose just one or at least introduce all terms that we use in one place so it is clear that we are speaking about the same thing → changed "power good line" to "status line"

Section 6.3:

in this section I was confused at which point we speak about power supply modules vs power supply units. At least the counting is confusing as the text says "the two low voltage channels (analog and digital)" which is two per PSU, while later talks about the 4603 having 4 HV channels which is per PSM. Maybe the two LV channels talks about the two `_types_` but also then I find it confusing. → OK. Let's spell this out clearly. In line 984, after the last sentence, we add: "One A4603D power supply module consists of two identical power supply units. Each unit features two low voltage channels and two high voltage channels."

L1026: "a path for the leakage current" I assume "to ground"? If so, can we just add this? → Yes.

L1048: "starts to condensate thanks to the incoming" this should be rephrased in my opinion. Maybe "starts to condensate due to the heat exchange with the incoming...." or "starts to condensate by virtue of transferring heat to the incoming..." → ok

L1055: "far away from the detector" maybe we can simply say "in USC" or "in the service cavern" or sth like this to make it clear what this means? → ok

L1065: "Therefore +15°C is set as the maximum coolant temperature". "Set" for me implies already what we do during operation, so maybe "defined" or "chosen" could be better here? During commissioning +15°C in principle is _minimum_ temperature we want to have to avoid condensation. → ok

L10170: "than 10K" the is the only time Kelvin is used, later °C is also used for temperature differences -> suggest to change to °C → ok

L1080: trailing line should be avoided in final version → ok

L1083/84: USC55 and UXC55 should be introduced much earlier if we want to use them. → removed. Also from Fig. 28 and glossary.

L1091: "through eight cooling loops" eight -> 8, also: it says "to the detector through 8 loops" but is this the total (as the phrasing suggests) or 8 per sub-detector? → per subdetector. Made explicit.

L1095: "cooling plant module" this term has not really be used or introduced so far → changed to "accumulator"

L1104: "to keep the vacuum level low enough to ensure" -> can we give an indicative value? → just a qualitative statement

L1113: suggest to remove parenthesis around "on each system" → ok

L1148: the goal for installing -> the goal of installing ? → rephrased

L1202: "was completed with four non-working modules". This seems like a good result, so we could say "with only four non-working modules"? → ok

L1220 "one section could be tested at a time" was the concept of "sector" introduced? → no, but now added in Section 5.

L1224: can we say where the LV came from? Probably not yet the DCDC converter but simply e.g. a Keithly, maybe we can mention that it's simply a direct delivery of LV to the modules? → this is with DCDC converters. We say on L1218 that this discusses tests of all final components.

L1244: "In the end" maybe "After all tests" → ok

L1309:

"muon detector and the calorimeter endcaps had to be moved" this seems a bit

convoluted. Maybe simply "the CMS endcaps had to be opened" ? → ok

L1311: both sides of CMS -> both ends of CMS (in L1342 you also use "end")
→ok

L1330: "were inserted from the same end of the CMS detector" could maybe add
"as for the old detector" → ok

L1351 "power, control and readout cables" suggest to change to "power cables,
control and readout fibers" → ok

L1361: "has been measured, relatively" -> relative → ok

L1376: During the course of data taking -> suggest to add "the course of collision
data taking" → ok

L1376: I wonder if we should say "is usually fully recalibrated" as some
calibrations might be performed more frequently? → prefer to keep as is.

L1416: "in the center of the working region and is found" -> was found? →
presence used before

Figure 41: this figure does not have a CMS label or anything like it → see answer
above

L1470/71: avoid line break between 45000 and electrons → ok

L1484: "exposed to the first collisions" suggest "exposed to the first stable proton-
proton collisions" → ok

Figure 42: this figure does not have a CMS label → see answer above

L1526: "added in the TBM" could maybe add "described in Section..." → ok

L1527: "will correct this feature" -> maybe suggest "will allow to correct this
feature" → ok

L1533: "before the start of collisions" this sounds a bit colloquial to me. Maybe
"before the start of collision data taking"? → ok

L1534: "from the center of CMS" suggest "from the interaction point" → ok

L1536: about a week -> suggest "about one week" → ok

L1565: "negative elementary charge" we have freely used "electrons" before,
seems odd to become so formal at this point. If you want to use this, you can use
it at the first occurrence but otherwise suggest to just use "electrons" → ok

L1567: (neighboring pixels form a cluster) suggest to drop this parenthesis, this
has been discussed/introduced earlier → rephrased

Caption Figure 44: "were operated at 100 V at [...]" suggest to simplify: "were
operated at 100 V at this point" as "the beginning of 2017" has been said in just
the previous sentence → ok

L1591: affected by the radiation damage -> suggest to drop "the" →ok

L1607: "one of the two most important" this leaves the reader to wonder for the whole sub-section which is the other one, not sure I have a good suggestion for a changed text → yes, but explained in the next paragraph.

