


IDRC MEETING

IHEP

14-16 April 2025



Introduction

The IDRC met at IHEP on April 14–16 to review the draft version of the Reference Technical Design Report (Ref-TDR) for the CEPC detector. The committee congratulates the CEPC team for their outstanding work and recognizes the remarkable progress made since the IDRC's first meeting in October 2024.

During the meeting, the CEPC team presented detailed responses to all comments and recommendations submitted by the committee following the initial meeting. The IDRC commends the team for their thorough and thoughtful replies.

The committee engaged with the CEPC team to review progress in mechanical and electronics integration, the design of the machine-detector interface, and the magnet system. In parallel sessions, panel experts evaluated updates on specific detector components, including tracking detectors, calorimetry, the muon system, as well as TDAQ and software and computing.

The committee carefully assessed the current status of the R&D program and discussed what additional work is needed to ensure the reference detector will be a robust tool capable of fully exploiting the CEPC's physics potential if the project is approved. While the presentations were excellent, the Ref-TDR requires further editing before it can be released publicly. The IDRC issues the following general recommendations:

- **Appoint an overall editor** to enhance the document's coherence, eliminate duplications, and ensure consistency across chapters.
- **Restructure chapters** to focus clearly on the baseline design for each detector component, including specific details showing that the required physics performance can be achieved. Clearly outline the R&D and prototyping still needed to demonstrate the feasibility of the selected baseline technologies.
- **Describe backup solutions** in detail, showing that they are available and could meet performance requirements if needed. The proposed use of HTS cables for the superconducting magnet, for instance, will require significant additional R&D.
- **Significantly shorten discussions** of alternative technologies.
- **Eliminate descriptions** of the basic working principles of well-established technologies.

The CEPC team has selected very innovative technologies for the Ref-TDR. Demonstrating sufficient technological maturity will require an intensive program of prototyping and beam testing. The committee strongly encourages the CEPC team to now focus on this critical phase.

Finally, we recommend beginning detailed evaluations of detector performance, including the impact of expected beam backgrounds, at least for key physics channels. The IDRC believes these are essential next steps on the path to realizing this exciting project.

Machine Detector Interface and Luminosity

Measurement

The committee would like to congratulate the group on the significant progress made since the last review. Most of the comments and recommendations raised at that time have been addressed or at least partially resolved. However, some open issues remain, as detailed below.

Findings

A layout of the interaction region—including the anti-solenoid, final focus quadrupoles, IP beam pipe, and Luminosity Monitor—has been presented. Although not yet final, it is sufficiently advanced to proceed with detailed component design.

The Be beam pipe has an inner diameter of 10 mm and a length of 220 mm. It consists of an inner Be layer 0.2 mm thick and an outer Be layer 0.15 mm thick, separated by a 0.2 mm gap, which seems quite aggressive. The default coolant is water, with paraffin as a backup. The cooling technology is well-established, and the cost estimates are considered reliable. The coolant flow remains in the stable laminar regime. Mechanical analyses for cantilever support during installation have been performed, and no issues have been found. With an inlet water temperature of 15°C, the Be pipe temperature remains below 20°C, which is within acceptable limits.

Beam background sources—including synchrotron radiation (SR) photons, pairs, and off-energy particles—have been studied in detail, along with other minor contributions. For SR photons, important interactions such as the photoelectric effect and Rayleigh scattering have been simulated. It was found critically important to employ high-Z SR masks to protect the IP beam pipe region. Pair backgrounds are mitigated by designing the IP beam pipe to stay outside the high-density region. The hit rate on the first layer of the vertex detector (VTX) is estimated to be around 1 hit per cm² per bunch crossing. Backgrounds from off-energy particles are suppressed by placing collimators at strategic points around the ring and heavy metal masks near the IP to absorb secondary showers.

Experience from BESIII was presented, showing that real background levels were about 1/5 of those predicted by simulation.

An updated design of the LumiCal, using a silicon detector and LYSO crystals for pile-up event veto, is included in the Reference TDR. Mechanical design and optimization are complete, and construction of a large prototype is planned for the EDR phase. The selected technology appears

feasible, although achieving the required luminosity precision will demand electron impact position measurements at the level of better than ten microns.

