LHCC Questions & Comments on SND LoI (LHCC-I-037), 27 October 2020

Answers by the SND Collaboration, 13th November 2020

LHCC referees:
Francesca Di Lodovico, Katja Krüger, Michelangelo Mangano, Frank Simon, David Waters

Physics Case (General)

1. How much of the physics programme would be possible with just the ECC bricks (i.e. without the active part of the detector)?

Electronic detectors fulfil three different tasks: provide the energy measurement and the muon identification for neutrino physics, provide the time stamp to the interaction for both neutrino physics and new particle searches. The measurement of the neutrino energy is essential to disentangle the two components in the muon neutrino source (heavy versus light flavour). The muon identification is essential for the neutrino flavour identification (done in charged-current interactions). Therefore, without electronic detectors the neutrino program would be significantly spoiled. The time stamp is an important validation for neutrino interactions (from the IP) but it is essential to search for neutral massive particle interactions in a wide mass range.

2. Would it be possible and/or desirable to have some rapidity overlap with FASERnu?

Yes, it would be desirable as well as an overlap with LHCb, where charm has been studied with large statistics and to also look at neutrinos from W decays. Nevertheless, at this stage overlapping with FASERnu could only be made at the expense of performance, given the constraints on civil engineering (detector length) and, therefore, it is not worth it. Moreover, we underline that the on-axis location is not optimal for neutrinos from the physics point of view (neutrinos from light and heavy quarks have similar energy spectra) and also not from the experimental point of view (events are not uniformly distributed on the detector surface, but rather accumulated on the edge).

3. Several aspects of the physics programme are severely limited by statistics. Is there an intention to continue SND beyond Run 3? Has the possibility been considered, from all relevant perspectives (technical, e.g. compatibility with HL-LHC event rates and frequency of emulsion replacements, financial, collaboration longevity etc.)?

The Collaboration is now focussed on the efforts of constructing the detector to run it in Run3. Nevertheless, we clearly see the possibility and interest to continue beyond Run3. As an example, we envisage the possibility to upgrade the detector with a fully active version in order to extend the pseudorapidity range and improve its performance in the HL-LHC environment. We will add an “Outlook” section in the TP.

Neutrino Programme

1. Please clarify the primary goals for the neutrino programme. If the goal is to measure heavy quark production in the forward region based on a certain number of observed neutrino events, then this would need to assume a neutrino interaction cross-section. If the goal is to measure the neutrino cross-sections, then this would rely on a calculation of the flux (with associated QCD production cross-section uncertainties etc.). Explain how you propose (with the use of flavour ratios etc.) to disentangle these different measurements. Are there detailed studies

which show how such an analysis will be performed, what are the expected uncertainties and how will the constraints (e.g. on sterile neutrinos) improve on current knowledge?
Observed electron and tau neutrinos are expected to come only from heavy quark decays. In particular, electron neutrinos show a relatively large statistics. The neutrino cross-section has been measured at energies up to 350 GeV for muon neutrinos and it agrees with the standard model with an accuracy of a few percent. One can assume that the electron neutrino cross-section agrees with the standard model too since it would be difficult to argue differently for the first lepton generation. In this case we can use the lower energy part of the spectrum to constrain the heavy quark production and thus obtain a prediction that we can apply to the higher energy part in order to derive the νN cross section beyond 350 GeV.

Muon neutrinos also come from pion/kaon decays and this produces a different energy spectrum. The energy measurement can be used to disentangle the two components. The comparison with the electron neutrino yield at lower energies (below 350 GeV) can be used to derive the pion/kaon contribution in the region where the cross-section is well known.

From the ratio between electron and tau neutrino events, one can derive the tau neutrino cross-section (i.e. a universality test) with about 30% accuracy.

