DUNE PDR: Cold Electronics WIB and System

Questions / comments received from the committee on Wednesday 25 March. The answers from the TPC electronics consortium are clearly marked in *italic*.

The questions are grouped in categories following the way the presentations are listed on the agenda page.

Grounding & Shielding

The presented scheme looks good to the committee and it has mainly been validated in protoDUNE. A question arose concerning the move of the DAQ hardware on the cryo mezzanine. Although this equipment is referenced to the facility ground system and should not interfere with the front-end electronics the team should comment on this.

All the connections between the TPC electronics and the DAQ will be made via optical fibers and there should be no direct electrical connection between the detector mezzanine and the DAQ hardware on the cryogenic mezzanine. As in ProtoDUNE, an impedance monitor will be used to continuously measure the impedance between detector ground and building ground. We do not expect the DAQ hardware to generate noise on the TPC electronics, given that both the Warm Interface Electronics Crates (WIECs) and the chassis of the PCs used for the DAQ (and even the hut housing the DAQ hardware) act as Faraday cages. Still, we think it is valuable to make tests at ProtoDUNE to ensure that a PC running close to a WIEC does not generate noise on the TPC electronics.

WIEC

The design is mainly unchanged wrt to the protoDUNE design. Some problems with the QSFP+ transceivers were observed during protoDUNE, likely due to loss of connection because of vibrations. Has this been analysed and are there some design improvements foreseen?

Thanks for reminding us of this problem. We had this problem only during the tests in the cold box, and only when the cold box was at room temperature. We did not observe this problem during the normal ProtoDUNE data taking. We will nevertheless revisit whether there are any improvements we can make on the mechanics of the WIEC and of the WIB. We note that in the long term the QSFP+ will be replaced by two SFP+ to simplify the cable plant.

How accessible are the warm flanges for regular maintenance of the air filters?

The warm flanges are easily accessible on the top of the cryostat, and we do not think that the maintenance of the air filters (i.e. replacing the old filter with a new one) will require more than a few minutes of down time for each WIEC. The following picture shows the spool piece (the cross) attached to the TPC electronics cryostat penetration (in purple) with the two WIECs and the PD flange (only one of the WIECs is clearly visible), installed closely to a photon detector system mini-rack (in blue). The ribs of the cryostat support structure are shown in red and the false flooring installed on the top of cryostat is shown in light brown. The WIECs are at a height

of ~20 cms above the false flooring, which means that the personnel replacing the air filters will have to kneel or bend over during the operation. The area on top of the cryostat will be accessible without any restriction during the operation of the detector (we actually have to plan for a fence around the TPC electronics cryostat penetrations and the photon detector mini-racks to prevent people from running into them). It should be noted that in this drawing the filter is at the bottom of the WIEC, with the air flowing upward. Given the small distance between the WIEC and the false floor it has been decided to put the air filter above the WIEC and have the air flowing downward. This change has yet to be propagated in the CAD models

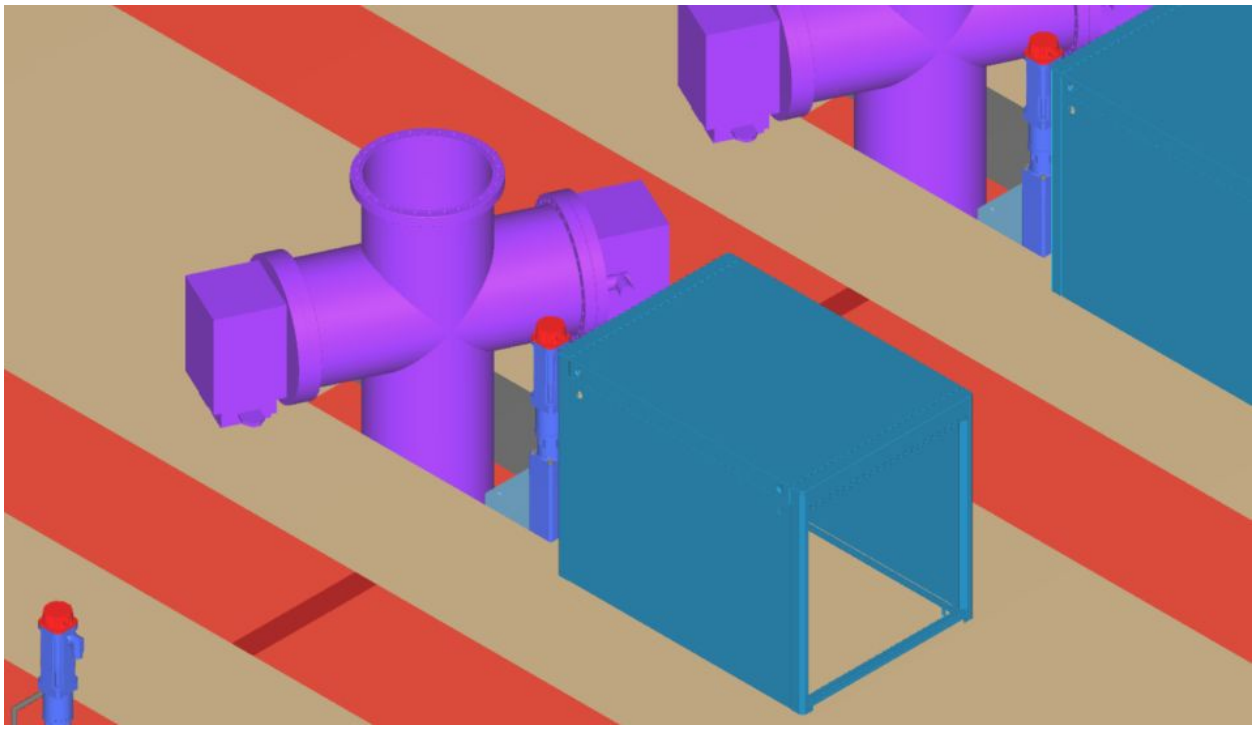

Is there a plan to have a complete working feed-through assembly including WIEC (aside from the ProtoDUNE detector) for testing and long-term maintenance?