Comments

From the presentations and draft Ref-TDR, it is unclear whether a gold (Au) coating is applied to the inside of the Be beam pipe, and if so, what thickness is used. Further study on this issue may be necessary.

The current SR mask configurations require further study. In particular, changes in beam orbit—such as during commissioning—must be considered to ensure continued protection of the Be beam pipe region.

While heavy metal masks can effectively reduce backgrounds from off-energy particles, they may also generate secondary backgrounds. This study suggests the effect is moderate, but further investigation is needed, especially concerning the tungsten (W) masks near the IP.

The discrepancy between real background measurements and simulations observed at BESIII is a significant concern. Although the energy scales differ between BESIII and CEPC, understanding this discrepancy is critical for evaluating the reliability of CEPC background predictions.

Although background rates (after shielding) are presented, detailed histograms and numerical data characterizing these backgrounds (e.g., energy spectra, multiplicity, polar and azimuthal angle distributions) are often missing. Such information is essential to assess the impact on the entire apparatus, particularly on the first layers of the vertex detector and LumiCal.

Recommendations

- Conduct studies to finalize the decision on the Au coating inside the Be beam pipe, including its thickness and the possibility of omitting it.
- Continue detailed studies of SR mask configurations and materials, considering the effects of beam orbit steering in collaboration with the accelerator group.
- Further investigate the role of heavy metal masks in absorbing particle backgrounds, including the possibility of operating without them.
- Pursue deeper studies of BESIII backgrounds to understand the discrepancy between simulation and data.
- Continue dedicated studies on LumiCal readout electronics at high rates, followed by beam tests with detector prototypes at later stages.
- Study the effects of beam backgrounds on LumiCal reconstruction for realistic configurations.

Vertex Detector

The design and technology choice for the vertex detector is highly ambitious and at the forefront of technological innovation. The baseline option—thin, bendable silicon CMOS sensors that can be stitched together to create a large, lightweight vertex detector—is certainly a very challenging yet promising approach. The backup solution, which also employs thin CMOS sensors mounted on ladders, remains technologically advanced. The project has made significant progress since the last review.

The vertex detector is expected to operate for 10 years during Higgs and low-luminosity Z running, after which it will need to be replaced for high-luminosity Z, WW, and $t\bar{t}$ operations. The ability to run at high-luminosity Z mode imposes additional demands on an already extremely challenging design (e.g., larger data rate capabilities, shorter charge collection times, lower noise, and tighter power consumption constraints).

There are several major challenges for the vertex detector. Achieving the desired position resolution has been demonstrated with the TaichuPix-3 chip in TJ180 nm technology. However, this process has limitations, particularly its high-power consumption (>100 mW/cm²), which is too high for air cooling. Transitioning to the TPSCo65 process is a way forward, reducing power consumption to approximately 40 mW/cm². The required position resolution and uniform hit efficiency have already been demonstrated in this new process.

The real challenge now lies in the development of large-area stitched sensors of ~ 40 μm thickness that can be bent to radii as small as 11 mm, achieving a material budget of just 0.06% X_0 . Successful bending of dummy 40 μm wafers has been demonstrated, but bending fully processed wafers (including metal layers for data and power lines) remains to be achieved. Securing access to TPSCo65 technology, including a modified process optimized for particle detection, is crucial. An alternative approach under early exploration is the HLMC 55 nm technology, although it has yet to be demonstrated.

The mechanical design is quite advanced for this stage, with many demonstrators and mock-up parts already tested, including a full ladder demonstrator. The simulations shown for mechanical and thermal performance, efficiency coverage, and background hit rates are detailed and of high quality. The impact of background hit rates on physics performance has been found to be negligible. The planned laser alignment system is a very positive development, but its effective operation must still be demonstrated.

The fifth double layer of the vertex detector employs a more conventional ladder design, achieving approximately 0.25% X_0 . At the larger radius of the fifth layer, this becomes necessary but also requires parallel development. Importantly, this ladder design offers a reliable fallback option in case insurmountable problems arise with the stitched sensor technology for the inner layers.