2. Table 2. Please indicate the uncertainties on these numbers, preferably separated into production and interaction uncertainties. What impact will these uncertainties have on the final physics measurements?
Electron neutrino interactions with energies below 350 GeV, where the cross-section is known, are the basis of our data-driven measurement technique, as explained above. They constitute ~20% of the total electron neutrino sample, i.e. about 80 events for 150 /fb as reported in Table 2, with a corresponding statistical uncertainty of ~11%. Therefore, accounting also for detector effects, we expect an accuracy at the level of 15-20% in the measurement of the neutrino production in pp collisions. With 3000 /fb of HL-LHC the statistical uncertainty of this method will improve to ~3%, while experimental (systematics) effects will be more relevant.

Electron neutrinos in the SND rapidity region mainly originate from the charm decay chain. The dominant process is gg—> ccbar (90.8%), other processes are gg—> bbar (7.4%), qqbar—>ccbar (1.7%), qqbar—>bbar (0.1%). The reaction is strongly boosted forward: <x1> ~ 0.2, x2 in the range from a few 10^{-7} to a few 10^{-6}. Reaction <Q^2> is ~4 GeV^2. See plots below. The neutrino energy at the end of the charm decay chain is dominated by the boost; it has lost a direct relation to the hard scattering process. The parton PDF is the same for both energy regions, below and above 350 GeV. Thus, from the ratio of these two samples we can derive the νN cross section at high energies. The bin sizes may be chosen so that the error is dominated by the absolute normalisation uncertainty of the low energy sample, i.e. <20%.

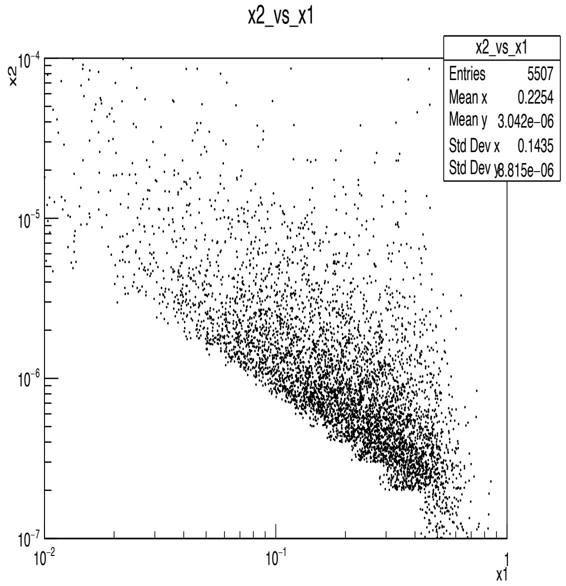

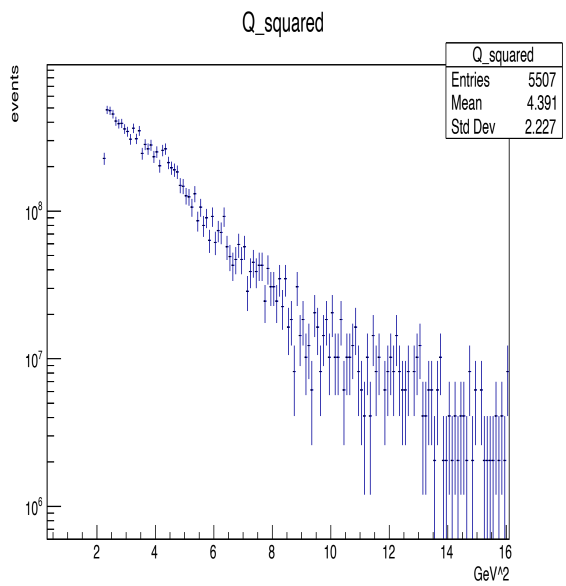

A second measurement can be extracted from the ratio (ν_μ at high-E)/(ν_e at low-E), that will demonstrate that the ν_μ sample at high energy is dominated by the charm contribution. Moreover, from the ratio (ν_τ / ν_e) we could measure the ν_τ cross-section, with an uncertainty in Run 3 of about 30%. It can improve to ~5% with 3000/fb.