For firmware development purposes we will have multiple WIECs operating. These may not be fully populated (i.e. have multiple WIBs), but the goal is to have one of them operating in each institution involved with the WIB/PTC firmware development. Eventually, during operations, we could think about having one of these test stands on the surface at SURF, in such a way that the personnel responsible for operation on site will be able to exercise the maintenance tasks there, prior to performing them on the systems used for data taking.

PTC

The proposed modification of the PTC makes sense. However the implementation and the relationship with the current WIB prototype is not clear. Can we have the new PTC for protoDUNE run2 with the intermediate WIB prototype?

While the idea of having an FPGA inside the PTC and of having the PTC communicate via I2C with the five WIBs in the same crate is very recent, we plan to use existing lines on the backplane of the WIEC (PTB) for this communication. Ideas on how to implement the communication between the WIBs and the PTC are shown in the following transparency prepared by Jack Fried. We think that the design of a new PTC which can communicate with the WIB can be achieved on the timescale of ProtoDUNE run 2 (taking also into account that we may start operating with an old PTC and replace the PTC during data taking).

- Each WIB connection to the PTB Include
  - Fixed 4 bit slot address
    - Used for geographical addressing
  - Two differential point to point connections to the PTC
    - Used to pass timing system signals
  - Six differential bussed signals OR 12 single ended signals.
    - For ProtoDUNE they were used as single ended signals
      - 4 lines used as a crate address for geographical address
      - 8 lines used as a priority encoder to select which WIB is selected to use the PTC SFP TX link
    - For DUNE they can be used as I2C links so that the smart PTC can communicate with the WIB
      - Slot address will be used for WIB I2C address
      - 2 lines used as I2C (PTC master)
      - 5 lines can be used as alarms from WIB to PTC where the slot address will either tri-state the line or use it for the alarm
      - Crate address will be delivered over I2C
      - Priority encoder function now controlled by smart PTC/I2C
      - 5 spare lines TBD

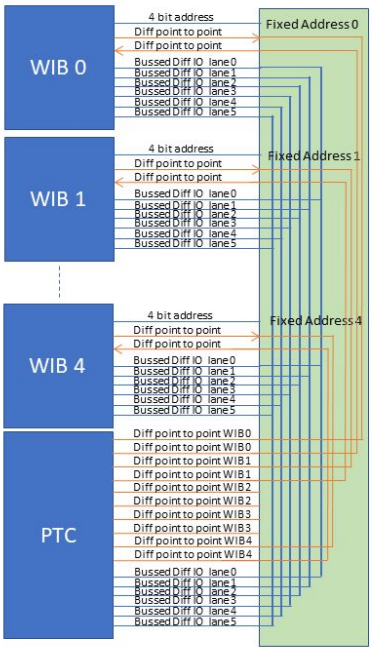

The PTC and its role wrt SC is quite confusing. What will be controlled/monitored via the WIBs and what via the PTC?

The addition of the FPGA on the PTC opens up multiple possibilities. At the very minimum we could turn on/off the power to individual WIBs as well as monitor the voltage and the current sent to each one of them. In the current PTC the power to individual WIBs can be enabled / disabled via a dip switch, and there is no monitoring of voltages / currents. It would also be appropriate to add the possibility of issuing a reset command to each individual WIB via the I2C path through the PTC, in case the communication between the WIB and SC is stuck. At the opposite end we could decide that all the slow control connections between SC and the WIB are replaced by a connection to the PTC and then by I2C communication between the PTC and the WIB. We do not think that we have to fully constrain now how the system will behave, and it would be good to keep some flexibility in the design of the firmware for both the WIB and the PTC.

Are requirement and specifications documents available or in preparation for the PTC?

The decision to add an FPGA to the PTC is very recent and we have not yet fully developed a requirement document for the PTC (hardware and firmware). We expect to have this document completed before the end of June.

WIB

The WIB requirements requirement document and the firmware/software document are extremely vague and it's not clear how real specifications will be defined/derived from it. For instance, the allowed power dissipation, the allowed weight should be clearly defined. The 1% downtime required should translate in allowed data transmission errors, allowed loss of synchronisation and time to recover.

We agree with this suggestion and we will try to improve both the hardware and the firmware/software requirement documents. Many of the comments of the committee on the WIB requirements document highlight the fact that these documents are preliminary and need to go through a joint review by the TPC electronics and DAQ consortia. The firmware/software document in particular is a planned extension of the interface document between the two consortia. We have tried to conform to the interface document as much as possible (the two documents were being prepared at the same time, and in some cases the DAQ interface document was finalized after some sections of the WIB interface document were written and made available to the review committee). We note that for both a fully detailed description of the planned functionality and requirements of both the CCM and SC is missing at this time. We expect to continue updating the WIB requirement documents as the design of the CCM and SC evolves, to ensure that the systems are fully compatible. Despite the few remaining uncertainties in the interface between the two systems, we think that we can start with some parts of the design of the WIB firmware/software.

It is true that neither the allowed power dissipation nor the allowed weight values are specified. For the total power dissipation preliminary numbers based on ProtoDUNE have been used to define the overall air cooling requirements in the cavern. The allowed weight has an impact on the thickness of the tubing used for the TPC electronics spool pieces (the "crosses"). This has been studied for ProtoDUNE and we need to document those studies in the requirement document.

Several points are unclear wrt to DAQ/SC interface. Detailed comments on these documents are given below.

As the specifications are not yet fully defined it's feared that this opens a path to have multiple design runs. Already the intermediate WIB design is not fully compatible with the new TPC.

As mentioned above the design of the intermediate WIB is compatible with both the current (ProtoDUNE) PTC and with an updated PTC with an added FPGA. While we think that the intermediate WIB is a significant step forward relative to the WIB used for ProtoDUNE, we also do not think that it is the final WIB prototype.