However, adopting this backup would come at the cost of a slightly higher material budget, and thus a somewhat degraded p_t resolution at lower energies.

Comments:

- The chapter in the TDR is very long and detailed. In sections covering previously published R&D, the level of detail could be reduced. However, we found it very positive that past efforts from various groups were acknowledged, and that valuable experience and knowledge have been incorporated into the current vertex detector design.
- Although pursuing the HLHC 55 nm process could dilute focus and resources, it remains a worthwhile avenue, especially if difficulties arise in securing access to TPSCo65.
- While the mechanical design, simulations, and testing are advanced for this stage, maintaining mechanical stability and achieving efficient cooling remain significant challenges.
- The foreseen cost of sensors (~4.0 MCHF) appears reasonable. The relatively higher cost per unit area compared to other silicon detectors likely reflects the fact that the baseline sensor producer is not domestic.
- The estimated costs for mechanics, electronics, and the alignment system are of the correct order of magnitude but constitute only a fraction of the sensor cost

Recommendations:

- Re-evaluate the required performance specifications, focusing on operation during ZH and low-luminosity Z runs over the first 10 years.
- We reiterate our previous recommendation to expand the current manpower dedicated to design and R&D efforts. A close collaboration with ALICE-ITS3 on stitched sensor development is strongly encouraged.
- Explore constructing a mock-up with dummy heaters for thermal performance tests, to be used also for simulation validation.
- While pixel functionalities are similar for the ladder and stitched sensor options, significant differences in chip-to-chip (RSU-to-RSU) connectivity could impact performance. Address potential challenges (e.g., power and data distribution across large, stitched sensors) as early as possible through simulation studies.

Silicon Trackers

ITK Findings

The Inner Tracker (ITK) of the CEPC silicon tracking system adopts a baseline technology based on HV-CMOS monolithic active pixel sensors (SMIC 55), offering excellent spatial resolution (pixel size: $34\ \mu\text{m} \times 150\ \mu\text{m}$), moderate time resolution (3–5 ns), and low material per layer ($<1\% X_0$). As an alternative for the endcap, a HV-CMOS strip sensor technology is also under development.

Since the last review in October 2024, the project has made notable progress in ASIC development. The COFFEE3 chip (SMIC 55 nm)—the second-generation HV-CMOS prototype—was successfully submitted for fabrication in January 2025 and is expected to be delivered for testing in May 2025. COFFEE3 integrates two readout architectures, supports in-pixel time stamping, and includes significant improvements in power optimization and data-driven readout. In parallel, the CSC1 chip, which integrates a passive CMOS strip sensor and an analog front-end placed in the periphery (CMSC 180 nm process), is scheduled for tape-out in April 2025. Together, these developments demonstrate strong and steady progress in the ITK technology program.

Significant system-level advancements have also been achieved, particularly in mechanical design, cooling, and thermal management. The ITK now incorporates a well-defined multi-loop water-cooling system capable of maintaining stable operation at a power density of ~ 200 mW/cm², with thermal gradients controlled below 4°C. Updated mechanical simulations confirm the structural integrity of staves and validate the mechanical design for integration and prototyping.

ITK Comments

- The alternative ITK endcap technology based on CMOS strips, while offering slightly better intrinsic spatial resolution (~ 4 μ m), presents greater challenges compared to the HV-CMOS pixel baseline. Achieving 3D tracking requires stereo configurations, increasing the material budget and mechanical complexity. Furthermore, the integration of a CMOS readout circuit in the strip sensor periphery, although innovative, may introduce additional design and integration risks.
- The COFFEE3 chip successfully consolidates the sensor, analog front-end, coarse-fine TDC for time-of-arrival and time-over-threshold measurements, and data-driven digital readout into a single device. Achieving these capabilities within a power density of ~ 200 mW/cm² is an impressive technical goal. A successful validation of COFFEE3 would represent a major milestone, demonstrating the technological feasibility of the CEPC ITK concept.