One can subtract the electron neutrinos at low energy (<350 GeV) from the muon neutrino sample in the same energy range (80% of the sample), thus providing the spectrum of muon neutrinos from pion/kaon decays with an accuracy of ~4%, which can be used for a cross-check of the FLUKA/DPMJET LHC description.

3. The various codes available to model pp interactions in the forward region show important quantitative differences. This is particularly true of the charm production processes, which for SND are a dominant source of e and tau neutrinos. For example, we are aware of warnings from the authors of DPMJET (used as default by SND), that its charm-production predictions have not been validated and are likely to significantly overestimate production rates. Aside from affecting the precision in the extraction of the absolute neutrino flux and thus cross sections etc, this could also bias the overall estimates of the expected statistics. This is a particular concern in the case of the tau neutrinos, where Table 2 predicts at most a score of events. A complete comparison of available event simulation codes, and their impact on the SND measurements, must be documented.
In the paper N. Beni et al., “Further studies on the physics potential of an experiment using LHC neutrinos.”, published on J. Phys. G: Nucl. Part. Phys. **47** (2020) 125004, https://iopscience.iop.org/article/10.1088/1361-6471/aba7ad/pdf, the neutrino energy spectra and eta distributions were investigated in the region relevant for SND.
The authors used PYTHIA (8.226, with Monash 2013 tune, for b and c quark production and decays) and DPMJET (minimum bias + charm production). DPMJET was embedded in the FLUKA package of CERN EN-STI (F.Cerutti, M.Sabaté-Gilarte), with detailed simulation of the charged particle propagation through the LHC beam line elements.
For the SND rapidity range, it is shown that:
-- electron neutrinos (essential for our measurement strategy) come from charm decays (plots at bottom of fig. 6 and 7)

-- the flux expectations of PYTHIA and DPMJET are in good agreement (fig. 8).
-- muon neutrinos from charged pion and kaon decays populate lower energies (upper plots in fig.6 and 7).

4. Page 31. What is the final answer regarding the tau neutrino backgrounds. Table 6 is presented but there are some comments that “backgrounds can be further reduced”. What is the current best estimate of backgrounds, and signal/background ratio, for tau neutrino detection?
The main background for the tau neutrino search is due to muon neutrino CC interactions with charm production. The current background estimate takes into account only the muon identification efficiency to reject the charm background. As demonstrated by the OPERA experiment, a further reduction (a factor of about 3 with a cut-based analysis) can be achieved by exploiting kinematical features such as the fact that lepton and the hadronic jet at the neutrino interaction are expected to be back-to-back in the transverse plane. See for instance the plot below where the distribution of the angle in the transverse plane is reported. We are confident that, with a multivariate approach one can further improve the rejection power.

