Question 1

- It is understood that the variance analysis is not yet being posted to IPD. It would, however, be useful to see what you have collected thus far (in any convenient format)
  - The June CPR is the first one for which we are writing variance reports
  - These initial reports have not been reviewed by the PO yet, and for many CAMs, this is the first one written

➡ It is premature, and not helpful, to share these outside the project team

  - However, here are the variance drivers (SPI < 0.90 **and** SV > $50k):
    - LAr FEE: submission of ADC pre-prototype 3 is delayed
      - If submission happens in August, PDR date likely to be maintained, so no significant impact
    - LAr BEE: delay in engineer hire (person will start this August) at SBU, extra time being spent by AZ team on FW specifications
      - Extra effort planned in next few months to recover, better FW specifications should make FW development more efficient

Question 1, ct'd

- Muon sMDT: Module 0 material and precision jigging for construction has not arrived yet
- Trigger HTT SV: Effort on the RTM was delayed due to deadlines on other projects, and the RTM/TFM work was deemed acceptable to delay short term
  - Effort will be available in the near future to compensate and recover the schedule.

Question 2

- The June monthly report says that the EV report has bugs. Please explain the origin of the problem. Is it e.g., related to the tool being used or the data.
  - The use of bug may have been too strong. It was intended to reflect the fact that it is still early days as we exercise the system and there may be issues yet to be resolved.
    - Some of the historical actuals from FY17-18 were not captured correctly in the May CPR. This was fixed in the June CPR.

Question 3

- The EAC reported in the June report is not same as the BAC. What is the primary reason for this?
  - The BAC reflects the current “baseline” cost, i.e. the planned cost for the RLS against which we report
  - The EAC is the BAC + cost variance + impact of escalation (due to schedule variance)
  - (To see trends, a stable “baseline” is needed over a certain period)

Question 4

- In the subsystem presentations two contingency numbers were usually given, for 70% CL and for 90% CL. When adding up the contingency for the total project which of these CL's were used?
  - We simulate the project, typically for 1000 “runs”, and integrate the expected project cost distribution to extract the xx% CL TPC
  - In Brooijmans’ talk, sl 32, the 70% and 90% CL cost contingency numbers for the full project are given ($16.1M and $22.0M)
    - NSF requires that the TPC falls into the 70%-90% CL range
  - In Brooijmans’ talk, sl 33, the 90% CL cost contingency numbers are given for each deliverable
  - The available contingency, given a $75M TPC, is $20.1M
    - This corresponds to the 85% CL
    - In other words, according to the simulation we will complete the planned scope within $75M at 85% CL
      - We also have 15% scope contingency, bringing the CL to ~100%
  - Contingency for individual deliverables is not meaningful
    - We show the numbers to show that less mature items “score worse”, i.e. as a check that things are consistent

Question 5

- A possible one year slippage of the overall ATLAS upgrade schedule was mentioned as a possibility. This might cause a standing army cost increase. Even though the US responsibility is defined as delivery of subsystems, and thus insensitive to standing army costs, a delayed overall schedule might delay the completion of US deliverables due to delays of prerequisite parts from overseas collaborators. How has this possibility taken into account in the contingency estimation?
  - At this time, the contingency calculation only includes the costs of delays caused by uncertainties and risks captured in the RLS and risk register
    - The risk register does include risks that items we depend on are delayed (lpGBT, ELMB2, GCM hardware)
    - The risk register does not (yet?) include a risk that the CERN schedule would be delayed

Question 5

- CERN delay risk?
  - To estimate the impact of a CERN delay risk, can look at Phase-I, as LS2 was delayed by 6 months **after** we baselined
    - Both NSF and DOE Phase-I projects were governed by DOE 413.3b
  - In LAr in Phase-I, we used up all the CD-2 schedule contingency + the added amount from the LS2 delay
    - Available time influences decisions on how to address features found during integration
    - However, while we used up ~18 months of schedule float (in a 4 year project), we only used 9% contingency, of which 0 went to "standing army" costs, but maybe 2% can be assigned to extra checks we would not have done if the extra 6 months had not been available
  - In TDAQ in Phase-I, similar situation
    - Of ~25% contingency drawn, none to "standing army" but maybe ~2-5% can be linked to extra time available

