TECHNICAL PROPOSAL

FASER

FORWARD SEARCH EXPERIMENT AT THE LHC

Akitaki Ariga,[1] Tomoko Ariga,[1,2] Jamie Boyd,[3,*] Franck Cadoux,[4] David W. Casper,[5] Francesco Cerutti,[3] Salvatore Danzeca,[3] Liam Dougherty,[3] Yannick Favre,[4] Jonathan L. Feng,[5,†] Didier Ferrere,[4] Jonathan Gall,[3] Iftah Galon,[6] Sergio Gonzalez-Sevilla,[4] Shih-Chieh Hsu,[7] Giuseppe Iacobucci,[4] Enrique Kajomovitz,[8] Felix Kling,[5] Susanne Kuehn,[3] Mike Lamont,[3] Lorne Levinson,[9] Hidetoshi Otono,[2] John Osborne,[3] Brian Petersen,[3] Osamu Sato,[10] Marta Sabaté-Gilarte,[3,11] Matthias Schott,[12] Anna Sfyrla,[4] Jordan Smolinsky,[5] Aaron M. Soffa,[5] Yosuke Takubo,[13] Pierre Thonet,[3] Eric Torrence,[14] Sebastian Trojanowski,[15] and Gang Zhang[16]

[1] *Universität Bern, Sidlerstrasse 5, CH-3012 Bern, Switzerland*
[2] *Kyushu University, Nishi-ku, 819-0395 Fukuoka, Japan*
[3] *CERN, CH-1211 Geneva 23, Switzerland*
[4] *University of Geneva, CH-1211 Geneva 4, Switzerland*
[5] *University of California, Irvine, CA 92697-4575, USA*
[6] *Rutgers, Piscataway, New Jersey 08854-8019, USA*
[7] *University of Washington, PO Box 351560, Seattle, WA 98195-1560, USA*
[8] *Technion—Israel Institute of Technology, Haifa 32000, Israel*
[9] *Weizmann Institute of Science, Rehovot 761001, Israel*
[10] *Nagoya University, Furo-cho, Chikusa-ku, Nagoya 464-8602, Japan*
[11] *University of Seville, Seville, Spain*
[12] *Institut fr Physik, Universitt Mainz, Mainz, Germany*
[13] *KEK, Oho 1-1, Tsukuba, Ibaraki 305-0801, Japan*
[14] *University of Oregon, Eugene, OR 97403, USA*
[15] *University of Sheffield, Sheffield, S3 7RH, UK*
[16] *Tsinghua University, Beijing, China*

Executive Summary

FASER is a proposed small and inexpensive experiment designed to search for light, weakly-interacting particles during Run 3 of the LHC from 2021–23. Such particles may be produced in large numbers along the beam collision axis, travel for hundreds of meters without interacting, and then decay to standard model particles. To search for such events, FASER will be located 480 m downstream of the ATLAS IP in the unused service tunnel TI12 and be sensitive to particles that decays in a cylindrical volume with radius $R = 10$ cm and length $L = 1.5$ m. FASER will complement the LHC's existing physics program, extending its discovery potential to a host of new, light particles, with potentially far-reaching implications for particle physics and cosmology.

This document describes the technical details of the FASER detector components: the magnets, the tracker, the scintillator system, and the calorimeter, as well as the trigger and readout system. The preparatory work that is needed to install and operate the detector, including civil engineering, transport, and integration with various services is also presented. The information presented includes preliminary cost estimates for the detector components and the infrastructure work, as well as a timeline for the design, construction, and installation of the experiment.

* Contact email: Jamie.Boyd@cern.ch

† Contact email: jlf@uci.edu

CONTENTS

I. INTRODUCTION

FASER is a proposed small and inexpensive experiment designed to search for light, weakly-interacting particles at the LHC. Such particles are dominantly produced along the beam collision axis and may be long-lived particles (LLPs), traveling hundreds of meters before decaying. To exploit both of these properties, FASER is to be located along the beam collision axis, 480 m downstream from the ATLAS interaction point, in the unused service tunnel TI12. FASER will be sensitive to particles that decay in a small cylindrical volume with radius $R = 10$ cm and length $L = 1.5$ m. Despite its small size, FASER will significantly extend the LHC's discovery potential to a host of new, light particles. We propose that FASER be installed in TI12 in Long Shutdown 2 in time to collect 150 fb^{-1} of data from 2021-23 during Run 3 of the 14 TeV LHC.

The basic physics of the FASER experiment and the concept of the detector have been described in FASER's Letter Of Intent (LOI) [1]. This document gives much more detail about the technical aspects of the experiment. In Section II we give a brief overview of the physics and experiment. We then describe the detector environment in Section III. In Sections IV–IX, we describe the components of the detector in turn: the magnets, the tracker, the scintillators, the calorimeter, the detector support structure and the readout and trigger. In Sections X–XIII we give details about the civil engineering required to install the detector, installation and integration, commissioning, and safety aspects. We discuss off-line software and computing in Section XIV and summarize the overall costing and schedule in Section XV.

II. OVERVIEW OF FASER

A. Physics Goals

For decades, the leading examples of new physics targets at particle colliders were particles with TeV-scale masses and $\mathcal{O}(1)$ couplings to the standard model (SM). More recently, however, there is a growing and complementary interest in new particles that are much lighter and more weakly coupled [2]. Among their many motivations, such particles may yield dark matter with the correct thermal relic density and resolve outstanding discrepancies between theory and low-energy experiments [3–5]. Perhaps most importantly, new particles that are light and weakly coupled can be discovered by relatively inexpensive, small, and fast experiments.

If new light and weakly-interacting particles exist, they are typically produced parallel to the beam line and may travel hundreds of meters without interacting before decaying to visible particles, such as electrons and positrons. The existing detectors at the LHC, such as ATLAS and CMS, are therefore not well-matched to these particles, as they have holes along the beam line to let the proton beams in, and new light, weakly-interacting particles would escape through these holes undetected.

The goal of FASER is to target this "blind spot" by being located along the beam collision axis, far downstream from the ATLAS interaction point (IP). At this point, the proton beams are bent by magnets and so pass by FASER unhindered. However, new light, weakly-interacting particles will travel in straight lines and can decay to visible particles in FASER. Moreover, such particles are highly collimated. For example, new particles that are produced in pion decays are typically produced within angles of $\theta \sim \Lambda_{\rm QCD}/E$ of the beam collision axis, where E is the energy of the particle. For $E \sim$ TeV, this implies that even ~ 480 m downstream, such particles have only spread out ~ 10 cm in the transverse plane. A small and inexpensive detector placed in the very forward region may therefore be capable of highly sensitive searches.

FASER's physics reach has now been investigated for a host of light, weakly-interacting particles. In these studies, the most typical signals are those of LLPs that are produced at or close to the IP, travel along the beam collision axis, and decay visibly in FASER:

$$pp \to \text{LLP} + X, \quad \text{LLP travels } \sim 480 \text{ m}, \quad \text{LLP} \to \text{charged tracks} + X \ (\text{or } \gamma\gamma + X) \ . \quad (1)$$

These signals are striking: two oppositely charged tracks (or two photons) with very high energy ($\sim$ TeV) that emanate from a common vertex inside the detector and which have a combined momentum that points back through 10 m of concrete and 90 m of rock to the IP. Studies of dark photons, dark Higgs bosons, heavy neutral leptons, light $B-L$ gauge bosons, axion-like particles, and others [6–14] have demonstrated that FASER and a possible follow-up experiment, FASER 2, have a full physics program, with significant discovery potential in a variety of models.

In addition to searches for new physics, FASER may also provide interesting probes of standard model physics. As an example, in Run 3, the number of muon neutrinos passing through FASER with energies above 100 GeV is $\sim 10^{13}$, with roughly 600 interacting in the 10 cm-thick block of lead that is near the front of FASER (see Sec. II B). Furthermore, a few tau neutrinos with $\sim$ TeV energies are expected to interact in FASER. Although more study is required, these event rates imply that FASER may also provide interesting information about SM particles by detecting the first neutrinos at the LHC and, for example,

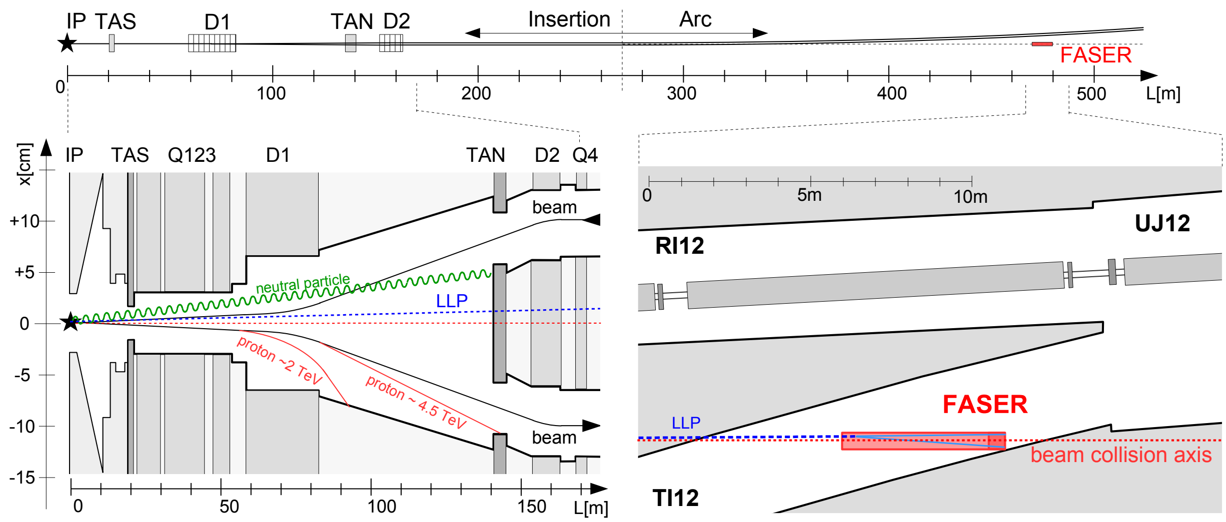

FIG. 1. Schematic view of the far-forward region downstream of ATLAS and various particle trajectories. **Upper panel**: FASER is located 480 m downstream of ATLAS along the beam collision axis (dotted line) after the main LHC tunnel curves away. **Lower left panel**: High-energy particles produced at the IP in the far-forward direction. Charged particles are deflected by LHC magnets, and neutral hadrons are absorbed by either the TAS or TAN, but LLPs pass through the LHC infrastructure without interacting. Note the extreme difference in horizontal and vertical scales. **Lower right panel**: LLPs may then travel $\sim$ 480 m further downstream, passing through 10 m of concrete and 90 m of rock, and decay within FASER in TI12.

constraining neutrino interaction rates in the energy range $E_\nu \sim 400$ GeV $-$ 4 TeV, where they are currently unconstrained. (See, for example, Appendix 1 of Ref. [1].)

B. The Experiment

The proposed location of FASER and the LHC infrastructure between the IP and the detector are shown in Fig. 1. A sketch of the proposed detector to be installed in the TI12 tunnel[1] is shown in Fig. 2, and a brief overview is given below.

At the entrance to the detector, two scintillator stations are used to veto charged particles coming through the cavern wall from the IP, primarily high-energy muons. Each station consists of two layers of scintillators. Between the stations is a 20-radiation-lengths-thick layer of lead that converts photons produced in the wall into electromagnetic showers that can be efficiently vetoed by the scintillators.

The veto stations are followed by a 1.5 m long, 0.6 T permanent dipole magnet with a 10 cm aperture radius. This is the decay volume for LLPs decaying into a pair of charged particles, with the magnet providing a horizontal kick to separate the decay products to a detectable distance. The decay volume is not foreseen to be under vacuum.

[1] In the LOI we were expecting the TI18 tunnel to be the best place for FASER, but additional measurements made by the CERN survey team during LHC Technical Stop 2 show that the TI12 tunnel on the other side of the LHC interaction point, IP1, will allow for a longer detector on the beam collision axis. TI12 is the same distance from IP1 as TI18, but is on the side towards the LHCb experiment, rather than towards ALICE.

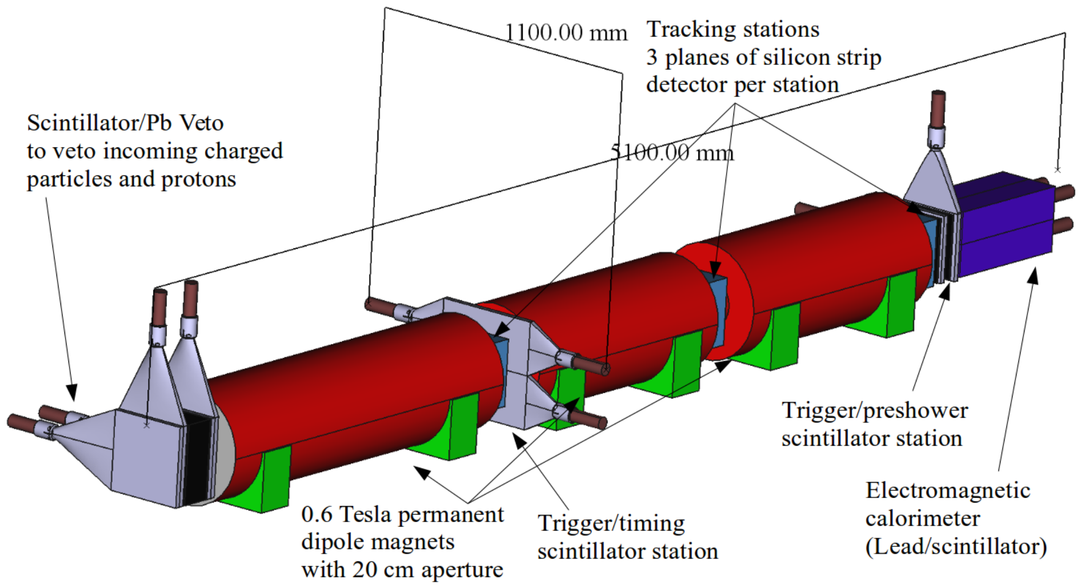

FIG. 2. Layout of the proposed FASER detector. LLPs enter from the left. The detector components include scintillators (gray), dipole magnets (red), tracking stations (blue), and a calorimeter (dark purple).

After the decay volume is a spectrometer consisting of two 1 m long, 0.6 T dipole magnets with three tracking stations, which are located at either end and in between the magnets. Each tracking station is composed of layers of precision silicon strip detectors. The three magnets will have their fields aligned to give the maximum separation for charged particles in the bending plane. Scintillator stations for triggering and precision time measurements are located at the entrance and exit of the spectrometer. The primary purpose of the spectrometer is to observe the characteristic signal of two oppositely-charged particles pointing back towards the IP, measure their momenta, and sweep out low-momentum charged particles before they reach the final layer of the spectrometer.

The final component is the electromagnetic calorimeter. This will identify high-energy electrons and photons and measure the total electromagnetic energy. The primary signals are two close-by electrons or photons with too-small separation for the calorimeter to resolve individually.

III. DETECTOR ENVIRONMENT

A. Detector Length Constraints

Detailed measurements from the CERN survey team have mapped out the beam collision axis or line of sight (LOS) in both the TI18 and TI12 tunnels. This LOS assumes no crossing angle between the beams at IP1. In reality, the LHC runs with a crossing angle in the range of about 150 μrad[2], the effect of which is discussed below. The TI18 and TI12 tunnels connect the LHC to the much shallower SPS, and they therefore slope steeply upwards as they leave the LHC tunnel. Because of this geometry, the LOS is below the tunnel floor as it enters the tunnel, and then emerges from the floor. Given this, to maximize the detector length that can be centered on the LOS, it will be desirable to lower the floor. Measurements from the CERN survey team show that, with the allowed digging that can be done in LS2 (limited to 460 mm, as discussed in Sec. X), a significantly longer detector can be installed in TI12 than in TI18. In particular, based on our design, we could fit a roughly 5 m-long detector in TI12, compared to a roughly 3 m-long detector in TI18.

A more detailed modeling of the tunnel and the detector shows that the length of the detector becomes limited by the back part hitting the wall of the tunnel, as shown in Fig. 3.

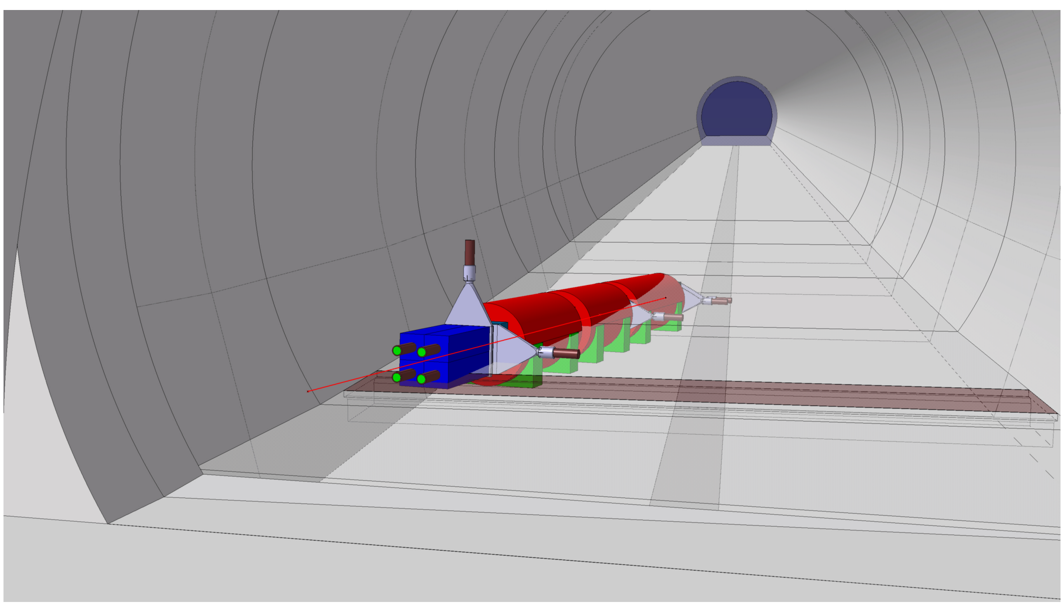

FIG. 3. A model of the FASER detector situated at the proposed location (centered on the nominal LOS) in the TI12 tunnel.

[2] This value is the half crossing angle which is typically what is quoted, and what is used for the rest of this document

B. Beam Configuration Effects

1. *Effect of IP1 beam crossing angle*

To avoid long range beam-beam effects and parasitic collisions inside the common beam pipe, the LHC runs with a half crossing angle at IP1 of about 150 μrad. In the LHC design, and for all running to date, the crossing angle is in the vertical plane. For Run 2, it was decided to flip the crossing angle direction (sign) periodically (e.g., once per year) to distribute the collision debris more evenly and prolong the lifetime of the focusing magnets, which are a potential limitation on the total integrated luminosity that can be delivered before LS3. In addition, since the start of 2017, the crossing angle was reduced during the physics fills from an initial value of 160 μrad to 120 μrad at the end of the fill. At the FASER location (480 m from IP1) a half crossing angle of 160 μrad corresponds to a shift of the collision axis (compared to the nominal LOS, assuming no crossing angle) of 7.7 cm, and 120 μrad will shift the axis by about 2 cm less.

For Run 3 the optics that will be used in the LHC has not been finalized. Two options are under consideration: (i) *round beams*, in which the β^* is the same in the horizontal and vertical planes, and (ii) *flat beams*, in which the β^* is larger in the crossing plane and which can give higher luminosities. In the case of *round beams*, the crossing angle will be similar to that used in Run 2, being in the vertical plane and with the sign possibly changed each year of running. With *flat beams*, the crossing angle will be changed to be in the horizontal plane, and the sign will not be changed (this will be fixed to the sign for which the LOS points towards the outside of the LHC ring). In all cases, the crossing angle values will be similar to those used in Run 2 and will be reduced in a similar way during the physics fills.

For FASER the considerations above mean that we need to be ready for the LOS to be displaced from the nominal LOS by up to 8 cm in either the horizontal plane or in the vertical plane (and switching between the displacement being up or down every year). The baseline strategy to deal with this would be to keep the detector centered on the nominal LOS and sacrifice some signal rate from the fact that the actual LOS is shifted with respect to the nominal LOS by the crossing angle. Simulation studies show that this leads to a minor loss in physics reach; for example, for dark photons, it leads to a loss in signal acceptance of roughly 25%, corresponding to an almost imperceptible change in sensitivity reach in the (mass, coupling) plane. It may be possible to align the detector to be more centered on the LOS taking into account the crossing angle (especially in the case of flat beams when this will not be changed between years), but for now we consider that we will center FASER on the nominal LOS.

2. *Effect of beam divergence*

The beam divergence is a measure of the intrinsic transverse momentum spread of the collision system due to the machine optics. It is given by $D = \sqrt{\epsilon/\beta^*}$, where ϵ is the transverse emittance of the beam, and β^* is the value of the β function at the IP. Whereas the crossing angle shifts the LOS, the effect of the divergence is to spread the per-collision collision products out around the nominal LOS. For the typical values of parameters expected to be used in Run 3 ($\beta* \approx 30$ cm and $\epsilon \approx 3\times10^{-10}$ m), we expect $D \approx 30$ μrad. In the case of flat optics with $\beta^*_x \approx 60$ cm and $\beta^*_y \approx 15$ cm, this will give $D_x \approx 20$ μrad and $D_y \approx 40$ μrad. Given the LHC machine parameters, then, the value of the divergence is usually of the order of 10% of the crossing angle, and so we expect the effect of the divergence to be very small for FASER. Of course, the effect of the divergence can be taken into account in the signal simulations.

3. *Effect of filling scheme*

The effect of the LHC filling scheme on FASER is expected to be minimal. The background particle flux coming from the IP will be related to the instantaneous luminosity, but these rates are much lower than the bunch crossing rate, and so the effect of pileup is completely negligible. The signal rate is also proportional to the instantaneous luminosity, but again, for the same luminosity, it will not be dependent on the filling scheme.

C. Particle Flux

FLUKA simulation [15, 16] studies have been carried out to estimate the expected particle fluxes entering the FASER detector [17]. These studies include particles arising from collisions in IP1 and from beam-related backgrounds in the LHC. In addition, detectors have been installed in the TI18 and TI12 tunnels during LHC Technical Stops in 2018 to measure the particle flux and validate the simulations. These fluxes are important as they will determine the trigger rate and the number of particles that can be used for detector alignment and calibration. The current simulations and most of the *in situ* measurements are for TI18. The preferred FASER location is now the TI12 tunnel, but the expectation is that the particle flux will be the same in TI12 and TI18, as the collision system and LHC infrastructure is symmetric around IP1. First *in situ* measurements from TI12 demonstrate that this is the case, and future simulations will be carried out for TI12.

1. *FLUKA simulation*

FLUKA simulations have been carried out by the EN-STI CERN group as part of the Physics Beyond Colliders effort. The simulations were made for the TI18 tunnel, normalized to Run 3 conditions (we assume a constant luminosity of 2×10^{34} cm^{-2} s^{-1}). The construction of the TI12 tunnel simulation model is envisaged, although its symmetry a priori suggests no major impact on the results.

The simulations show, as expected, that muons and neutrinos are the only particles that traverse the FASER detector locations with an appreciable rate. The muons and neutrinos are either produced directly from proton-proton interactions at the IP ,or in upstream showers from particles produced at the IP hitting machine components (for example, high energy neutrons hitting the TAN 140 m from the IP). The FLUKA simulation contains a realistic geometry of the LHC tunnel and the LHC optics, and is therefore expected to model such processes well. Similar simulations using this setup have been used for understanding particle fluences and radiation levels in different areas of the LHC complex, and in many cases have been validated with measurements.

The simulated particle fluences as functions of energy for muons and neutrinos are shown in Fig. 4 for an instantaneous luminosity of 2×10^{34} cm^{-2} s^{-1} and collision energy of 13 TeV. The figure only shows negatively charged muons, as the biasing technique used does

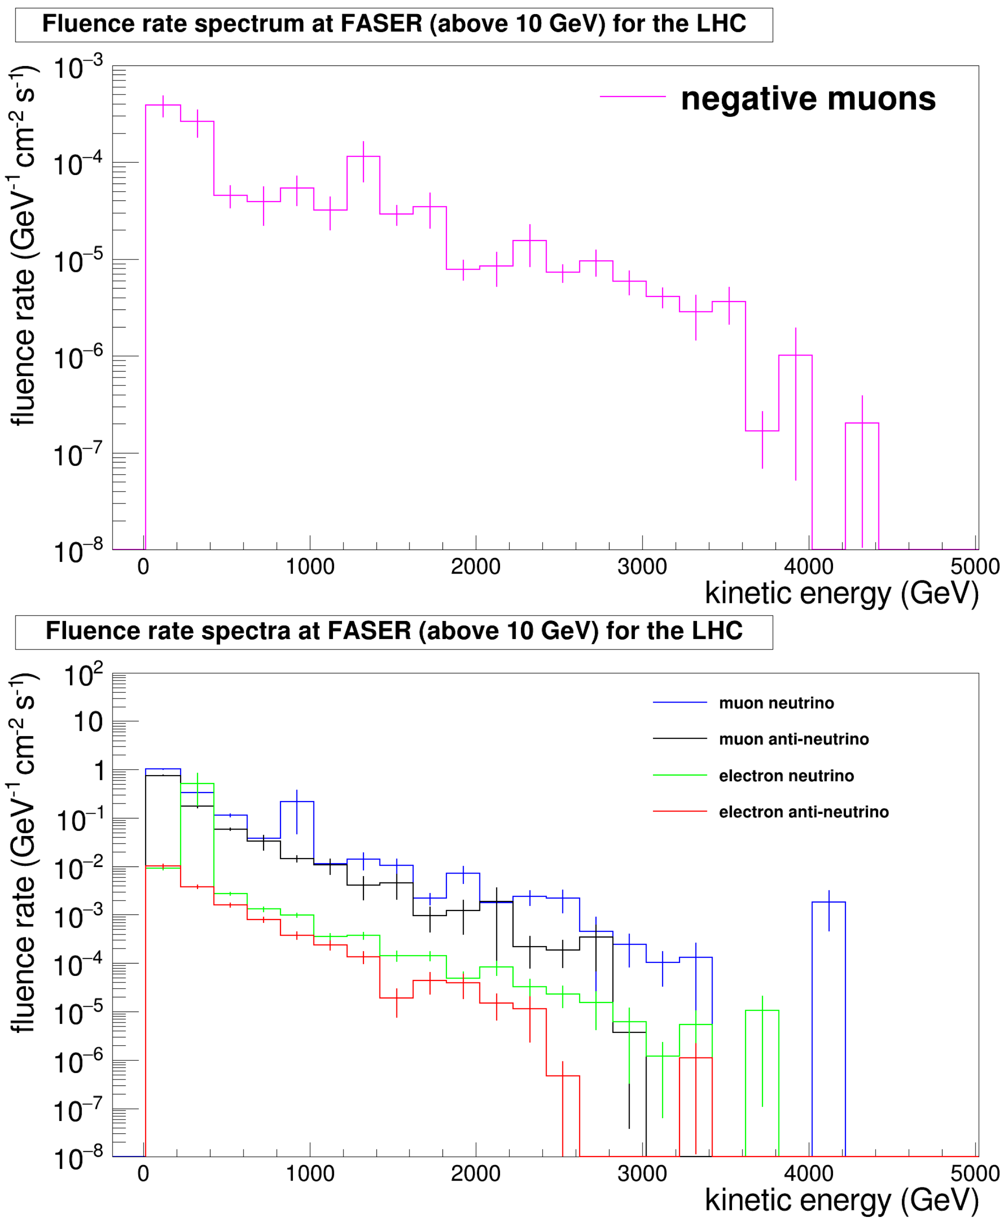

FIG. 4. FLUKA simulation estimation of the particle flux as a function of energy at the FASER location: (top) for negative muons; (bottom) for different neutrino species. These are normalized to a luminosity of 2×10^{34} cm^{-2} s^{-1}. Due to limited statistics, the positive muon spectra is not included but this is of the order of 10-20% of the negative muon flux above 100 GeV.

not give reliable spectra for positive muons at the moment. The simulations indicate that above 100 GeV the positive muon flux at FASER is between 10-20% of the negative muon flux, while this fraction increases below 100 GeV. The rate for all muons with energy above different thresholds is presented in Table I. For particles with energy above 10 GeV, this gives a charged particle rate through the 10 cm-radius detector volume of around 100 Hz.

FLUKA does not include the rate of particles produced in neutrino interactions in the rock, but this is estimated to be completely negligible (less than 0.01 Hz).

The FLUKA estimations of the particle flux have an uncertainty of order a factor of a

Energy threshold [GeV]	Charged particle flux [$\mathrm{cm}^{-2}\ \mathrm{s}^{-1}$]
10	0.40
100	0.20
1000	0.06

TABLE I. The expected charge particle flux at the FASER location from FLUKA simulations for different energy thresholds, normalized to the expected Run 3 luminosity of $2 \times 10^{34}\ \mathrm{cm}^{-2}\ \mathrm{s}^{-1}$. The rate is entirely from muons.

few, dominated by statistical effects in the current simulation samples. With larger samples there will still be sizable systematic uncertainties at the tens of percent level.

Figure 5 shows the muon flux as a function of radial position around FASER. It can be seen that the FASER detector is in a region with reduced particle flux, with significantly higher rates expected 1–2 m on either side of FASER (for positive and negative muons separately), due to the bending of the muons in the LHC magnetic field.