Certainly not wasted work but it might be useful to finalise the specs in a very detailed way before rushing into modifications. What are the plans for this?

We are not rushing to finalize the hardware requirement document and we are not rushing to make design changes to the WIB at this point. At the very least we need to complete system tests with the new WIBs and new FEMBs, including optimizing the power distribution between the WIB and the FEMBs, taking into account the constraints on the number of power and signal connections. We should also incorporate in the requirement document lessons learned from the operation of the intermediate WIB. We do not expect to finalize the hardware requirement document until the Fall of 2020. We think that there will be sufficient time to make changes in the design, even if it is decided that a new version of the WIB is required for ProtoDUNE Run 2, since we do not expect the design changes will be major.

The firmware is to be designed by several institutes and the way this will be efficiently done is not yet fully defined. How will the different blocks be defined (including detailed interfaces between blocks),
An initial definition of the various blocks and their interfaces, including the separation between firmware and software is included in the slides and needs to be reflected in the WIB requirements document. We have now identified an engineer with significant firmware expertise, Alex Madorsky from Florida, to lead the firmware development for the WIB. Alex has expertise from CDF and CMS, where he has been working on the L1 forward muon trigger. One of his first tasks will be to review the definition of the various blocks and of their interfaces.

What will be the design flow, what will be the verification flow, etc.? Also, the separation between firmware and software is not fully defined yet. Can we get clarifications on this?
General criteria for the separation of tasks between firmware and software running in the Zynq CPU are discussed in the presentation of Josh Klein. Some details are given in the discussion of the individual firmware blocks, but we do expect that there will be some changes as the design of the firmware and of the software evolves. Alex Madorsky will decide what design approach to follow and how to perform all firmware verifications.

Is there a formal assignment for a System Engineer for Firmware? To what extent are bit-wise simulations of firmware required?
We now have a system engineer for the WIB, as mentioned above. We think that there are blocks of firmware where bit wise simulation of firmware is required. Clearly the FEMB block, that handles all the communication with the FEMBs requires that, and the same is true for the block where the frames sent to the DAQ back-end are built.

It is not clear whether the new firmware includes enough debugging tools which proved to be missing for protoDUNE (i.e. needed to change firmware for running low level diagnostics).
One of the firmware requirements (mentioned in the slides) is that all of the functionality, including the diagnostic and debugging capability, can be done without changing firmware. The Debug/Diagnostic block in the firmware will be the place for these functional pieces.

The qualification of the high speed links is to be detailed. In the slides the links were tested up to a BER of 10^-13. Is this enough? How does it translate in data loss and hence down time?

*For each time sample (there are 1.95M samples per second, 1.7*10^11 samples per day), each WIB transmits 7,616 bits to FELIX. This corresponds to a total of ~1*10^18 bits per day. Assuming a BER of 10^-12, a total of 10^6 samples would be corrupted in the entire detector, i.e. a loss of 6*10-6 of the samples, which translates in a downtime of 0.44%. The BER of 10^-13 was used during the quality control of the links for ProtoDUNE after installation (this corresponds to testing the links for 20 minutes). When designing all the links in the system we are aiming for numbers that are much smaller than this. For the links between the FEMBs and the WIBs tests have been performed to the 10^-15 level (those links are much slower, and the tests we performed in very pessimistic conditions). For the links between the WIBs and FELIX we can increase without any problem the testing period during the design phase to achieve a BER of 10^-15. We will try to do the same during the final installation at SURF (we may start using some of the links and then decide to replace transceivers / fibers if we notice problems with the BER of specific links).*

Why are so many 125 MHz oscillators required on the WIB in Jack's talk. Why not use a common oscillator and distribute it through the Si5344 (or larger Si5345)?
Most of the oscillators are not required for normal WIB operation. They were added as backups for this initial version of the WIB.

Today there are two options for the front-end (3-ASIC solution or CRYO) and the interfaces to the WIB are not the same for the 2 solutions. The requirement documents should reflect these differences and there should be a detailed description of the interfaces for the two solutions (data format, protocol, control of synchronisation, resynchronisation process). When will this happen? Is there a priority defined between the two solutions (I.e. is the firmware developed in parallel for the 2 solutions or is there one before the other)? Are there enough resources to be ready for the two solutions in due time? It is understood that the hardware will be the same for the 2 solutions. Is that true?
MV: when we prepared the WIB hardware / firmware requirement documents and the slides for the review we were making the assumption that WIB should be compatible from the hardware point of view with FEMBs housing the 3 ASICs solution (i.e. communication via COLDATA) and FEMBs with the CRYO solution. The powering scheme of the current intermediate WIB prototype is compatible with both types of FEMBs. From the firmware point of view we expect to need two different versions of firmware and we have tried to divide the firmware blocks in a way that minimizes the amount of firmware that differs between the two solutions. We think we can limit the differences between the two versions of firmware to the block that communicates with the FEMB. This block controls both the communication from the WIB to the FEMB, which uses I2C in the case of COLDATA and SACI in the case of CRYO, and the decoding of the data from the FEMB. This block of firmware needs to be able to read back asynchronous data with 8b/10b encoding (COLDATA) and synchronous data with 12b/14b encoding (CRYO). The FEMB block also needs to handle the different synchronization mechanism for the two types of FEMBs.
We plan to develop the COLDATA version of the firmware first, followed by the version for CRYO. When developing the CRYO version of the firmware we plan to reuse blocks of firmware library (communication with SACI, 12b/14b decoder) that have already been developed by the

SLAC group. We should also note that it is not clear that we will pursue the CRYO option for DUNE beyond this Summer.