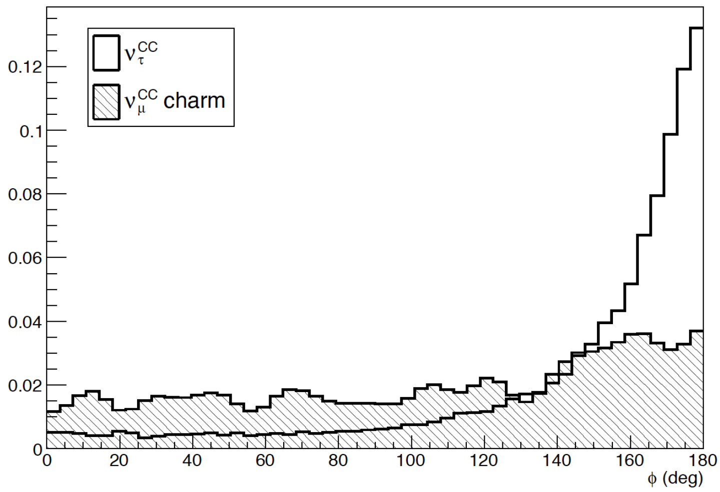

Feebly Interacting Particle (FIP) Programme
1. The case for making FIP searches “in an unexplored domain” is not currently made. The only direct comparison with other experiments (Figure 27) does not seem to show that SND extends the reach of NA64. Figure 28 shows regions of kinematic phase space for a time-of-flight analysis but there is no estimate of physics sensitivity – can the TOF-based FIP search be combined with the analysis summarised in Figure 27? In general, what is the status of the detailed simulation studies into the FIP sensitivity of SND?
Indeed, the model chosen is not competitive, although it is based on the direct search rather than on the missing energy as for NA64. Results reported in Fig. 27 do not exploit the TOF since the background was estimated to be negligible. TOF-based analysis could play a key role in confirming an evidence. We are now studying other channels and plan to report the corresponding sensitivity in the TP.

2. Section 7.3 argues that the dominant background comes from CCDIS interactions, which can be rejected very efficiently. Other backgrounds are considered to be too low to be relevant. Can you clarify the estimate of the neutrino-electron scattering background – is it truly

negligible even with the full exposure, or can it give rise to a small number of irreducible background events (with exactly the same signature as the LDM elastic scattering signal)?

The contribution from the different processes integrated over the full exposure (150 fb^{-1}) was estimated by applying a topological selection (requiring that only the electron is reconstructed in the final state). As a result, we expect 0.1 events from neutrino elastic scattering and 0.16 events from quasi-elastic interactions. By applying a kinematical analysis one can make this contribution negligible. All the other processes (deep inelastic and resonant scattering) provide a negligible contribution already at the topological level. As an example, one can see the analysis reported in the paper https://arxiv.org/abs/2010.11057

Technical Questions

1. The timing goals for the different sub-detectors are not quite clear. Section 2 (p6) mentions a particular need for excellent timing in the target region, but in fact the SciFis used there reach ~250 ps, while the scintillator bar system in the muon detector has ~100 ps or better. Are these goals just adapted to the capabilities of the respective technology, or is there a clear physics motivation for a particular choice?

The quoted time resolution in the LoI is per-plane, limited by the technology. The goal here is to have the TOF resolution of the experiment dominated by the fluctuations of the interaction time. This is true if the detector time resolution is O(100 ps). The event time resolution is expected to be considerably better by combining measurements in multiple planes.

2. It would be expected that the event time for an interaction that results in significant activity in the detector is substantially better than the single particle, single plane time resolution. Since the length of the bunches and the size of the overlap region limits precision of the interaction time to 200ps, it is not quite obvious why 100ps or better are needed for individual detector planes. Would a relaxation of the specifications here result in savings in terms of cost, or power budget (and thus cooling needs)?

The collaboration has recently re-discussed the use of the SAMPIC waveform sampling for the muon timing system. The conclusion was that a significant simplification can be achieved by using the SciFi DAQ system, based on the TOFPET2 timing ASIC for the scintillator bars as well (VETO and muon stations). This also leads to a read-out system based on a unique board for all detectors. Event building, timing alignment, synchronisation will be simplified. In this scheme power consumption for the muon timing system is a minor concern. Using the SciFi DAQ boards reduces significantly the cost for the readout electronics.

3. Section 4.1 Upstream veto: What size do the SiPMs have / what fraction of the surface area of the bar ends are covered by photon sensors? Is this a significant cost factor? If yes, what would be the consequences in terms of energy, time resolution if the area or the number of SiPM were to be reduced?

The SiPMs are 6 x 6 mm^2, while the cross section of the bar is 6 x 1 cm^2 so that the total coverage considering 8 SiPMs per each bar end is 48%. The SiPMs represent about half the material cost, although a reduction in the number of SiPMs does scale linearly with the cost/piece. The cost for the SiPMs is actually dominated by the muon system, where there are 880 SiPMs vs. 112 for the Upstream veto. The time resolution scales as sqrt(n), where n is the number of SiPMs. So reducing the number of SiPMs by 2 would degrade the resolution of the Upstream veto by a factor of sqrt(2) (~70 ps -> ~100 ps).