Question 5

- CERN delay risk?
  - It's probably a good idea to add such a risk, with cost impact range 2-8% of base cost or so
- (This would have a slightly larger impact than the "MREFC delay" risk we have been asked to retire)

Question 6

- The HTT project builds on FTK is key ways. Please reflect on the aspects of FTK that were successful and aspects that were less than successful. Which lessons are appropriate to HTT? How will these lessons help you manage risk to cost and schedule in HTT?
  - Next two slides

Lessons learned from FTK

- Simplify the hardware
  - FTK has 7 board types; HTT has 3
  - Data transfer is simpler (from DAQ rather than directly from detector)
- Integration of multiple board types should be started early
  - Integration tests of HTT are in the RLS including joint demonstrator tests of the TP and TFM, and slice tests at CERN with all prototype board types.
- Start writing firmware long before final FPGA decision is made; in FTK some of the FPGA resources turned out to be marginally sufficient.
  - HTT: Two prototype rounds with a substantial period of firmware development/optimization prior to the final FPGA decision.

Gus

Lessons learned from FTK

- Have professional oversight of scientific personnel writing the functional firmware.
  - HTT: There will be firmware engineers at each of the firmware writing institutions. They will set the overall structure of the firmware, write some of the most challenging functions (I/O, memory access, etc.), and oversee the work of the scientists.
- There should be system-wide oversight of the firmware
  - There is an HTT firmware coordinator whose responsibilities include system verification procedures.
  - All HTT board reviews will include reviews of the firmware.
- When testing with simulated data, include dropped data and data errors.
  - A comprehensive test suite which includes data corruption is part of the RLS.
  - QA/QC specifically includes inserting data with errors.

Gus

Question 7

- The key to the proposed change management plan is the CCB. Please explain how this committee functions, e.g., by consensus, majority vote, unanimity? How are conflicts of interest among its members managed/mitigated?
  - The CCB is constructed to serve as a forum and clearinghouse to openly air and discuss all issues relating to changes to the project plan
  - It is intended to foster broad project engagement and ownership of the project -- scope, cost, schedule, risk, etc. The PO takes this function of the CCB quite seriously, both in imparting our views and carefully taking into consideration those of the L2 principals.
  - Central to the discussion is contingency usage, which is owned by the PO, and is always subject to its final discretion. We have made clear from the outset, and often, that while we evaluate contingency at the deliverable level, it is all held as one lump sum by the PO: some systems may ultimately receive a far larger fraction than their "share" might imply, or none at all, depending ultimately on the evaluation of the global needs of the project determined by the PO.

Question 7

- Each request is evaluated on its merits at the time of evaluation. The PO has made clear that positive or negative votes do not influence future decisions on their systems.
- Requiring or expecting a CCB-wide consensus (or unanimity) would be unrealistic and inconsistent with the temporal and other pressures associated with effective project management and execution. Issues of fundamental disagreement will be decided by the PM, Deputies and NSF PI (factoring in CCB input); if this proves inconclusive, it will be decided by the PM (CCB Chair). We have not yet confronted such a situation, but the process is in place and understood by all of the principals. The process has worked smoothly so far; there has not been a case of fundamental disagreement.
- We are in the relatively early stages of the process; many of our change requests to date have been associated with truing up the RLS and base plan, and have been reasonably straightforward. The PO has been, and will continue to, regularly evaluate whether modifications to the process would be advantageous as the project moves through its various stages.
- Funds do not cross the DOE/NSF boundary – contingency requests are each treated independently for each funding line/scope. The process is identical regardless of funding source, and the full CCB participates in the evaluation of each BCP.

Question 8

- The core management team has used the US-ATLAS Phase-1 upgrade project cost and scheduling data to inform your expectations for the MREFC HL-LHC project. Could you make available to us any documentation you have on “lessons learned in Phase 1”?
  - The Phase-I CD-4 Lessons Learned write-up is posted on the indico page under the Tuesday homework response entry
  - (Phase-I’s similarity to the current project makes it very useful in this context; few projects have something like that.)