In Shanshan's presentation a quite high number of WIBs and PTC had to be reworked when. Coming back from the assembly house. Have the reasons for this been analysed and understood?
Three out of the 35 WIBs required rework. The first WIB had an equalizer chip installed in the wrong direction. The QSFP on the second WIB did not function properly, which was fixed by reflowing the QSFP socket. The third WIB, the connector between WIB and flange board was damaged during the installation (improper alignment and excessive force). There are two types of PTC cards, one using the Linear LTM8064 DC-DC modules, the other one using Vicor PI3456 modules. Both PTC cards were designed by UC Davis in cooperation with BNL. The PTC cards were assembled in an assembly house chosen by UC Davis which is not the assembly house that we (BNL) usually work with. All ten PTC cards with the LTM8064 pass QC test, no failure. Seven out of the ten PTC cards with PI3456 didn't pass the BNL reception test. We figured out that there were one or more broken PI3456 modules in these seven boards but we didn't replace the module since PTC cards with LTM8064 were sufficient for ProtoDUNE. The reason for the failure of the PI3456 was not investigated. It is clear that in the long term we need to understand any failures during the assembly of the WIBs and PTCs and that the overall yield has to be very close to 1.

"Formal" comments on WIB and FW requirements documents:
- It would be useful to use the typical requirements language using "shall", "should", "may" depending on the importance/priority of requirements.
We agree with the committee and we will edit the requirement documents to use a more formal language. In this first version of the requirements documents we focused mostly in capturing all the aspects of the design.

- Requirements should have a unique identifier, that will be usable at specification/design level to indicate which feature satisfies (or contributes to satisfying) which requirements.
We agree with this suggestion of the committee. We will be doing this for a future version of the document.

- There is a mixture of DAQ, SC and DAQ/SC used in the document, without a clear logic. DAQ/SC is the consortium and should be used every time a generic reference is made to the components covered by the consortium. DAQ is specifically the data acquisition part (including CCM) and SC is specifically the slow control subsystem.
We will edit the requirements document as suggested by the committee.

Specific comments on WIB requirements doc:
2) Functional requirements:
5. Provide real-time diagnostic data… it is not clear what is meant here

We think that describing the requirement is "provide diagnostic data continuously". This can be achieved in two ways: first of all for every data frame sent from the WIB to FELIX there will be flags indicating whether the data is corrupted. This could include flags indicating errors in the FEMB timestamps, flags indicating data corruption in the transmission between the FEMBs and the WIB. The WIB should have a way of informing the CCM that it does require actions to be taken to correct error status, after failing to perform internal error recovery. The CCM should have a way of requesting the internal status of the WIB and of the FEMBs in case of problems. All the procedures for error handling need to be determined jointly by the TPC electronics and DAQ/SC consortia.

9. Should the data readout be added here?
In the WIB requirements document we have an item 4. "Transmit data to the DAQ system". In the WIB firmware requirement document we have an item 4. "Send data to DAQ without corruption but with checks on the local data integrity" and an item 7. "Pass the reformatted data continuously to the DAQ over the high-speed links". The two documents need to be aligned.

Fig 2. To be correct, the yellow line from CCM to Timing should be green (ethernet connection)
Yes, in fact those details are not entirely relevant to the WIB requirements (the Timing System talks to the WIB via its own endpoint and its own clock-encoded commands). We will correct the figure.

4) Interlocks
"The WIB must provide hardware interlock(s) …" which HW interlocks?
Is the DDSS only acting through the PTC?
"Multi-level interlock system": can you add the logics? Which interlock acts when? Which has precedence? What is the hierarchy?
We need to revisit the text in Section 4 of the hardware requirements document. The current text confuses interlocks, hardware protections, and software monitoring. The WIB needs to provide input to the DUNE detector safety system (mostly information on the status of the low voltage power to the FEMBs) which can then be used to form hardware interlock signals to be sent to the bias voltage supplies. It can also receive inhibit signals from the DDSS. In both cases the communications between the WIB and the DDSS is handled through the PTC. There are no direct links between the WIBs and the DDSS, to reduce the number of optical to electric conversion units in the DDSS. The hierarchy and the precedence of the interlock system are defined in the action matrix of the DDSS, that at the moment has not been defined. In the design of the interface between the TPC electronics and the DDSS we have focused, for the moment, on ensuring that we can provide as much information as possible from the detector to the DDSS and that we can apply inhibit signals on individual channels of the detector.
Even with a fully functional DDSS we still plan to add fuses to the power lines for the FEMBs. These will be the highest priority protection system for the FEMB. Next comes the DDSS, and at lower priority monitoring will allow us to raise alarms on the status of the WIB/FEMBs.

5) Mechanics

This section is very vague. What is “low weight” what is a small enough form factor? Is there no mechanical requirement on the WIB? (Insertion, extraction, locking, resilience to vibration, ….)
We agree that this section is vague and needs to be expanded. The requirements need to be clearer. The weight has to be compatible with being supported from the TPC electronics spool piece (the “crosses”). The form factor needs to be consistent with not interfering with other elements of the detector on the top of the cryostat, including the false flooring. We need to add mechanical requirements on the WIB as you suggest.

7) Hardware architecture
What is the requirement? What is the aim? Here a hardware implementation choice is described not the requirement.
The requirement is to have a programmable system that can handle the data throughput from 4 FEMBs (~15 Gbit/s), check for data integrity and reformat the data for transmission to FELIX. This is most easily implemented with one (or more) FPGAs on the board. The choice of using a Xilinx Ultrascale+ with a Zynq CPU is an implementation choice that was taken because of cost reasons and because it adds flexibility in the system, allowing the migration of some tasks from firmware to software. We need to ensure that both the hardware requirement document and the firmware/software requirement document are consistent on this point.

8) Environment
Is it possible to be more quantitative? What is the power budget? What is the requirement on air quality?
We will try to be more quantitative and have an estimate of the power dissipation (mostly in the FPGA, in the transceivers, and in the converters) that we can deal with on the WIB. We can only allow for air cooling in the WIEC. The problem with setting an air quality requirement inside the WIEC is that we do not have a good estimate of what will be the air quality inside the experimental cavern at SURF and how it will evolve with time. We do not know whether there will be chemical components in the dust that may affect the lifetime of the WIBs and PTCs, and what is the expected distribution for the size of the dust particles. We have to admit that we are trying to derive requirements for a currently existing design that has started with the choice of a cooling system, and that has been retrofitted to use the best possible filter compatible with keeping the temperature on the WIB in a good range for operation.