L1618: "dynamical" inefficiency -> isn't "dynamic inefficiency" more commonly used? → ok

L1637: ", namely the" maybe we could say "notably the" which means the list does not claim to be fully exhaustive? → ok

L1659: " additional innermost layer" this phrasing is at odds with what has been said earlier, see comment on L87 → rephrased

L1668: "tracking at CMS" -> suggest "tracking in CMS" → ok

Type B:

Throughout the paper you are using eta of 2.5 while from Figure 1 one can see that the detector provides 4-point coverage well beyond. Also, CMS is starting to use pixel tracks up to eta ~ 3.0. Maybe we can at least spend one sentence on why we consider only $\eta < 2.5$ to be relevant here? → Danek

Eta of 2.5 was the requirement set in the TDR. All performance figures in the TDR are for 2.5 because this was considered the limit of track reconstruction. It is also not clear that we cover $\eta=3$ also for all PVs spread in Z. The paper motivates the cut of 2.5 with the design choice of optimizing for that region. For higher eta the momentum resolution is poor since there is no coverage from the strip.

Table 2: the upper fluence for the inner ring FPIX modules is about the same as for the L1 modules. In the text, L142, it is said that the modules in FPIX will stay operational while L1 has to be replaced which is only partly obvious from the bare numbers. → The fluence for L1 in this table is for 300fb and takes the replacement into account.

L183: "trapping caused by radiation damage in the silicon" I would somehow add "after high irradiation" or sth like this as charge trapping is not the dominant effect at low doses. → ok

L205: "here the spatial resolution relies on charge sharing [...]"
it is unclear from what has been said why FPIX would not rely at all on charge sharing to improve spatial resolution, notably when later it is said that the modules are tilted for exactly that reason. It might actually relate to charge drift, but this is not at all clear from what is written → sentence removed.

L285: the concept of the TMB readout token has not yet been discussed at this point. Maybe one can add a "(cf. Sec...)" as forward reference? → changed to "a"

Figure 7: The region at the highest rates and high efficiencies is not really visible in this plot, I assume this is intentional as the low- and high-rate inefficiency should be highlighted and minor differences at high rates for the new chip should

not be the focus point? →yes. This is also different chips at different times, so some fluctuations are expected.

The data taking mode with a reset preceding every trigger seems peculiar. I don't recall: was this done explicitly to not have the inefficiencies or was it something that masked these effects in tests? In case this was specifically done to test the efficiency without being affected by the DC desynchronization one could explicitly add this. → Resets are used for many tests/calibrations to have a clean state. Its use in efficiency measurement as a mistake, it just was not taken out in 2016, so I would not go into much detail here.

L531: this might be a language comment, but we speak about a 2-dimensional scan here and then say that we choose the working point for CALDel to be in the center of the CALDel scan range which sounds like we would not need to perform any scan at all for it. → should say. Center of the valid range for caldel. fixed

Figure 15, right: there is an interruption in the module production around week 30 that is not mentioned in the caption and -- unless I missed it -- also not in the text. Could it be worthwhile to include sth to explain this? → do not want to get into these details. Few weeks interruptions can happen.

L654/55: the values quoted here are with respect to (0,0), I assume. Could it be worth to discuss briefly a possible further effect coming from the offset beam position in CMS? → This would require additional simulations and I am not sure what would be the added value of such comment to the paper. In principle there should be no beam offsets.

L670: the number 175 mm for the inner radius is identical to the one quoted for the outer radius of the FPIX service cylinders in line 761. As they do not touch/coincide, one of them ought to be different, I would say? → No, this is correct. We checked on the CAD drawings. The reason that it fits is that the rails of the BPIX service cylinder are further out and the BPIX service cylinder is not a 180 degree cylinder, but less.

L822: "up to 16 peripheral interface adapter (PIA) channels" this should be "parallel interface adapter" and reference 46 does not agree with the number 16 for the PIA channels. Therein it is said that the CCU25 supports 16 I2C channels and 4 PIA channels. Maybe you can still steer up to 16 peripherals, but at least the channel number seems incorrect. → ok, for parallel. The counting is correct. It is 4 8-bit registers, so 32 channels. 13 are used to enable/disable DCDC converters (some sectors have less DCDC converters) and 3 used to send reset signals. The remaining ones are configured as input and are used to read the power good bit of the DCDC converters and the status of the QPLL locking.

L823: "generate reset signals" it is a bit unclear who these signals are generated for → added "readout electronics on the service half-cylinders and the detector modules"

Section 7.2/7.4

I am very confused by the numbers in these sections, but this could be my fault. The power mentioned is $6+3+2 = 11$ kW and it is said that the cooling system (not one

cooling plant) was designed for 15 kW. The maximum power for one plant is then given as 11kW for one plant, so 22kW for two plants, in normal mode and 18kW in backup mode. It is not really obvious from what is written why the cooling power would be so much higher in the backup case for a single plant.