Question 9

- Provide a pointer to NSF project milestones that are linked to international milestones and indicate how this is tracked in the RLS
  - The majority of these are labeled with “EX” in the task name (there are many)
  - They include the ATLAS reviews (SVR, PDR, FDR, PRR), needed at CERN dates, as well as required deliveries from international collaborators (but there are not so many of these in NSF scope)
  - In our RLS these milestones are tracked as all other tasks and milestones
    - We plan to emphasize the need to status their expected completion dates, not just actual completion
    - Expected completion information is obtained from regular ATLAS working meetings, as well as the International ATLAS schedules (which are statused quarterly)

Question 10

- Top level talks emphasized the difficulty of controlling the risks of international contributions to the US project. We would like to a specific list of those contributions, if any, and the plans for mitigating these risks (where possible). Integration aspects such as power, cooling, space (rack and cable), and common projects such as lpGBT or bPOL are of particular interest
  - In phase-I international contribution delays/unstable specifications generated the majority of cost and schedule contingency draws, but the international schedules were coarse and essentially not revised or statused after the MoUs were finalized
    - The Phase-II ATLAS approach is completely different, a world of difference
      - Not a guarantee of anything, but very significant mitigation
  - In particular, the international schedules clearly identify when prototypes are needed to allow integration tests to proceed

Question 10

- The risk register identifies the impacts of delays in lpGBT, ELMB2, … availability
- Power and cooling density specs are already set; the question is rather whether we can meet these specs (for ATCA blades)
  - ATLAS has set up a dedicated test crate at Point 1 and encourages team to use it for thermal measurements
- Similarly, rack space has been identified and cable chain inventory indicates there is sufficient space
- (One unknown is when exactly USA15 - the counting room - will be available for installation, but that only affects us in the I&C phase)

Question 11 (Muon+Trigger)

- Point 5j in the NSF charge reads "Performance verification and acceptance test policies for all deliverables are defined and complete. Documentation describes how acceptance tests will verify that deliverables meet design performance specifications and safety requirements. (i. QA plans and activities are integrated into the RLS. ii. QA and radiation exposure policies are applied consistently across the project.) Can you tell us what the status of this documentation is and point us to it? In particular is the SMDT integrated in RLS (this question applies to both trigger and muon).
  - QA testing-related tasks precede ATLAS FDR in schedule (often several tasks)
    - ATLAS FDR is final sign-off on the design, requires integrated testing

Question 11 (Muon)

Example from the sMDT

Tubes: The FDR/PRR for the tubes was 14 Nov 2018 (https://indico.cern.ch/event/773575/)

https://edms.cern.ch/document/2048104/1

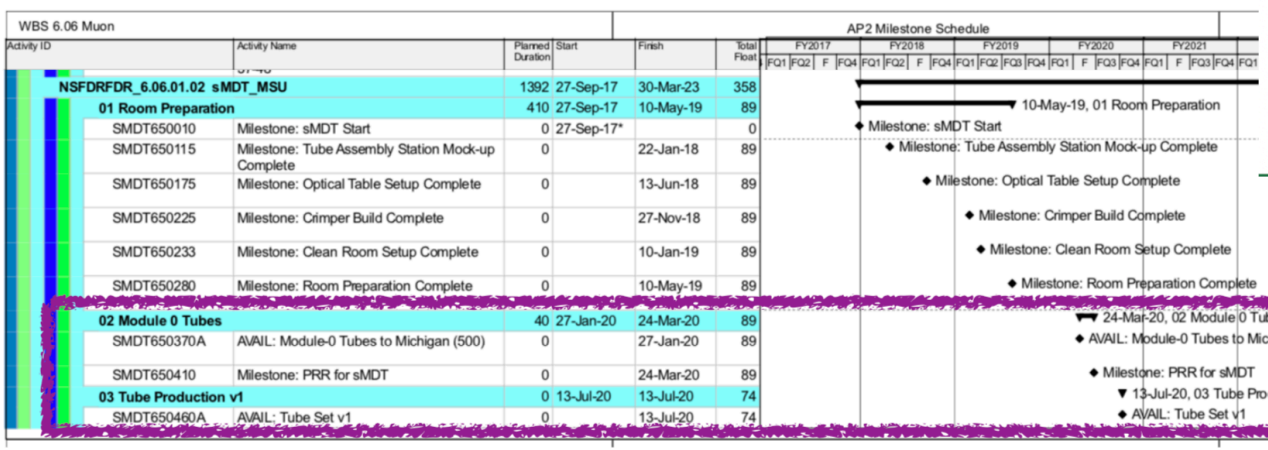

Tasks leading to FDR

1.02 SMDT650115	Milestone: Tube Assembly Station Mock-up Complete		1/22/18
1.02 SMDT650140	Review of Tube Assembly Mock-up and Tube Handling	3/21/18	3/27/18
1.02 SMDT650175	Milestone: Optical Table Setup Complete		6/13/18
1.02 SMDT650220	Setup Tube Construction Clean Room (MSU)	10/1/18	11/27/18
1.02 SMDT650225	Milestone: Crimper Build Complete		11/27/18
1.02 SMDT650230	Install Clean Room	11/28/18	1/10/19
1.02 SMDT650233	Milestone: Clean Room Setup Complete		1/10/19