9) Uptime
These requirements need to be quantified. What are acceptable configuration times? What means “quickly”? What is a “low hardware failure rate”? Etc
This is a fair point, but we think that strictly speaking the “requirement” is first and foremost that configuration time will be short enough to satisfy the DUNE exposure requirements. Getting more quantitative will require use cases: do we mean cold start configuration, or configuration of one new WIB that just got replaced, or configuration of one WIB because we’re debugging it…?

Is hot-swapping of WIBs in a WIEC going to be supported/needed to satisfy the uptime requirements?

At the moment we do not have any plans to support hot-swapping of a WIB. This is not consistent with the planned design of the inputs from the WIBs to the DUNE detector safety system.

11) Data integrity

CRC checking is planned. What else? What is acceptable/target BER? Specify that data integrity of WIB frames must be guaranteed even if one of the FEMBs is failing.

Additional details are in the firmware document. In particular, we plan to continue providing data from the other FEMBs and guarantee the integrity of the WIB frame, even when one of the FEMBs has a problem. We do need to explicitly mention this case in the requirements document.

12) Latency

Indeed latency is not an issue. In order to put a reasonable target aiming at <1ms is sufficient

13) Communication latency

It is not really clear what is meant here. Again a quantitative number is needed.

Communication latency means the time between the receipt of a command from CCM or SC and its execution and acknowledgment. Ultimately this should be the same as any DAQ subsystem that interfaces to the CCM. If SC is used to power off the board in case of emergencies, then there does need to be a distinct requirement there.

14) Geographical addressing

Indeed each WIB must have a unique identifier and the geographical addressing is a good proposal. Nevertheless the downstream connection should not be part of the identifier, i.e. you don't care which FELIX you are connected to. Right?

Each WIB will have a unique identifier given by an hardware address on the PTC plus the geographical address of the WIB in the WIEC. Each FEMB will also have a unique hardware identifier hard-wires in COLDATA (or in CRYO) with E-fuses. The hardware identifiers in the ASICs will be established at the time of testing the FEMBs, while the hardware identifiers in the PTC and the WIBs will be established during installation.These hardware identifiers will be used during commissioning to verify that all the data transmission links have been installed correctly. While it is true that the WIB does not care what FELIX card it is connected to, for online and offline purposes we do need to have all the WIBs from a single APA to be connected to the same FELIX card.

15) Monitoring

Really real-time readout of temperatures and voltages? What is realtime here? There shall be a clear separation of monitoring data for SC and for DAQ. Anything that needs to be monitored at all times irrespective of data taking shall go through SC. Quantities that are relevant for the data taking process (data corruption counters, synchronisation status, …)shall go to DAQ.

Here real-time means continuous, compatible with the rate that can be handled by SC. We will modify the language in the requirements document.

16) Failure and recovery
While the recovery procedures may be infinitely complex, the basic functions to be carried out by the WIB are always the same. Those should be part of the requirements. E.g.: flag that a FEMB is failing, disable a FEMB from data taking, resynchronise a FEMB, re-include a FEMB into data taking, ….
These details sound like they go deeper than requirements. They are part of a table of use cases. We agree that they have not been worked out yet.

18) Installation and integraty
Is “remote connectivity” the 1G ethernet link? No reason to be vague, if to be vague if you know it.
We agree that we need to specify “via the 1 Gbit/s Ethernet link”

19.3) Power
Noise should be subdominant… at what level?
We are requesting that all the noise sources that are not internal to the front-end amplifier are at a level of 100 e-. This includes the noise from voltage ripples on the power rail of the front-end amplifier. It is difficult to specify a value for the voltage ripple on the power line going from the WIB to the FEMB, because that depends on the power supply rejection ratio of the LDO regulators on the FEMB. The text of the requirement should be made more explicit, saying that voltage ripples on the power line to the FEMB should not contribute more than 100 e- to the overall noise of each readout channel.

21.2.1) Option open for a second 1G ethernet link. With 750 WIBs this is not really a detail….
Had this not been already discarded?
Yes the option of a second 1 Gbit/s Ethernet link has been discarded. The requirements document needs to be updated to reflect this.

23) Cost
This requirement means nothing.
We agree and we will remove this requirement.

24) Timeline
Can requirements on verification/validation/integration be added in the text and then also in the timeline?
We will do this for a future version of the requirements document.

Specific comments on WIB FW requirements:
1.1) Background
- what monitoring data goes from WIB to PTC and what from WIB directly to SC?

As discussed below, the addition of an FPGA on the PTC and the addition of the I2C communication between the WIBs and the PTC opens up the possibility of having some of the WIB/SC communications going through the PTC. For the moment, the default is to have all the WIB/SC communications going directly through the 1 Gbit/s Ethernet connection on the WIB, but we will build in the system the possibility of re-routing some of these communications through the PTC. Information required by the DUNE detector safety system to build interlocks will be transmitted to the PTC. All the communications with the CCM will go exclusively through the 1 Gbit/s Ethernet connection on the WIB.

2) Functional requirements
8. Real-time diagnostic data: not clear what is meant, really. Real-time in what sense?
Real time should be restated as "continuous"---these are things that will always be monitored and in some cases cause action to be taken.

3.1.1 The use of DQM in this context is misleading. Is it possible to find a better acronym?
We agree with this suggestion. We should leave DQM for the online/offline data quality monitoring, and use instead Data Integrity Monitoring (DIM). We would also like to point out that we could perform some operations usually meant for the online DQM (calculating a rolling mean and width of the pedestal) in the Zynq.