FLUKA can also be used to estimate the flux of particles entering FASER that are not produced in interactions at IP1. These include particles coming from showers in the LHC

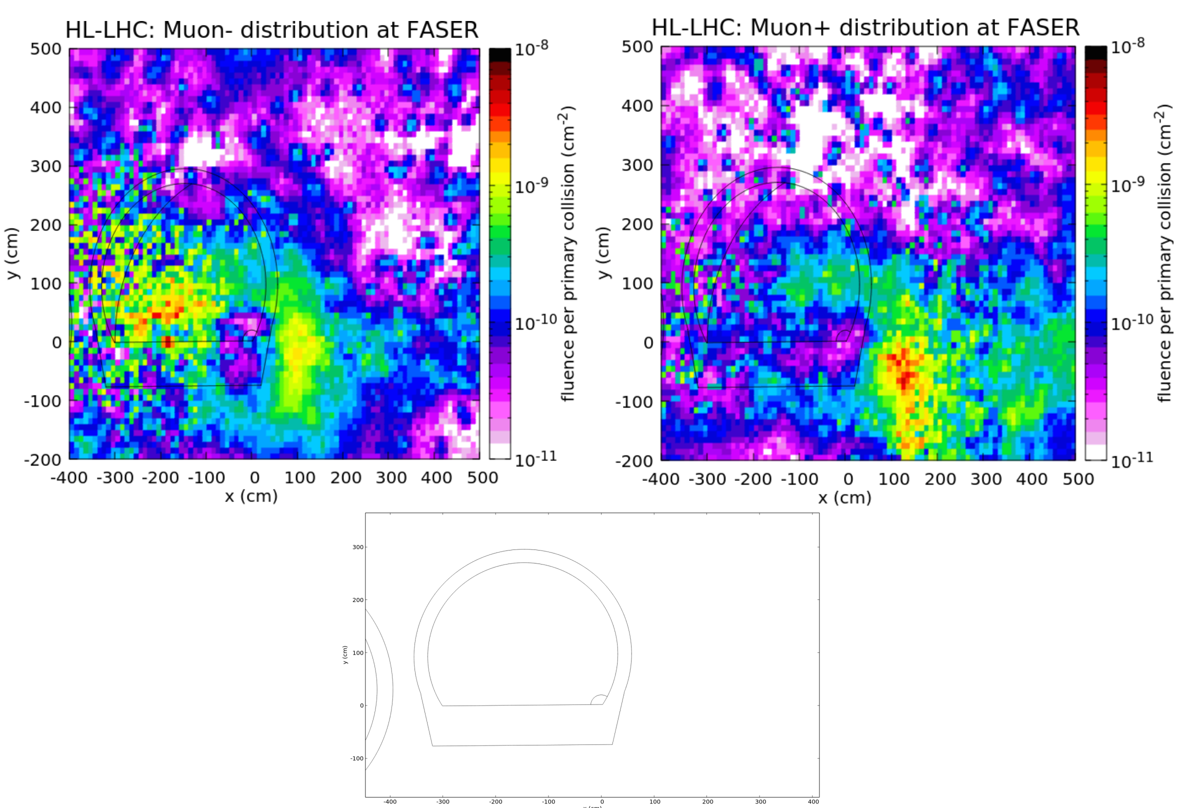

FIG. 5. FLUKA simulation estimates of negative muon fluxes (top left) and positive muon fluxes (top right) in the transverse plane at the FASER location. These results assume the TI18 location, 485 m from the IP. The diagram at the bottom shows the geometry used in the simulations. The FASER detector is visible as a small, partially-cut circle of radius 20 cm at the bottom right of the tunnel.

	beam [fb^{-1}]	observed tracks [cm^{-2}]	efficiency	normalized flux, all [fb cm^{-2}]	normalized flux, main peak [fb cm^{-2}]
TI18	2.86	18407	0.25	$(2.6 \pm 0.7) \times 10^4$	$(1.2 \pm 0.4) \times 10^4$
TI12	7.07	174208	0.80	$(3.0 \pm 0.3) \times 10^4$	$(1.9 \pm 0.2) \times 10^4$
FLUKA simulation, E>100 GeV				1×10^4	

TABLE II. Measured fluxes from emulsion detector data. The fluxes in the main peak (within 10 mrad) should be compared with the FLUKA simulation.

beam pipe in the dispersion suppressor (which arise from off-momentum protons following diffractive interactions in IP1, hitting the beam aperture, and causing particle showers), and particles produced in beam-gas interactions in the LHC beampipe. Both of these sources are highly suppressed: the dispersion function close to the FASER location minimizes proton losses in this region, and the excellent vacuum in the LHC means beam-gas interaction rates are extremely small. The FLUKA simulations predict the background of high-energy particles entering FASER from these processes is negligible.

2. In situ *measurements*

To measure particle fluxes at FASER's location, emulsion detectors were placed in the TI18 and TI12 tunnels. Photos and maps of the installation and the angular distributions of particles detected are shown in Fig. 6. The angular distributions show clear peaks of charged particles entering the detector from directions compatible with the ATLAS IP (the peak population of tracks at close to (0,0) on the figures). The width of the main peak was measured to be 2.3 mrad as shown in Fig. 7, which is as narrow as the angular resolution of the emulsion films, which is 2 mrad. This implies that the particles in this peak are sufficiently energetic not to be affected by multiple Coulomb scattering through about 100 m of rock and concrete shielding. Note that the expected transverse momentum due to the multiple Coulomb scattering through 100 m of rock is 0.54 GeV, corresponding to a scattering of 2 mrad for 270 GeV particles without including their ionization energy loss. In addition, the emulsion detector measurements were taken both with and without 0.5-mm-thick tungsten plates, which effectively imposed different energy cutoffs and provided insights about particle energy as a function of angle. As shown in Fig. 8, the tracks at the peripheral part are mostly composed of low-energy particles below 1 GeV.

The observed number of charged particles and rates are summarized in Table II. The particle fluxes in the main peak (within 10 mrad in angular space) were measured to be $(1.2 - 1.9) \times 10^4$ fb cm^{-2} when corrected for detection efficiency. These can be compared with the expected flux from the FLUKA simulations of 2.0×10^4 fb cm^{-2} for particles with the energy above 10 GeV. The *in situ* measurement therefore agrees well with the simulation given the uncertainties in the detection efficiency and uncertainty in the simulations.

An interesting feature in the emulsion detector data is small secondary peaks visible in the angular distributions, for example, at $(-0.75, 0)$ in the TI18 data shown in Fig. 6, with $\sim 1\%$ of the total number of tracks. This corresponds to tracks entering the detector with an angle consistent with originating at the LHC beamline at the bottom of the TI18 tunnel and, therefore, entering the detector without passing through any rock. The coordinate system and the tunnel geometry can also be seen in the center panels of Fig. 6. The FLUKA

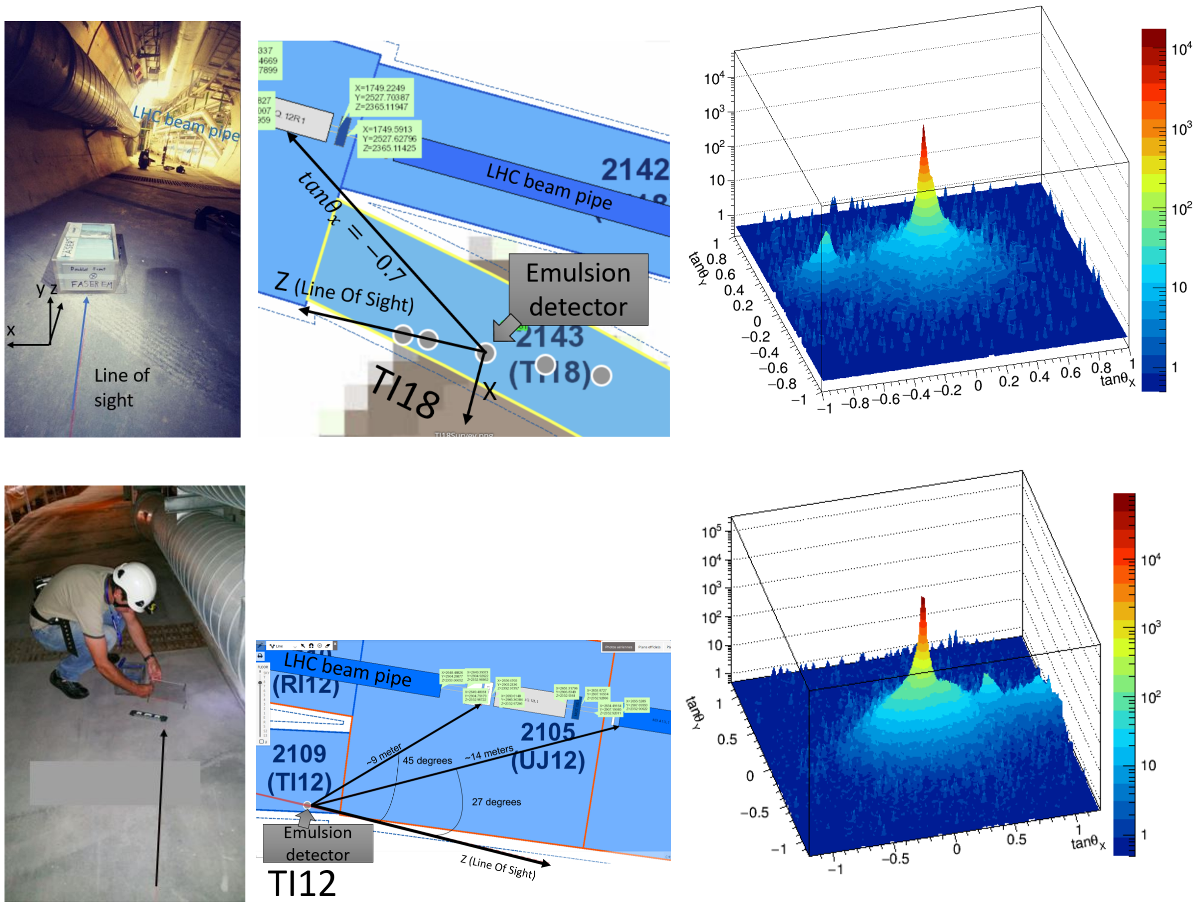

FIG. 6. *In situ* measurements by emulsion detectors at the TI18 location (upper panels) and TI12 location (lower panels). We show photos of the installed detectors (left), maps of the installation locations (center), and angular distributions of the detected particles (right).

simulations do not show such a population of tracks, but for the beam-gas simulations they only included high energy particles with $E > 100$ GeV, whereas the emulsion detectors are sensitive to particles with much lower energies. It is therefore likely that these are low energy particles and will not be problematic for the experiment. A more detailed analysis of the emulsion detector data is ongoing, and it should provide a determination of the rate of high-energy electromagnetic objects.

An active monitoring device (a TimePix3 Beam Loss Monitor [18]) was installed in the TI18 tunnel on the LOS during LHC Technical Stop 2. This device has the capability to correlate the rate of detected particles with the beam conditions. In particular, it can separately determine the particle detection rate during periods with high-luminosity collisions, periods with high-energy beams but no collisions (for example, during the 'squeeze' beam process), and periods with no beam in the machine. During periods with high-luminosity collisions, the rate can also be correlated with the instantaneous luminosity at IP1. However the device and the reconstruction algorithm are not calibrated, so it cannot currently provide absolute measurements of the particle flux. A first analysis of the TimePix detector data shows a clear correlation between the observed cluster counts and the instantaneous luminosity in IP1. It also shows slightly larger rates when there is beam in the LHC, but

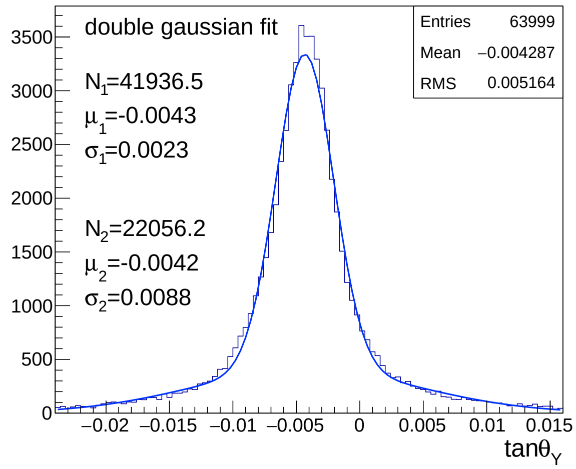

FIG. 7. A magnified view of the main peak of the angular distribution of particles detected at TI12, projected into the y-axis. The width of the main peak is 2.3 mrad, which is nearly consistent with the emulsion detector's angular resolution of 2 mrad.

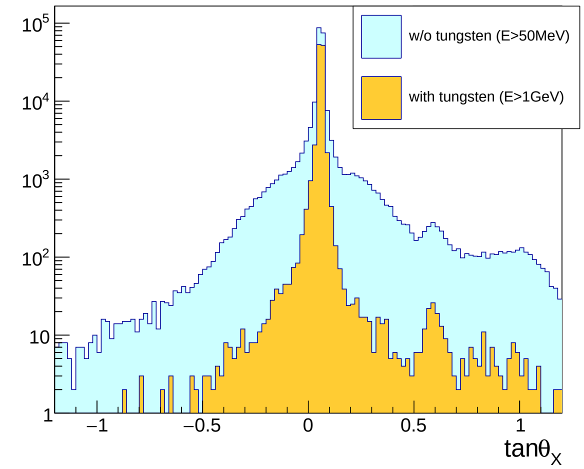

FIG. 8. Angular distributions, projected into the x-axis, by three emulsion films with and without tungsten plates, corresponding to energy cutoffs of about 1 GeV and 50 MeV, due to multiple Coulomb scattering, respectively.

Period	**Luminosity** $[10^{34}\ \mathrm{s}^{-1}\ \mathrm{cm}^{-2}]$	**Counting Rate** $[\mathrm{s}^{-1}]$	**Counting Rate/Luminosity** $[10^{-34}\ \mathrm{cm}^{2}]$
No beam	-	0.16	-
Beam (no collisions)	-	0.55	-
Collisions	1.8	7.0	4.0
Collisions	1.3	4.8	3.8
Collisions	0.8	3.3	4.2
Collisions	0.6	2.7	4.3
Collisions	0.5	2.2	4.1

TABLE III. Preliminary results from the TimePix detector installed in TI18, indicating that the main particle rate is proportional to luminosity in IP1. This also shows a small, but significant, increase in rate with non-colliding beam, compared to no beam in the machine. Beam (no collisions) corresponds to a full machine (2556 bunches) at the start of a physics fill, providing a total intensity of 2.7×10^{14} protons per beam.

with no collisions, compared to no beam in the machine. Example results are shown in Table III.

D. Radiation Level

The radiation level in the TI12 tunnel is an important input for determining what electronics can be operated in and close-to the detector. Both simulation studies and *in situ* measurements have been carried out to assess the radiation level.

1. *FLUKA simulation*

The FLUKA setup described above was also used to estimate the radiation level in the TI18 tunnel. For such locations, the radiation field is usually driven by proton showers in the dispersion suppressor, and, as mentioned above, these losses are very low for the LHC cell around FASER. Given this, radiation from showers due to beam-gas interactions in the incoming beam, where the particles from the beam-gas interaction showers can enter FASER without passing through any rock, can also be relevant. The simulations estimate a dose less than 5×10^{-3} Gy per year and a 1 MeV neutron equivalent fluence of less than 5×10^{7} per year. These numbers are roughly estimated to be comparable to or smaller than the dose in the ATLAS underground area USA15, where non-radiation-hard commercial electronics are used.

2. In situ *measurements*

Four BatMon battery operated radiation monitoring devices were installed in the TI18 tunnel by the EN-SMM CERN group during LHC Technical Stop 1. Two were tuned to measure the high-energy hadron flux, and two to measure the thermal neutron level. Given the expected low radiation in TI18, they were deliberately installed at the entrance to TI18 from the LHC, so closer to the LHC than any FASER electronics would be, as otherwise they may not have been able to make meaningful measurements. This means the measurements should represent a conservative estimate of the radiation field that FASER electronics will be exposed to. The devices were read out after 3 fb^{-1} of 13 TeV *pp* collision data was delivered to IP1. The readings show that the high-energy hadron fluence is below the device sensitivity (corresponding to 10^{6} cm^{-2}), completely consistent with the expectation from the FLUKA simulation studies. For thermal neutrons the measured flux is 4×10^{6} cm^{-2}, to be compared with the simulation estimate of 3×10^{6} cm^{-2}. In general, the measured radiation level is consistent with the simulations, and it is generally low and dominated by beam-gas interactions in beam-2 (the incoming beam for TI18). The BatMon detectors have now been installed in TI12, but the detector data has not been analyzed yet.

In summary, in general we expect to be able to use non-radiation-hard electronics in FASER without problems.

E. Other Environmental Factors

A temperature and humidity sensor (of type TAND TR-72wf) was installed in the TI12 tunnel on the LOS as part of the emulsion detector installed there to measure the charged particle flux. The measured temperature and humidity as functions of time are shown in Fig. 9. During this period the temperature in TI12 was constant at 18°C, whereas the humidity varied between 40% and 60%, with a value around 55% for most of the time. The figure also shows the variation over a longer timescale of about a year, but using the LHC environmental monitoring system, with the reading for the sensors closest to TI12. The temperature in the tunnel is very stable also over the longer time scale.

Vibrations and ground movements in the LHC tunnel are carefully monitored, as even small movements can cause the beam to be lost. For FASER, movement can effect the detector alignment. Studies by LHC experts [19] show that, for wavelengths up to 500 m,

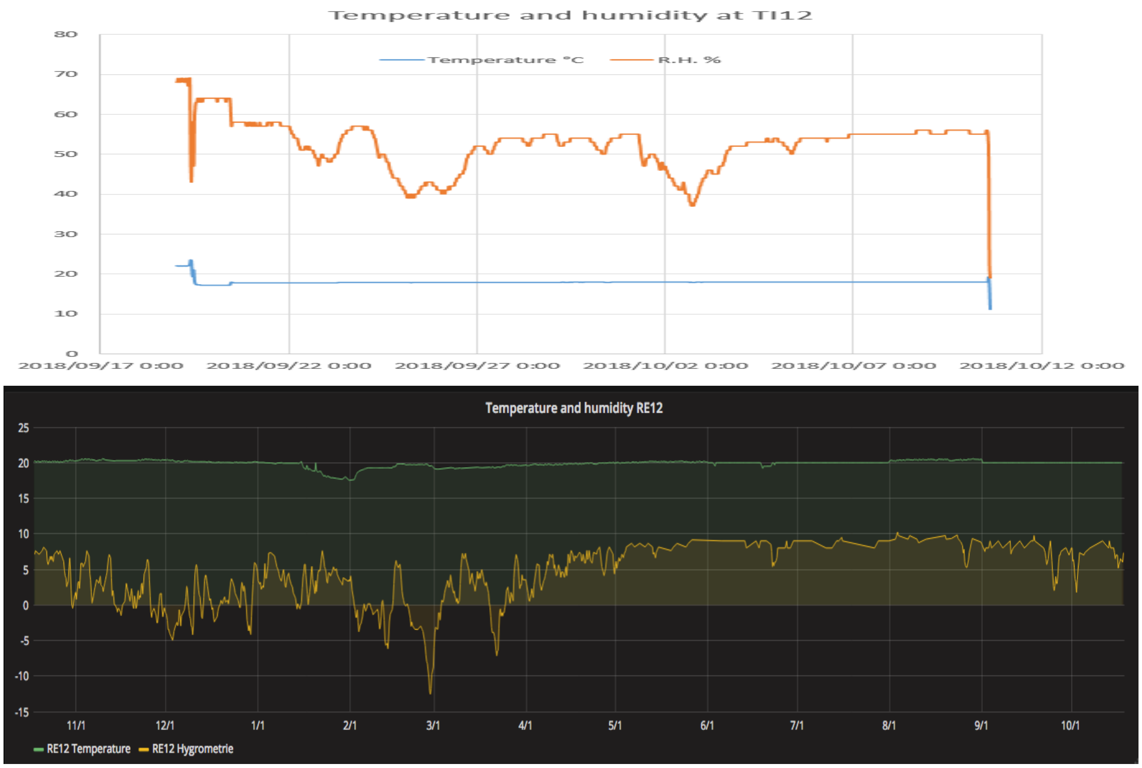

FIG. 9. Measured temperature and humidity. Top: during the period when the emulsion detector was installed on the LOS in TI12 (between mid-September and mid-October). Bottom: in the LHC tunnel close to TI12 (in the UJ12 region) from LHC monitoring, over about 1 year.

the tolerable range of seismic wave amplitudes is in the range of a few micrometers or less, with large uncertainties due to the exact wavelength or frequency content of the earthquake and due to its orientation relative to the LHC ring plane. For wavelengths above 500 m, wave amplitudes of a few tens of micrometers are acceptable and should not provoke beam aborts. This suggests that any ground motion would be a problem for the LHC before it becomes problematic for FASER.

IV. MAGNETS

A. Requirements and Magnet Design

The main requirements on the magnets needed for FASER are:

- The dipole field should be large enough to sufficiently separate pairs of oppositely-charged, high-energy particles originating from a high-energy, low-mass particle decay inside the detector decay volume;
- The magnet should be sufficiently thin to be able to fit into the tunnel (after lowering the floor by up to 46 cm) to allow a 5 m-long detector centered on the LOS (this suggests a magnet thickness of 15 cm would be best);
- The services needed for the magnet should be minimized (for example, minimizing power and cooling requirements);
- The stray field must be small enough that the PMTs used for the scintillator and calorimeter readout can operate correctly;
- The magnet dimensions and weight must be compatible with transport to the TI12 location (including being lifted over the LHC machine).

The above requirements suggest that a permanent magnet with a field of ≈0.5 T would be a good solution. For such a magnet there are no services, and the magnet can be thin. A conceptual design of such a magnet, based on a design made for the NTOF experiment, has been made by the CERN warm magnet section. This is based on a Halbach array of permanent magnets arranged to produce a dipole field. Since the radiation level in TI12 is very low, the NTOF design can be slightly improved for FASER by replacing the magnetic material of SmCo with NdFeB, which allows a slightly higher field of 0.6 T to be achieved. The magnet design is shown in Fig. 10, and Table IV shows the main parameters.

The stray field outside the magnet openings can be quite large, as is demonstrated in Fig. 11. It maybe that shielding is needed to reduce the field sufficiently to be able to operate the PMTs in the detector. Because of the field between the magnets, the magnets should be separated by at least 20 cm, and no magnetic material should be placed between them.

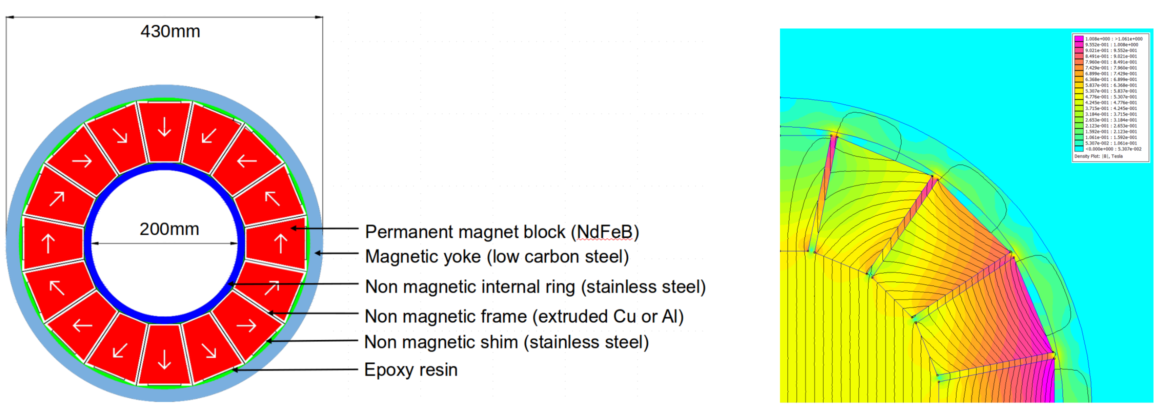

FIG. 10. Left: a sketch of the magnet design. Right: the field lines in the magnet.

Parameter	Value	Unit
Magnetic material	NdFeB	
Central field	0.6	T
Aperture	200	mm
Outer diameter	430	mm
Field homogeneity	± 2	%
Temperature dependence	-0.12	%/K
Weight/length	≈ 1000	kg/m

TABLE IV. Main magnet parameters for the FASER magnet design.

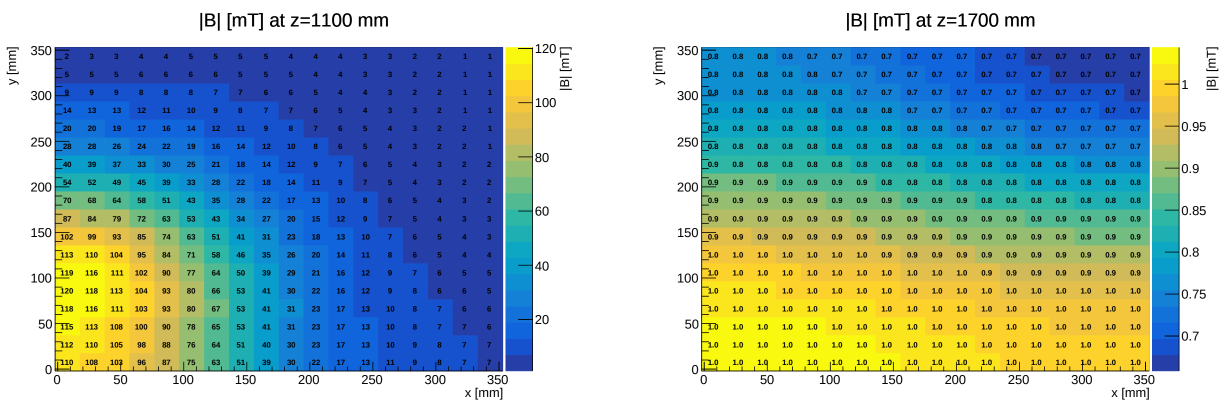

FIG. 11. The stray field (in mT) outside of the magnet estimated from the magnet design model at positions 10 cm (left) and 70 cm (right) in front of the magnet opening, corresponding to the locations of the closest scintillator (at 35cm from center) and calorimeter PMTs (at 7cm from center), respectively.

B. Cost and Schedule

The cost of producing the FASER magnets has been estimated by the CERN magnet group to be 420 kCHF, including 100 kCHF for personnel costs. The cost breakdown is outlined in Table V.

The timeline for producing the FASER magnets has been estimated by the CERN magnet group to be 15 months, with the breakdown outlined in Fig. 12.

1.5m Magnet Component	**Cost [kCHF]**
Permanent Magnet Blocks	90
Machining of Magnetic and Non-Magnetic Parts	30
Magnet Assembly	10
1.5m Magnet Total (1 Unit)	**130**
1.0m Magnet Component	**Cost [kCHF]**
Permanent Magnet Blocks	60
Machining of Magnetic and Non-Magnetic Parts	25
Magnet Assembly	10
1.0m Magnet Total (1 Unit)	**95**
Magnet Engineer	**100**
Total (1 1.5m Magnet + 2 1.0m Magnets + Engineer)	**420**

TABLE V. Budget for magnet construction for the FASER experiment.

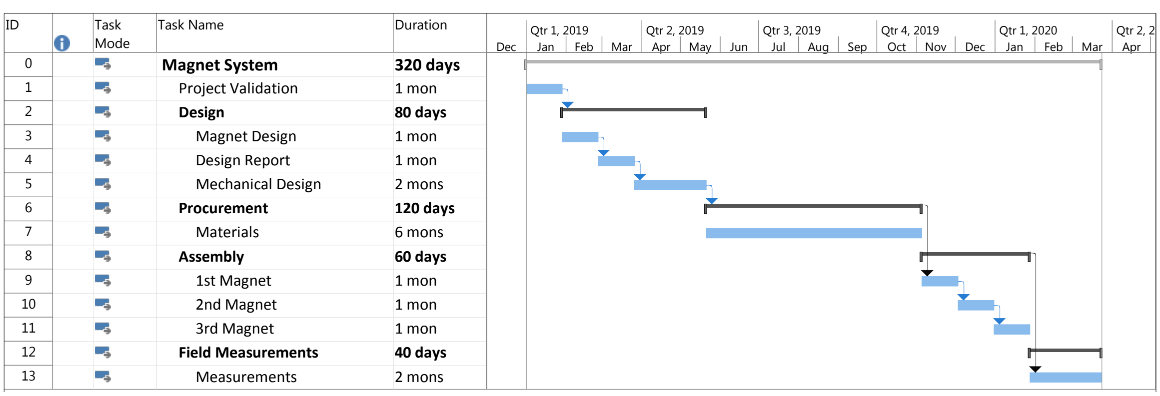

FIG. 12. Timeline for magnet construction for the FASER experiment.

V. TRACKER

A. Overview

The FASER tracker's purpose is to reconstruct the trajectories of energetic charged particles from the direction of ATLAS. It must have sufficient hit resolution to efficiently and reliably identify the signature of two closely-spaced, oppositely-charged tracks from decay of a very high-energy, low-mass particle, where the tracks are separated by the detector's magnetic field. The Letter of Intent [1] showed that requiring tracks separated by at least $300\,\mu$m results in good physics performance.

Given the detector length and magnetic field that can be practically achieved, the high momentum of signal particles, and the uncertainty in the relative alignment of tracking stations in the magnetic bending plane, precise momentum measurement is not expected, and is not required by FASER's searches for new physics. Nevertheless, the detector is designed to achieve the best momentum resolution possible, which will be important for understanding backgrounds and for any auxiliary physics measurements (for example, neutrino studies with FASER).