ITK Recommendations

- Given the limited resources available for R&D and the increased complexity associated with the CMOS strip-based solution for the ITK endcap, we recommend a careful evaluation of the merit and timing of continuing this development. The CMOS strip approach remains a promising technological direction; however, it may be appropriate to continue work at a lower priority to preserve long-term potential, while focusing current resources on advancing the baseline HV-CMOS pixel system.
- We strongly encourage prioritizing the COFFEE3 validation campaign. As a comprehensive integration of position sensing and timing capabilities within a monolithic HV-CMOS technology, successful testing of COFFEE3 would constitute a critical milestone for the CEPC silicon tracking system, helping to de-risk the ITK concept and validate the soundness of the baseline detector design.

OTK Findings

The Outer Tracker (OTK) of the CEPC is a large area tracking system designed to provide precision timing and position measurements, complementing the inner tracking systems and the central Time Projection Chamber (TPC). It plays a critical role in improving momentum resolution for high-momentum tracks and mitigating performance degradation of the TPC in high-luminosity operations, where ion backflow-induced space charge can distort the drift field and impair tracking accuracy. The OTK is based on AC-coupled Low-Gain Avalanche Detectors (AC-LGADs) arranged as microstrip sensors, capable of delivering $\sim 10\ \mu\text{m}$ spatial and $\sim 50\ \text{ps}$ timing resolution, covering approximately $85\ \text{m}^2$ across the barrel and endcap sections.

Since the October 2024 IDRC review, the CEPC OTK system has made significant progress in response to the committee's recommendations. The AC-LGAD sensor design was updated, reducing the baseline size from approximately $8 \times 5\ \text{cm}^2$ to $4 \times 5\ \text{cm}^2$ to improve timing performance and manage higher particle rates. A new prototype sensor layout was submitted for tape-out in February 2025, featuring a variety of strip lengths, pitches, and electrode widths to optimize capacitance, noise, and spatial resolution. In parallel, the LATRIC readout ASIC was finalized and submitted for tape-out in April 2025, integrating all main functionalities. Mechanical and thermal simulations were also updated, confirming stable operation under a heat flux of $300\ \text{mW/cm}^2$.

OTK Comments

- The latest submission of the redesigned AC-LGAD sensor marks a key step toward determining the optimal strip geometry and layout for the OTK system. By systematically varying strip length, p-well doping, electrode width, pitch, and implementing isolation structures (e.g., trenches) to reduce capacitance, this prototype will provide crucial data to balance timing resolution, position resolution, power dissipation, and occupancy, ensuring the final design meets the stringent performance and rate requirements of the CEPC detector.
- The LATRIC ASIC is a critical demonstrator, consolidating all essential functionalities required in the OTK readout chain into a single chip. It includes a low-jitter analog front-end, clock generation via PLL, coarse-fine TDCs, internal calibration circuits, and a data readout architecture with serializer and control interfaces. Its successful validation will be essential to confirm the feasibility of a compact, low-power, fully integrated readout solution for the AC-LGAD-based OTK system.

OTK Recommendations

- To fully exploit the potential of this design iteration, it is strongly recommended to conduct exhaustive performance characterization of the newly submitted AC-LGAD sensors in test beams. These studies are critical to experimentally validate the effects of strip length, pitch, electrode width, n^+ -well doping, and isolation structures on timing and position resolution, power dissipation, and occupancy. Given the importance of timely feedback and the

current unavailability of the LATRIC ASIC, it is advised to perform these measurements using fast discrete amplifiers until the LATRIC ASIC becomes available.

- The collaboration should continue evaluating other LGAD sensor options and should closely follow novel developments, particularly the trench-isolated DC-LGADs currently under investigation by the team.
- It is recommended to ensure that the LATRIC ASIC is made available for sensor bonding as early as possible to enable realistic and fully representative characterization of the AC-LGAD sensor performance. Integrating the final readout ASIC with the sensor is crucial to obtaining accurate estimates of timing, position resolution, and power consumption.
- Given the critical role of the OTK in correcting space charge distortions in the TPC, it is essential that the OTK be robustly designed with service granularity in mind, exploring all known failure modes and developing strategies to mitigate them.