3.1.3) If requirements cannot be formulated how can you move to design phase?
In writing the WIB firmware requirements document and preparing the slides for the review we have made the assumption that the WIB should be compatible both with the 3 ASIC solution (i.e. COLDATA) and with the CRYO solution. At the firmware level this requires that the FEMB block be able to communicate with the FEMB both via I2C (COLDATA) and via SACI (CRYO) and then be able to read back asynchronous data with 8b/10b encoding (COLDATA) and synchronous data with 12b/14b encoding (CRYO). We believe that we can proceed to the firmware design by isolating the blocks of firmware (possibly only the FEMB block) that need to be different between the two solutions. In the meantime we are also reconsidering whether to pursue the CRYO option at all, which would simplify the design.

3.2.3) If a 62.5 MHz clock is not sufficient for cable length compensation, is it relevant at all to make this compensation, given the 2 MHz ADC sampling rate?
We have a requirement (SP-ELEC-3) that all the wires in the TPC are sampled simultaneously within a window of 50 ns, with a goal of reaching 10 ns. The difference in length between the cables for the upper and lower APAs is approximately three clock ticks in units of the 62.5 MHz clock (16 ns). If we can use a 125 MHz clock or a 250 MHz clock, we can improve this number to 8 or 4 ns respectively. This would allow us to fully exploit the precision of the timing system (which has a target of a 1 ns precision).

3.2.4) lost?

No text was lost. A section was deleted on purpose and the numbering of the following section was not changed to reflect this.

4) For clarification: the commands are not limited to initialisation and synchronisation, though it may be that the WIB doesn't need any others.
It is not anticipated that more will be needed by the WIB, although that could certainly change.

5) The communication interface with CCM depends on what is available on the WIB. It would be important to include in the req. document a clear interface of SW and FW and then, eventually to add a SW requirements section.
We agree with this but the FW/SW line will evolve as the design does. Both have to satisfy the high-level requirements (e.g., uptime, functionality, etc.), independent of whether the implementation is via FW or SW.

5.2) I think that this section indicates that the interface of CCM (which is external to the WIB), WIB software and WIB firmware is unclear.
This is a reflection of the current status of all the protocols for the communication between DAQ/SC and the various detector components. We will revisit the list of requirements once these protocols are better defined. From the point of view of the WIB development, we think that the interface with the CCM will be entirely in software.

5.3) CCM and SC are distinct systems with a different scope. It is not the case that one is faster than the other. Any parameter that needs to be continuously monitored irrespective of the data taking status shall go through SC.
This is a useful point, and we will make that change to reflect this.

6) A soft reset is mentioned here. What are the requirements on different types of reset at the WIB level? If the SC does the soft reset, why does the CCM do the FW reloading? There seems to be some inconsistence in the separation of duties.
We agree with your comment, it may be better for SC to be responsible for loading new versions of the firmware on the WIB.

8) It is mentioned that there will be no buffering at the WIB. At NP04, the problem on one of the FEMBs would require some buffering, in order to align data an do correct event building to construct the WIB frames. How will this be handled?
We certainly have the capability for a minimal amount of buffering (a few frames) inside the FPGA. We do not plan to use this feature extensively, unless needed.

9) There is no reason to send commands with timestamps in the future. The DAQ(CCM) will send a command to start the calibration. As soon as the calibration will start data will be marked as such. Also, the data unit that is received at the DAQ side are WIB frames, not single FEMBs. Wasn't it agreed that pulser runs would be done at overall APA level?

We will rewrite this requirement. We agree that this is not a feature that is required for operating the WIB in DUNE.

10.1) It is not clear what is meant with “WIB-only path”. Could you elaborate? In no case should the WIB care about the data buffering scheme downstream…
There is not going to be a WIB-only path. For specific configurations of the DAQ, the WIB will have the capability of tagging event frames with specific bits that will instruct the data selection system to write the data to a specific stream. This is a feature, agreed with the DAQ consortium, that will allow us to write the data from the TPC electronics in a separate data stream, either for internal calibrations of the TPC electronics or for debugging purposes.

10.2) This is not agreed. What was said is that this may be a possibility if a clear use-case is identified.
We agree and we will change this to be an “optional” or “possible” feature.

10.3) Can this data stream path be elaborated on? We should be careful, in the integrated system, to not have any high bandwidth data stream saturating the 1G link which is used for control and monitoring.
The WIB will have the possibility of recording some data in memory. This data could then be sent to slow control (with a throttle on the maximum speed) for specific debugging purposes. This feature is not intended for normal data taking, and is expected to be used mainly during the testing of ASICs and FEMBs, and during the commissioning of the experiment.

11) Is there really an intention of having high level data quality monitoring inside the WIB?
We think that we have resources inside the WIB (in the Zynq) to keep running averages of the pedestals and of their widths. This is the only data quality monitoring that we are considering for the WIB, to complement what is going to be done with the monitoring data sent from the FELIX to the online DAQ. The WIB is the only place that will have continuous access to the entire raw data stream.

13) Not clear what information goes through the 1G connection on the WIBs and what goes through PTC
As discussed below, the addition of an FPGA on the PTC and the addition of the I2C communication between the WIBs and the PTC opens up the possibility of having some of the WIB/SC communications going through the PTC. For the moment, the default is to have all the WIB/SC communications going directly through the 1 Gbit/s Ethernet connection on the WIB, but we will build in the system the possibility of re-routing some of these communications through the PTC. All the communications with the CCM will go exclusively through the 1 Gbit/s Ethernet connection on the WIB.

Power supplies

The proposed scheme makes sense. The statement "Both the WIENER PL506 and the ISEG HV modules that are in use in ProtoDUNE have ripple voltages that exceed the requirements for DUNE" should be documented with measurements. It is said that the last LDO reduces the input ripple at the level of 10^-4 which should be enough to reach the few μV ripple needed at the input of the front-end chips. Is it planned to make measurements (or has it been done) with a full chain including the different converters and LDOs and the filters and to give the ripple vs frequency at each stage?