For cost and schedule reasons, the tracker will be constructed from spare modules of the ATLAS semiconductor strip tracker (SCT) [20, 21]. The SCT Institute Board has kindly agreed to allow FASER to use 80 of the available spare barrel modules [22, 23]. Figure 13 shows a barrel module with 6 on-detector ASICs per side, which are integrated into the module. These "ABCD" ASICs [24] are the first stage of the detector readout and control the module's hit threshold and overall configuration. The SCT modules have proven their reliability in ATLAS during LHC Runs 1 and 2, operating without problems and delivering excellent physics performance.

The basic tracker design is three tracking stations, located at the front, center and back of the spectrometer, where each station contains three planes of SCT modules. An SCT module comprises two pairs of silicon strip detectors glued back-to-back to a central thermal pyrolytic graphite (TPG) baseboard. The modules will be oriented to optimize measurement precision in the magnetic bending plane; a stereo angle of $\pm 20\,$mrad between the two sides allows track position measurements in the non-precision direction as well. The strip pitch is 80 μm allowing track position resolution in the precision coordinate of order 17 μm per plane (assuming 4 SCT planes, and perfect detector alignment). The resolution in the non-precision coordinate is about 580 μm. [25]

The number of tracking planes per station is chosen to give high efficiency and sufficient redundancy, while limiting the amount of material inside the tracking volume. To simplify the system, all three tracking stations have an identical design. The design is also matched to the number of available SCT spare modules. Each SCT module is 6 cm $\times$ 12 cm and the tracking plane is made of 8 modules, giving a square of about 24 cm $\times$ 24 cm covering the full active area of the detector (a 10 cm-radius circle determined by the magnet aperture). Each SCT module has 768 readout channels per side, giving the full tracker system a total of 111 $\times 10^3$ channels.

With the low radiation levels in the TI12 tunnel, the silicon itself does not require cooling to give acceptable efficiency and noise, but the on-detector ASIC chips (12 chips per module) must be cooled. Based on experience during commissioning of the ATLAS SCT, the silicon should be operated at a temperature below 30°C. A water cooling system with inlet water temperature of 5 to 10°C is sufficient to achieve this, but the tracking station will need to

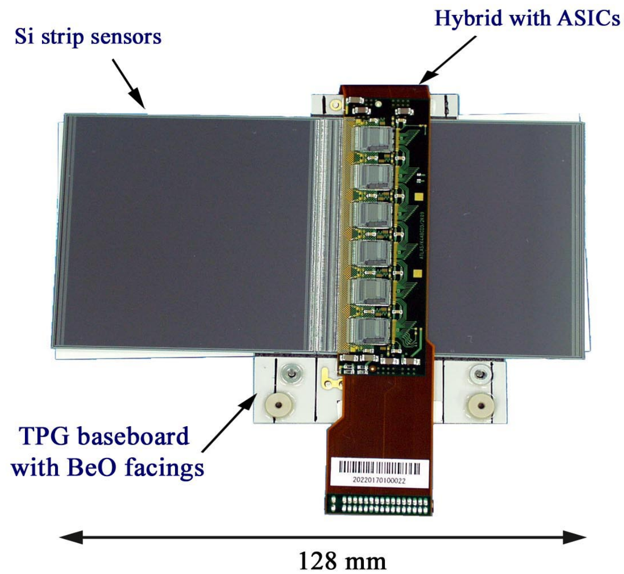

FIG. 13. Illustration of the ATLAS SCT barrel module to be used in FASER.

be enclosed in dry air with a dew point low enough to prevent condensation.

The low trigger rate (expected to be $\sim 600\,\mathrm{Hz}$, as described in Sec. IX), the very small occupancy of the tracking detector (generally a few hits per plane from a single muon traversing the detector, along with a small number of noise hits), and the low radiation level all combine to greatly reduce the challenges of reading out SCT modules in FASER, compared to ATLAS. A simpler readout architecture than ATLAS, based on shipping the data to a custom made board via flex cable, has therefore been chosen. The readout board contains a single FPGA, which carries out simple data processing and error handling before sending the data to an event builder process running on the surface. Whether one readout board per tracker plane (nine total) or per tracking station (three total) will be required is still under study. Either way, the readout system consists of only a small number of boards; these need to be situated about 1 m from the detector to minimize data cable lengths, while remaining out of the stray magnetic field.

The readout system must also collect calibration data from the SCT modules during beam-off periods, when all strips are read to measure noise and efficiency and to tune thresholds. Preliminary calculations show that the proposed architecture can complete the most bandwidth-intensive calibration for the full detector in short dewntimes (≈ 30 minutes).

B. Module Selection and Quality Assurance

SCT modules for the FASER tracker will be selected based on a series of Quality Assurance (QA) tests. For module QA, a so-called Chimaera digital board attached with a Tengja trigger card will be used (see Fig. 14). These were originally developed for a test system of the LHCb RICH upgrade at Cambridge University and were later used for cosmic muon tomography with eight spare SCT modules [26].

The Chimaera digital board is FPGA-based. It sends commands and reads data to/from the module, operated by a PC connected via Ethernet. The Tengja card is an interface between the digital board and an adapter card for the SCT module. The trigger board is directly connected to the digital board via an AUX connector and provides the trigger signal to issue a calibration pulse to the ABCD readout ASICs on the SCT module.

The adapter card connects to the SCT module via a 36-pin connector and provides electrical lines for the clock, command, data, thermistor readout and several voltages to the SCT module. A low-voltage supply is connected with the adapter card to power the ABCD chips, while high voltage for the sensors is supplied via a LEMO connector on the card.

The digital board and trigger board are operated via a GUI-based operating software. The user can set the scan parameters via the GUI window.

The SCT module will be operated at room temperature during the QA and placed on a cooling jig connected with a chiller. The temperature on the module is monitored by using a Negative Coefficient Thermistor (NTC) on the module.

During the QA, the performance of the modules will be checked by a set of tests that includes scans of threshold and trigger latency for qualification of the ABCD chips, and an HV scan for qualification of the sensors themselves. Modules will be ranked according to the test results and those with the highest scores selected for use in the FASER tracker.

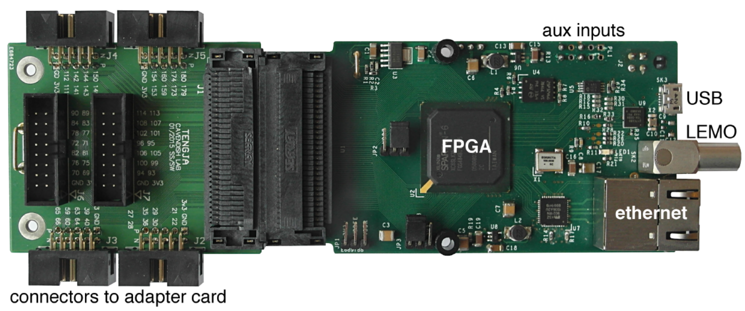

FIG. 14. Chimaera digital board (right) with the Tengja card (left).

C. Mechanics

The tracker consists of three stations deployed between magnet sections. A station comprises three planes of eight SCT modules, arranged over two rows as shown in Fig. 15. The tracker as a whole thus uses 72 of the 80 SCT modules provided by ATLAS, with the remaining 8 available as spares.

1. *SCT modules*

Mechanical assembly of tracker modules into planes re-uses most of the existing parts of an SCT module (Figs. 16 and 17). The aim is to retain the module's primary assembly (two tilted silicon planes + TPG backbone + hybrid) apart from the cooling pipe. Attachment will be made through the three existing kinematics mounts (shown in Fig. 16) so that no major modifications to the SCT modules themselves are necessary, saving time.

Overlap between adjacent modules (a minimum of five silicon strips) is optimized to guarantee full geometric efficiency for near-normal tracks over the entire area of a plane. The arrangement of the two rows is chosen to minimize the material in the central region

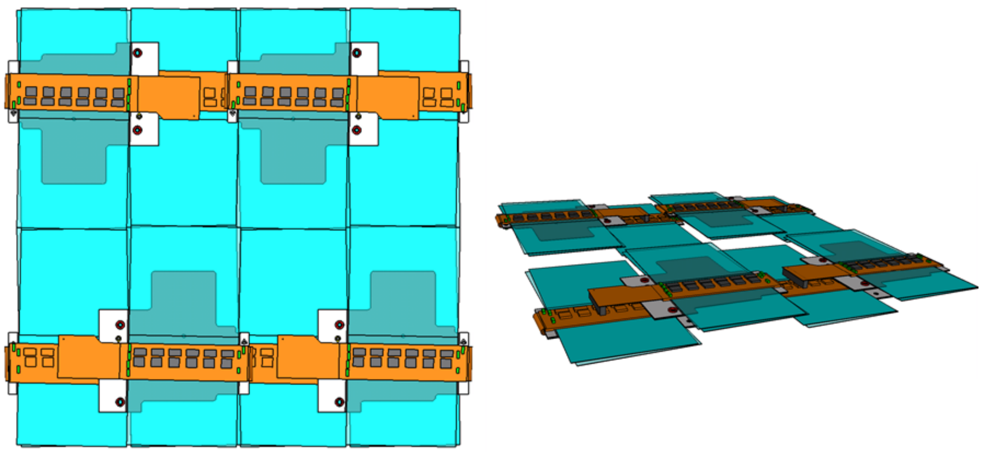

FIG. 15. Left: A full FASER tracker plane, with two rows of four SCT modules. Right: To allow the edges to overlap, modules are staggered in and out of the plane by several millimeters.

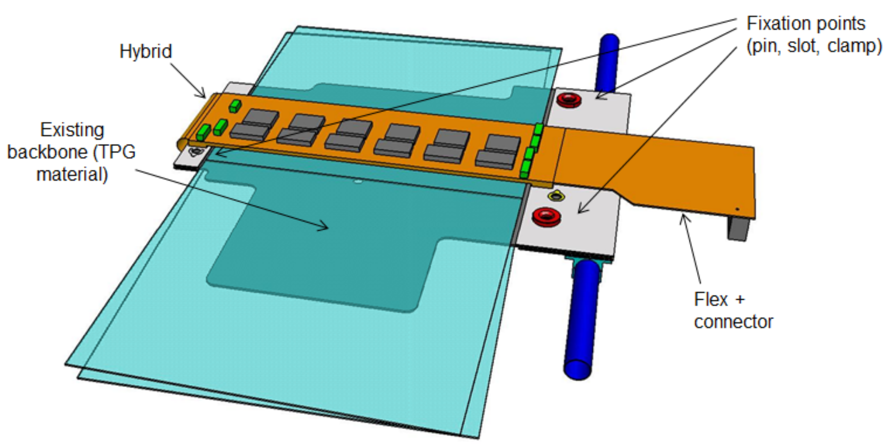

FIG. 16. SCT module with its cooling pipe (dark blue).

(no hybrids, connectors, TPG). This mechanical design also allows replacement of any individual module of a plane, if necessary.

2. *Module frame*

Tracker planes are mounted in a module frame (see Fig. 18) that serves several purposes:

- Provide stable attachments for the kinematic mounts of each of the eight modules in a plane,
- Provide attachment and alignment within 30 microns for each of the three planes within a station,
- Provide the main thermal path for module cooling.

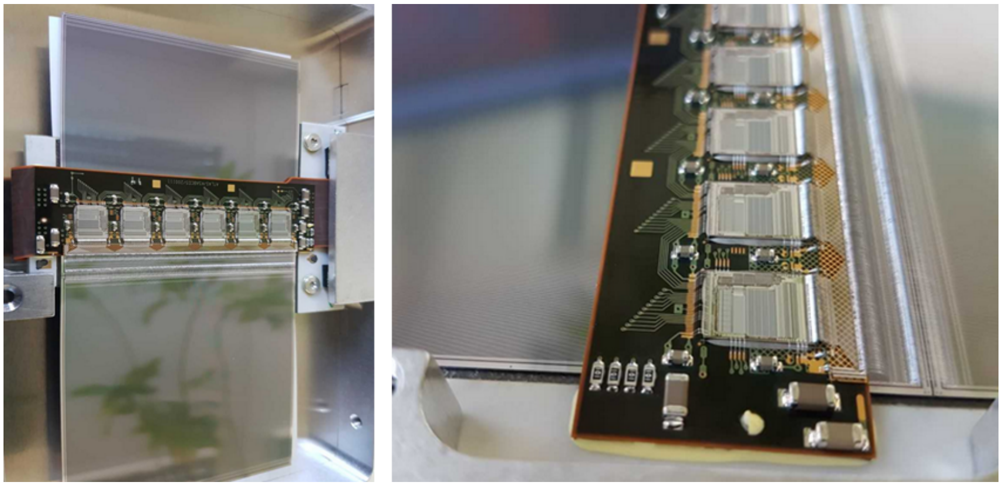

FIG. 17. An SCT module in its holding frame and close-up of the hybrid part.

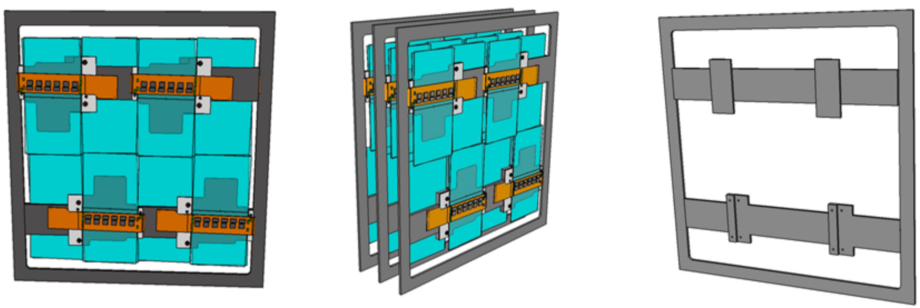

FIG. 18. Left: Module frame with a tracker plane mounted inside. Center: Three module frames assembled into a tracking station. Right: The module frame itself (holes for alignment and attachment are not shown).

Aluminum is a candidate material for the module frame.

The spacer of each module frame (Fig. 19) determines the separation between planes and allows connections for cooling. It is also the main interface to the outer box that allows each station to be aligned with respect to the others and isolated from the tunnel environment. All service connections (power, data, cooling water, dry air flushing) are located on the same side of a station due to tunnel access constraints.

D. Cooling and Humidity Control

Each tracking station is housed in an outer box (Figs. 20 and 21), which provides:

- An attachment for the enclosed planes,
- A light-tight and air-tight environment for the modules inside, to isolate them from the detector tunnel,
- An interface for services (cables, water pipes, dry air pipes),

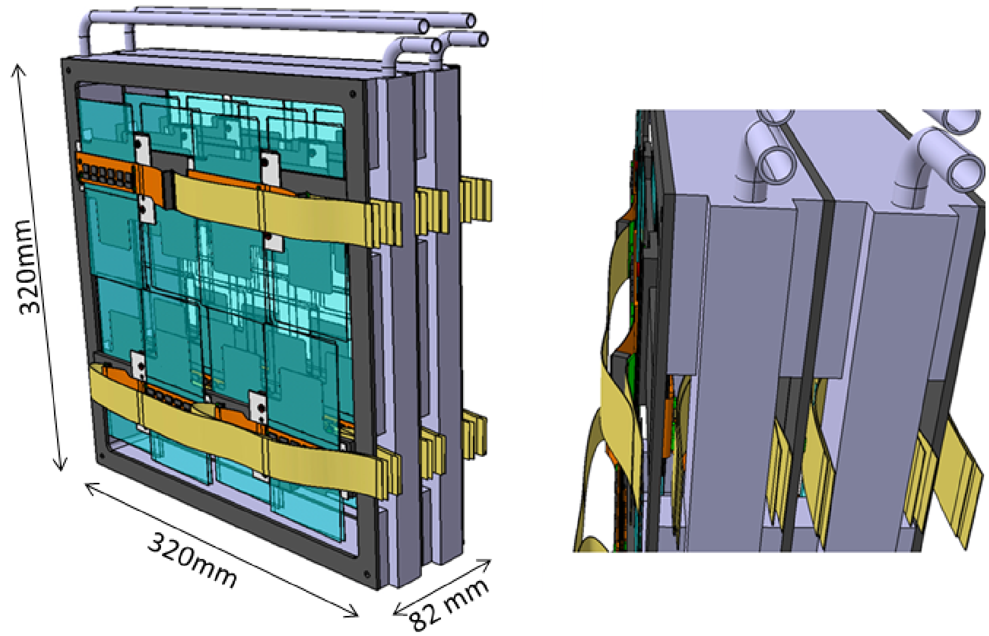

FIG. 19. A tracking station with module frame spacers, showing cooling and electronic connectors.

- An interface to the FASER support system (with an adjustable mechanism for alignment).

The outer box has a patch panel for connectors and fittings. The material will be aluminum to enhance EMI performance by providing a Faraday cage of sorts. It also has fiducial points to establish the position of the station as accurately as possible via survey.

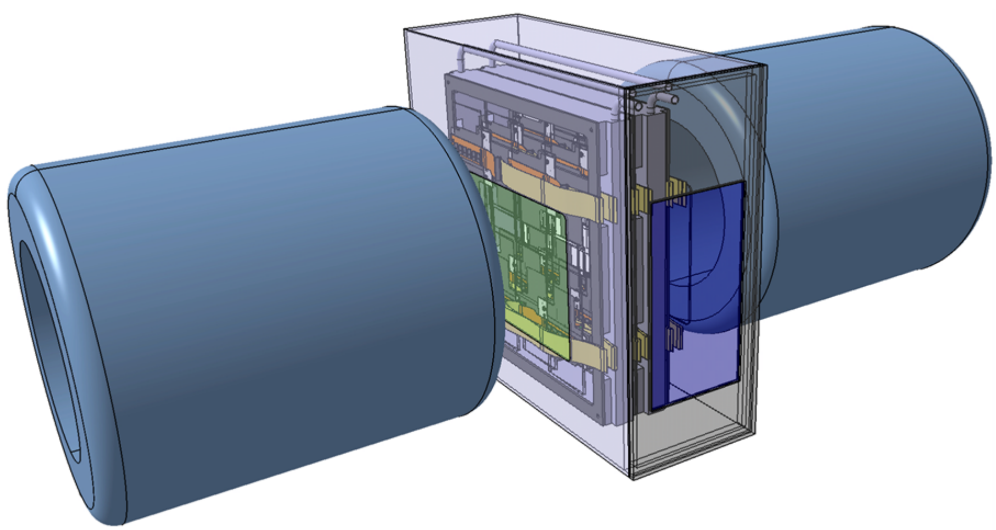

FIG. 20. A tracking station in its outer box between two magnet segments. Nearby scintillators are not shown.

The tracking stations will be cooled by chilled water at an inlet temperature between 5 and 10°C. To provide this chilled water, the CERN cooling and ventilation group (EN-

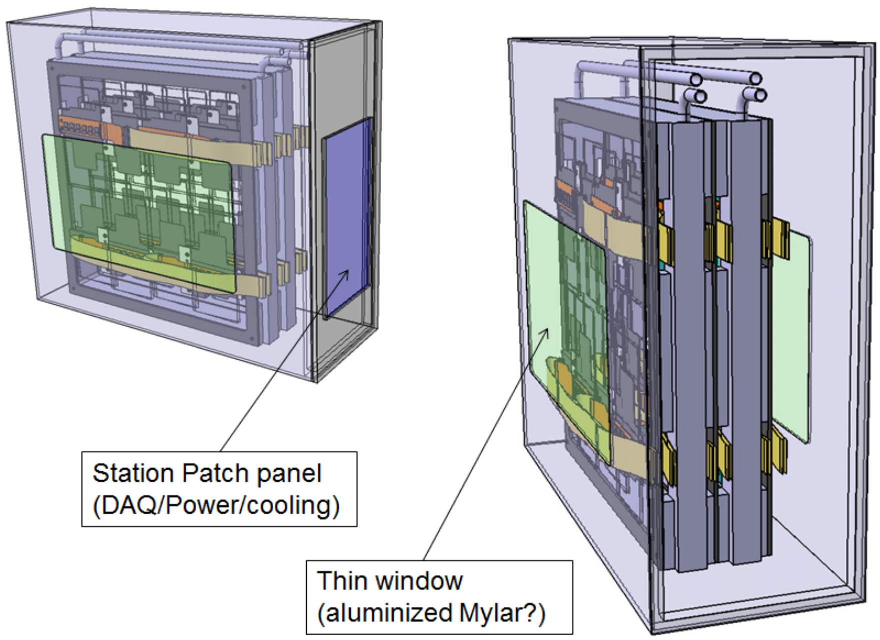

FIG. 21. The tracking station outer box showing the locations of the patch panel and aluminized mylar windows (to minimize material and multiple scattering).

CV) recommends a small chiller installed close to the experiment in the TI12 tunnel. An SMC chiller (model HRS030-AN-020 or equivalent) with a 5-liter water tank can provide cooling capacity up to 1500 W (Figs. 22 and 23) at 25°C ambient temperature. The total cooling capacity required for the three tracking stations is about 600 W. The chiller would be controlled and monitored using an RS485-to-Ethernet converter (Fig. 24). Some insulation (Armaflex) will be required along the water lines to the tracking stations (including the patch panels for the stations). The chiller will be placed in the tunnel at least 30 cm from any obstacle, to allow free air ventilation, but limiting the length of the cooling lines as much as possible.

Model			HRS012-A□-20	HRS018-A□-20	HRS024-A□-20	HRS030-A□-20	HRS050-A□-20	HRS060-A□-20
Cooling method			Air-cooled refrigeration					
Refrigerant			R407C (HFC)				R410A (HFC)	
Control method			PID control					
Ambient temperature/humidity Note 1)			Temperature: 41 to 104°F (5 to 40°C), High-temperature environment specifications (option): 41 to 113°F (5 to 45°C), Humidity: 30 to 70%					
Circulating fluid system	Circulating fluid Note 2)		Tap water, 15% ethylene glycol aqueous solution Note 4)					
	Set temperature range Note 1)		41 to 104°F (5 to 40°C)					
	Cooling capacity Note 3) (50/60 Hz)	W	1100/1300	1700/1900	2100/2400	2600/3200	4700/5100	4900/5900
	Heating capacity Note 3) (50/60 Hz)	W	530/650			600/640	1100/1400	1000/1300
	Temperature stability Note 5)		±0.18°F (±0.1°C)					
	Pump: Rated flow Note 6) Note 7) (50/60 Hz)		1.8 gal/min (7 L/min)(19 psi (0.13 MPa))/1.8 gal/min (7 L/min) (26 psi (0.18 MPa))				23 (0.24 MPa)/28 (0.32 MPa)	23 (0.21 MPa)/28 (0.29 MPa)
	Pump: Maximum flow rate (50/60 Hz)		7.1/7.7 gal/min (27/29 L/min)			34/40	31/42	29/38
	Pump: Maximum pump head (50/60 Hz)		46/62 ft (14/19 m)				50	
	Pump: Output	W	200				550	
	Tank capacity		Approx. 1.3 gal (5 L)					
	Port size		Rc1/2					
	Fluid contact material		Stainless steel, Copper (Heat exchanger brazing), Bronze, Alumina ceramic, Carbon, PP, PE, POM, FKM, EPDM, PVC					
Electrical	Power supply		Single-phase 200 to 230 VAC (50/60 Hz) Allowable voltage range ±10%					
	Circuit protector	A	10				20	30
	Applicable earth leakage breaker capacity Note 8)	A	10				20	30
	Rated operating current	A	4.6/5.1	4.7/5.2	5.1/5.9	5.2/6.0	8/11	8.9/11.5
	Rated power consumption Note 3) (50/60 Hz)	kVA	0.9/1.0	0.9/1.0	1.0/1.2	1.0/1.2	1.7/2.2	1.8/2.3
Noise level Note 9) (50/60 Hz)		dB	60/61			62/65	65/68	66/68
Accessories			Fitting (for drain outlet) 1 pc. Note 11), Input/output signal connector 1 pc., Power supply connector 1 pc. Note 11), Operation manual (for installation/operation) 1, Quick manual (with a clear case) 1 Note 11), Alarm code list sticker 1, Ferritic core (for communication) 1 pc. Power supply cable should be ordered the option (sold separately) or prepared by the user.					
Weight Note 10)			95 lb (43 kg)			104 lb (47 kg)	52 lb (69 kg)	161 lb (73kg)

FIG. 22. Chiller data sheet.

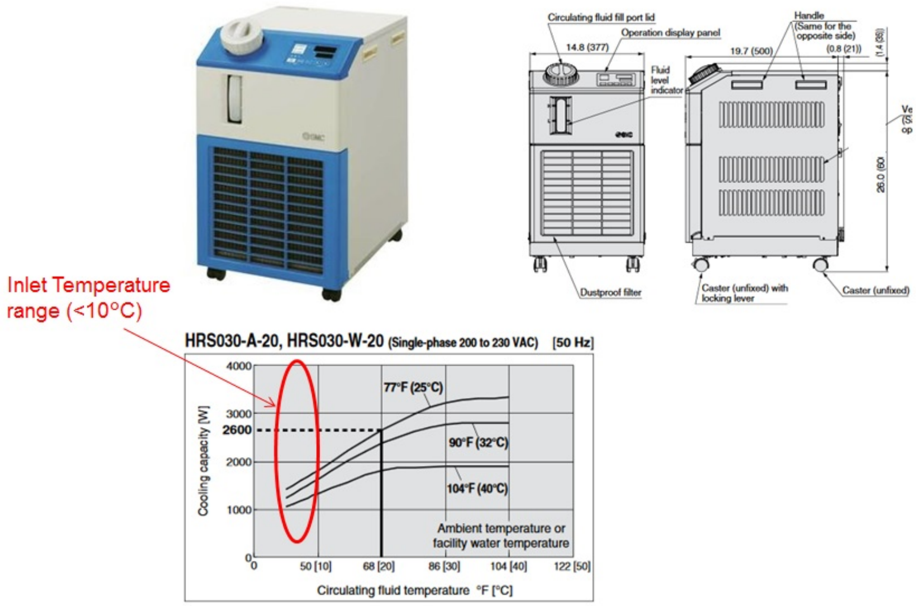

FIG. 23. HRS030 with cooling capacity and dimensions.

Humidity inside the tracker stations will be controlled using compressed air provided by the EN-CV group, which can be installed during LS2. This air has a guaranteed dew point below −40°C and thus meets FASER's requirements perfectly.

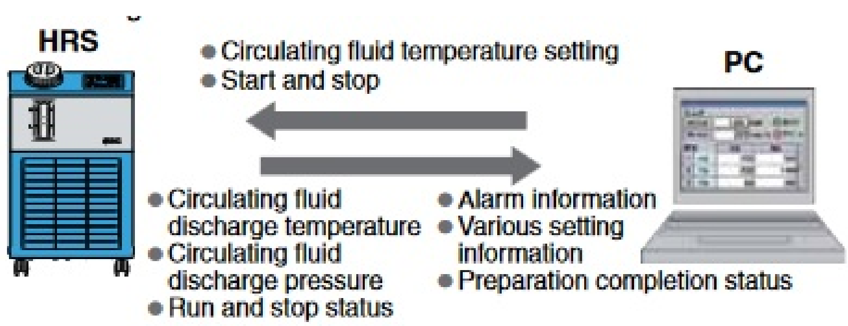

FIG. 24. HRS030 Monitoring via PC.

E. Power, Control and Interlocks

Each SCT module requires three separate power lines. For a non-irradiated module the requirements are typically 900 mA at 3.5 V for the analogue low voltage, 400 mA at 4 V for the digital low voltage, and 150 V reverse bias voltage with less than 0.5 μA leakage current. Several powering schemes are under consideration; the most conservative (baseline) design would provide dedicated power supply channels for each module, i.e., one 3.5 V line, one 4 V line, and one 150 V channel per module. Having separate powering for each module minimizes the impact on the tracker performance of a power supply channel or module failure. It also allows for individualized monitoring and tuning of each module for best performance and reliability.

This powering system can be implemented using eighteen CAEN A2517A 8-channel floating low voltage power supply modules and three A1542 24-channel 500 V floating high voltage boards. These can be housed in two SY4527 power supply crates and controlled via Ethernet. The PS modules should allow for individual channels to be enabled/disabled via an external input signal. Both analogue and digital low voltages are sensed (four wires) for a precise control of the supply voltages.

One of the most problematic issues for the SCT modules is overheating. This could occur due to a cooling failure (e.g., chiller failure, chiller power-cut, cooling circuit clogging), and therefore temperature monitoring is essential.

The tracker's detector control system (DCS) will control and monitor the power supplies, chiller, module temperatures, and environmental sensors inside the stations. Each SCT module has two (one per hybrid side) NTC thermistors ($R25 = 10\,\mathrm{k}\Omega \pm 1\%$) that provide two temperature measurements. In addition, the environmental conditions will be monitored by other, redundant sensors. Each tracking station will have at least four NTCs for temperature measurements (two located at the exhaust pipes to measure the temperature of the spent cooling water and two to measure the air temperature inside the station) and two humidity sensors (e.g., Honeywell) for relative humidity measurements. The complete tracker will include 144 hybrid thermistor channels and 18 environmental sensors to be monitored. The DCS system monitors all supplied low and high voltages and return currents, hybrid temperatures, and environmental conditions. The power supply crate and chiller are controlled by the DCS software, allowing for remote operation of the power and cooling systems. All conditions data will be regularly archived in a database.