Next step is production of the Module-0 tubes to get the MSU construction site approved for the PRR

Chambers: The PDR chambers, PRR jigging was 3 May 2019 https://indico.cern.ch/event/813217/

https://edms.cern.ch/document/2145116/1

FDR for chambers will be in September 2019

1.01 SMDT630200	Validation of Final Precision Chamber Assembly Jigging	6/4/19	7/30/19
1.01 SMDT630210	Milestone: Setup of Final Precision Jigging Complete		7/30/19

The Preliminary Design Review (PDR) of the BIS1-6 sMDT chambers is passed with a small number of recommendations, mainly to clarify the drawing approval process given the very large number of chamber variants and design changes with respect to BIS7/8, which are minor but should be fully documented and signed-off. At the FDR a solid and tested QA/QC plan is expected.

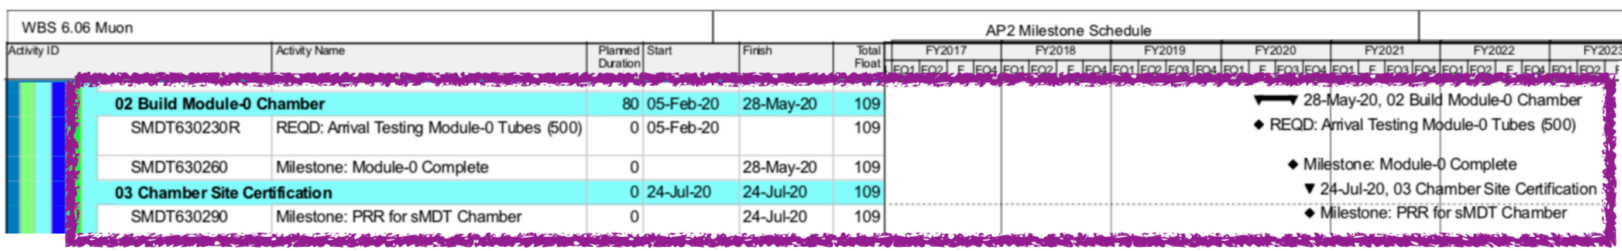

Next step is production of the Module-0 chamber to get the UM construction site approved for the PRR

Question 11 (Trigger)

- Example from HTT TFM board

TFHW10450R	REQD: Full Data Flow and Processing Firmware Complete	0	20-Jul-21	
TFHW10400A	AVAIL: Prototype Board Manufacturing Available	0		26-Jul-21
TFHW10410R	REQD: Main Board Available for Testing (Pennsylvania)	0	12-Aug-21	
TFHW10420	Test Control and Loading of the Board	23	12-Aug-21	14-Sep-21
TFHW10460	Test I/O on the Prototype Board	19	15-Sep-21	12-Oct-21
TFHW10470	Test Full Processing on the Prototype Board	64	13-Oct-21	13-Jan-22
TFHW10480R	REQD: Modified Firmware	0	24-Mar-22	
TFHW10490	Retest Board	44	24-Mar-22	24-May-22
HTTMS1120	MILESTONE: TFM ATLAS Final Design Review	0	25-May-22	