We have reached the same conclusion as the review committee. We will measure the ripple attenuation factor provided by the LDOs on the WIBs and on the FEMBs, and also for the filter cards used for the bias voltage (this measurement will actually be done in ProtoDUNE, to check at what point we start seeing noise if we increase the ripple on the bias voltage).

The way the interlock is implemented on the power supplies should be documented. Which signal, where does it come from, how is it transmitted?

For DUNE we plan to use a version of the PL506 low voltage power supply and a version of the ISEG bias voltage modules that feature the possibility of using an individual inhibit signal on each channel (this is an additional option for these supplies). Both the PL506 and the ISEG require TTL signals (i.e. in the range 0-5 V) for the inhibits. These signals will come from level translation cards that will convert the 0-24 V range used by a PLC system (industrial standard) to the 0-5 V required for TTL. Assuming that the PLC system has digital output modules with 8 or 16 channels, the translation boards are also required to ensure the appropriate mapping of the connectors. This scheme allows us to use the full power of a PLC system to form interlocks, possibly including inputs from other systems, including the cryogenics. In order for this scheme to work, the signals from the WIEC needs to be transmitted optically to the DUNE detector safety system, translated into electrical signals, and eventually decoded prior to being fed into the action matrix of the PLC system. While this system may appear to be complicated, the only custom boards that are required are the translation boards used for the conversion from the 0-24 V range of the PLC to the 0-5 V required by the PL506 and the ISEG modules. All other components would be commercially available off-the shelf components. We should note that at this point this is far from being a conceptual design. In order to reach that stage we do need to have more information on the actual implementation of the DUNE detector safety system.

The power dissipated in the warm cable trays is quite high. Is there a need for cooling?

In the slides we say "O(1 kW) of power dissipated in the warm cables that run below the false floor on top of the cryostat". This is the power dissipated over the entire surface of the cryostat (i.e. 150 power cables), not the power dissipated by a single cable. Let's look at the math: the AWG10 power cables have a resistance of 3 mΩ per meter. Assuming 20 m of cables the total voltage drop on the power plus return would be 120 mV per A of current. The current estimate of the power required for 1 crate is 7.5 A, which gives a voltage drop of .9 V, and corresponds to a power dissipation per cable (i.e. per WIEC) of ~ 7W. The total power dissipated by all the cables providing power to the WIECs is obtained by multiplying this number by 150, which gives the estimate of O(1 kW). In obtaining this estimate we have overestimated the length of cables. There are 25 cable trays (one per row of APAs), and therefore the power in each cable trays is

~40W, distributed over a length in the range of 10-20m. While we do not think that there is any need for cooling of the cable trays that run under the false floor of the cryostat, we will be performing studies to check this assumption.

All the power lines (Bias, HV, LV to the front-end) are potential single point failures. Are there any redundancy plans?
We have to distinguish between the top of the cryostat and inside the cryostat. On the top of the cryostat we will be able to replace power supplies and cables if needed. We are making the assumption that a failure on the power supplies or on the bias voltage supplies will result in a downtime of up to 2-3 days, until a spare unit (available at SURF) can be installed. Components inside the cryostat cannot be accessed or replaced. Unfortunately, we also don't have sufficient room to increase the redundancy of the power or bias voltage distribution system. For the bias voltage lines it could be conceivable to add some redundancy, but for all other cables inside the liquid argon this is not possible. This stresses the importance of having reliable connections on the FEMBs and on the flange (this has been a sour point in ProtoDUNE, in particular for the signal cables, and we believe we have significantly improved the reliability of the connector on the FEMB). We have not yet decided on the final scheme for the power distribution between the WIB and the FEMBs, but it is likely that some redundancy could be built in the system. In the design presented in the TDR we argued for having a single LV being sent from the WIB to each FEMB, over multiple AWG20 wires. If a single wire was disconnected, this would lead to having the same current being transferred over a smaller number of wires, resulting in larger ohmic losses between the WIB and the FEMB. This could be compensated by increasing the voltage on the WIB, such that the final input voltage on the FEMB is unchanged.

BY: For the field cage termination bias voltages, there are failsafe terminations inside the cryostat such that in case of a bias cable connection failure, the field cage module is terminated to a voltage not too different from the optimal bias value to maintain near normal operation.

Interfaces to DCS, DDSS and interlocks
As admitted it is still a bit fuzzy. In particular the interlocks should be more clearly defined: which signal level, where are they formed, etc.
Many issues around the so called DDSS have not been understood by the committee. Several HW interlocks seem local to a single WIEC: thus the action is taken and it is sufficient to inform SC. Which external hardware signals would act on the WIBs/PTC? Also, why would the DDSS need an optical link for (ethernet) connection?
We think that there is some misunderstanding here. The DDSS system needs inputs from all the low voltage power supplies, the bias voltage supplies, the supplies for the heaters and the fans, but also from the WIBs (because the only way to know whether there is power on a FEMB is to check at the WIB level). There are 45 bias voltage modules in 5 MPOD crates (for a total of 450 channels of APA wires and 208 field cage termination electrodes), 25 low voltage power supplies (for a total of 150 channels), 25 power supplies for the fans (with a total of 150 channels, or a multiple thereof), and 25 power supplies for the heaters (again with a total of 150

channels). There are then 750 WIBs. All these channels can provide input to the DDSS and should be controllable by the DDSS. We do not want to use the existing Ethernet connections in the WIBs and on the power supplies for the communication with the DDSS. Also for the communications coming from the WIEC (either from the WIBs or the PTCs) we cannot use a copper connection because that violates the grounding rules of the experiment. There are appropriate interface modules between PLC systems and optical fibers, and therefore we could decide to have 750 optical connections between the WIBs and the DDSS. To reduce this number, we propose to concentrate the information at the PTC level reducing the number of connections by a factor 5. This could imply that we transmit all the information from the WIBs to the PTC and later to the DDSS, or that we already form interlocks (like "turn off the bias voltages for this APA") in the PTC and transmit that interlock to the DDSS. In the first case the information would be decoded in the DDSS and the interlock signal would be formed there. In the second case the interlock signal would be transmitted optically from the WIB to the DDSS where it would be converted to an electric signal. In both cases the electrical signal from the DDSS sent to the appropriate interface board to a bias voltage supply, and used to turn off the bias voltage (or prevent the bias voltage from being turned on). In both cases (interlock signal formed on the PTC and transmitted to the DDSS, or interlock signal formed inside the DDSS from data sent from the PTC) the communication never goes through the Ethernet interface of the WIB or of the PTC or of the DDSS. We still expect the central DDSS node to have an Ethernet interface for monitoring purposes, that would communicate with the Slow Controls (if something turns off the bias voltage supplies, you would want to know what it is and act remotely rather than going down in the cavern and look at a LED switch).