Ensuring safe operation of an underground experiment with very limited access is mandatory. Several safety mechanisms at both software and hardware levels will be implemented. The power supply cards contain programmable current and voltage protections, so that a given output channel will be automatically switched off if it exceeds a safe maximum value. The DCS software will also be programmed to take automated protective actions based on monitored sensor readings (e.g., HV-channel ramped-down in case the measured leakage current goes above a certain threshold). Finally, a purely hardware-based interlock system (IS) will provide a final line of protection for the detector modules in the event that software-based protection mechanisms fail. The IS acts as an "ultimate emergency switch," and as such it must be autonomous, simple, and reliable (without relying on software). The baseline is to have custom circuit boards with simple analogue circuitry perform a comparison of temperatures, voltages, and currents to fixed values. Some logic circuits (e.g., OR-ing the comparator outputs for a given module) will be implemented. If thresholds are exceeded, a logical signal (typically TTL-logic) is sent to the interlock inputs of the PS cards to trip the corresponding channel. User intervention would be required to re-enable the channel through the DCS software after diagnosing the fault.

F. Readout

The hardware readout architecture of the FASER tracker is illustrated in Fig. 25.

Each plane is composed of 8 SCT modules and has 2 patch panel boards, one for the top and bottom row, respectively. The connection between the patch panel board and the four modules of a row is made with a module flex circuit having one flex for the two modules with the connectors on the front and another for the two modules with connectors on the back. The patch panel board merges signals and preserves the light-tightness of the tracking station, having only one connector on the external face of the station. The SCT modules will be controlled and read out using a "UniGe USB3 GPIO" board described in Section IX C. This board requires an additional interface board (the "FASER station GPIO" board) which will be designed to allow communication with the SCT modules. The architecture of this interface board is relatively simple: it is composed of LDVS drivers and receivers for the clock, command and two data signals of each module connected (eight per plane) and is connected to the two patch panel boards per plane by cables.

Figure 25 presents an optimized architecture with one "Unige USB3 GPIO" board per station, each driving three planes. If the firmware cannot handle three planes but only one, then nine "FASER station GPIO" boards and nine "Unige USB3 GPIO" boards will be required in the readout box for the overall experiment. The detailed functionality of the readout boards is described in Section IX E.

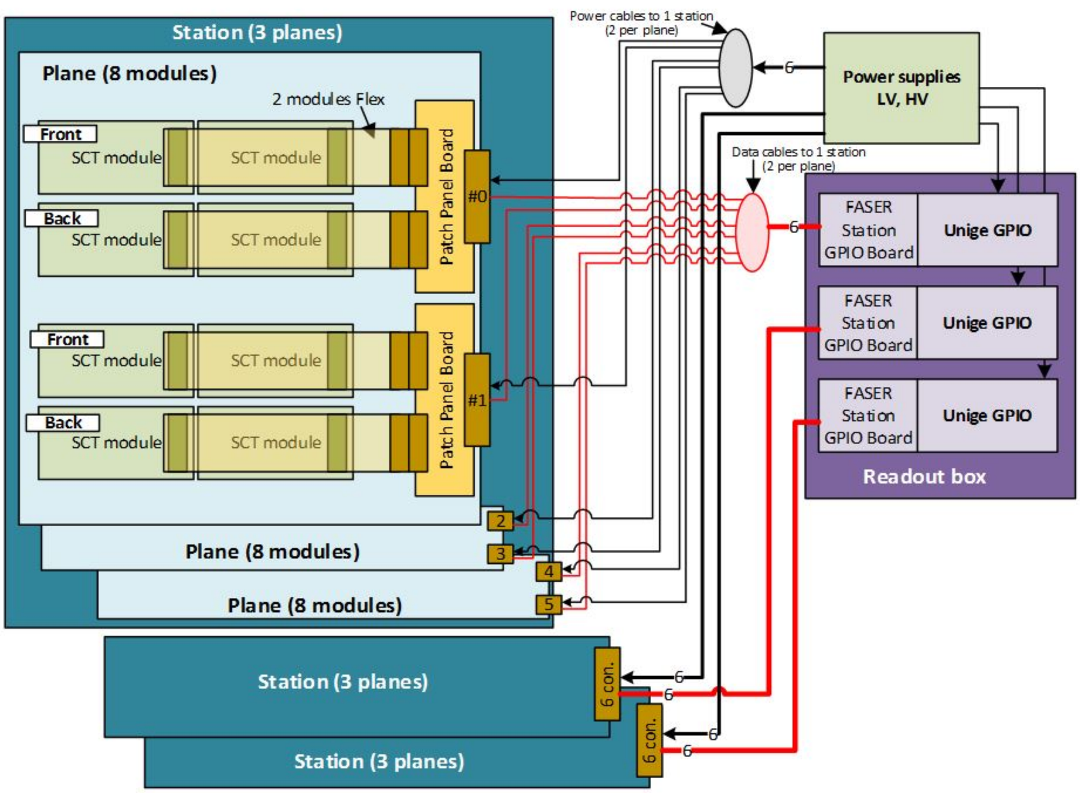

FIG. 25. Tracker readout hardware architecture.

G. Alignment

Preliminary studies of detector performance have assumed perfect alignment, but the precision with which the positions and orientations of the 72 tracker modules can be determined will be an important limitation on the detector's momentum resolution.

The modules in each tracking station will be measured carefully when assembled, and surveyed upon installation. A crucial tool for constraining the module positions will be the sizable flux of high-energy muons during proton-proton collisions (see Section III C). The overlapping geometry of the modules in each station will facilitate precision measurement of their relative positions with these nearly straight tracks.

Without, for instance, a detectable resonant particle whose mass can be used for absolute momentum calibration, certain "weak modes" of misalignment [27] that cannot be constrained by geometrical single track fits may remain and would limit the momentum resolution achievable. Use of calorimeter measurements to help resolve these degeneracies will be explored. Fortunately, FASER's ability to discover a dark photon signal does not rely on precision momentum measurement.

H. Calibration

Calibration of the SCT modules is performed to optimize the threshold and voltage settings and ensure stable performance by measuring the properties of the sensors and their electronics. The calibration procedure also includes basic tests to identify and flag dead or noisy channels.

The most important chip parameter probed by the calibration is the threshold of the discriminator. Each front-end chip has an 8-bit DAC, which allows the threshold to be set globally across that chip. To assure uniformity of threshold across the chip, each channel has its own 4-bit correction DAC (TrimDAC). The DAC steps can themselves be set to one of four different values. This means that uniformity of threshold can be maintained even with uncorrected channel-to-channel variations. Each channel also has an internal capacitor used to inject test charges. By injecting various known charges the analogue properties of each channel can be measured.

In normal operation, the target threshold is set to 1 fC, which should be compared to the 3.6 fC deposited by a minimum ionizing particle at normal incidence. With this threshold, the noise occupancy should be below 5×10^{-4}.

Since radiation damage is not a serious issue in FASER, it should be unnecessary to tune the thresholds frequently. Calibrations will likely be performed during technical stops and machine development periods.

I. Cost and Schedule

The expected cost for the tracking mechanics is detailed in Table VI, and the costs for other Tracker related items are details in Table VII. To efficiently deal with hardware failures, we will need a spare chiller, a spare power supply crate, one spare HV PS and two spare low-voltage power supplies for a total cost of spares of 22 kCHF. The schedule for the production of the Tracker is shown in Fig. 26.

Item	Cost [kCHF]
9 Tracking plane frames	15
6 Spacers frames	20
3 Tracking Station Outer boxes	20
3 set of cooling piping	1
3 Fixation to the base frame	10
Total	**66**

TABLE VI. Budget for tracker station mechanical assembly.

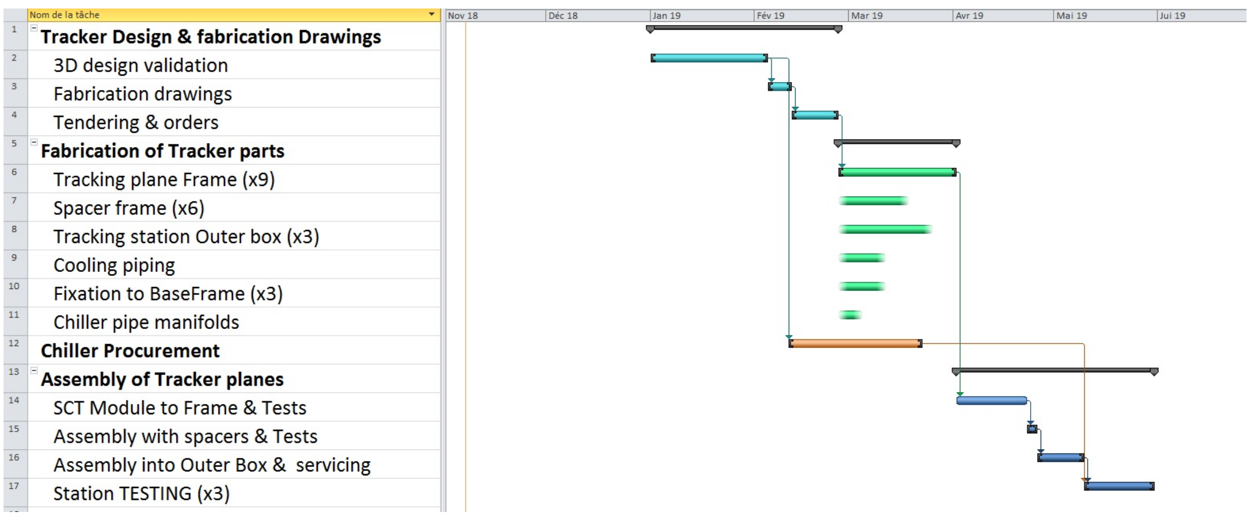

FIG. 26. Schedule of Tracker stations

Item	Cost [kCHF]
Chiller	7
Connection to EN-CV compressed air	1
Power supply crate	10
HV Power supplies	17
LV Power supplies	36
Power cables	6
Flex cables	10
Patch panel and connectors	8
Interlock and environmental monitoring	10
Total	**105**

TABLE VII. Budget for other tracker-related hardware.

VI. SCINTILLATOR VETO AND TRIGGER LAYERS

The FASER experiment has four scintillator stations as shown in Fig. 2. These serve several different purposes in the experiment. The stations are constructed from one or two scintillator layers.

A. Veto Stations

The first two stations, the veto stations located in front of the dipole magnets, are primarily used to suppress events with incoming particles, mostly high-energy muons. To avoid energetic photons from muon bremsstrahlung before the detector entering undetected, an absorber block of lead is placed between the two veto stations, which will either contain the photons completely or generate a shower that is detectable by the second station. The block is foreseen to be 20 radiation lengths (about 11 cm) thick. High-energy muons passing through the absorber will also radiate photons, but in this case the muons will be detected by the front station. To fully suppress background related to muons from the interaction point, each station should detect more than 99.99% of the incoming muons.

Each of the two veto stations consists of two identical scintillator layers. The basic design of each layer is shown on the left in Fig. 27 and consists of a 2 cm-thick, 30 cm $\times$ 30 cm plastic scintillator (Bicron BC 408) connected through a light guide to a PMT (for example a Hamamatsu H3178-51). The PMT will be encased in mu-metal to ensure its operation is not affected by the magnetic fringe fields (up to 5 mT). To minimize the required space in the horizontal direction, the two layers will be installed with a 90° angle between them. The transverse size of the scintillator layer is larger than the magnet aperture to further ensure no charged particles can enter undetected, and the 2 cm thickness is chosen to provide a very high single layer detection efficiency (well above 99%). The two independent layers provide redundancy and ensure a very high veto efficiency, which can easily be measured *in situ*, as there should be no correlated inefficiencies.

B. Trigger and Timing Station

The trigger/timing station located after the first magnet and first tracking station is used to detect the appearance of a charged particle pair from the decay of a LLP in the decay volume of the first dipole magnet. It provides the primary trigger signal and will also be

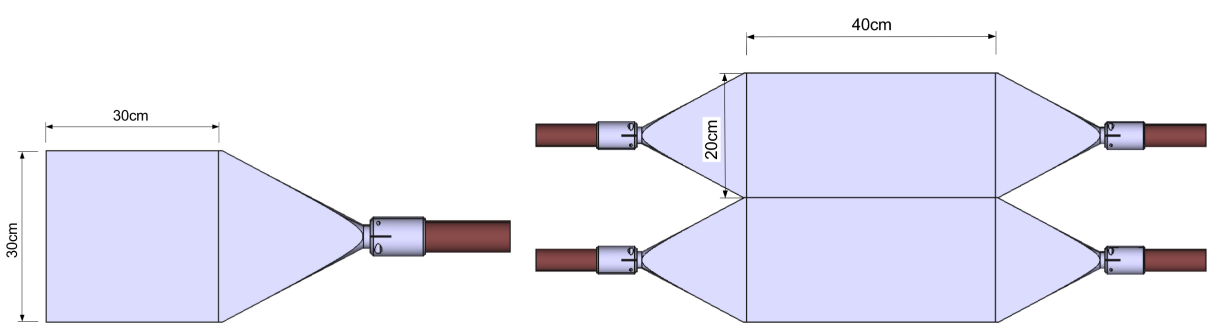

FIG. 27. Design of the two types of scintillator layers.

used to precisely measure the arrival time of the signal with respect to the *pp* interaction at the IP. The precision should be better than 1 ns to suppress any non-collision backgrounds. Additional signal to background suppression is provided by measuring the light yield since the prospective signal has two charged particles unlike most background types, but this is not expected to be a decisive discriminator due to the large event-to-event Landau fluctuations in the light yield. Since this layer is in the tracker volume, the material should be minimized while still maintaining an efficiency above 98%. Finally the active area of this station will also be made large enough to cover most of the magnet front surface. This is to detect muons coming in at an angle, missing the veto stations and the first two tracker stations, but causing an electromagnetic shower in the magnet material and thus a detectable shower in the downstream calorimeter that mimics a photon-only signature.

The trigger/timing station consists of a single scintillator layer made from two 20 cm $\times$ 40 cm scintillator blocks with a thickness of 1 cm (2.5% radiation length) shown on the right in Fig. 27. The layer is split in to two blocks to reduce the size of the light guide, keep high detection efficiency, and improve the timing precision by minimizing the vertical time-walk. The horizontal time-walk is compensated by using the average signal time of the two PMTs, hence the requirement to have a PMT on each side of the scintillator layer. With this setup, the timing resolution will be limited by the precision of the readout electronics.

C. Preshower Station

The final trigger/preshower station is located just in front of the calorimeter. This station provides an additional trigger signal which, if needed, can be used in a coincidence with the first trigger station to reduce the rate of non-physics triggers. It will be preceded by a thin layer of radiator (about 2 radiation lengths of tungsten or lead) to create a simple preshower detector that with high efficiency can detect a physics signal of two close-by energetic photons. The preshower helps distinguish this physics signal, which would otherwise leave only large energy deposition in the calorimeter, from deep inelastic scattering of high energy neutrinos in the calorimeter. This is needed because the calorimeter does not have any longitudinal segmentation; see Sec. VII.

To reduce the number of different scintillator configurations, the preshower station will use two layers in the same configuration as the veto stations with half of the radiator material placed in front of the first layer and the second half in between the layers. A double arm-layer as used in the timing station, which would have provided another precise timing measurement, is not possible due to physical constraints from the tunnel wall on one side of the detector.

D. Powering

Each PMT requires 1.5-2.5 kV. To avoid correlated inefficiencies, each PMT should be supplied by its own power supply channel. The baseline is to provide this with two 6-channel high-voltage power supplies (CAEN V6533). These will be installed in the VME crate that is also used by the PMT readout described in Sec. IX B. An alternative option is a separate high voltage power supply module in the power supply crates used for the tracker stations.

Item	Cost [kCHF]
8 scintillator plates	4
10 light guides	8
Assembly	5
10 PMTs	20
2 HV power supply	8
HV Cables	1
Lead absorber block	1.5
2 Tungsten plates	4.5
Total	**52**

TABLE VIII. Budget for scintillator system for the FASER experiment.

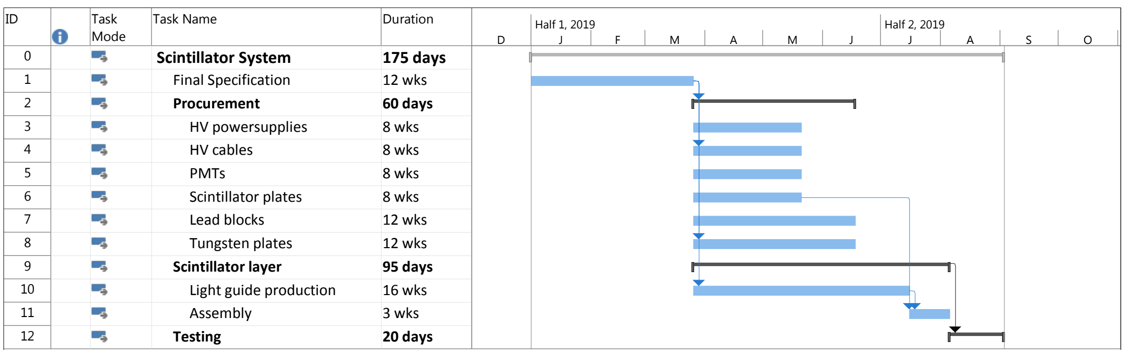

FIG. 28. Timeline for the construction of the scintillator system.

E. Cost and Schedule

The cost of the different scintillator components apart from the readout is estimated in Table VIII. One spare PMT and a spare HV power supply module at a total cost of 6 kCHF is also foreseen to be purchased.

The expected schedule for the construction of the FASER scintillator system is shown in Fig. 28. The schedule assumes that the light guides and PMT mounts are produced by an outside company, while the assembly is done by the CERN scintillator lab.

VII. ELECTROMAGNETIC CALORIMETER

The electromagnetic calorimeter provides strong identification of high-energy electrons and photons over muons and hadrons and allows to measure their energy. Since for most signal events the e^+e^- or photon pair is separated by less than a few millimeters, it is not feasible to measure the individual particle energies. The main calorimeter requirement is, therefore, to measure the total electromagnetic energy with good accuracy for multi-TeV deposits in a compact detector.

A. LHCb ECAL Modules

The calorimeter will use four spare LHCb outer ECAL modules [28], shown in Fig. 29. The LHCb Collaboration has kindly agreed to allow FASER to use six of these modules on indefinite loan. These are so-called Shashlik-type calorimeters with interleaved scintillator and lead plates. Each module contain wavelength shifting fibers penetrating the full calorimeter, which extract the signals to a single PMT on the back. The modules are 754 mm-long, including the PMT, and have transverse dimensions of 121.2 mm $\times$ 121.2 mm. The full FASER acceptance is therefore covered with these four modules. The calorimeter contains 66 layers of 2 mm lead and 4 mm plastic scintillator, for a total depth of 25 radiation lengths. The baseline is to use the same PMT as LHCb (Hamamatsu R7899-20), but it might be necessary to use either a different model or attenuate the light in order to not saturate at the highest energy. The high voltage for the PMTs will come from the same type of power supply as for the scintillator system and the same readout will be used as well. The energy resolution for TeV deposits in such a calorimeter is expected to be around 1%, although this will be degraded at the highest energies, since 25 radiation lengths will not fully contain all such showers. It is not foreseen to have a leakage detector behind the calorimeter as the impact on the signal efficiency is expected to be negligible, and there is very limited space.

B. Cost and Schedule

The cost of the different calorimeter components, apart from the readout, is estimated in Table IX. One spare PMT, with a cost of 2 kCHF, is planned, but no additional HV modules are planned, as the spare is shared with the scintillators. As noted above, six calorimeter modules will be on indefinite loan from the LHCb collaboration.

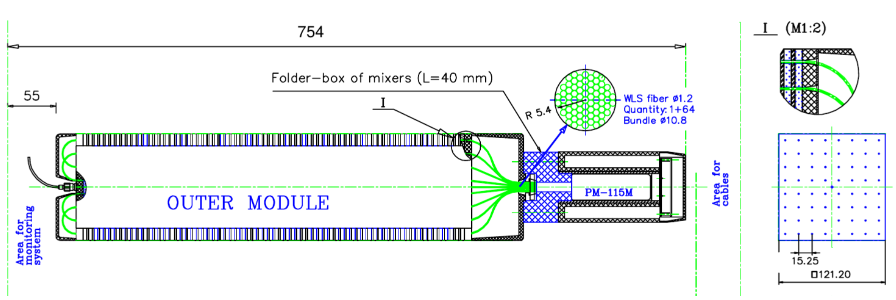

FIG. 29. Design of the LHCb outer ECAL modules [28].

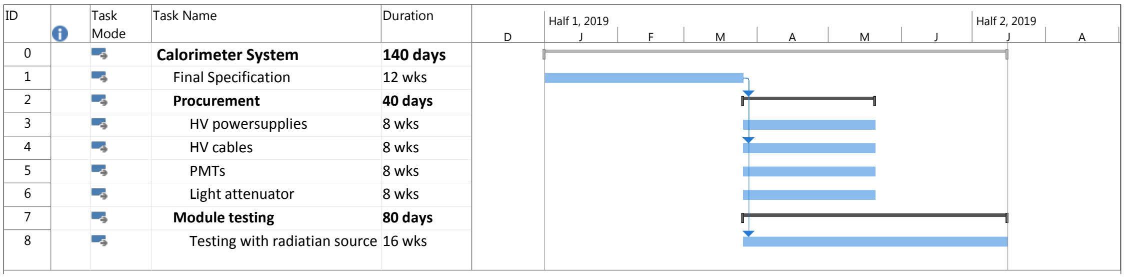

FIG. 30. Timeline for the preparation of the calorimeter system.

The expected schedule for the preparation of the FASER calorimeter system is shown in Fig. 30. Tests of each module will be done in collaboration with LHCb ECAL experts.

Item	**Cost [kCHF]**
4 PMTs	8
HV power supply	4
HV Cables	1
Total	**13**

TABLE IX. Budget for the calorimeter system for the FASER experiment.

VIII. DETECTOR AND MAGNET SUPPORT

The detector is mechanically supported by a common structure (the detector support structure or main cradle), which also allows fine alignment of the different components. At this stage a preliminary conceptual design has been carried out, which will be refined over the next months including adding the sub detectors specific needs (for the calorimeter, scintillators, and Tracker stations).

A. Requirements

The main requirements for the detector support structure are:

- Providing a common mechanical support for the detector components and the magnet segments (limiting also the number of interfaces needed for the detectors with respect to the floor);
- Providing a guaranteed safe distance between the magnets of 20 cm;
- Allow for fine alignment between the components *in situ*;
- Provide stable alignment between the components (especially the tracking stations) as precisely as possible given the detector design. A placement alignment precision of 100 μm and stability of the alignment of 100 μm would be desirable to provide good spectrometer momentum resolution for high energy charged particles.

B. Design

The main support structure as shown in Fig. 31 is a conceptual design to describe the key features. It is a common cradle which positions each part independently and acts as the main interface to perform the necessary alignment fine tuning with respect to the theoretical beam axis. The base plate (or base frame) shown in Fig. 32 will be the first element to be

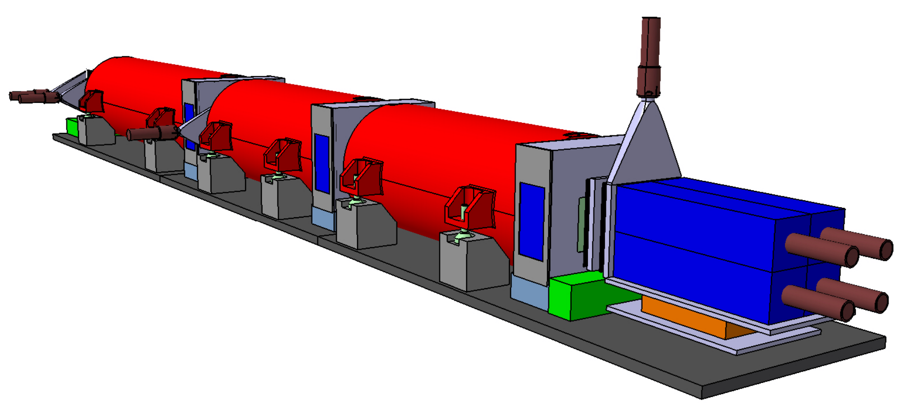

FIG. 31. Detector support overview

installed in the tunnel prior to any detectors assembly. It will be surveyed in 3-dimensions and aligned by means of jack screws (or equivalent) before being clamped to the ground to provide a stable interface for the detectors. The conceptual approach is to limit the civil engineering work by removing individual fixations per detector. Instead the fixation is only done once at the main support structure level. Due to size limitation (in particular for CNC machining, and transport to TI12), this base plate or frame has to be build in segments to be precisely re-assembled using calibrated pins when installed in the tunnel. The overall precision (both for machining and assembly) that can be expected here is at the order of 100-200 microns (along the beam axis).The lateral precision over 50 cm (plate width) could be lowered to 50 microns at best. This strongly depends on the CNC machine type. The in-plane alignment between plate segments is performed and frozen by the machining itself (no fine tuning). Only the out-of-plane alignment (perpendicular to the base pate) is needed to provide a common theoretical axis to be aligned with the beam axis. This will be carried out with the CERN survey group by means of a laser system and targets attached to the base frame. The precision of the measurement is of the order of 100 microns (to be confirmed with CERN survey group).

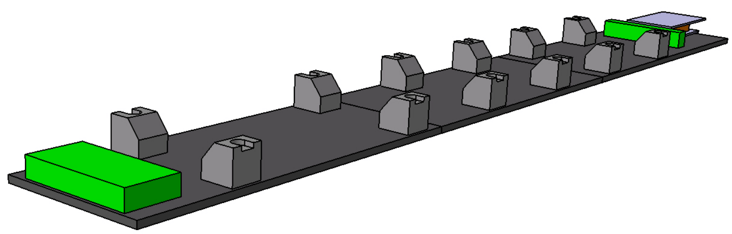

FIG. 32. Base plate with detector resting blocks.

The material choice is based on stability requirements, magnetic compatibility and cost (procurement, and fabrication). Stainless steel is a good candidate since it is less sensitive to temperature changes (as compared to aluminum). As an example, over a 5 m long steel part (the overall base plate length), and for 1 degC change in temperature, the length variation is of the order of 70 microns. The design is still open and will be driven mostly by the final clearance between the Magnet outer cylinder and the floor.

Fig. 33 shows the basic principle of the magnet vertical tuning by means of jack screws (4 per magnet segment). The in-plane alignment is done by inserting the jack screw into machined resting blocs (acting as a pocket), but such as system cannot guarantee any better than 1 mm precision. If a greater precision is requested from the physics requirements (to be confirmed), lateral jack screws could be added to solve this issue. An example of the support structure for a Tracker station is shown (with the block in light blue), that could either be a precisely machined block or some linear stages (in 2 axes, perpendicular to the beam axis) to allow for fine tuning. A similar support structure would be used for each detector component. All of the interfaces to the base plate are aligned (in-plane) precisely by the CNC plate machining and will comply with the above mentioned level of precision.

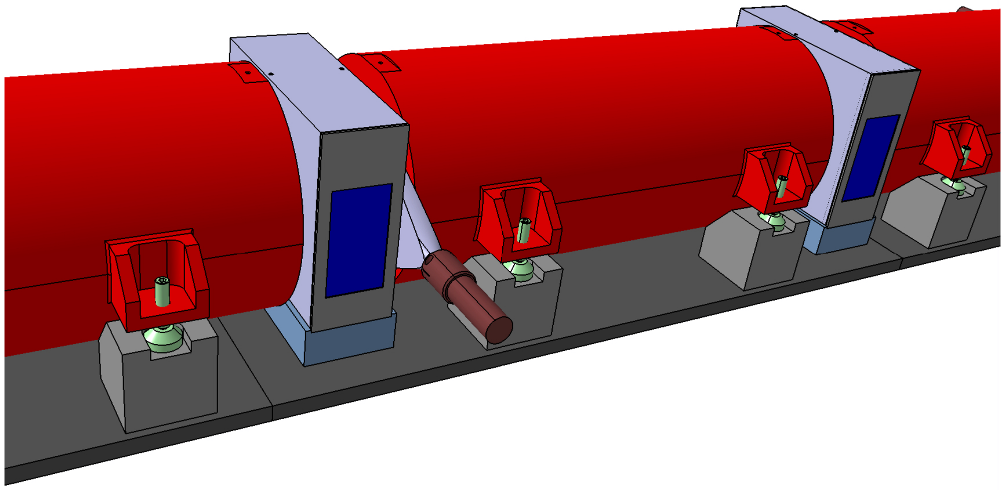

FIG. 33. Close up of magnet and Tracker station alignment.

C. Cost and Schedule

The expected cost estimate for producing the detector support structure is presented in Table X, based on the preliminary conceptual design.

Item	**Cost [kCHF]**
Stainless steel Plate or frame	30
Detector interfaces	5
Magnet adjustable feet	25
Total	**60**

TABLE X. Budget for Detector support for FASER experiment.

Fig. 34 shows the schedule for producing the FASER support structure. This is interconnected with the other sub-system timelines, for instance the design completion of the Tracker stations. Such interconnection will apply also for the other sub-parts since the design will be a common effort. The starting date is given as an example.