4. Page 24. You claim that this experiment demonstrates that you can reconstruct vertices in environments with up to 5´104 tracks/cm2, but you show only 5´103 in figure 20, which would mean a factor of 20 in extrapolation to the densities expected for SND in 25 fb-1. Is this correct?

Data reported in Fig. 20 refer to a lower density run of the SHiP-charm exposure. Higher density runs, as the one reported in the picture below, show the capability to perform track and vertex reconstruction with track densities up to $4\times10^4/cm^2$.

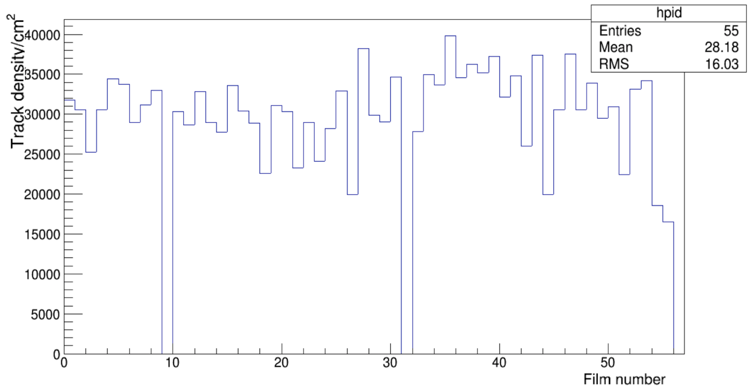

5. Figure 21, there is 30-50% difference between data and MC in the first half of the z distribution. If this is really mis-modelling, what would be the consequences for the SND measurements?

The analysis was recently finalised and a very good agreement was observed between data and simulations. An example from one of the SHiP-charm runs is reported in the picture below, for vertex position along the beam direction and for track multiplicity.

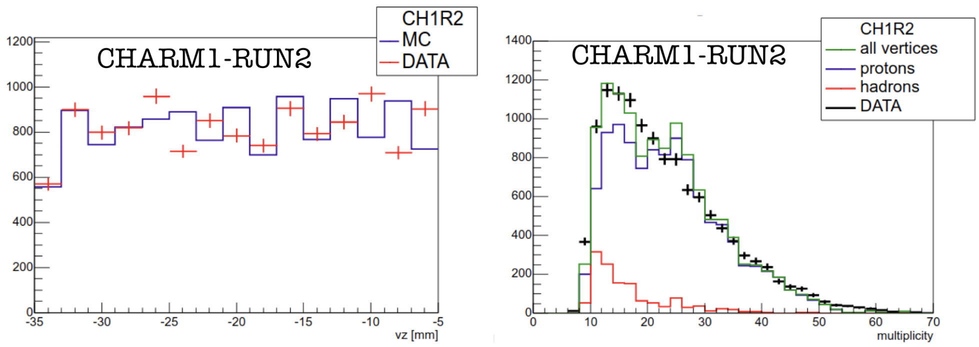

6. Do you need any signals from the LHC machine? If yes, are the corresponding cables/fibres included in the networking costs?

The SND detector will need the LHC clock and orbit signals. It will be provided by a fibre from the Beam Synchronous Timing (BST) system via the BST distribution point in SR1 to the shared FASER/SND rack in SR1. This fibre is part of the six fibres that are included in the cost presented in the LoI. Four will be used for Ethernet and one for BST. In reality, twelve fibres will actually be installed during the preparatory works in 2021, while SND will pay pro rata for six.

7. We understand that you cannot install the emulsions until right before the run starts, but do you plan to exercise the installation of the ECC bricks (162kg each) with a mockup?

We plan to perform a test of the assembly procedure on surface. A real-size prototype of the ECC made will be used.