There is a statement that SC accepts only ethernet interfaces: this is a misunderstanding. With several other systems there are directly analog/digital signals treatment.
We agree that this is a little bit of a misunderstanding (the statement in the interface document with Slow Control was that systems with Ethernet interfaces were preferred). We agree that SC can receive analog and digital inputs, however we do not think that SC would like to interface with systems as complex as the low voltage supplies and the bias voltage supplies through interfaces that are different from Ethernet.

If WIBs are used both for SC and DAQ traffic, probably different levels of reset are needed, in order to keep the SC communication up and running at all times.
We agree. In principle we can have two Ethernet interfaces on the WIB and keep the traffic separate. Or we could decide to channel all SC communications through the PTC and then use the I2C communication between the PTC and the WIB. Or we can agree to have different levels of reset inside the WIB. This needs to be specified in detail in the design documents for the WIB and in the procedures for the communication between the CCM, SC, and the WIB.

While it is clear that DDSS and SC are themselves still unclear, we would expect a clear list of needs in terms of interlocks, hardware signals, software control/monitoring from the CE, irrespective of which system(s) will implement those.

We have already provided (backup slides) a list of all the quantities that are going to be monitored and controlled, or simply monitored from the SC. We have also made a list of all the channels that provide inputs to the DDSS. We also state clearly that the DDSS should be able to turn off or disable all the power in the detector (i.e. power to the 3,000 FEMBs, 450 bias voltages for the APA wires, bias voltages for 258 field cage termination electrodes, low voltage power to the PTCs and WIBs, low voltage power to heaters and fans). We understand most of these connections (see for example the need for an optical interface between the PTC and DDSS, as well as the need for level translation boards to control the low voltage power and the bias voltage supplies). The part that at this point we understand the least is the interface between the DDSS, SC and the power supplies for heaters and fans. These could be entirely under the control of DDSS or they could be controlled by SC and have interlocks with the DDSS. This depends on the actual design of DDSS and this is why we do need somebody from Technical coordination to take the lead on the design of DDSS, such that we can fully specify also the design for the controls of the fans and heaters supplies, in addition to making progress with the design of the interfaces between the DDSS and the TPC electronics supplies.

Concerning the fans control, it is currently based on the current drawn to decide whether they are in rotation or not. Is it really sufficient, i.e. cannot we have a broken fan drawing current? Either an air flow sensor or a movement sensor would be safer.
We agree that the current ProtoDUNE fans control is very limited. For ProtoDUNE there is a design (not implemented) of the control of the fans supplies that actually checks the fans speed. The design of DUNE will be an evolution of that design, and it could also include feedback from the temperature sensors on the WIBs and on the PTC (which would be transmitted via the DDSS).

Timeline and production
The production for protoDUNE run2 requires additional thinking. New PTC? Modification of the new WIB prototype to be compatible with the new PTC? How many boards of each types are needed for protoDUNE run2?
The idea of having a PTC with an FPGA dates from February, when we already had prototypes of the intermediate WIBs in hand. It turns out that we can implement the communication between the PTC and the WIB on a bus that was already available on the backplane of the warm interface electronics crate (the PTB). This bus was originally meant only to assign an address for each WIB (and therefore there is already a connection on the printed circuit board of the WIB between that bus and the FPGA), but it can be repurposed for the I2C communication between the WIBs and the PTC. Therefore, we will be able to test the communication between the new PTC and the current WIB as soon as a new PTC prototype is available. For ProtoDUNE run 2 we need a total of 20 WIBs (five per APA, with 4 APA total), and 4 PTCs (one per APA), plus spares. It is too early to decide that the current WIB prototype will be the one used for the ProtoDUNE run 2. It is possible that we may have an updated version (with a simpler power distribution and a reduced number of voltage regulators) by the time of the review. We will also include any modification based on lessons learned from testing the current version of the board.

The full review process requires clarification. This is a PDR however we don't have yet a full specification of the prototype to be delivered. It is assumed we'll have prototypes of the WIB and PTC tested in March 2021 so that an FDR could be made in view of the production of the boards for protoDUNE run2. Is this correct? How much of the prototypes will have been tested at that time (with the front-end)?
As mentioned above, we will be able to test a new "intelligent" PTC with the current WIB prototype. The current WIB prototype will be used with prototype FEMBs with the current version of the ASICs in Summer 2020. It will be used in system tests at the Fermilab ICEBERG test stand and in the cold box at CERN using the first prototype DUNE APA. The FDR for all the TPC electronics components is planned (are planned, we will again split this into multiple reviews) for Spring of 2021, before we fabricate the components for the ProtoDUNE run 2. We would like to remind the committee that the WIBs and the PTCs are elements that we can replace during the ProtoDUNE data taking. This would allow us to test newer prototypes, should they become available.

The PRR would then happen in the course of 2022 for a start of production end of 2022. Is there enough time for validation of the design on protoDUNE?
We do not need to have the PRR for the WIB to take place early in 2022. The WIBs are not needed at SURF until late in 2025. We could test additional WIB/PTC prototypes in ProtoDUNE in 2022, held the PRR in early 2023, and complete the production well ahead of the need-by date at SURF.