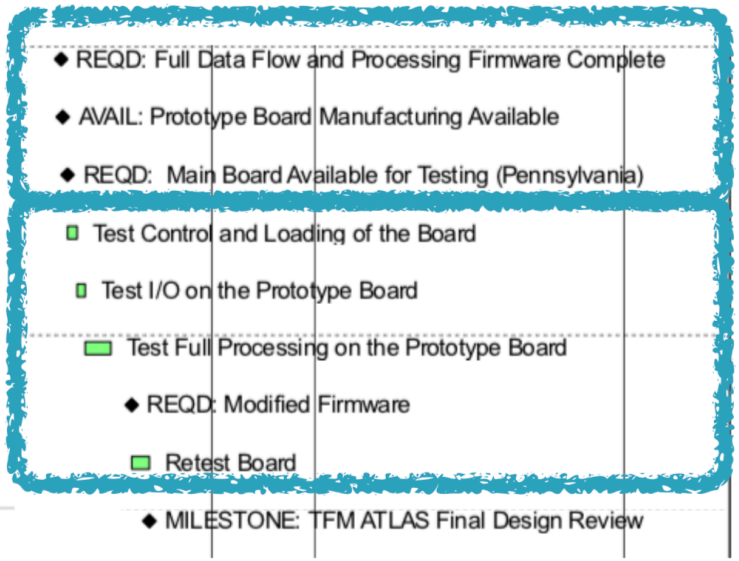

TP board and firmware needed for QA testing

TFM QA Testing

- Example from Global firmware

GEPTFR8160	Document V1 Trigger Framework Firmware	15	21-Apr-22	11-May-22
GEPTFR8170	Coordinate Integration of Algorithm Snapshots	28	29-Jun-22	08-Aug-22
GEPTFR8180	Coordinate Integration of Trigger Framework Firmware	36	03-Nov-22	23-Dec-22
GEPTFR8190	Prepare for Trigger Framework Final Design Review	6	27-Dec-22	04-Jan-23
GEPTFR8200T	Travel for Trigger Framework Final Design Review	6	05-Jan-23	12-Jan-23
GEPTFR8200	Trigger Framework FDR	8	05-Jan-23	16-Jan-23

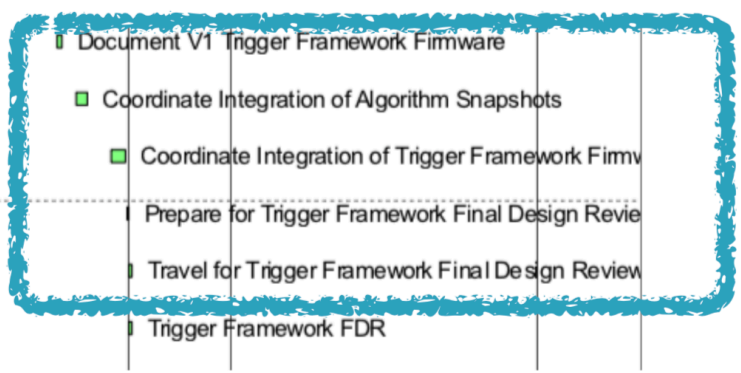

Documentation, and Integration test for FDR

Question 12

- In the "NSF review tracking_2017 v7" excel spreadsheet, there many cells which indicate the response is "underway" or recommendation status is "in progress". Indicate which of these have been completed and which are still in progress.
  - Indeed, there are a number of items there that are complete but have not been marked that way; we will address this

REVIEW # *year-xx*	REVIEW TITLE *name of review committee*	DATE PERFORMED *start of review*	RECOMMENDATION # *(Rxx)*	RECOMMENDATION DESCRIPTION	OWNER	RESPONSE	RECOMMENDATION STATUS per L2 Manager
2018-01	Review	16-Jan-18	R10	Proceed with finalizing technical specifications for all … perform radiation aging and SEU studies to assure proposed technical solutions satisfy all specifications.	T. Schwarz	According to the current schedule for our … time for the FDR for all deliverables.	In progress
2018-01	NSF Preliminary Design Review	16-Jan-18	R10	Assure that environmental specifications are part of the technical specifications for all elements of the muon system.	T. Schwarz	This is being done (see below).	Done
2018-01	Review	16-Jan-18	R11	Make a detailed list of all, including environmental, specifications for L0MDT and CSM boards and keep track of …	T. Schwarz	The L0MDT and CSM specifications are currently being defined. This is work towards … the coming fiscal year. Environmental specifications are part of the required parameters.	In progress
2018-01	NSF Preliminary Design Review	16-Jan-18	R12	Add more technical milestones in the R&D and MREFC construction phases.	T. Schwarz	This has been implemented.	Done
2018-01	NSF Preliminary Design Review	16-Jan-18	R13	For TDC develop a more extensive technical specifications, including power budget, input hit rate, radiation survival and SEU/SEE requirement, INL/DNL requirement, maximum data loss, and more detailed functionality and interface descriptions.	T. Schwarz	Most of this has been implemented. A specifications review has taken place. We will continuously add detail to the as-built documentation.	Done
2018-01	NSF Preliminary Design Review			Optimize the number of spares, based on the experience with prototypes, to minimize production of boards which might …	T. Schwarz	We haven't deined the numbers yet as we are still in early prototyping phase, but it will be … of the iterative process of defining our total production after prototyping.	In progress

Done - Spec Reviews passed for all deliverables except L0MDT, which will have a spec review in Sep 2019.