Gaseous Tracker

Findings

The Time Projection Chamber (TPC) with pixelated micro-pattern gaseous detector readout has been chosen as the baseline for the central tracker. It offers continuous 3D tracking with minimal material budget over a large volume and provides particle identification (PID) capabilities. This is a solid and well-justified choice for the Reference TDR, building on extensive past TPC experience at lepton colliders and more recent developments from ALICE, the LCTPC, and DRD1 collaborations. Several technological advances have recently been achieved; notably, a novel ultra-light QM55 carbon fiber material is now used to construct the TPC cylinder.

However, realizing the TPC as the tracker will require further significant efforts, particularly in mechanical design, structural integrity, and achieving uniform magnetic field conditions (see also the comments and recommendations in the Magnet chapter). Full simulation and digitization tools, based on the software framework and reconstruction algorithms, have been developed and must now be used to produce an updated set of performance plots (e.g., separation power plots).

The possibility of using a drift chamber as an alternative tracking option is also presented in the TDR.

Comments

Since the IDRC 2024 review, significant progress has been made in the simulation of space-charge distortions. For CEPC Higgs and low-luminosity Z runs, the maximum distortions after 3 meters of drift were found to be approximately 10 μm and 150 μm , respectively. Nevertheless, these figures require further verification, and careful evaluation of beam-related background effects must continue.

The choice of the pixel readout chip (TEPIX) with a $500 \times 500 \mu\text{m}^2$ pixel size is currently driven by power consumption considerations. A TPC module prototype with 10×300 readout channels (using 24 TEPIX chips) has been developed and is planned for beam testing at DESY in 2025. Larger prototypes will need to be developed and tested during the EDR phase.

If the alternative drift chamber option is pursued, it will be necessary to identify a hydrocarbon-free gas mixture that maintains the required performance characteristics.

Recommendations

- The structure of the TDR chapter needs better definition and further consolidation; some figures presented during the review must be updated in the document.
- Ion backflow during operation in low-luminosity Z-mode remains the primary concern. Further simulations of TPC space-charge distortions induced by beam-related backgrounds are necessary to optimize the Machine-Detector Interface (MDI). New module prototypes implementing advanced ion suppression techniques (e.g., double-grid structures or graphene layer coatings) should be studied.
- Optimization of pad size and gas mixture must be finalized, clearly distinguishing between the digital pixel readout case and the pad readout case. These two approaches lead to distinctly different electronics requirements, particularly regarding the need for analog-to-digital conversion in the pad case.
- In the longer term, dedicated simulation efforts are needed to develop and refine a "data-driven" approach for extracting space-charge corrections, based on the alignment and correction models developed for the ALICE TPC.

Electromagnetic Calorimeter

Findings

The high-granularity crystal ECAL is a recently proposed concept designed to be compatible with particle flow algorithm (PFA) reconstruction of jet energy within a homogeneous structure. The calorimeter is modular, with the fundamental detection units consisting of long orthogonal BGO crystal bars, read out at both ends by SiPMs.

Achieving high granularity at an affordable cost requires making strategic choices and compromises, guided by specific performance benchmarks. The team has made steady progress in understanding the performance and optimizing the performance-cost balance, using the ECAL standalone energy resolution and PFA jet resolution as primary benchmarks.

The baseline granularity was recently updated to $15 \times 15 \times 40 \text{ cm}^3$, resulting in a significant reduction in the number of readout channels and associated power needs. Simulations indicate that, even with this updated configuration, the calorimeter meets the target requirements for

boson-mass resolution and standalone electromagnetic energy resolution. The overall performance remains excellent, despite some degradation in π^0/γ identification and two-photon separation; the latter may potentially be recovered through improvements in offline reconstruction methods.

A full-scale prototype with the updated granularity is planned. Current simulation studies and experimental results from similar granularities provide confidence that the ECAL performance and component requirements are sufficiently understood, despite limitations in previous test beams due to electron beam spread. The team is also aware of the need for further progress beyond the reference TDR, particularly in defining QA/QC aspects for the components.