The cooling power requirements for the full detector are indeed 11kW (6 kW/BPix + 3kW/FPix + 2kW/Heat leaks) so we have designed a cooling system with 2 15kW capacity cooling plants to be able to operate with sufficient margin the full detector on 1 single plant in case the second one is down for whatever reasons.

Then we tested our cooling system.

In nominal operation (when both plants are running independently on their own sub detector circuits) we need to provide max ≈ 7.5 kW on each plant (15kW in total). We have pushed up to 11kW on each plant (22kW in total) and the limiting factor here was the main chiller (that we have in common with the TK C6F14 system) that could not provide more cooling power because the full capacity was reached.

In back-up operation (when one system is running on both sub detector circuits) we need to provide max 15kW on one plant. We have pushed up to 18kW on one plant and the limiting factor here was one component of the system (on the freon circuit side in the plant at the exit of the heat exchanger) giving too high pressure drop at this higher load.

So in the two cases the limiting factors were different, this is why we have different numbers and not 2 x 18kW in nominal operation...

Another thing that confused me is the statement that each plant runs at about half capacity when cooling one subsystem while at first glance this is very asymmetric at 6 vs 3 kW and the 3kW would be far away from 50% capacity. Is this because of the ambient losses are largely on FPIX?

This is let's say a shortcut to make it more simple... Ambient losses are the same with about 1kW per system. So Bpix runs with ≈ 7 kW load (47% of the total capacity) and Fpix runs with ≈ 4 kW (27% of the total capacity).

L1120/21: "which highlights the importance of a careful design for the general experimental services" while this might be a true statement, it can likely also be said in other places and is not only applicable from the viewpoint of the cooling system, so not sure we want to highlight it so much here. → ok, removed the whole sentence.

L1131/32: "along the cooling loops for nominal power load" what is this nominal power load, e.g. where in the life cycle of the detector? startup? towards end-of-life? the current state one would like to compare with? After 300\fb of data-taking. Added to the text.

L1136/37: "to improve the temperature uniformity along the loops" -> and? Can we make any statement if this succeeded? We expect it from what we saw in the mock-up, but did it also happen in the real system? Is the effect comparable? Can this be quantified? If not, why not? → yes, worked as expected. Rephrased.

L1299: The statement about the loss pattern comes rather out of nowhere. The next section of text suggests that this might have to do with a too rapid cooldown as the ramping speeds etc are mentioned. Even if understand that we do not want to dwell on this too much, can we include a statement along the lines of "were lost in patterns at the corner of the modules, likely due to a too rapid cool-down during the test procedure" ?

→ changed to: "For the first half-cylinder tested, 0.15% of pixels were lost in patterns at the corners of modules due to mechanical stresses."

Section 10.2: while you describe the bias point selection, I was I was missing how it is decided which gain setting is chosen → we do not optimize the gain. It is kept constant as said in L1398.

L1422++: this maybe comes down to wording, but the difference between "relatively non-uniform" to "very uniform" is only a factor of two → rephrased

L1457/58: "with significantly lower thresholds than before." this statement somehow very much invites to give a guestimate of the achievable new threshold, but I understand if you would like to avoid this → do not want to quantify, but still like to keep the statement.

L1488: "with the clock distribution in one sector" I don't know if the reader can use information from before to infer how many disks in this sector would be affected by such a problem, maybe it is worthwhile to state this explicitly? → don't think this is needed. It is clear from the sentence that it is less than 3.9%.

L1507++: the question I asked myself reading these numbers (and not remembering) was whether there was any significant overlap between the inactive detector regions as this could have an impact on tracking. Obviously we do not want to discuss tracking performance in details but could indicate (non-)overlaps. → lead to overlap as it affected layers 3&4 in the same sector.

L1580: "resulting from split clusters" with only the information in the text it might be difficult to understand what these are and what their cause is. Maybe we can add a reference which provides more explanation or add a bit more explanation? → We do not have a reference. The sentence says what the cause is: "lost pixels due to the PROC600 inefficiencies"

L1611++: one aspect that is not mentioned here is if and how the efficiency measurement is biased by the fact that the layer is(?) included in the initial pattern recognition. Is there something simple we can say about this? The other question is I had was that inefficient pixels can be a source of poor resolution, but in principle these are known and could be taken into account (e.g. excluded to get a more unaffected measurement)?

It think presenting results of such a study is more appropriate for the planned performance paper.

Here we just present simplified results.

Inefficient/dead pixels are dominated by SEUs and this cannot be taken into account in any way.

L1651: "and often on the radial position of the layer" could we add a bit more information as to why this might be the case? "Often" obviously means "not always" we seems to exclude a clear-cut effect, but could we give any hint why this is the case?

The radial dependence enters because of the radiation damage and also because layer 1 & 4 measurements rely on extrapolation and not interpolation. Removed "often".