8. What alignment precision is needed for the installation of the ECC bricks? Is this guaranteed by the mechanical structure, also for the replacement?
A reference corner and pins on the mechanical structure will guarantee an accuracy better than 1mm in the positioning of the ECC, also during the subsequent replacements.

9. Are there any surface radioactivity requirements for the tungsten plates in the ECC bricks, in order to avoid unnecessary film exposure?
A compatibility test between emulsions and tungsten plates was already performed in 2019 and showed no effect due to the radioactivity.

10. How will the data be brought out from the tunnel? Where will the surface rack be located and what other requirements for surface infrastructure will there be? Are there any potential conflicts with ATLAS?
In discussion with ATLAS (Michael Raymond) and FASER (Jamie Boyd), it has been agreed that the FASER rack in SR1 has sufficient space to host the SND control, monitoring and DAQ server and the fibre patch panel, and additional equipment if needed. A fibre cable will connect the rack in SR1 with the detector rack in TI18, providing a network connection between the server and the network switch in the detector rack. The FE DAQ boards are connected directly to the switch and send data to the server over Ethernet. No other requirements within the ATLAS area.

11. How will the detector be operated – will there be a “control room” for the operation of the detector?
There is no need on site for an experiment control room. It is envisaged that an office at CERN is equipped with a few computers for controlling and monitoring the detector and the readout, and for performing data checks via remote connection to the server in SR1.

Schedule & Planning

1. Are there any general updates to your planned schedule since the LoI preparation?
The tentative periods that have been identified together with operation, based on LHC schedule v3.1, for the preparatory works and installation are:
· 2021, July - September: Bulk of preparatory works and start of detector installation.
· 2021, November - December: SND installation.
· 2022, January: Installation of emulsion (one week) during the cryo recovery after the YETS.
This will be revisited together with operation in the coming days.

2. Is the commissioning foreseen to take place completely in the tunnel, or do you foresee a "dry run" on the surface?
We plan a full commissioning of the detector on the surface, including DAQ with cosmics, cooling tests, time alignment, etc.

3. How much do you rely on being able to exchange ECC bricks after exactly 25 fb-1? We assume that the time of LHC technical stops will need to be agreed between many interested parties, so you might not get the times that would fit exactly to the 25 fb-1.

We do not rely on the extraction of ECC bricks exactly after 25 fb-1. We have some flexibility and we will decide according to the actual planning of the Technical Stops. We will be more conservative in the first run. In addition, a detailed plan between the transport, the radiation protection and the detector teams has to be done to fit the exchange of emulsion bricks within Technical Stops.

4. How long does it take to take apart an ECC brick and re-assemble it with new emulsions? Is that feasible between two technical stops?
The construction of five ECC bricks requires less than one week in a dark room. In the baseline plan, this operation will be done right before the technical stop, having at CERN already a second set of tungsten plates for each swap. Exposed ECC bricks will be disassembled and their films will be developed soon after the extraction from the tunnel.

5. In figure 29, it looks like you only plan an extraction at the end of 2022. Depending on the run plan of the LHC, you might have to replace the emulsion once already during the first year of run. Are you ready to do that if required?
Yes, we are. We expect to have enough passive material and emulsion films to replace ECC bricks during 2022.

Costs and Resources
1. Do you need a crane to install/replace the ECC bricks in the tunnel (in TI18)? Is the cost included in the "transport" cost item?
A hydraulic manual folding workshop crane supported by a fixed crane is foreseen. This handling equipment is already available at CERN. A rail in TI18 was not considered due to other constraints in the area.