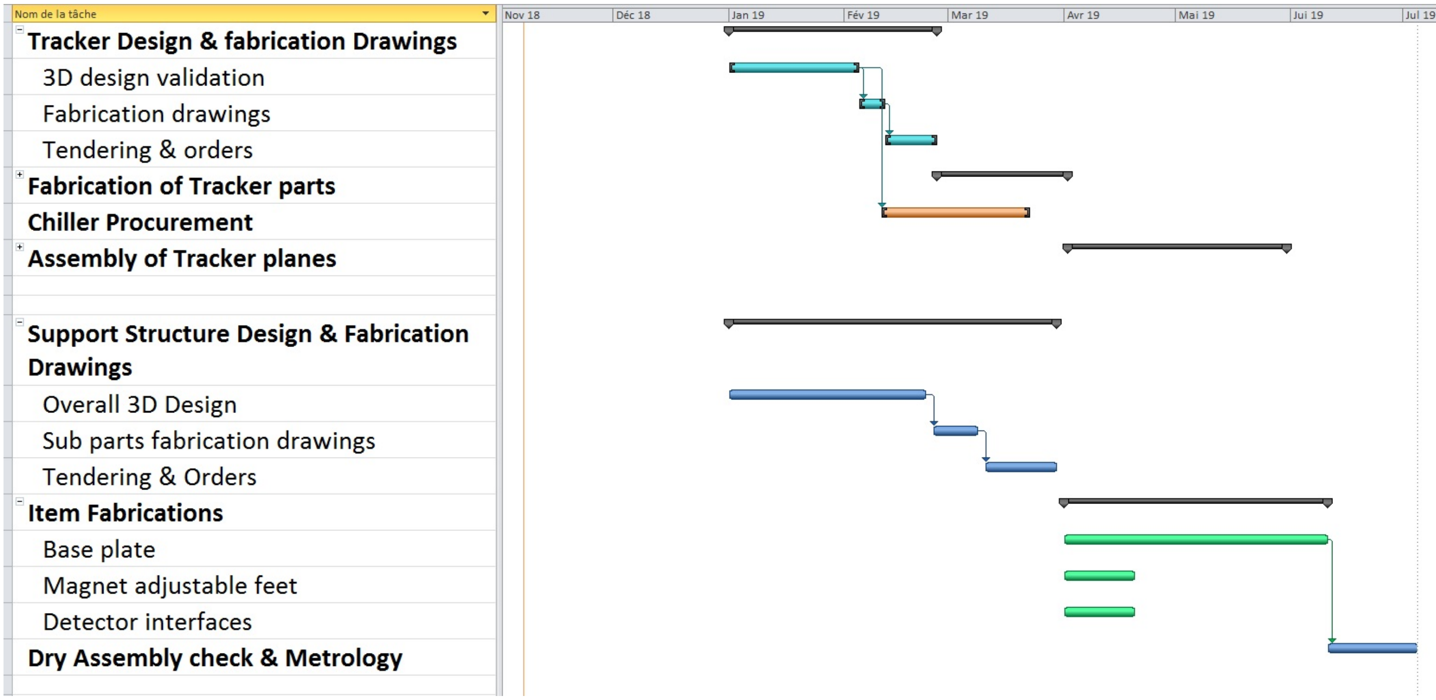

FIG. 34. Schedule for FASER Support structure.

IX. TRIGGER AND DATA ACQUISITION

The FASER detector is located in a low-background area of the LHC. The aim is therefore to trigger the read out of the full detector on every detectable high energy particle passage. For redundancy and to measure the detector efficiencies, all scintillators layers and the calorimeter will provide trigger signals. The scintillator trigger threshold will be below that of a single minimum ionizing particle, while the calorimeter threshold will be set to trigger on any significant electromagnetic shower. The expected trigger rate per source is shown in Table XI under the assumption that the rates will be completely dominated by high-energy muons, as shown by simulation studies and *in situ* measurements (see Sec. III). The rates are therefore highly correlated, and the total physics rate is expected to be approximately equal to the highest rate source. In case the rate of a trigger source is too high for the available readout bandwidth, the trigger system will allow triggering only on combinations of certain (anti-)coincidences, such as a calorimeter signal or signal in preshower scintillators with no signal in the veto scintillators. In this case a fraction of the individual trigger sources will still be recorded in the form of "pre-scaled" triggers for calibration and alignment purposes.

To be able to identify a rare non-SM physics signal, as much information as possible about each triggered event will be read from the detector. For the tracker, the information available is limited by the pre-existing readout chip to three 25 ns-wide time bins (hit/no hit) per strip. The trigger signal will be adjusted so that signals originating from highly relativistic particles from the IP will be seen in the central bin. For all PMTs, the full waveform will be read out in a much wider time window (about 1 μs) to accurately reconstruct the charge and time of each pulse and to ensure there is no anomalous activity near the time of a possible physics signal. Basic information from the trigger system, such as the source of the trigger firing, event number, etc., will also be read out.

Source	**Rate [Hz]**
Veto scintillators	360
Timing scintillators	640
Preshower scintillators	360
Calorimeter ($E > 100$GeV)	< 5 Hz
Random trigger	10
Total	**650**

TABLE XI. Expected trigger rate per trigger source at a instantaneous luminosity of 2×10^{34} cm^{-2} s^{-1}. Note that since the primary source of triggers is high-energy muons passing through the full detector, the rates are highly correlated. The random trigger rate is meant for pedestal calibration and noise monitoring.

Source	**Size [kB]**
PMTs	20
Tracker	5
Trigger Logic Board	< 1
Total	**25**

TABLE XII. Average expected event sizes. The PMT size assumes a 1 μs readout window (500 samples of 2 ns).

A diagram of the trigger and date acquisition system is given in Fig. 35. The system is designed to minimize the amount of equipment underground where access is highly limited, while having the minimal number of signal cables coming into or out of the experiment, as those have to be more than 500 m long. The trigger system will run synchronously to the 40.08 MHz clock of the LHC. All of the PMTs will be read out by a high-speed digitizer, which also provides a trigger signal when the signal on a pair of PMTs exceeds a preset threshold. The trigger signals are combined in the Trigger Logic Board (TLB) to form the final trigger signal (L1A) after possibly applying coincidence and/or veto logic, which combines information from the scintillator and calorimeter PMTs. The L1A is sent to the PMT digitizer and the tracker readout boards to initiate the readout of all of the detectors. The readout is done over Ethernet to a PC (the Event Builder) located on the surface. The expected event size is listed in Table XII; given the trigger rate, a single Gbit/s connection is more than sufficient. The event builder merges the data fragments into events that are buffered and undergo basic data quality checks on a second PC before being copied to CASTOR for permanent storage. The various Detector Control Systems (DCS) in the experiment will be connected over the same Ethernet connection to a separately-controlled PC on the surface, which will also be used to monitor and control the trigger and readout system. The Ethernet connection to the surface will use a single optical fiber pair connecting the TI12 experimental area to the SR1 surface building, close to the ATLAS control room.

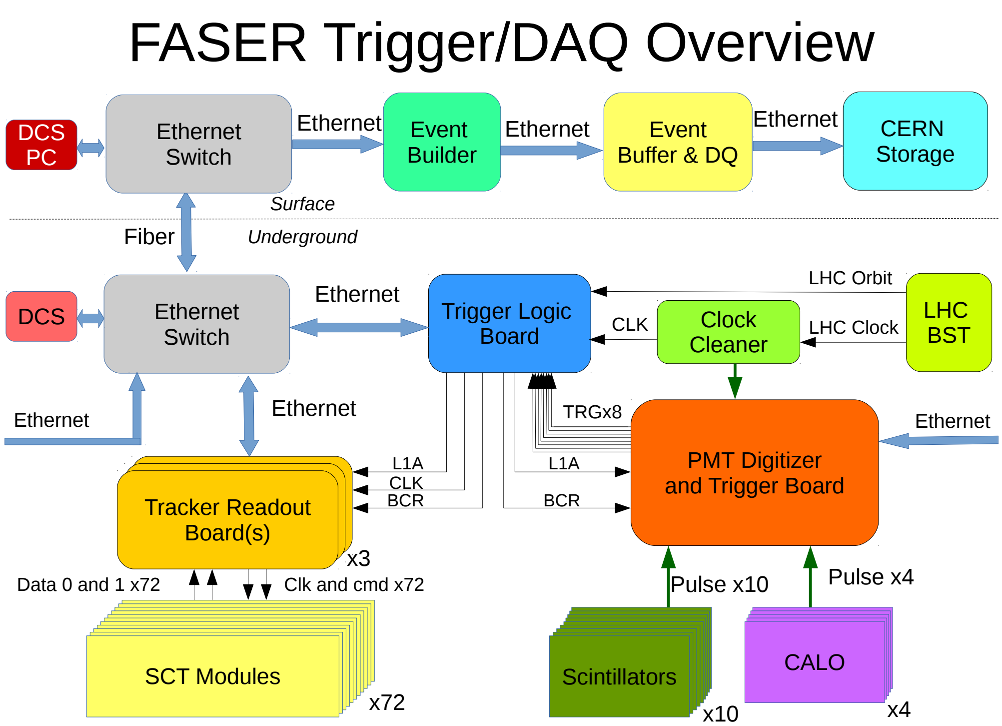

FIG. 35. Schematic diagram of the Trigger and DAQ system.

A. LHC Signals

The LHC clock (40.08 MHz) and orbit signal (11.245 kHz) are part of the Beam Synchronous Timing (BST), transmitted to beam instrumentation equipment around the LHC over optical fibers using the TTC system [29]. For FASER, this signal will be received over a single optical fiber by a BOBR VME card [30] from the LHC beam instrumentation group, which extracts the LHC clock and orbit signal to two front panel LEMO connectors. It also provides additional information, such as the global time and accelerator mode over the VME bus, but this is not foreseen to be used.

The LHC clock provided by the BOBR has a non-negligible jitter, changes during the ramp of the LHC, and is not guaranteed to be continuous. The clock will therefore pass through a Si5344 jitter attenuating clock multiplier chip [31] mounted on an evaluation board in the VME crate. This provides a ultra low-jitter clock (90 fs), which tracks the LHC clock at a programmable speed, provides a holdover clock in case the input clock disappears, and provides a zero delay mode that ensures the clock is always at the same delay with respect to the beam interactions. Daily and seasonal variations of up to several ns in the phase of the clock is expected, due primarily to ground temperature variations, and the absolute timing offset will therefore need to be monitored and calibrated with high-energy muons from the IP.

The orbit signal provides a 25 ns pulse once per turn of the LHC. It will be sent directly to the input of the TLB and used to reset the bunch counters in all of the readout.

B. Scintillator and Calorimeter Trigger and Readout

The readout and triggering of the scintillator and calorimeter PMT signals will be done by a single, 16-channel CAEN digitizer VME board, VX1730 [32]. The digitizer samples each channel at 500 MHz using 14-bit ADCs and stores the samples in a circular buffer of programmable length. For FASER, a 1 μs long buffer (500 samples) will be used to capture the full PMT waveform and preceding data for offline analysis.

Trigger thresholds can be set separately on the ADC value of each channel and a trigger signal generated for each pair of input channels. For the FASER scintillators, the threshold will be set well below that of a MIP, but both PMTs in a station will be required to be above threshold to keep noise triggers to a negligible level. For the calorimeter, a larger signal is required, but any of the four channels can raise a trigger signal. The board supports using a digital constant fraction discriminator logic for generating the trigger signal, but since the rise time of the PMT is shorter than the 2 ns sampling time, this is not foreseen to be used. The trigger logic runs at 125 MHz and provides eight trigger signals transmitted as 32 ns-wide pulses to the TLB. The 32 ns-wide pulse is needed since the signals are not synchronized to LHC clock and therefore could arrive at the TLB in separate bunch-crossings and not trigger the coincidence logic. The digitizer trigger signals will not on their own trigger the digitizer readout, which instead is triggered on the signal from the TLB.

The digitizer does not run synchronously with the LHC clock. Instead it provides a trigger time in units of 16 ns since the last LHC orbit signal. This does not provide sufficient precision to determine if any signal is in time with the collisions, which have an intrinsic spread of around 180 ps. Therefore the LHC clock from the clock cleaner is also sampled on one of the input channels to measure the relative timing with respect to a bunch crossing at a precision better than the 1 ns precision targeted for the timing scintillator station. The absolute timing offset will be calibrated using high-energy muons.

The VX1730 board is able to buffer up to 1024 events while waiting to be read out. The board does not support direct readout over Ethernet, but instead has to be read out over the VME backplane. The maximum readout speed over the VME exceeds 50 MB/s, well above the needs of FASER. The readout to Ethernet will be done with an SIS3153 Ethernet to VME interface [33], eliminating the need for a single-board computer in the VME crate.

C. "UniGe USB3 GPIO" Board

The Trigger Logic Board and the Tracker Readout Boards require custom firmware in an FPGA interfaced to the hardware signals and Ethernet readout/control. For both cases, this will be based on a general purpose board developed by the University of Geneva for slow control and readout of particle physics ASICs, detectors tests, and qualification. This board, named the "Unige USB3 GPIO" board, can also be used for small experiments using several detectors like the FASER experiment.

For each application, a dedicated interface board needs to be developed to interface to the hardware. This interface board will be connected with the connectors of the analogue and/or digital inputs/outputs (the latter only, for the case of FASER).

Figure 36 describes the architecture of the "Unige USB3 GPIO" board, and a photo of the board is shown in Fig. 37. The core of the board is composed of a CYCLONE V A7 FPGA having 150K Logic Elements or 56.5K Adaptive Logic Modules (1 ALM = 8 inputs LUT and 4 registers), i.e., 226K registers, 6.86Mbits + 836Kbits of RAM, 156 double 18x18 DSP blocks and 7 PLLs. The FPGA is driven by a 40MHz local clock or an external clock thanks to a LVDS clock input through a dedicated 2 pin LEMO connector.

Concerning the experiment interface of the board (located on the left side of Fig. 36), the FPGA is connected to a 40MSPS 12-bits 8 channels ADC and 2MSPS 16-bits 8 channels DAC, providing a versatile analogue interface through a dedicated 20-pin connector. For the digital part, two connectors dispatch 117 single-ended inputs/outputs, which can be configured through the FPGA firmware. In addition, two NIM inputs and one NIM output are available through LEMO connectors.

For the readout and the slow control, two interfaces can be used. Already fully implemented is USB3, where the communication is handled by a FX3 Cypress micro-controller providing a 32-bit parallel interface clocked at 100MHz with the FPGA having a maximum potential data throughput of 3.2Gb/s, but finally limited to 1.8 to 2.5Gb/s depending on the host PC performance when connected with the board. A full and very versatile FPGA library for the readout and the slow control is already available for this USB3 link; see below. Additionally an Ethernet PHY IC is also embedded on the board and should be able to provide a 10M/100M/1000Mbps connection. This interface is currently under development (Ethernet MAC and UDP integration into the FPGA firmware library) and should be available in the 1st quarter of 2019.

Finally a S9S12G64, G96 or G128 (for 64, 96 or 128KB of Flash memory) 16-bit micro-controller can also be mounted on the board. This device is typically used for slow control measurements (e.g., temperature, humidity, or voltage monitoring). This micro-controller has a CAN bus interface, 12 10-bits channels ADC, and a couple of digital I/Os which can also be configured in SPI or UART mode. All of those I/Os are accessible through a 2x20-pin header connector.

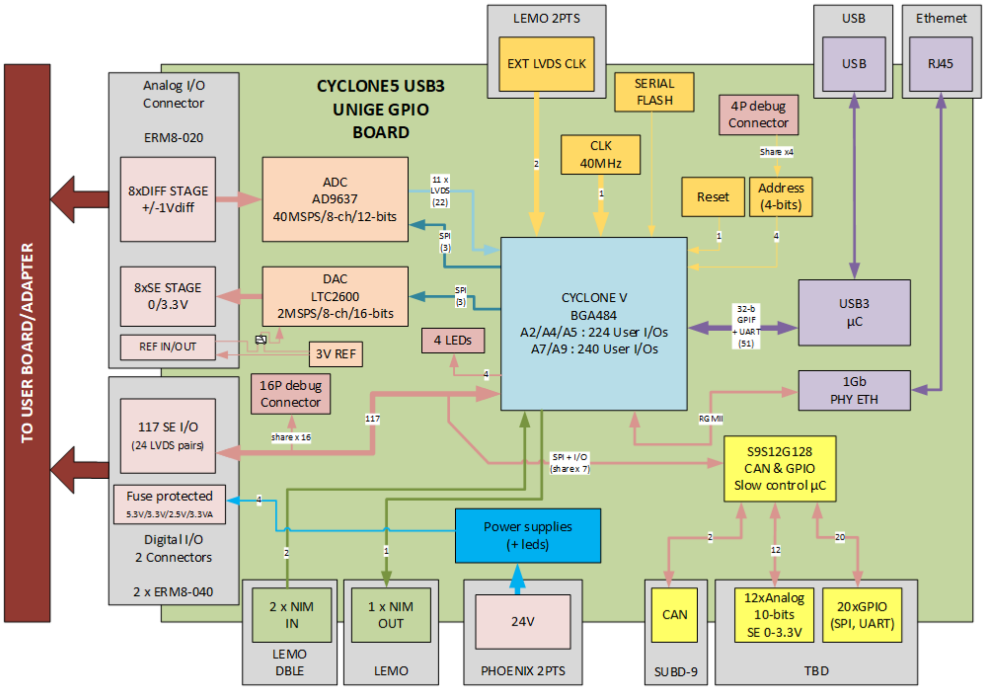

FIG. 36. Architecture of the "UniGe USB3 GPIO" board.

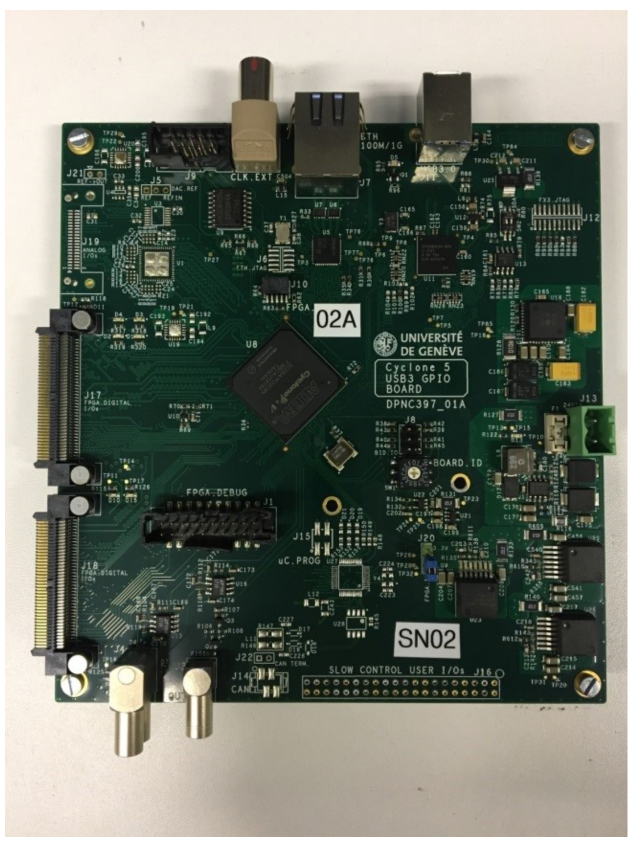

FIG. 37. The "UniGe USB3 GPIO" board.

One single 24 V power supply input is required to supply the board and 5.3 V, 3.3 V and 2.5 V is accessible through the user connectors for the interface board.

1. DAQ framework

Figure 38 describes the architecture of an experiment DAQ using the "Unige USB3 GPIO" board, which should be based on the existing generic UniGe Library and have two specific application parts: a VHDL application firmware code in FPGA using the UniGe VHDL library, and a Host application code using the C# monodev library for Windows/Linux.

The UniGe Library is composed of a VDHL firmware library in FPGA, the generic UniGe FX3 Firmware for the USB3 Cypress FX3 C, the Windows/Linux FX3 cypress USB driver and the UniGe Windows/Linux C# monodev library in host PC.

As an example, the Baby-Mind experiment (neutrino experiment installed in January 2018 at J-PARC, Tokai Japan) uses the UniGe Library and works with 4000 channels connected on 48 boards, 6 boards chained to 1 USB link, having 1 DAQ PC and 8 USB links for the overall experiment.

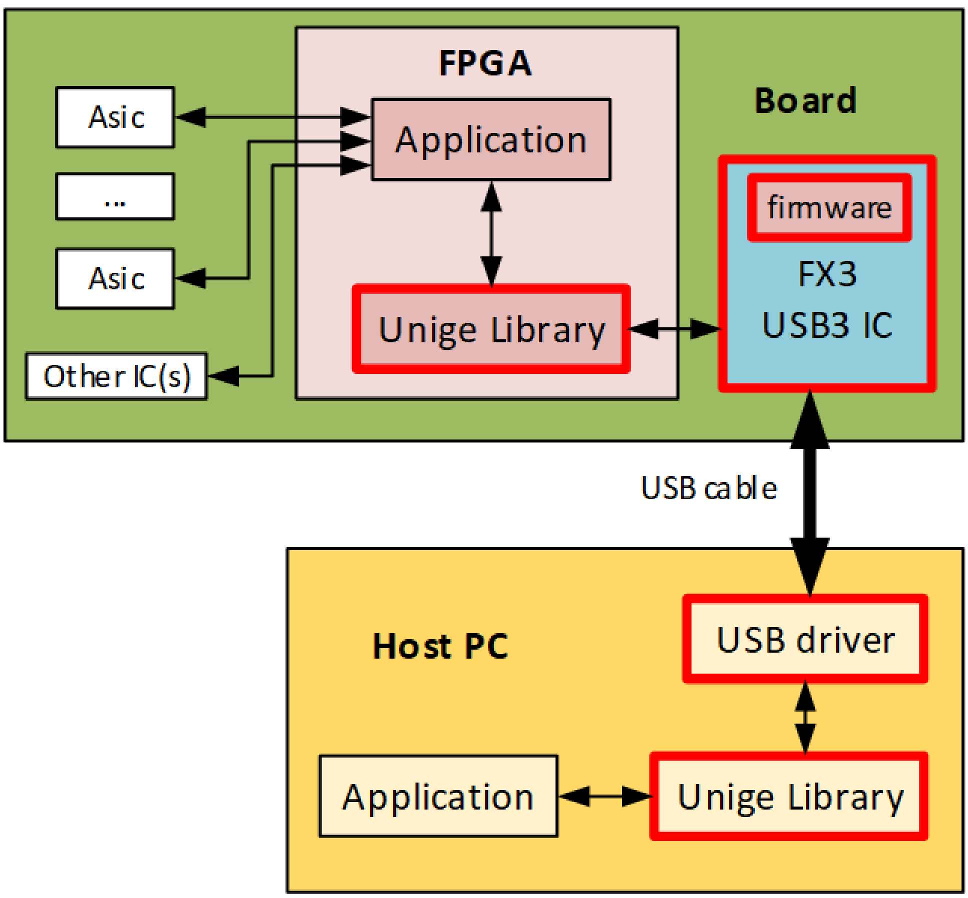

FIG. 38. DAQ framework for USB3 GPIO board.

2. *VHDL library*

The Unige VHDL library (Fig. 39) is linked with the C# Host Library for the slow control protocol handling, for status, read and write parameters, for the usage of Structured configuration objects defined by a JSON file (hardware descriptor, no C# code to be written), for the USB3 Readout and is compatible with single or multiple boards in a chain.

This means that USB3 is transparent from both sides, providing 1 bi-directional slow control channel and 1 fast readout FPGA to PC channel. A similar interface is planned for the Ethernet-based interface.

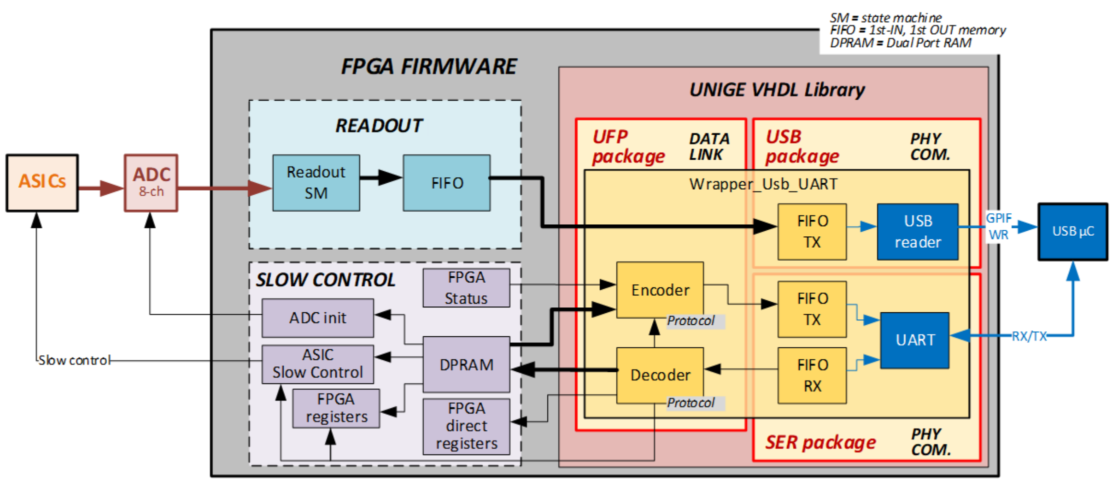

FIG. 39. UniGe VHDL library and application firmware overview.

D. Trigger Logic Board (TLB)

The TLB is the central triggering module of the FASER TDAQ system. It receives trigger signals from the scintillators and the calorimeter, digitized in the CAEN board. The TLB processes these signals to create the trigger decision (L1A) which is transfered to the detector readout boards. The TLB provides the FASER detector readout with a common clock (CLK), envisaged to be synchronous with the LHC clock, as well as bunch counter reset (BCR) signals on every LHC orbit signal. Finally, the TLB receives slow control signals (such as system reset) and sends out trigger data and monitoring information over Ethernet.

The main functionality of the TLB blocks, shown in Fig. 40, is described in what follows. The input trigger lines coming from the CAEN digitizer board are aligned to the 40 MHz clock, which is either coming from the LHC clock cleaner or generated internally. The eight synchronized trigger lines are merged to four trigger items[3] based on user-defined coincidence logic. User-settable prescales can be applied to each of these trigger items, as well as two additional trigger items resulting from a pseudo-random generator (providing triggers up to 1kHz) and a software trigger accept send by slow control. The logical OR of the six

[3] There could be more than four items, but from initial evaluations, four appear to be sufficient.

trigger items is created and unless vetoed, the signal becomes the L1A. A simple dead time requirement ensures that detector read-out is vetoed if it comes within 5-10 μs from the previous one; the dead time requirements are driven by the tracker readout transmission time (more detail is given in Sec. IX E) and are user-settable. An event counter will provide a count of the total number of accepted events and a bunch counter will calculate the number of bunch crossings since the last LHC orbit signal. Additional counters will provide monitoring for the board: Monitoring will consist of counters for each trigger item before prescale, after prescale and after veto. The counters will be pushed out over Ethernet and the counters reset at a fixed rate of the order of 1 Hz.

The output data from the TLB in normal running will be very small. Besides the event and bunch counters, the TLB data fragment to be sent to the Event Builder will contain just a copy of the trigger lines at the different stages of the trigger logic.

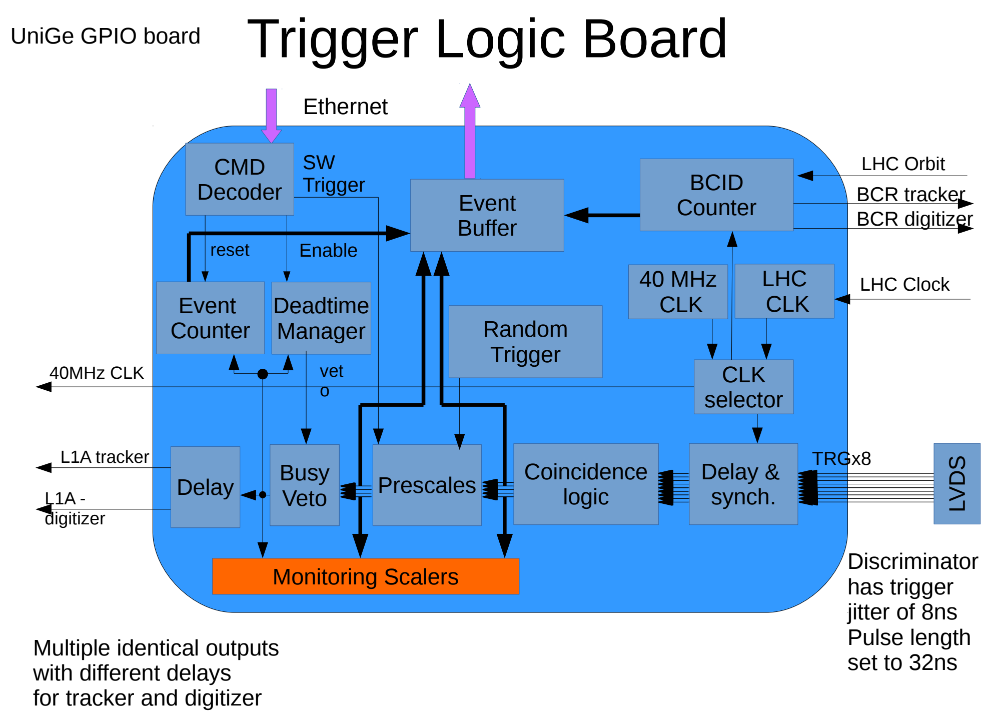

FIG. 40. Schematic diagram of the functionality implemented in the Trigger Logic Board

The TLB will be implemented using one “UniGe USB3 GPIO board” described in Sec. IX C. The board has sufficient I/O lines (LVDS and Ethernet) and the functionality will be implemented by firmware in the Cyclone V FPGA. A detailed evaluation of the required FPGA resources to implement the above described trigger system still has to be done. However, the FPGA contains more logic cells than the combined sum of the FPGAs used to implement the original ATLAS CTP [34] which was significantly more complex (160 incoming trigger lines and 256 trigger items). An adapter board will be built to provide the needed input and output connectors as well as holding the board in place in the VME crate next to the scintillator digitizer.