Done - Spec reviews include environmental specifications. Several reviews scheduled as milestones to track R&D.

Done - Spares defined and updates integrated into project plans/estimates.

Question 13 (LAr)

- It seemed from the plots showing performance of the ADC in terms of energy resolution, the ultimate goals were not met for the latest, v2, prototype. Please be more specific on the progress of the ASIC development and what known issues were addressed between v1 and v2 and what issues, so far, need to be addressed in a v3. As homework, can you tell us what improvements to the performance, either in hardware, software, or firmware can we expect with the current v2 chip before the FDR?
  - With the v2 pre-prototype the primary goal was to integrate the individual blocks into a 2 channel ADC, testing it as a single LAr channel system, with on-chip digital processing and control and an e-link for lpGBT testing. As far as precision, this was an incremental step toward the ultimate precision goal and, though it would have been great to meet the ultimate goal in v2, we anticipated 3 pre-prototypes, before the final prototype (4 iterations).

Question 13 (LAr)

- V1->v2 was a major step in integrating the DRE and SAR in a 2 channel design, with many changes. Looking forward, for v3 we had a 2 day workshop at UT Austin in April to address the challenges remaining. In v3 the issues are the noise limiting the performance at the lowest energies, and the calibration of the 4x DRE gain, limiting the performance at the highest energies. In the DRE improvements will come through adding an on-chip gain calibration using a DAC ladder, and improvements in the sampling network to reduce kickback. In the SAR the improvements come through adding an additional physical bit in the second stage of the SAR, to improve the overall resolution to 11.5b, at the cost of some additional power.
- For the v2 chip we don't anticipate any changes or improvements, testing the v2 on it's own is effectively concluded (though still ongoing as part of the analog test board). The ATLAS FDR in December 2019 will focus on the results of the v3.

Question 14 (LAr)

- Can you be a little more specific on your concerns of the rad hardness of the Amplifier shaper? Given what is known about radiation damage in similar circuits, what are the potential impacts of what is known/unknown about potential radiation damage to this external (DOE) circuit on the NSF part of the project?
  - There is no reason to believe there will be a radiation issue with the PA/S. The 130 nm CMOS process has been demonstrated by several others ASICs to be radiation-hard even at tracker levels, around a factor of 1000 beyond our levels. However, the PA/S shaper needs to be irradiated to validate that the challenging analog performance, in particular low ENI and high dynamic range of the preamp, is not degraded due to radiation. A test has been scheduled for end of October.
  - The FEB2 in NSF scope has a risk in the Risk Register (RN-06-04-02-001) that accounts for a delay in FEB2 preproduction due to a variety of possible reasons, including late delivery of the PA/S.

Question 15 (LAr)

- The improvement to the Higgs>gamma gamma mass resolution seems an important benchmark. It was unclear from the presentations whether it was just an impact on the trigger or if it affects the offline resolution. It is also not clear what algorithm for digital filtering was used, and how that corresponds to the currently envisioned algorithm, and what portion of the improvement comes purely from the dual gain ADC. Can you clarify?
  - This is discussed in detail in Sec. 4.1 and Ch. 8 of the LAr TDR. The offline resolution is impacted because the only the online system sees the extended time history needed to establish the base line. The Phase I upgrade has already demonstrated an implementation of a real-time online determination of base line subtraction for pileup which must be detector location and bunch number dependent.