2. Is the cost of the hoist and the protection of the LHC machine included in the cost estimate?
Yes, it is.

3. How much lab & office space do you need at CERN? Do you need access to a dark room at CERN for the development of the emulsion films and the assembly of the bricks? If yes, is there an overlap with FASERnu? Has this been agreed?
The core collaboration of Opera that has now formed the SND collaboration is still installed and using the dark room and neighbouring facilities that have been transferred from WA75 and CHORUS in building 169. The Dark Room was cleaned, refurbished, and safety inspected for the SHiP-Charm measurement. It was also offered for use by the DsTau collaboration. We are currently in contact with FASERnu about their requirements in order to draw up a plan for shared use, and determine the changes that will need to be done to the lab to support transport, brick assembly, development, and handling and storage/disposal of chemicals for both experiments.

The SND collaboration is interested in a small lab at CERN for installation of an emulsion scanning station but this is to be confirmed.

A possible limited request for office space needs further discussion.

4. Infrastructure costs (Table 8): are these costs limited to the construction and installation phase, or do they cover the needs and M&O for the full Run 3? Are these costs implicitly assumed to be borne by CERN, or are they partly covered by the collaboration?
Table 8 shows the preliminary costs for the works to prepare the experimental area, services, and the transport channel, as well as the installation of the detector. They do not include M&O, this is in preparation for the TP.

It is expected that the costs associated with preparing the experimental area, installation and M&O are shared between Host Lab and the Collaboration, including the interested members from CERN EP. Discussions are currently ongoing.

5. In general, what is the overall situation regarding funding: what fraction is already pledged by collaboration institutes?
More than 60% of the budget is already pledged by Collaboration Institutes. An additional 15% is expected to be pledged at the end of November.

Organisation & Management

1. Do you have an organisation chart for the SND collaboration, with names assigned to all of the key managerial and technical roles? Please include the key CERN contacts.
The Collaboration has formed an Institution Board (IB) consisting of a representative from each Institute. The IB has elected a chair (N. Polukhina, Russian Academy of Science, RU), the Spokesperson (G. De Lellis, Naples University and INFN, IT), and appointed an interim Technical Coordinator (iTC) and contact person at CERN (R. Jacobsson, CERN). A Technical Board has been formed with appointed project leaders for each subsystem. iTC is currently acting as EXSO, and safety contacts are L. Di Giulio (CERN/EP) and M. Andreini (CERN/HSE).

Specific Comments on the LoI:

1. Figure 1: what is the “Missing Momentum” curve (it may be in the original references, but the caption should be self-contained).
The “Missing momentum” curve refers to the ultimate projected performance of the LDMX experiment.

2. Page 14: the number of SciFi places mentioned here is 4, but it seems to be 5 in Figure 5. Please clarify.
The current design foresees 5 planes.

3. Page 15/16: what are the requirements on spatial accuracy for muon identification/isolation? Can you show that a tile upgrade would be justified in terms of expected physics performance?
The current algorithm, used to classify isolated tracks in the muon identification system, assumes that the tracks are firstly identified in the emulsion brick and then projected onto the different muon identification planes. In the emulsion films, tracks are reconstructed with an angular accuracy of a few mrad. This results in an uncertainty better than 1 cm. Therefore, assuming 1cm bars in X-Y, the search window corresponds to three bars on both X and Y axes. Concerning the isolation requirement, this is satisfied if no hit is found in the bars immediately outside this search window.

Replacing the scintillating bars of the three most downstream planes with tiles would improve the isolation performance, increasing the muon identification by about 10%.

4. Page 17: what defines the “noise” data rate? Is a MIP threshold (thousands of photons) applied at the front-end? Would this not dramatically reduce the data-rate without any loss of physics?
Since the LoI, several things have changed:

- significantly lower noise rate using the TOFPET2 chip rather than the StiC (dynamic range better adapted to SciFi).
- Using the same readout system for all the detectors (no waveform sampling).