E. Tracker Readout

The tracker readout is largely defined by the pre-existing ABCD chips mounted on the tracking modules. Each module requires a minimum of two input lines, a clock and a command line, and two output data lines (one per side of the module), all running synchronized to the 40.08 MHz LHC clock. On receiving a L1A over the command line, the location of all strips with a hit are read out to a buffer and transmitted over the data lines. For FASER, a strip is considered hit if the strip fired in the previous, current or next bunch crossing with respect to the L1A, though the trigger timing will be adjusted so that for collision signals the L1A should be aligned to the current bunch crossing. Besides the hit data, the output data contains the lowest 4 bits of the event counter and 8 bits of the bunch crossing counter as seen by the module which allows for synchronization checks at the module level. The minimum data size is 228 bits if no hit is present while a typical size of around 260 bits is expected for most events in FASER. The typical readout transmission time is therefore about 6.5 μs. Up to eight events can be buffered in a module while it is read out, allowing for close-by triggers.

One "UniGe USB3 GPIO" board has sufficient logic cells and I/O connections to handle the readout of at least one out of the three planes located in the tracker station and possible all three planes. Each tracker plane has eight modules. Nine (or three) such boards will therefore be used for the three tracker stations. A special interface card will be constructed to provide the connections from the "UniGe USB3 GPIO" board to the tracker stations and the TLB. To keep the cables short and maintain a good signal integrity, the boards will be located in a small mini-crate about one meter away from the center of the middle tracker station, just outside the fringe field of the dipoles. The clock, L1A and BCR from the TLB will be received over a Cat5e cable and distributed from a fan-out to the nine read out boards. These will convert the L1A and BCR signals to the appropriate commands for the tracker modules. Additional commands for configuring and calibrating tracker modules can be received over the Ethernet connection and transmitted to selected modules. The data received by the board from the different modules will be merged and appropriate identifying headers, event and bunch counters added, before being transmitted to the Event Builder over Ethernet. No detailed processing of the data in the readout board is foreseen.

The resource requirements for implementing the tracker readout can be estimated by comparing the resources needed in the Chimaera board used in the module QA in Sec. V B with those of the "UniGe USB3 GPIO" board. The Chimaera board has a different, smaller FPGA (Spartan-6 XC6SLX25 FPGA vs Cyclone V A7), but only reads out four modules and has somewhat less functionality than planned for the FASER readout. A schematic of the circuit blocks is given in Fig. 41 and the comparison of resources is shown in Table XIII.

Table XIII shows that the readout of 24 modules would consume less than 50% of the Cyclone V resources. This leaves 50% of the resources for the Ethernet PHY + UDP implementation. While this indicates that a single "Unige USB3 GPIO" can control the three planes of a tracker station, the additional requirement for the FASER readout could put a strain on the FPGA resources. Until the preliminary design phase has been completed successfully, it will be assumed that one board will be needed per plane.

Resource	Chimaera design (4 SCT modules) Spartan-6 XC6SLX25	UniGe USB3 GPIO board Cyclone V A7	Ratio (UniGe/Cambridge)
Logic cells/ Logic elements	24k	150k	
Equivalent registers	30k	226k	7.5
RAM	936 kB	6860 kB + 836 kB	8.2
18×18 DSP blocks	38	156×2	8.2
PLLs	2	7	3.5
Chimaera design usage (4 modules)	45% Logic 50% RAM	6.0% Logic 6.1% RAM	
UniGe USB3 library		1% Logic 0.5% RAM	
Total for 24 modules + USB3 library		36.0% Logic 37.1% RAM	

TABLE XIII. Comparison of Xilinx XC6 and Intel Cyclone V. The usage of Logic and RAM was estimated from ratio between XC6 and Intel Cyclone V.

F. Event Building, Data-Quality and Storage

The Event Builder collects the data fragments from the TLB, the tracker readout boards and the PMT digitizer. The first two will push their fragments to the event builder when receiving the L1A, while for the digitizer, the event builder will have to pull the data whenever it receives a fragment from the TLB. The fragments are merged into a single event based on their event counter number. A time-stamp and a run-number will be added and the full event sent to the Event Buffer. The Event Builder will be implemented as a process on a commercial Linux PC located on the surface.

The same process also controls the overall data acquisition system and will send the configuration to all of the readout boards, enable the trigger and also reconfigure the system in case data corruption is detected. Data corruption is checked by comparing the bunch counter values in the different fragments and inside each tracker module sub-fragment. In case of continuous mismatch or missing data fragments, the system will reconfigure itself and in case this does not help, it will contact the expert on-call. The continuous rate of random triggers ensures that the system is maintained in good running condition even when there are no collisions and therefore no physics triggers.

The Event Buffer is a separate process, normally on a second PC, which collects events with the same run number into files on disk with a convenient size for archiving and analysis (about 1 GByte per file). As the event size is dominated by the PMT readout and most of that data should just be pedestal values given the large readout window, each event will be compressed using for example the Lempel-Ziv algorithm [35]. This should reduce the event size by at least a factor 2, allowing one file to hold about 5 minutes of data. The files will be copied asynchronously to the CERN CASTOR storage area and a backup at another FASER institute.

The latest event received by the Event Buffer will also be available for sampling by debugging and data quality tools running on the same PC as the Event Buffer. The data-quality will be monitored by a fast offline-based event reconstruction which produces a small

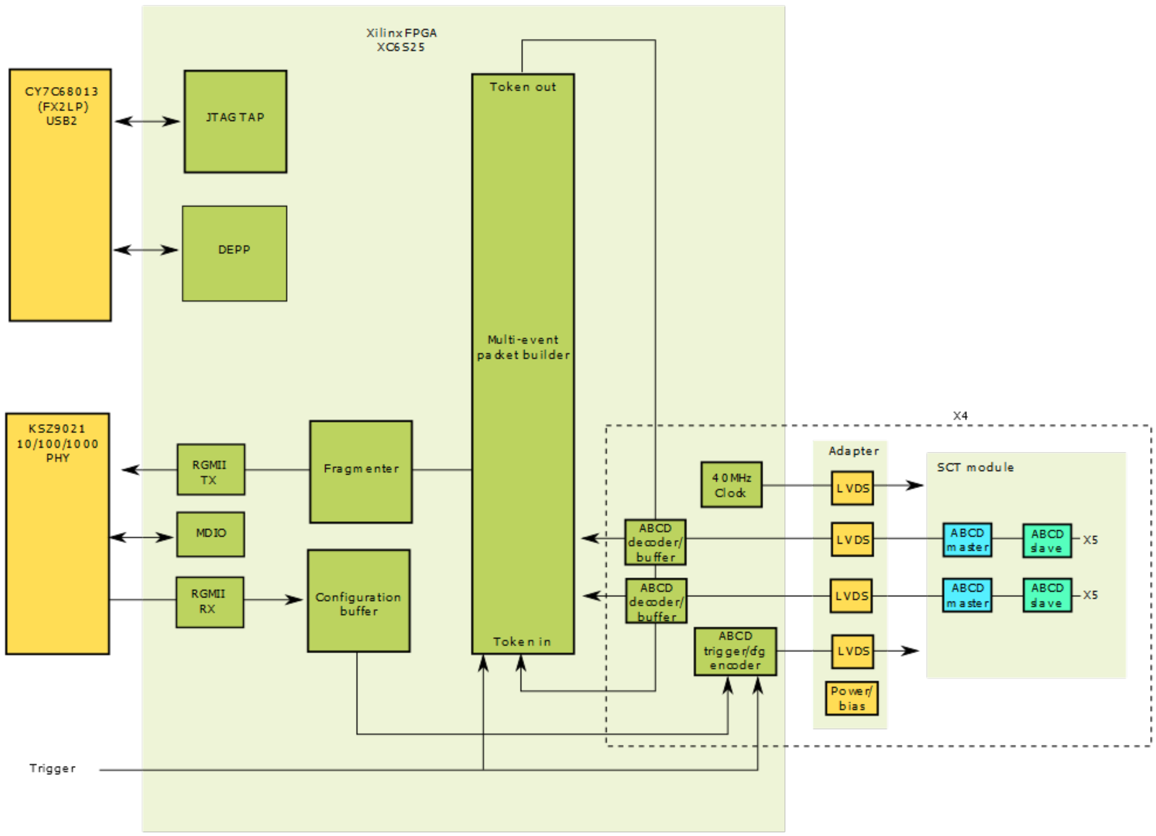

FIG. 41. Schematic diagram of the functionality implemented in the Chimaera board for reading out four SCT modules.

set of monitoring histograms for expert monitoring and automatic checks against reference histograms.

For redundancy and to minimize the risk of interference, the Event Builder and the Event Buffer runs on two separate Linux PCs located on the surface. In case of a hardware failure it will be made possible to run both processes on a single PC. Each PC will be server-grade rack-mounted PCs with redundant power supplies, at least 16 GB of memory, at least eight cores and 2 TBytes of raid-storage. The latter provides enough storage for one week of data taking at the nominal trigger rate.

Design of the data-acquisition and detector control software has not yet begun, but it is expected a simple web-based interface will be used for day-to-day monitoring and control.

For security reasons, the readout and control system will be on a private network, only accessible from the CERN GPN through a dedicated, locked-down PC.

G. Detector Control System (DCS)

The DCS is responsible for controlling and monitoring all of the different power supplies and electronics crates, as well as monitoring temperatures and humidity in the different detector parts. All devices will be connected through Ethernet to a dedicated DCS Linux

Trigger system	**Cost [kCHF]**
VME Crate	5
Digitizer	9
BOBR and fiber	0.5
Clock Cleaner	0.5
Trigger Logic Board	1.5
Trigger fan out	1
Trigger Cables	1
Trigger system Total	**18.5**
Data acquisition system	**Cost [kCHF]**
Mini Crate	2
Tracker readout boards	11
Data cables to tracker	1
VME to Ethernet converter	3
LV power supply	2
Four readout and control PCs	12
DAQ Total	**31**
Networking	**Cost [kCHF]**
2 Ethernet switches	3
Networking Total	**3**
Total	**52.5**

TABLE XIV. Budget for trigger/DAQ system for the FASER experiment.

PC running the standard LHC SCADA system WinCC OA (formerly known as PVSS).

H. Cost and Schedule

The cost of the different components is estimated in Table XIV. At least one spare card of each module of the system, one spare VME crate, a spare readout PC, and network switch will be needed in case of a failure. These spares will also be used to form a small test-stand for testing new firmware and software developments. The cost of the spares is estimated to be about 27 kCHF.

The expected schedule for the construction and commissioning of the FASER trigger and DAQ system is shown in Fig. 42.

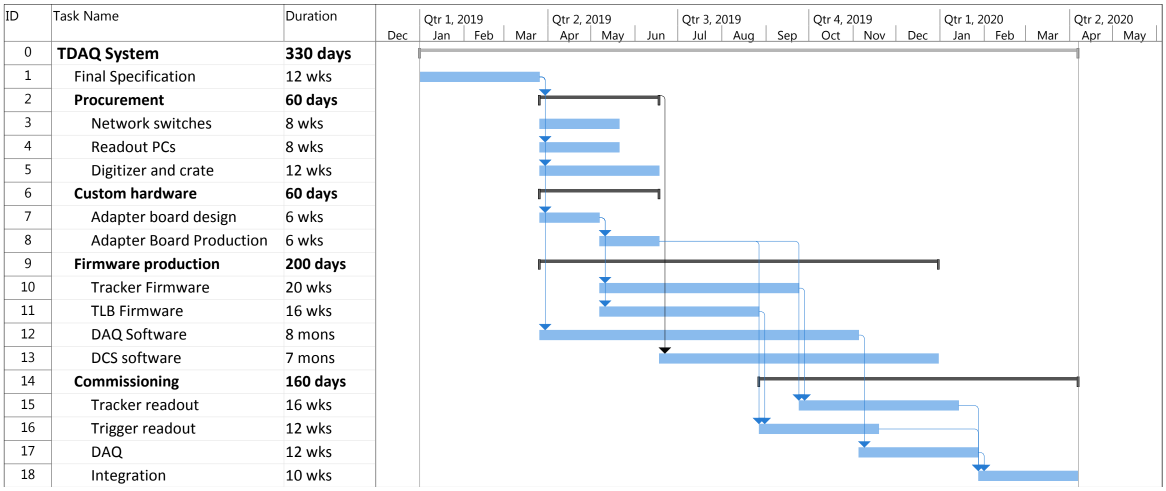

FIG. 42. Timeline for TDAQ construction and commissioning.

X. CIVIL ENGINEERING

The experiment is to be located within TI12, a former injection tunnel that connected the Super Proton Synchrotron (SPS) to the Large Electron–Positron (LEP) Collider at CERN's Meyrin site just outside Geneva. TI12 is located just off the LHC, entirely within Switzerland as shown in Fig. 43. The tunnel is currently disused due to the greater radii required for beam injection into the LHC. To accommodate the experiment within the chosen location in TI12, a certain amount of civil engineering (CE) work will be required. This section details the existing structure, the civil engineering requirements, the design process, the planned works methodology, and further work recommended to implement FASER.

A. Existing Structure

TI12 is a mined tunnel with a horseshoe shaped tunnel cross-section. The tunnel is 2.94 m in width at floor level and is 2.9 m tall from floor to crown. The tunnel structure consists of a 100 mm-thick spray applied shotcrete lining, rock bolts projecting into the surrounding Molasse bedrock, as well as a 250 mm-thick cast *in situ* invert and secondary lining as shown in Fig. 44.

The invert of the tunnel is formed of a shotcrete base 200 mm deep with a 250 mm deep first phase slab cast *in situ* beneath a final cast *in situ* concrete infill slab 500 mm deep. Within the first phase invert slab, there is a 200 mm diameter PVC drain surrounded by gravel infill running longitudinally down the centre of the tunnel. Between the outer layer of shotcrete and the cast *in situ* tunnel structure, there is a waterproofing membrane, allowing any groundwater seepage to drain into the longitudinal drain. Towards the bottom of tunnel TI12, there is a transverse drainage channel or caniveau, as shown in Fig. 45. The caniveau connects into the longitudinal drainage system.

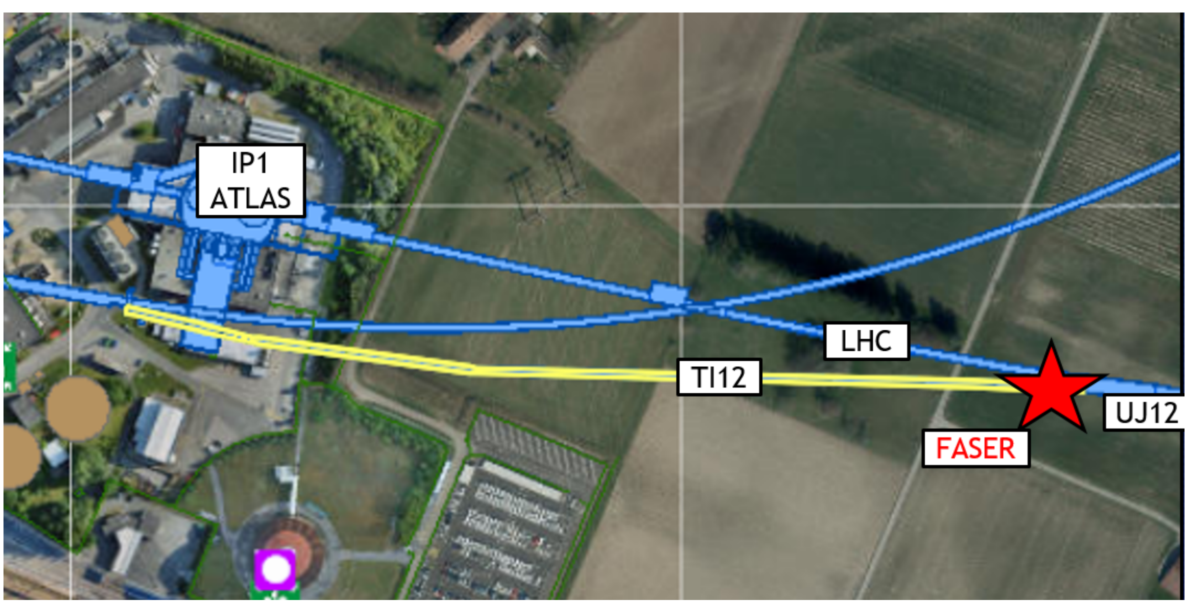

FIG. 43. Location plan showing the proposed situation of FASER in TI12

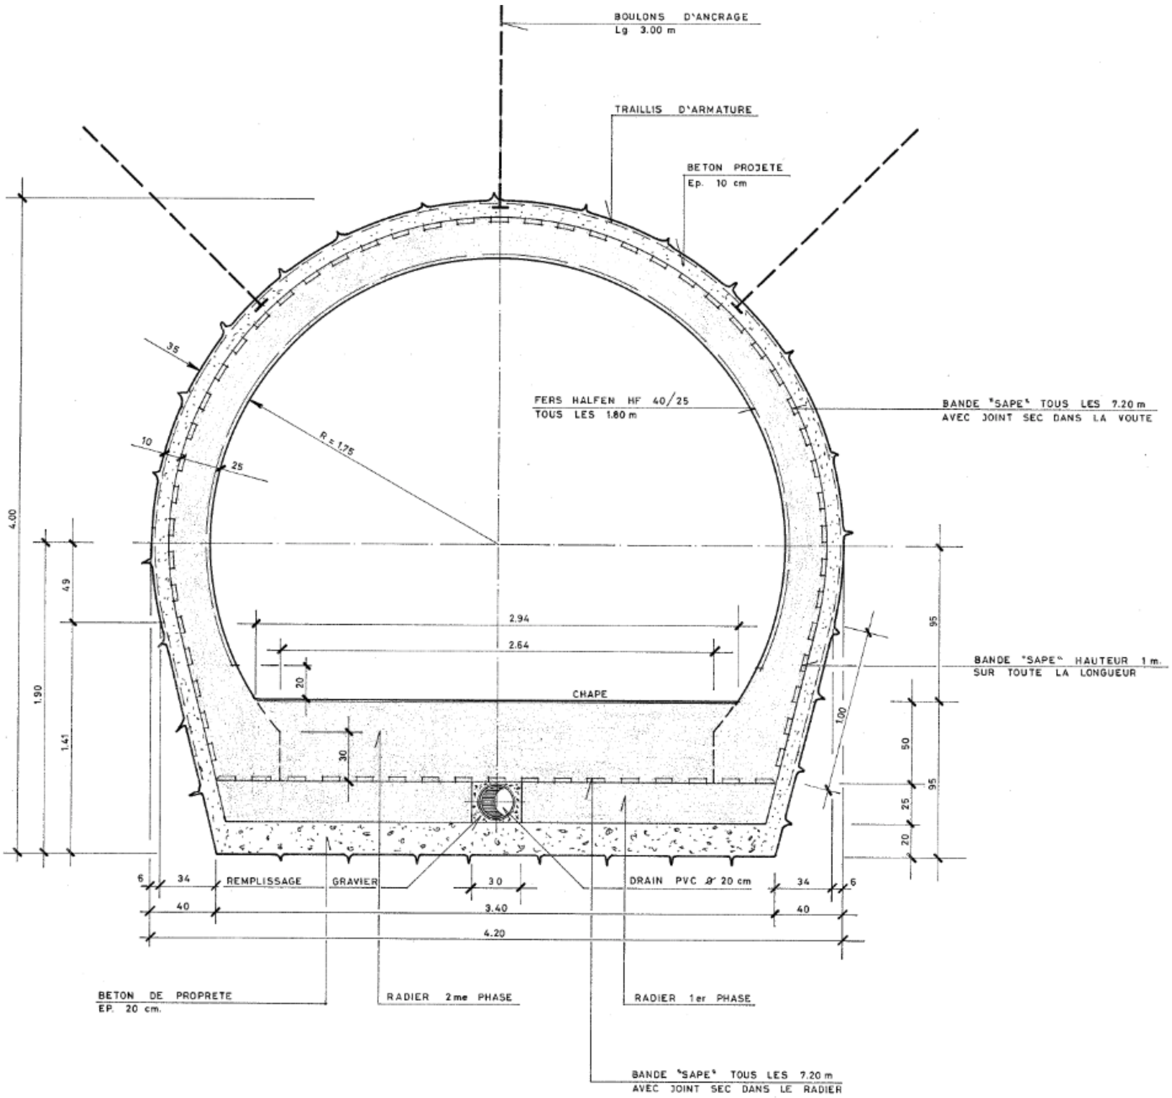

FIG. 44. As-built typical section of tunnel TI12 showing form of construction.

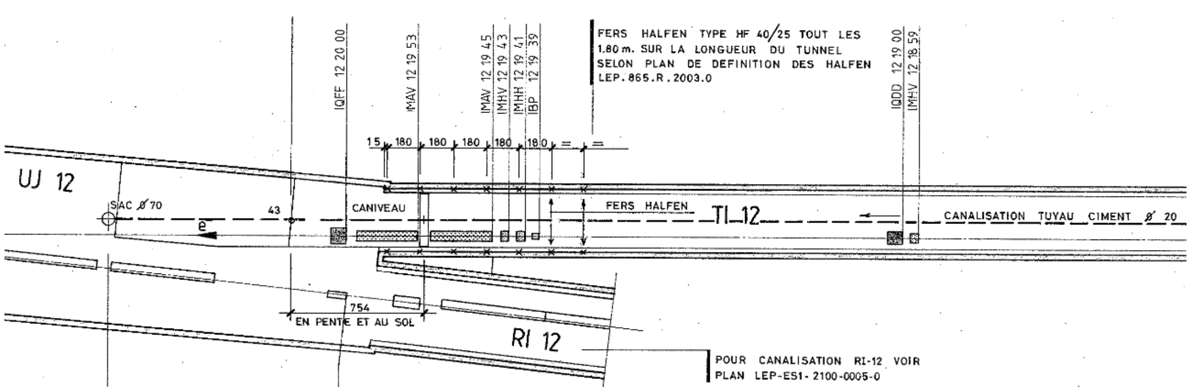

FIG. 45. As-built plan view of tunnel showing arrangement of existing drainage infrastructure.

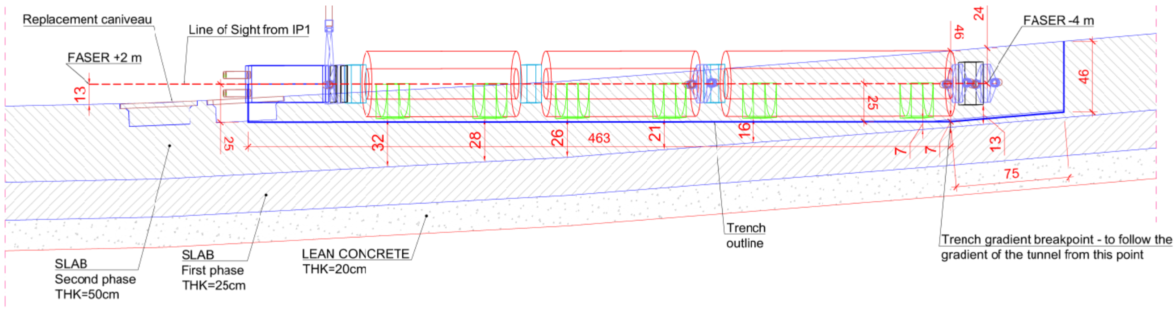

FIG. 46. Section showing trench and experiment in relation to TI12 floor and LOS.

B. Requirements

The axis of the LOS crosses TI12 at an oblique angle to the tunnel in plan and emerges through the tunnel floor close to TI12's junction with cavern UJ12. A trench will be required to create the space required for the experimental arrangement so it can be positioned centrally along the axis of the LOS. The dimensions and position of the trench in relation to the tunnel floor and the LOS are shown in Fig. 46.

The location of the trench and experiment interferes with existing drainage systems and will necessitate the removal of an existing transverse drain. This will need to be replaced with two new transverse drainage channels, above and below the experiment. The new drains will need to be connected into the existing tunnel drainage systems, which in turn outfall into those in UJ12.

In case of any unforeseen seepage through the existing tunnel structure, the trench for the experiment will include an outlet connection to the lower transverse drain to ensure no water can build up within the trench.

As part of the project, CERN's transportation team will also need to be involved, both during the works and for installation and maintenance of the experimental arrangement. To enable their works, lifting points will need to be installed in the tunnel soffit at regular intervals.

C. Design of Modifications

An outline feasibility design has been carried out to consider what CE enabling works are required for FASER. To enable the feasibility study, a 3D scan of the area has been carried out and matched with the existing 3D model of TI12 to corroborate the accuracy and position of the tunnel. A 3D model of the experimental arrangement has been positioned along the axis of the LOS (provided by CERN's Survey team) in combination with the 3D model of the tunnel shown in Fig. 47. This has allowed the interaction between experiment and civil engineering infrastructure to be studied at a preliminary level. Further study will be required: in particular, a detailed structural analysis will be needed to confirm the structural stability of the proposed arrangement. The proposed design will continue to be optimised as all constraints and requirements are clarified. The current preliminary design is shown in Fig. 48.

The proposed trench dimensions have been determined to allow for the installation and maintenance of the experimental arrangement. These dimensions have been reduced as far

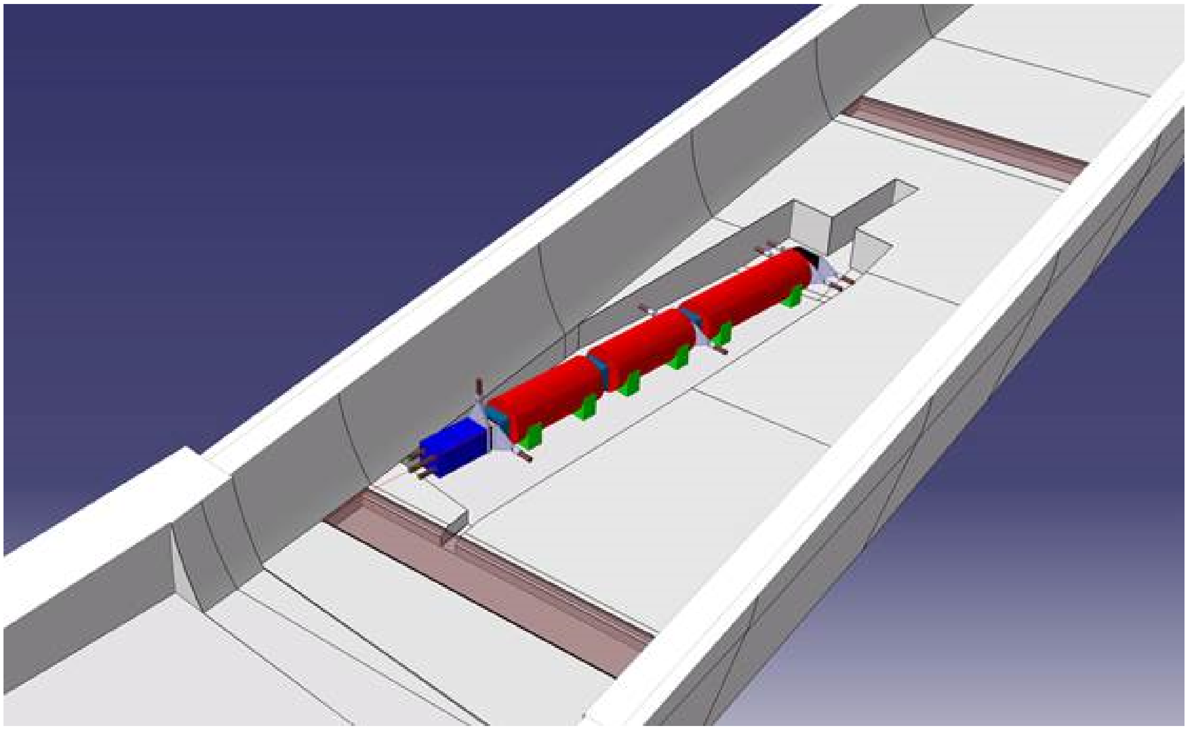

FIG. 47. View showing extents of trench and proposed drainage in relation to tunnel TI12.