Question 15 (LAr)

- The digital filtering implemented for the performance studies is the standard Optimal Filtering Coefficient (OFC) algorithm with 5 samples already being used by ATLAS. Fully simulated GEANT samples were generated with mu values up to 200.
- Improvements from extending the number of samples and other filtering algorithms have been studied. In particular a Wiener Filter with Forward Corrections indicates improvements are possible but imply need for a longer time sequence of data is needed prior to the collision than for the OFC

Question 17 (LAr)

- While measurements of the two-channel coherent noise made using the LAr Analog Testboard meets specifications, what confidence do you have that the fully populated board will also meet the coherent noise specifications? What is the remediation plan if the coherent noise level is too high?
  - The fact that the 2-channel Analog Testboard exceeds the coherent noise specification is very encouraging. Tests with the full channel density COLUTAv3 ADC will be important to demonstrate that the ASICs deliver low coherent noise on-chip, a requirement that has been taken into account in the design.
  - Having designed the original FEB, which exceeds the coherent noise specification, we understand how to apply techniques at the board level to combat coherent noise effects, and plan to employ similar techniques in the FEB2 design. The 32-channel Slice Testboard will be a critical step in validating the board design and performance. There is a risk in the Risk Register (RN-6-4-2-2) that system performance issues, including coherent noise, could require additional changes to the FEB2 design.

Question 18 (LAr)

- In the interest of high visibility, how well is the US contribution to the lpGPT working. Is the jitter measurement enough to demonstrate that the US contribution is solid?
  - The US contribution to the lpGBT, including the phase aligner at the ePort inputs, is working well. This has been extensively validated in the tests at CERN. Recent radiation results have identified some issues with parts of the lpGBT design that will be addressed in the next iteration. The jitter measurement does not directly validate the US blocks, but is an important performance mark for the application of the lpGBT on the FEB2.

Question 19 (LAr)

- Have any tests of the performance of the off-the-shelf ADC that will be used if the 65nm fails been done? If so, can you describe these tests? What are the considerations that go into deciding whether or not to prototype at FEB2 with this chip? Has it been radiation tested?
  - Yes, we made extensive performance tests of the COTS ADC with injected LAr pulseshapes, and successfully demonstrated that the required energy and time resolution could be achieved. The device was radiation-qualified to HL-LHC levels as part of evaluating it as a candidate for the Phase 1 LTDB.
  - The custom 65 nm ADC development is the baseline solution. The COTS ADC is a backup, should the custom ADC fail to meet the specifications. Given the progress so far, we feel this possibility is not very likely, but retain the risk for now until test results are available of COLUTAv3, which will have the full channel density and on-chip functionality. The decision will be made at PDR, scheduled for December 2019.

Question 20 (Tile)

- In your documentation, you describe simplifications of the Tile ELMB2 motherboard. Can you explain what allowed this simplification and if it has any impact on the performance of the circuits?
  - This is related to the single brick control. The simplification is the decision to use a tri-state voltage control instead of counters or registers that turn on/off individual bricks. This reduces the number of components needed. There is no impact on the performance of the circuits.

Question 21 (Tile)

- What is the magnetic field at that location in ATLAS? Have the components been tested in that field?
  - For the fingers, we have up to 20 Gauss in normal fingers and up to 50 Gauss in some special fingers.
    - The transformers we use on the bricks are rated for up to 100 Gauss. But we also have a few inductors in the bricks. We are planning to perform magnetic field tests on the entire LV box to make sure that brick performance would not be affected under magnetic field.
    - The Main Board has already been tested in magnetic fields 20 times higher than the ambient field.

Questions 22 and 23 (Tile)

- Does the long board get extensive temperature cycling? Are their hidden vias?
  - The Main Boards are burned-in at 60C for 5 days as part of the production process. We have not experienced any temperature related failures on the 20 boards produced so far. Yes, there are hidden vias on this 14-layer board.
- Are there really fuses in the low voltage power supply?
  - We do not have any fuses on the bricks themselves. But the 200 V input is being distributed to the bricks through a fuse board in the LV Box, with individual fuses for each brick. There are also some fuses on the output voltage of the bricks, which are located on the mainboards. These are supposed to disconnect a faulty mainboard side from one brick.

Risk Probability

- A complex issue; most people use categories
  - We use:

Rating	Description	Interval
Very Low	May occur in rare circumstances	0-10%
Low	Could only occur some time	11-25%
Moderate Low	Might occur some time	26-50%
Moderate	More likely to occur than not	51-75%
High	Is likely to occur	75-100%

  - For the moderate and high categories, mitigation that brings it down to moderate low or lower is required