The following assumptions are used:

- the expected muon rate is 0.1 Hz/cm2 (upper limit, should be conservative)
- a hit has a 128-bit size (16 bytes)
- a muon will produce:
    - 3 hits in each SciFi plane, for a total of 30 hits
    - 16 hits in the veto and in the upstream muon planes (8+8 SiPM per bar, assuming readout on both sides), for a total of 16+80=96 hits
    - 2 hits in the horizontal downstream muon planes, 1 in the vertical ones, for a total of 9 hits
- total of 135 hits, event size per muon = 2160 bytes

In the approximation that all planes are 60x60 cm^2, the muon data rate will be 2160 bytes x 3600 cm2 x 0.1 Hz/cm2 = 0.78 MB/s.

The computation of the instrumental noise rate uses the following ingredients:

- 512 channels * 6 boards * 5 planes = 15360 channels in the SciFi
- 7 bars * 8 SiPMs * 2 sides = 112 channels in the veto
- 11 bars * 8 SiPMs * 2 sides * 5 planes = 880 channels in the upstream muons
- 77 bars * 3 sides * 3 planes = 693 channels in the downstream muons
- total of 17k channels.

The noise rate with TOFPET2 at 4.5V overvoltage is ~10 Hz per channel with a threshold of ~4 pe and ~1Hz per channel with a threshold of 5 pe. Assuming 10 Hz noise rate, we get a data rate of 17k * 10Hz * 16B = 2.8 MB/s.

5. Page 17: why are there 13 modules (we could not easily reproduce this number)?
13 was wrong, the correct number is 12 = 1+5+3*2.

6. Figure 18: we are not sure how to interpret this picture - are fibres laid in existing conduits?
The picture is an illustration of the path of the fibres. A fibre panel will be installed in the detector electronics rack in TI18. From there to RE18, 12 optical fibres will be blown through 600m of blowing tube to be installed in 2021 (part of the preparatory works) and connected to the existing patch panel in RE18. From RE18 to the FASER/SND rack in SR1, the fibre infrastructure is already in place.

7. Page 21: what are the potential implications (technical/cost/schedule) if more detailed radiation modelling indicates a problem for “commercial” electronics?
The expected level of radiation is not a problem for commercial electronics. The location is considered to be safe after FLUKA simulation done by FASERnu (<0.01 Gy/year). The radiation protection team has performed a cross-check with the FLUKA simulations and no radiation levels constraints are expected in TI12 where FASER is being installed (equivalent to our TI18 location).

8. Page 23: what is the expected overall scanning time for a full set of ECC bricks (i.e. per 25/fb)?
We plan to have the scanning time approximately equal to the exposure time. Within the Collaboration, we have enough scanning power to fulfil this goal.

9. Page 27: the requirement for at least one track with a momentum larger than 1 GeV/c. Is this applied to data (simulated) based on energy estimators (discussed later) or is this a generator level fiducial requirement?
The requirement is applied to simulated data and it is based on true momentum information.

10. Page 32: to clarify, ECC information does not enter the reconstructed energy estimator, only SciFi and MuFilter. Does ECC information (number of hits/track multiplicity etc.) not add any further information?
Information from ECC can improve the energy estimation of the electromagnetic component of the event. Not exploited yet, data from a test-beam exposure in DESY are under analysis.

11. Figure 27: there are 3 lines on the plot; 2 in the legend. Please clarify.
The grey region is the one already excluded by previous experiments.

12. Page 34: regarding SHIP-charm. Is there a publication (or a talk) describing these results in more detail?
An analysis note has been circulated within the corresponding working group (SHiP-charm) and a paper is in preparation. We will share it as soon as it is available.

13. Page 32, figure 26: what is the energy for which this resolution is determined?
This plot includes the convolution of all the energy spectrum.

14. Page 35: 150 fb-1 looks too much for 2022. Do you mean all of Run-3 (also for the subsequent estimates that are provided)?
Yes, all of Run3.

15. Figure 28: what do the colours represent? Is emulsion production on the critical path for most of 2021?
Colours are only meant to distinguish different detector/infrastructure installation blocks. Emulsion bricks will be installed at the very last moment to reduce the integration of cosmic-rays before the exposure. The schedule reflects this requirement. It is not on the critical path.