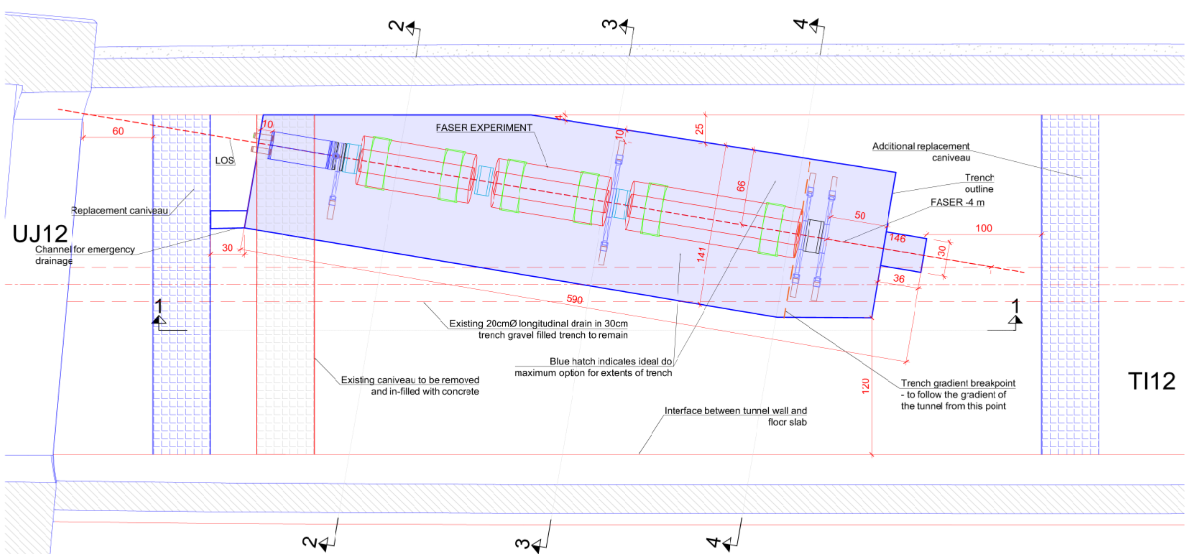

FIG. 48. Plan view showing proposed arrangement.

as possible to minimise the impact on the existing tunnel's structural elements. The trench will still have an effect on the stability of the tunnel, however, which may still require some local strengthening. The location of the experiment in relation to the structure of the tunnel is illustrated in Fig. 49.

The design has been carried out with a view to the construction works being undertaken during Long Shutdown 2 (LS2). This has meant aiming to minimise the complexity of civil engineering work required so it is feasible to plan, design, and implement the works within

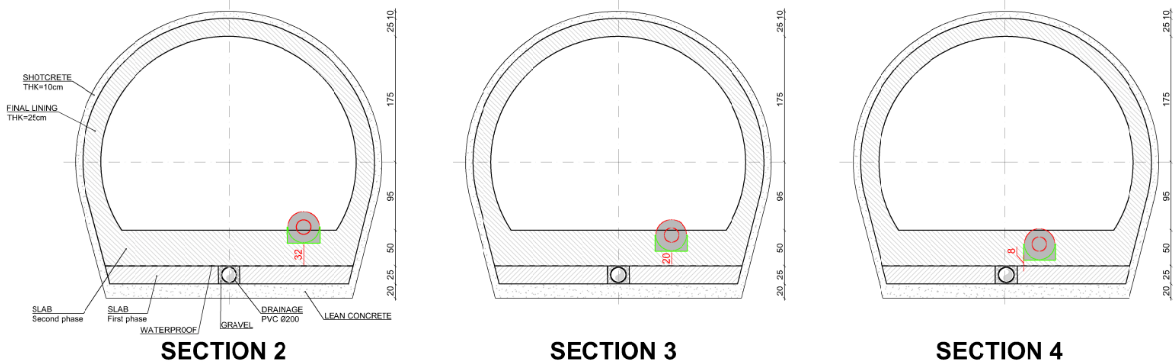

FIG. 49. Proposed sections showing depth of experiment along LOS in relation to the structure and drainage of tunnel TI12 (trench omitted).

a suitable timeframe. To avoid more significant works, for example, the depth of trench has been limited to avoid damage to the longitudinal drain and drainage membrane, 500 mm below the tunnel invert. The position of the experiment along the line of sight has also been optimised with the aim of avoiding the need for openings or cores into the tunnel arch barrel.

Following finalization of the experimental requirements, a full structural stability analysis will be undertaken. The analysis will be carried out in line with Eurocode design standards. The assessment will consider the removal of concrete as part of the trench and drainage along with the lifting points and associated loading required to transport, install, and maintain FASER's equipment. The analysis will determine whether there is any effect on stability and whether any modification is required to the design to enable the project. At this stage, it is considered likely that some local strengthening will be required. This may consist of additional rock bolts or installation of steel reinforcement in the tunnel invert, but analysis and design will be needed to confirm provisions.

The drainage systems will be designed in line with standard CERN specifications and will be of the same or greater standard as the existing system.

The CERN Radio-Protection (RP) group has been consulted and will advise on working methods and waste disposal as the project develops. The concrete in this location should be considered as potentially activated. At the beginning of LS2, representative sample(s) will be taken for analysis to measure the level of activation.

D. Construction Methodology

This type of construction work is not unusual at CERN, meaning that the work can be carried out by existing framework contractors, which will negate the need for lengthy procurement processes. The works will consist of the following main activities:

- Isolate works area (and provide temporary services and ventilation);
- Install transport lifting fixtures in tunnel soffit;
- Saw cut new trenches;
- Remove existing transverse drain and repair tunnel infill slab;

- Install transverse drains and connect to existing longitudinal drains;
- Undertake strengthening works;
- Remove concrete to 40 mm below finished level;
- Cast screed base to trench to level;
- Clear site removing all arisings and isolation.

Concerns have been raised about the works creating dust with the potential to affect existing CERN operations. This will be actively managed throughout the works to avoid any issues. As noted, the first activity will be to isolate the works area. This will be achieved by putting in place a double or triple skin polythene sheet wall system (or a 'SAS' as they are commonly known at CERN) to prevent any dust from escaping the works area. Filtration systems will be put in place to scrub dust from the air. A negative pressure will be created in the works area as necessary to ensure any small quantity of dust created remains within the confined zone. Suitable respiratory equipment will be utilised where required by construction workers. Framework contracts are in place at CERN with companies experienced in delivery of such confinement works. The methods that will be used are industry standard in the field of asbestos removal and treatment.

The dust created will also be kept to an absolute minimum. Dust suppression techniques will be applied to all activities with the potential to create fine airborne material. Water suppression will be used during diamond saw cutting to prevent mobilisation of dust, with arisings and run off contained and collected. Following saw cutting, typical concrete excavation/demolition techniques using a mechanical breaker will be substituted for rig-mounted diamond coring. The trench and drainage channels will be cored out to a suitable depth with drill-powered coring rigs in combination with water suppression, waste water collection and filtration systems, producing a minimum of airborne material. This change of technique will have the dual benefit of producing less dust and allowing more accurate control of the depth of excavation to ensure the drainage membrane remains intact. Transportation fixings will be installed by means of rig mounted drills with dust extraction/collection systems, followed by resin anchoring fixings into the tunnel soffit.

CERN's RP department have advised that the use of a sealed SAS will be sufficient to prevent any spreading of potentially activated dust. Further planning will be carried out in conjunction with the RP team.

Prior to leaving the site, the site will be thoroughly cleaned. Air testing to confirm results can be carried out to ensure no dust remains. All systems and procedures will be agreed with the relevant parties in advance of works.

Further study will be necessary to confirm the exact arrangement, particularly in terms of ventilation, but the techniques detailed are feasible, in widespread use within the construction industry, and have been used numerous times in similar situations at CERN.

E. Required Further Studies

Some additional studies will need to be carried out in advance of construction, including:

- A suitably detailed structural analysis of the existing tunnel should be carried out to confirm that the works will have no impact on tunnel stability or to allow the design of local strengthening measures necessary.

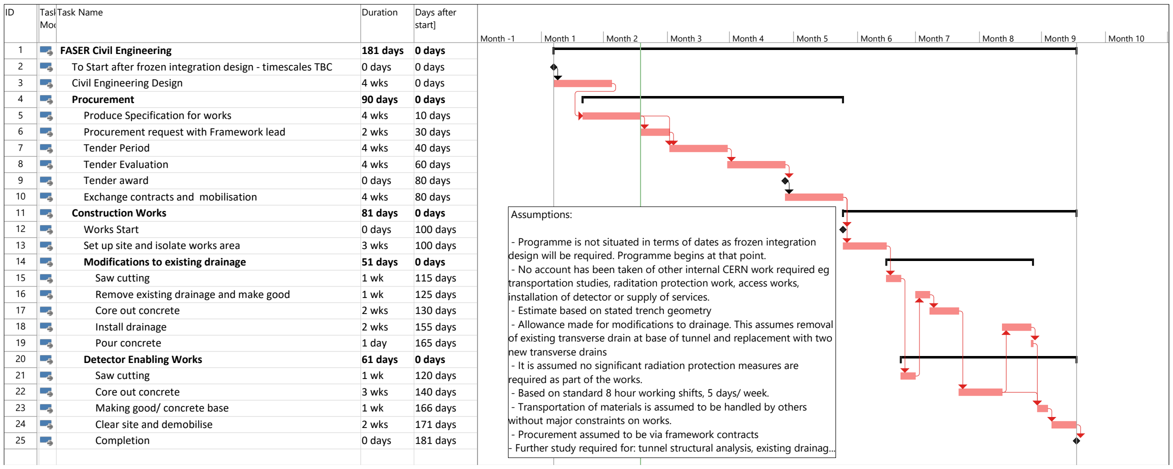

FIG. 50. Draft schedule for civil engineering work to be done in 2019.

- A ground penetrating radar survey of the existing tunnel invert structure should be undertaken to confirm the exact depth of the existing drainage and drainage membrane to facilitate works and confirm the design.

- A CCTV survey of the existing drainage systems should be carried out to confirm their condition and operation prior to connection.

- Representative samples will need to be taken to determine the level of concrete activation.

F. Cost and Schedule

The cost for the civil engineering work is estimated as 83.3 kCHF. This has been based on the layout presented in this chapter and includes costs for detailed design, construction, and construction management, excluding personnel costs for CERN resources. The estimate does not include development costs, materials or personnel costs in advance of detailed design and construction. The estimate and schedule do not include strengthening works, as these cannot be assessed until requirements are confirmed by structural analysis. Rates used have been based on existing framework rates with some tweaking based on comparative analysis of other similar schemes, as well as knowledge of the specific requirements of FASER.

The estimated schedule for the work is shown in Fig. 50. The cost and schedule have been estimated using the following assumptions:

- Costs and schedule are based on the design as detailed without allowance for significant tunnel strengthening works if required.

- The basis of the estimates is that the trench is no deeper than 460mm and does not interfere with drainage.

- The proposed drainage can be connected into existing tunnel drainage without significant repairs or maintenance of the existing.

- No allowance has been made in tunnel geometry for transport, logistics, fire compartmentalization, access staircases, or other alcoves.
- Provision of tunnel services, e.g., ventilation, electricity, etc. are not included.
- No allowance made for network service diversions.
- Costs are based on dust suppression methods as detailed earlier in this section.

The accuracy of the cost estimate is Class 4—Study or Feasibility which could be 15-30% lower or 20-50% higher (in line with Ref. [36], as has been used for previous CERN projects). Until the project requirements are further developed, it is suggested that the maximum band be adopted, i.e., −15% to +50%.

XI. INSTALLATION AND INTEGRATION

A. Transport

The biggest and most complicated detector component to transport is the 1.5 m-long, 50 cm-diameter, 1.5 tonne permanent magnet. The lift at point-1 can carry up to 3 tonnes, and the roof bracket at TI12 can support 8 tonnes. It is therefore not expected to be a problem to transport such an object to TI12. Analysis of the laser scan in the TI12 region indicates that there is sufficient room above the LHC machine to carry the magnet over the machine (and the corresponding QRL cryogenic line) using a simple crane attached to the roof bracket of the main tunnel. The magnet would be carried along the LHC tunnel using an electric tractor. For this, care must be taken as the trolley is made of steel, so the magnetic field must be taken into account in the detailed transport work planning.

The CERN transport group have studied the transport scenario; Fig. 51 shows two pictures from the 3D model of the transport work. This work requires the following tooling:

- A 2-3t hoist for the rail into UJ12 (the original ones have been removed and reinstalled on UJ62),
- Two rails into TI12 (design, production installation), and the transverse bar to make a simplified crane,
- Two 2t hoists to be installed on this rail.

This set of tools will be installed and tested at full load in collaboration with the CERN safety group (HSE). The whole of the process of installing and testing the tooling will take 2 weeks, and is estimated to cost 50 kCHF.

The transport of the FASER detector components is estimated to take 3 days, and cost 5 kCHF for personnel (this is assuming the transport of 10 pieces). This work needs to be integrated into the master LS2 schedule to avoid periods when work is ongoing in the LHC tunnel between IP1 and TI12.

For the transport and installation of the detector, unused ventilation tubes in TI12 need to be removed. This work will be done by the relevant CERN group (EN-CV) and is estimated to take about 1 week, and cost 4 kCHF. Unused cable trays in UJ12 and TI12 will also need to be removed, which can be done by the relevant group (EN-EL).

The access at point-1 is through the PM15 LHC shaft and the bypass tunnel, which joins the LHC beam line just after the triplet magnets, thus bypassing the ATLAS cavern. This means FASER work will not interfere at all with any LS2 work of the ATLAS experiment, which has its own access shaft. Transporting the digging tools for the civil engineering work, and removing the spoil from the excavation (estimated to be $\approx 3\text{m}^3$) is not included in the above cost estimates yet.

B. Services

The following services are needed for FASER:

- Cooling: As discussed in Sec. V D, cooling will be provided by a standalone water chiller situated in the TI12 tunnel close to the detector;

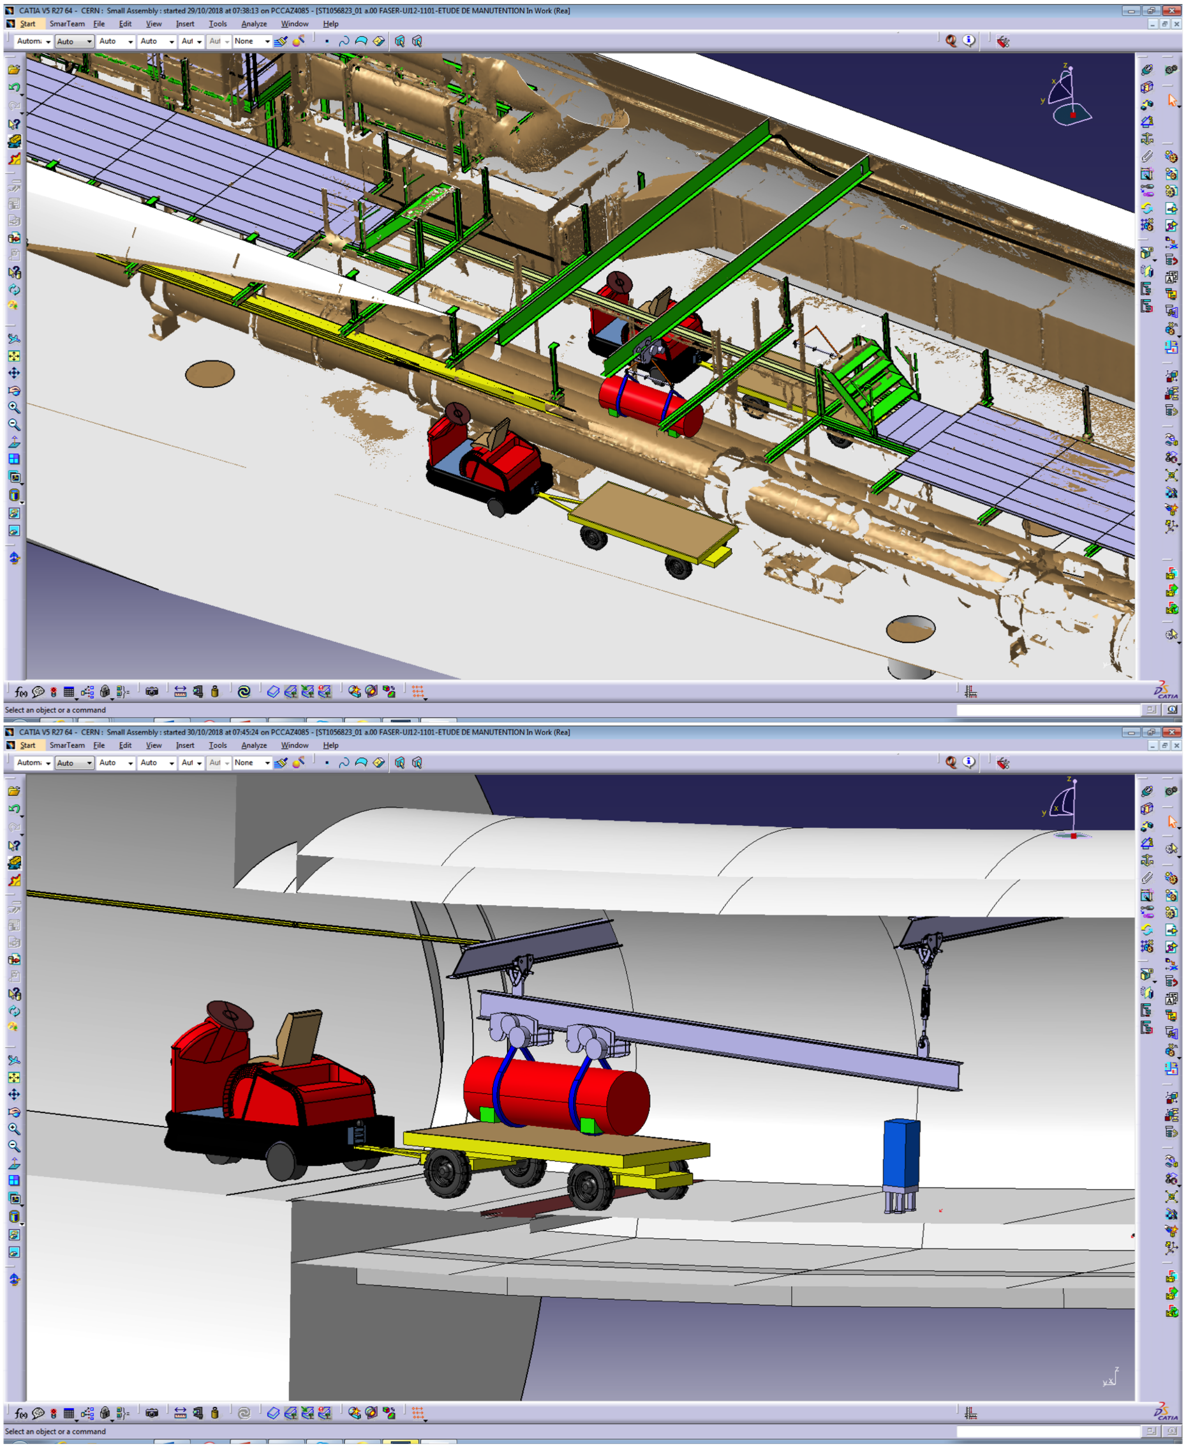

FIG. 51. Pictures from the study of transporting the detector components of FASER to the detector location. Both pictures show the transport of the largest magnet. The top picture shows how this will be lifted over the LHC machine from one tractor to another. The bottom picture shows how this will be transported into position in the TI12 tunnel.

Component	Power [W]
SCT module	540
SCT readout	450
SCT power supply	200
Chiller	1620
PMT HV	180
PMT Readout	60
Trigger logic	50
VME Crate	120
Network switch	180
Total	**3400**

TABLE XV. Expected power consumption of FASER experiment in TI12.

- Dry air: As discussed in Sec. V D, compressed air with a dew point below −40°C will be provided by EN-CV. The estimated cost of connecting the compressed air system into FASER at TI12 is 6 kCHF and it is one day's work, although this needs to be planned in the LS2 schedule;

- Power: The estimated power consumption for FASER is 3.4 kW (a breakdown is shown in Table XV). After discussion with the relevant CERN group (EN-EL), the following scheme for powering the experiment is assumed, although its feasibility remains to be confirmed. The powering would use the 400 V at 16 A outlet available in UJ12 about 50 m from FASER. EN-EL would lay 50 m of cable and install a switchboard with a master breaker and 3 breaker circuits separated for the cooling, readout, and powering. Two emergency cutoff switches (AUG) would also be installed at either end of the detector. This solution is estimated to cost about 10 kCHF, including installing lighting in TI12. The above solution would not provide UPS power for FASER, which would be significantly more expensive. The need for UPS power is still under evaluation.

- Signals and readout: Data will come in/out of the experiment via 3 dedicated optical fibers linking TI12 to the surface building in point-1 (SR1). These fibers will be installed during LS2 by EN-EL (TBC). Two fibers will connect Ethernet switches for data and commands (this gives sufficient bandwidth for the experiment readout as discussed in Sec. IX), while the third fiber provides the BST signal from the BE-BI group to the BOBR module in TI12. On the surface, rack space will be needed for an Ethernet switch and the four readout/controller PC's. These will need to be connected to the CERN-IT network in a safe manner for data-logging and expert access. The exact surface location and connection points are still under discussion.

C. Integration

The CERN integration team have started to develop a model for the TI12 area, including the FASER detector and its associated infrastructure (chiller, electronics, power supplies, etc.). Figure 52 shows a screen-shot from the integration model of the TI12 area.

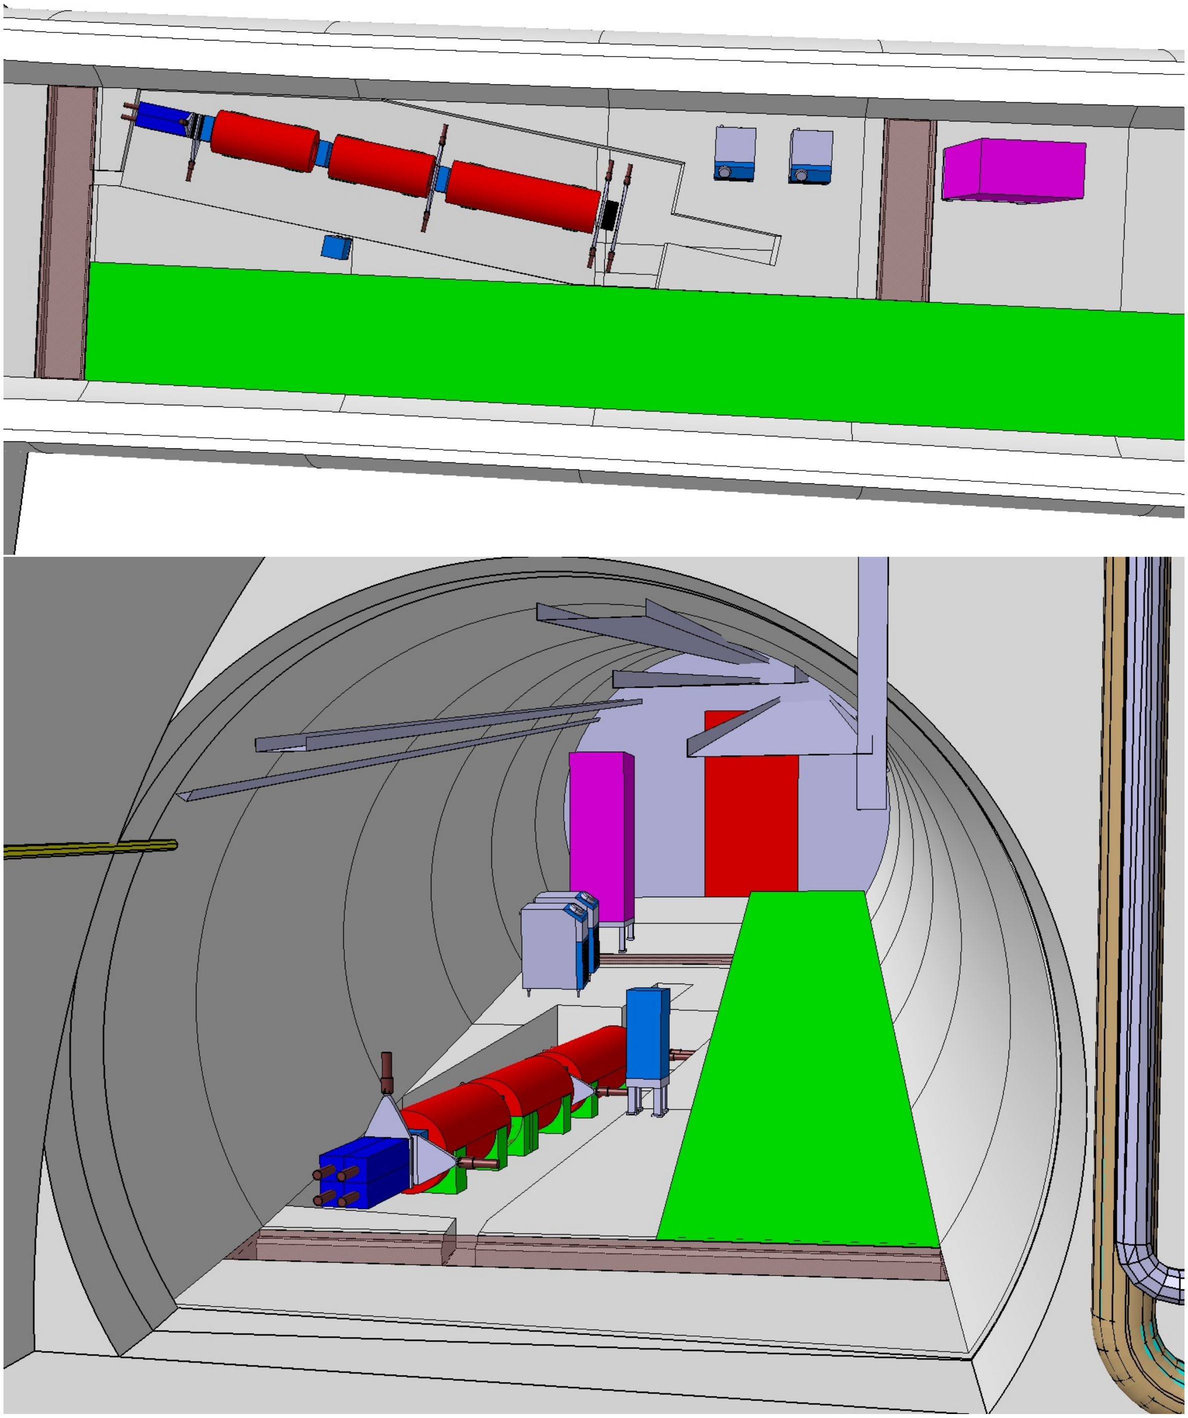

FIG. 52. Pictures from 3D model of integration of the experiment in the TI12 tunnel. The model includes: the detector, the mini-crate for the tracker readout electronics, a VME crate for the TDAQ and power supplies, and the chillers (one is the spare). The green space is 1.2 m reserved for access.

XII. COMMISSIONING

Four commissioning phases are planned for the FASER experiment:

1. **Individual component testing.** Every individual component of the FASER experiment will be separately tested in the place of construction or assembly. This phase includes quality assurance tests for existing components, such as the SCT modules, and testing of cables and boards with pulses. Individual components will need to be tested under magnetic field to ensure proper functioning in realistic conditions.

2. **Integration on surface.** All the components, except for the magnets, will be assembled for integration on the surface. The commissioning will proceed with cosmic runs in a setup where the assembled components will be placed at an angle with respect to the horizontal line. It is expected that the full detector and readout pieces will participate to such cosmic runs, with the exception of the magnets. Stability runs for extended periods of time are expected to take place.

3. **Commissioning in the tunnel, without beam.** Following the installation of individual components in the tunnel, basic testing similar to what was done for 'Individual component testing' will be performed to ensure that there is no damage from manipulation during installation. Once all components (including the magnets) are installed in their final position, combined runs will take place with test pulses. Integration with the external, BST-based, clock will also be done during this period.

4. **Commissioning with beam.** The final commissioning of the FASER experiment will happen with beams. Initial beam commissioning will be attempted with beam-gas interactions, which are expected to be higher after a long shutdown due to worse vacuum conditions, while the final commissioning will be done with the first collision data. Extended calibration runs for the detector are expected to take place.

XIII. SAFETY

The FASER project is followed by the Project and Experiment Safety Support (PESS) of the HSE-OHS group. An HSE correspondent has been named, and the following domains are concerned: civil/structural engineering including worksite aspects, asbestos detection and eventual removal, environmental protection, mechanical safety, HVAC, fire and chemical safety, non-ionizing radiation, and magnetic and electrical safety. In particular, it has been individuated that the civil and structural engineering domain presents major safety implication aspects and, therefore, a safety clearance shall be released by the HSE unit head before starting the civil works (see Ref. [37]). A team of HSE specialists is currently drawing the Launch Safety Agreement, listing the applicable rules and standards in the mentioned safety domains. The HSE correspondent, along with the EP DSO office, will support the project leader in gathering the required documentation to complete the safety file.

XIV. OFFLINE SOFTWARE AND COMPUTING

A. Detector Simulation

A GEANT4-based [38] model of the detector generates simulated data for detector optimization and preliminary performance studies. The model includes the nominal geometry of the ATLAS SCT modules and LHCb ECAL modules, along with a simplified approximation (0.6 T uniform dipole) of the magnetic field. No fringe fields, scintillator planes or mechanical support structures are included at present. For the purpose of simulating multiple scattering, the amount of material along the paths of particles through the tracker is accurate to within $\pm 15\%$.

In addition to the generic "particle gun" functionality provided by the toolkit, an internal generator can simulate decays of dark photons with user-specified mass, momentum, and daughter particles. Simulated dark photons arrive with a uniform solid angle distribution, assuming a source located at the ATLAS interaction point, and decay uniformly along the length of the detector. Figure 53 shows two views of a simulated dark photon decay recorded in FASER.

The tracker simulation produces digitized strip firing data using a simplified electronics model with Gaussian charge smearing and a uniform threshold, along with information to determine which Monte Carlo truth particles contributed energy to them. Each simulated ECAL module produces a single value corresponding to the total energy deposited in its scintillating layers. The time structure of the signal, which will be available for the real ECAL modules, is not yet simulated.

B. Data Reconstruction

1. Strategy

The strategy for reconstruction of tracker data will be similar to that used in ATLAS, although the much lower occupancy in FASER allows considerable simplification. Neighboring hit strips on a single side of a tracker plane are grouped into clusters. Then clusters on opposite sides of a plane are combined to form space-point candidates. Combinatorial pattern recognition will be used to identify track candidates passing through space-point candidates in different tracker planes. Track candidates could pick up additional compatible space-points in a road around their trajectory. The final track fit will use all available information, with extrapolation through a realistic model of the detector material (including energy loss and multiple scattering) and a detailed map of the magnetic field.

2. Performance studies

FASER has developed "proof-of-concept" clustering, space-point and track reconstruction algorithms to optimize and validate the detector design. Track identification and hit association (pattern recognition) is simulated, assuming perfect efficiency, by using Monte Carlo truth information. Misalignment effects are not included.

The tracker must resolve closely-separated, oppositely-charged tracks to identify the dark photon signal. Figure 54 shows the efficiency to reconstruct two isolated space-points in the

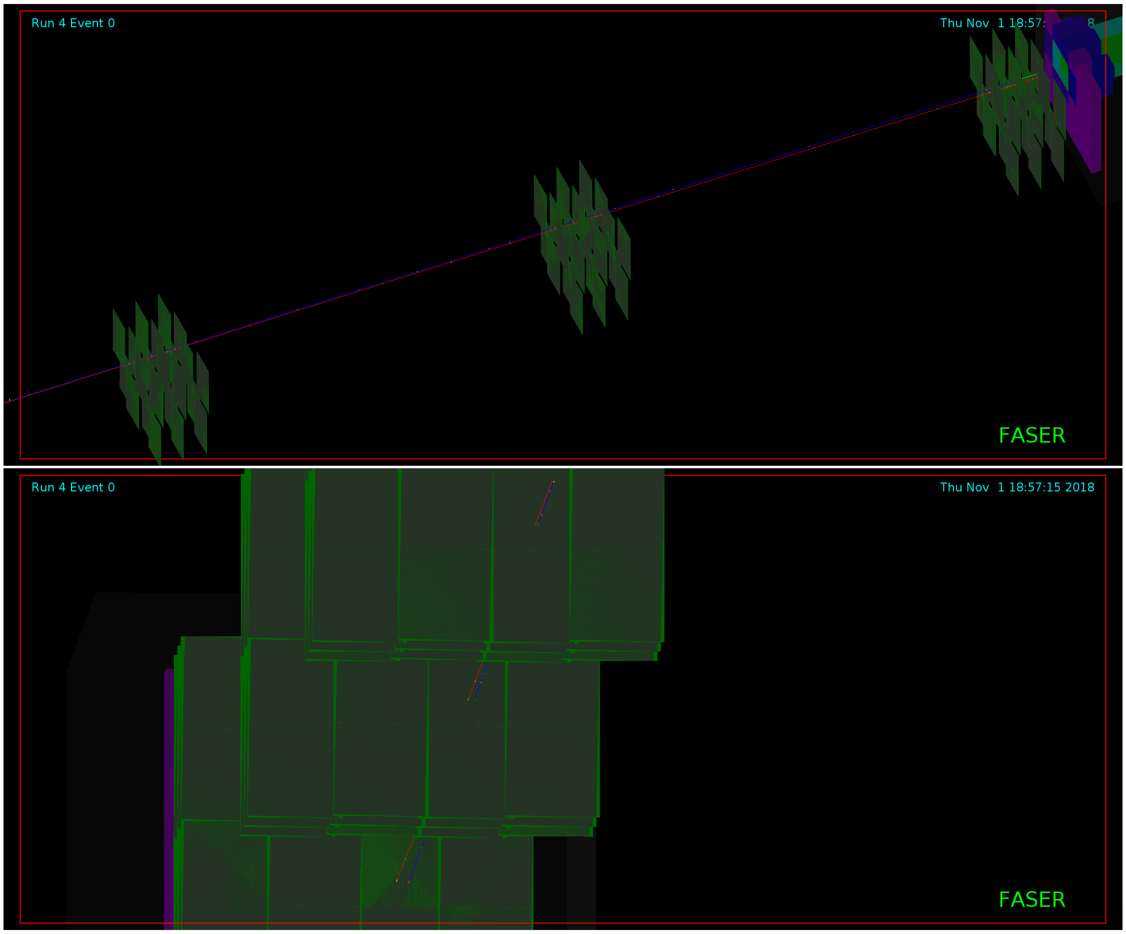

FIG. 53. Two views of a simulated dark photon decay to e^+e^- recorded in FASER. Above: Top view showing the two charged tracks separating in the magnetic field as they pass through the tracker and enter the calorimeter at top right. Below: Close-up showing the silicon strips fired in the first planes of the three tracking stations.

first tracker plane for several dark photon masses and momenta.

Figure 55 shows the reconstructed space-point resolution in the first tracker plane for simulated dark photon decays. The resolution achieved is comparable to that expected from the ATLAS SCT [25].

Karimaki [39] calculated the expected momentum and other resolutions of a magnetic spectrometer from first principles, under assumptions valid for FASER. Figure 56 compares this theoretical performance, including measured space-point resolution and expected multiple scattering, with the results of a global χ^2 fit to reconstructed space-points. The preliminary track fit does somewhat worse than the theoretical prediction at higher momentum, but still gives acceptable resolution above 1 TeV.

To summarize, in these and other studies, using relatively unrefined reconstruction software, the simulated performance of the detector closely tracks our design expectations.

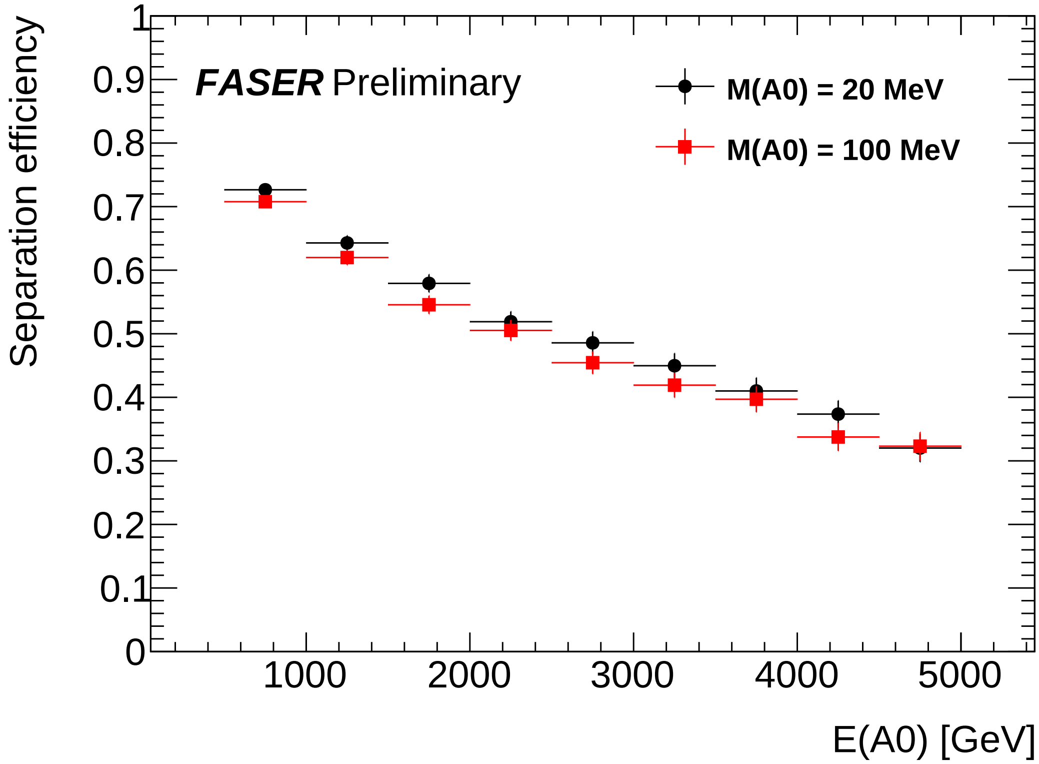

FIG. 54. Two-track separation efficiency based on isolated space-points in the first tracker plane for the indicated dark photon masses and momenta. Dark photon decays are uniformly distributed over the length of the 1.5 m decay volume. Because most of the separation comes from magnetic bending, rather than the transverse momenta of the decay products, the separation efficiency does not depend on the dark photon mass. Higher dark-photon energies, on the other hand, significantly degrade efficiency.

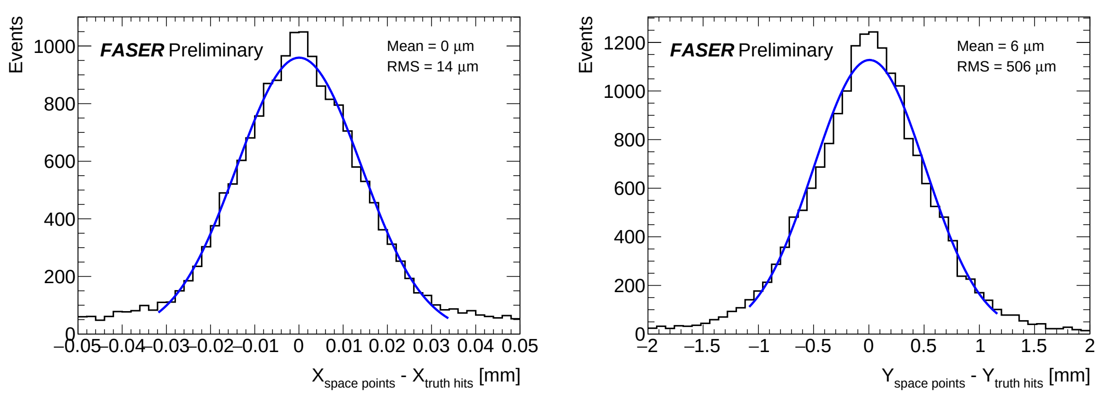

FIG. 55. Spatial resolution of reconstructed space-points in the magnetic bending direction (left) and along the strip direction (right), with respect to Monte Carlo truth, for the first detector plane.

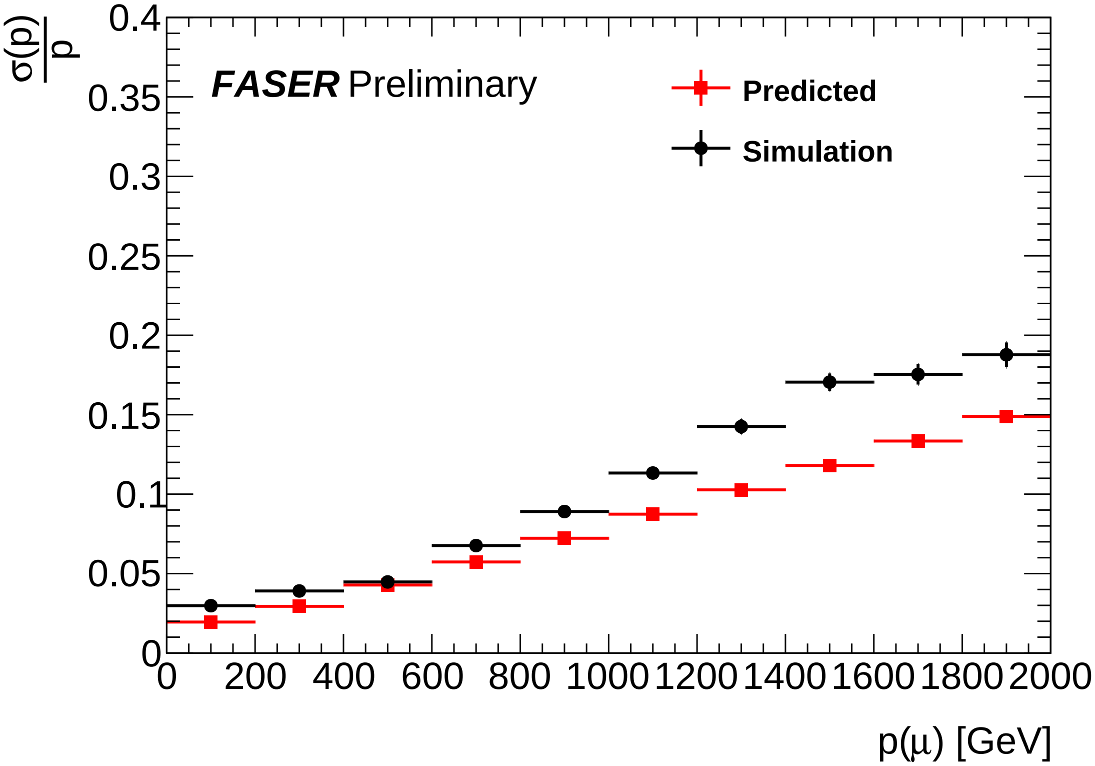

FIG. 56. Fractional momentum resolution (σ_p/p) for reconstructed muon tracks as a function of momentum, compared to the predicted resolution from Karimaki [39].

3. *ACTS for FASER*

The prototype reconstruction software used for validation and optimization of the detector design is inadequate for physics use because it lacks the ability to accommodate a realistic magnetic field or a detailed description of the detector's material. More sophisticated pattern recognition, extrapolation, and track-fitting will be required.

Moving forward, the ACTS ("A Common Tracking System") project [40] offers FASER the opportunity to leverage a fully-featured, open-source reconstruction framework based on the ATLAS tracking software that several FASER members are already familiar with. ATLAS plans to replace its original tracking software with this modernized version after Run 3, and so it will be commissioned and then supported during the lifetime of FASER.

FASER members have worked with the ACTS team over the last six months, and a proof-of-concept implementation of the FASER geometry in the ACTS detector description framework has already been achieved. The ACTS developers have embraced FASER as a useful milestone on their road to the more demanding requirements of ATLAS itself. Thus, we are confident FASER will have access to world-class, collaboratively developed and maintained track reconstruction software, with a far more realistic commitment of effort than would be required to deliver the same functionality ourselves.

C. Offline Software Infrastructure and Schedule

In keeping with its economical design and compact scale, FASER's offline software will be much simpler than the larger and more complex flagship LHC experiments. ROOT [41] and related technology will be sufficient for many of FASER's geometry, conditions, data quality and other database needs. The challenge will not be to find software capable of meeting our needs, but avoiding (as much as possible) the overhead of tools that far exceed them.

Planning for offline software development and commissioning is still in its earliest stages, but one of the first steps will be to implement a prototype geometry database in the first half of 2019, to provide a common reference for existing simulation and evolving ACTS reconstruction software.

As graduate students join the experiment, they will help integrate simulation and reconstruction tools with the geometry database. Demonstration of "full-chain" data simulation and reconstruction with early versions of the production software by the end of 2019 is a realistic milestone.

Conditions data and associated tools for calibration, alignment, data quality, and luminosity, along with physics analysis code, would be built on top of this foundation as it continues to mature in 2020.

D. Computing

FASER's computing and storage requirements will be orders of magnitude smaller than the flagship LHC experiments. The detector's 72 ATLAS SCT modules represent less than 2% of the corresponding ATLAS sub-detector, and the occupancy is lower by a similar factor. As discussed in Sec. IX, the bulk of FASER's raw data will consist of waveform information from the trigger/timing scintillator PMTs and LHCb ECAL modules, with an expected average event size of 25 kB. Pessimistically, at expected Run 3 luminosity, and with no trigger pre-scaling or data reduction, the expected raw data rate from the experiment is $16\,\mathrm{MB/s}$ or roughly $1\,\mathrm{TB/fb}^{-1}$. If raw waveform data from events firing the veto can be discarded or summarized, the event size shrinks to about 5 kB and the data rate falls to $\sim 200\,\mathrm{GB/fb}^{-1}$ even without further data reduction. This rate is unlikely to tax the capacity of the CASTOR data storage system.

Reconstruction of FASER data will consume negligible CPU time. Simulation time (for signal events) is completely dominated by the calorimeter in the present, unoptimized implementation. Because the calorimeter serves primarily as a tag for energetic showers, fast simulation or parameterization of its response can be used to avoid the much more costly full simulation.

TeV muons in FASER can be simulated on a single desktop PC at rates exceeding the realtime trigger rate at peak luminosity. Neutrino background simulations will likely be the most CPU-intensive: for an integrated luminosity of $150\,\mathrm{fb}^{-1}$, roughly 100,000 charged current neutrino interactions above 100 GeV are expected in FASER's magnets; scaling the single-muon simulation time by the expected multiplicity (30), this corresponds to approximately 250 CPU hours on a single-core PC to simulate twenty times the expected background sample. FASER should require only a token allocation of shared capacity on the CERN batch system to meet its computing needs.

XV. OVERALL COST AND SCHEDULE

A. Cost

The overall cost of the experiment hardware, to be borne by the FASER collaboration, is summarized in Table XVI. The costs for each major subsystem are itemized in the corresponding sections above. Spares are shown separately in the table, as in some cases these can be shared between different systems. To ensure adequate funding a contingency of 20% will be included on top of the costs shown. The cost estimates for the main infrastructure work that is assumed to be borne by CERN is detailed in Table XVII.

B. Schedule

Discussions on accommodating the various elements of FASER's project schedule within the master schedule for LS2 are currently ongoing and should be completed before the LHCC meeting in late November 2018.

Preparation of the TI12 tunnel (removing ventilation pipes, installing lights) must happen early in 2019 to allow the civil engineering work to proceed. Magnet construction (the

Detector component	**Cost [kCHF]**	**Detailed Table**
Magnet	420	Table V
Tracker Mechanics	66	Table VI
Tracker Services	105	Table VII
Scintillator Trigger & Veto	52	Table VIII
Calorimeter	13	Table IX
Support structure	60	Table X
Trigger & Data Acquisition	52	Table XIV
Total	**768**	-
Spares	55	-

TABLE XVI. Overall budget for FASER experiment hardware. The TDAQ system includes the readout for all detectors (including the Tracker).

Work	**Cost [kCHF]**
Civil Engineering	83
Transport	55
Optical Fiber & Network Connection	TBD
Power Connection	10
Compressed Air Connection	6
Preparation of TI12	TBD
Total	> 160

TABLE XVII. Budget for infrastructure work whose cost is assumed to be borne by CERN. Preparation of TI12 includes removal of ventilation tubes and cable trays, installation of lighting and installation of temporary ventilation and water. The cost for most of these are in the process of being determined.

other main schedule driver) will progress in parallel during 2019, along with many detector construction and commissioning tasks on the surface. The target for installation and testing of the transport tooling, followed by detector installation, is mid-2020, before the LHC is refilled with Helium. After installation, commissioning the detector *in situ* (first with access, and then without) will become the main activity.

ACKNOWLEDGMENTS

We are grateful to the ATLAS SCT project and the LHCb Calorimeter project for letting us use spare modules as part of the FASER experiment. In addition, FASER gratefully acknowledges invaluable assistance from many people, including the CERN Physics Beyond Colliders study group; the LHC Tunnel Region Experiment (TREX) working group; Rhodri Jones, James Storey, Swann Levasseur, Christos Zamantzas, Tom Levens, Enrico Bravin (beam instrumentation); Dominique Missiaen, Pierre Valentin, Tobias Dobers (survey); Caterina Bertone, Serge Pelletier, Frederic Delsaux (transport); Andrea Tsinganis (FLUKA simulation and background characterization); Attilio Milanese, Davide Tommasini, Luca Bottura (magnets); Burkhard Schmitt, Christian Joram, Raphael Dumps, Sune Jacobsen (scintillators); Dave Robinson, Steve McMahon (ATLAS SCT); Yuri Guz (LHCb calorimeters); Stephane Fartoukh, Jorg Wenninger (LHC optics), Michaela Schaumann (LHC vibrations); Marzia Bernardini, Anne-Laure Perrot, Thomas Otto, Markus Brugger (LHC access and schedule); Simon Marsh, Marco Andreini, Olga Beltramello (safety); Stephen Wotton, Floris Keizer (SCT QA system and SCT readout); Yannic Body, Olivier Crespo-Lopez (cooling/ventilation); Yann Maurer (power); Gianluca Canale, Jeremy Blanc (readout signals); and Ludovico Pontecorvo, Christoph Rembser (general support).

---

[1] A. Ariga, T. Ariga, J. T. Boyd, D. W. Casper, J. L. Feng, I. Galon, S. Hsu, F. Kling, H. Otono, B. Petersen, O. Sato, A. M. Soffa, J. R. Swaney, and S. Trojanowski, "Letter of Intent: FASER—Forward Search Experiment at the LHC," Tech. Rep. CERN-LHCC-2018-030. LHCC-I-032, CERN, Geneva, Oct, 2018. `https://cds.cern.ch/record/2642351`.

[2] M. Battaglieri *et al.*, "US Cosmic Visions: New Ideas in Dark Matter 2017: Community Report," `arXiv:1707.04591 [hep-ph]`.

[3] **Muon g-2** Collaboration, G. W. Bennett *et al.*, "Final Report of the Muon E821 Anomalous Magnetic Moment Measurement at BNL," *Phys. Rev.* **D73** (2006) 072003, `arXiv:hep-ex/0602035 [hep-ex]`.

[4] R. Pohl *et al.*, "The size of the proton," *Nature* **466** (2010) 213–216.

[5] A. J. Krasznahorkay *et al.*, "Observation of Anomalous Internal Pair Creation in Be8 : A Possible Indication of a Light, Neutral Boson," *Phys. Rev. Lett.* **116** no. 4, (2016) 042501, `arXiv:1504.01527 [nucl-ex]`.

[6] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, "ForwArd Search ExpeRiment at the LHC," *Phys. Rev.* **D97** no. 3, (2018) 035001, `arXiv:1708.09389 [hep-ph]`.

[7] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, "Dark Higgs bosons at the ForwArd Search ExpeRiment," *Phys. Rev.* **D97** no. 5, (2018) 055034, `arXiv:1710.09387 [hep-ph]`.

[8] B. Batell, A. Freitas, A. Ismail, and D. Mckeen, "Flavor-specific scalar mediators," *Phys. Rev.* **D98** no. 5, (2018) 055026, arXiv:1712.10022 [hep-ph].

[9] F. Kling and S. Trojanowski, "Heavy Neutral Leptons at FASER," *Phys. Rev.* **D97** no. 9, (2018) 095016, arXiv:1801.08947 [hep-ph].

[10] J. C. Helo, M. Hirsch, and Z. S. Wang, "Heavy neutral fermions at the high-luminosity LHC," *JHEP* **07** (2018) 056, arXiv:1803.02212 [hep-ph].

[11] M. Bauer, P. Foldenauer, and J. Jaeckel, "Hunting All the Hidden Photons," *JHEP* **07** (2018) 094, arXiv:1803.05466 [hep-ph].

[12] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, "Axionlike particles at FASER: The LHC as a photon beam dump," *Phys. Rev.* **D98** no. 5, (2018) 055021, arXiv:1806.02348 [hep-ph].

[13] A. Berlin and F. Kling, "Inelastic Dark Matter at the LHC Lifetime Frontier: ATLAS, CMS, LHCb, CODEX-b, FASER, and MATHUSLA," arXiv:1810.01879 [hep-ph].

[14] D. Dercks, J. de Vries, H. K. Dreiner, and Z. S. Wang, "R-parity Violation and Light Neutralinos at CODEX-b, FASER, and MATHUSLA," arXiv:1810.03617 [hep-ph].

[15] A. Ferrari, P. R. Sala, A. Fass, and J. Ranft, *FLUKA: A Multi-particle Transport Code (Program Version 2005).* CERN Yellow Reports: Monographs. CERN, Geneva, 2005. http://cds.cern.ch/record/898301.

[16] T. T. Böhlen, F. Cerutti, M. P. W. Chin, A. Fass, A. Ferrari, P. G. Ortega, A. Mairani, P. R. Sala, G. Smirnov, and V. Vlachoudis, "The FLUKA Code: Developments and Challenges for High Energy and Medical Applications," *Nucl. Data Sheets* **120** (2014) 211–214.

[17] CERN Sources, Targets, and Interactions Group, M. Sabate-Gilarte, F. Cerutti, and A. Tsinganis, "Characterization of the radiation field for the FASER experiment,".

[18] **Medipix3** Collaboration, C. Brezina, Y. Fu, M. De Gaspari, V. Gromov, X. Llopart, T. Poikela, F. Zappon, and A. Kruth, "The Timepix3 chip," tech. rep. https://indico.cern.ch/event/267425/attachments/477859/661149/Timepix3_final.pdf.

[19] M. Cabon, C. Charrondire, K. Develle, P. Fessia, M. Guinchard, and J. Wenninger, "LHC seismic network design, installation and operation," Tech. Rep. EDMS:1549343, 2017. https://edms.cern.ch/ui/file/1549343/1/EDMS1549343_Report_V11_docx_cpdf.pdf.

[20] **ATLAS** Collaboration, *ATLAS Inner Detector: Technical Design Report, Vol. 1.* Technical Design Report ATLAS. CERN, Geneva, 1997. http://cds.cern.ch/record/331063.

[21] **ATLAS** Collaboration, *ATLAS Inner Detector: Technical Design Report, Vol. 2.* Technical Design Report ATLAS. CERN, Geneva, 1997. http://cds.cern.ch/record/331064.

[22] A. Abdesselam *et al.*, "The barrel modules of the ATLAS semiconductor tracker," *Nucl. Instrum. Meth.* **A568** (2006) 642–671.

[23] **ATLAS** Collaboration, A. Abdesselam *et al.*, "The ATLAS semiconductor tracker end-cap module," *Nucl. Instrum. Meth.* **A575** (2007) 353–389.

[24] F. Campabadal *et al.*, "Design and performance of the ABCD3TA ASIC for readout of silicon strip detectors in the ATLAS semiconductor tracker," *Nucl. Instrum. Meth.* **A552** (2005) 292–328.

[25] **ATLAS** Collaboration, G. Aad *et al.*, "Operation and performance of the ATLAS semiconductor tracker," *JINST* **9** (2014) P08009, arXiv:1404.7473 [hep-ex].

[26] F. Keizer, A. Gorbatch, M. A. Parker, C. Steer, and S. A. Wotton, "A compact, high resolution tracker for cosmic ray muon scattering tomography using semiconductor sensors," *JINST* **13** no. 10, (2018) P10028, arXiv:1810.12174 [physics.ins-det].

[27] **ATLAS** Collaboration, "Alignment Performance of the ATLAS Inner Detector Tracking System in 7 TeV proton-proton collisions at the LHC," Tech. Rep. ATLAS-CONF-2010-067, CERN, Geneva, Jul, 2010. http://cds.cern.ch/record/1281342.

[28] **LHCb** Collaboration, *LHCb Calorimeters: Technical Design Report*. Technical Design Report LHCb. CERN, Geneva, 2000. http://cds.cern.ch/record/494264.

[29] "Timing, Trigger and Control (TTC) Systems for the LHC." http://ttc.web.cern.ch/TTC/intro.html.

[30] J.-J. Savioz, "The Beam Synchonous Timing Receiver Interface for the Beam Observation System," Tech. Rep. LHC-BOBR-ES-0001, CERN, Geneva, Sep, 2003. https://edms.cern.ch/file/406137/1.0/LHC-BOBR-ES-0001.pdf.

[31] Silicon Labs, "Si5345/44/42 Rev D Data Sheet." https://www.silabs.com/documents/public/data-sheets/Si5345-44-42-D-DataSheet.pdf.

[32] CAEN Electronic Instrumentation, "VX1730, 16/8 Channel 14-bit 500 MS/s Digitizer." http://www.caen.it/csite/CaenProd.jsp?parent=11&idmod=780.

[33] Struck Innovative Systeme, "SIS3153 USB3.0 and Ethernet To VME Interface." https://struck.de/sis3153.html.

[34] S. Ask, D. Berge, P. Borrego-Amaral, D. Caracinha, N. Ellis, P. Farthouat, P. Glln, S. Haas, J. Haller, P. Klofver, A. Krasznahorkay, A. Messina, C. Ohm, T. Pauly, M. Perantoni, H. P. L. Junior, G. Schuler, D. Sherman, R. Spiwoks, T. Wengler, J. M. de Seixas, and R. T. Teixeira, "The ATLAS central level-1 trigger logic and TTC system," *JINST* **3** (2008) P08002.

[35] J. Ziv and A. Lempel, "A Universal Algorithm for Sequential Data Compression," *IEEE Trans. Inf. Theor.* **23** no. 3, (Sept., 2006) 337–343. http://dx.doi.org/10.1109/TIT.1977.1055714.

[36] P. Christensen, L. R. Dysert, J. Bates, D. Burton, R. Creese, and J. Hollmann, "Cost Estimate Classification system-as applied in engineering, procurement, and construction for the process industries," *AACE, Inc* **2011** (2005) .

[37] HSE, "Safety Regulation SR-SO: Responsibilities and Organisational Structure in Matters of Safety at CERN." https://edms.cern.ch/ui/file/1389540/LAST_RELEASED/SR-SO_E.pdf.

[38] **GEANT4** Collaboration, S. Agostinelli *et al.*, "GEANT4: A Simulation toolkit," *Nucl. Instrum. Meth.* **A506** (2003) 250–303.

[39] V. Karimaki, "Explicit covariance matrix for particle measurement precision," *Nucl. Instrum. Meth.* **A410** (1998) 284–292.

[40] **ATLAS** Collaboration, C. Gumpert, A. Salzburger, M. Kiehn, J. Hrdinka, and N. Calace, "ACTS: from ATLAS software towards a common track reconstruction software," *J. Phys. Conf. Ser.* **898** no. 4, (2017) 042011.

[41] R. Brun and F. Rademakers, "ROOT: An object oriented data analysis framework," *Nucl. Instrum. Meth.* **A389** (1997) 81–86.