Comments

- The ECAL standalone energy reconstruction requires an energy threshold of 0.1 MIPs (as shown in Fig. 7.17), whereas the PFA jet reconstruction, based on fast simulation, adopts a higher threshold (Fig. 8.2). The ECAL team acknowledges the need to further develop and refine particle flow algorithms, photon identification at low energies, and π^0/γ separation to fully exploit the calorimeter's potential.
- The timing specification of 0.5 ns for MIPs was not strongly motivated. This corresponds approximately to the time spread across a full bar and is insufficient to provide significant benefits for event reconstruction. Additionally, the current time resolution analysis appears suboptimal. The team acknowledges that a deeper understanding of timing response and its potential use is needed, although this is not a priority at this stage.
- Crystal transparency variations are significant. While progress has been made toward developing a calibration plan using collision events, a quantitative demonstration that the precision and event rates are sufficient for monitoring response evolution across the detector is still missing.
- The non-linearity of the SiPMs poses a potential risk to the constant term in the energy resolution. Although compact photon detectors with more linear responses (such as APDs) were noted, SiPMs are preferred for cost and design uniformity reasons. However, SiPMs will require per-channel calibration, and possibly sorting during construction, along with continuous in-situ monitoring and corrections.
- Prototyping with close-to-final components is essential to confirm performance. Using existing readout ASICs can help decouple detector characterization from electronics debugging and allow parallel progress rather than sequential steps.
- The committee was pleased to see a substantial effort in understanding the impact of gaps between modules. However, the mechanical design and service layout within the gaps remains sketchy, and some components, such as the cooling plates, could be prone to underperformance.
- From a stylistic perspective, the TDR would benefit from an upfront presentation of the key performance benchmarks and detector specifications, followed by a focused discussion of the performance-cost optimization supported by R&D and simulation evidence. The detailed discussion of alternative options currently in Section 7.2 could be summarized briefly and moved to an appendix.

Recommendations

- Continue developing a full-scale prototype with the final geometry (using existing readout ASICs) and aim to confirm performance in an electron beam with low momentum spread.
- Further refine calibration strategies to ensure that necessary stability in transparency, linearity, and uniformity can be achieved in situ, without relying on a dedicated monitoring system.
- Advance preliminary engineering of the gaps between modules and fully assess their impact on reconstruction performance.

Hadronic Calorimeter

Findings

A key innovative feature of the Hadronic Calorimeter (HCAL) is the large sampling fraction achieved through the use of heavy scintillating glasses (GS). The proposed layout offers the potential to significantly improve the PFA resolution by enhancing the stochastic term of the single-particle energy resolution. The committee acknowledges that the team is highly motivated and is making steady progress toward this ambitious design.

The proposal is reasonable but aggressive. Detector specifications and performance benchmarks are clearly defined; however, considerable work is needed to bring the GS-HCAL baseline choice from its current R&D phase to a full-scale detector. For instance, during the review, the team was uncertain about the origin of the constant term observed in the hadronic energy resolution, a result obtained from idealized simulations with only limited hardware effects included. The committee is pleased to note that the next step is the construction of a large-scale prototype. In this prototype, each tile will be read out by four SiPMs, whereas in the final detector only one SiPM per tile is foreseen for cost reasons. Regarding costs, the committee notes that the current cost estimate does not include the cost of PCBs, which experience shows can become a significant item for granular calorimeters due to the complexity of design and production.

It was also acknowledged that a sampling calorimeter based on plastic scintillator (PS-AHCAL), which is more mature and better understood, remains a viable fallback option. The committee encourages the team to continue pursuing the GS-HCAL option, while maintaining the PS-HCAL as a backup.

Comments

- Scintillating glasses represent new territory for hadronic calorimetry. The material properties, such as radiation length and hadronic interaction length, are not yet fully characterized. Although the decision to adopt this technology is well justified, it carries significant risk. Therefore, extensive prototyping and simulation studies are mandatory to validate the concept. A deep understanding of the response to hadrons is essential,

including clarification of the constant term origin, study of the e/h ratio (software compensation), validation of GEANT4 physics lists, and accurate characterization of material properties such as quenching (Birks' law).

- The introduction of the TDR currently lacks references to important developments such as the CALICE AHCAL, built by German, Czech, and Japanese groups, which served as a foundation for the scintillator section of the CMS HGCAL.
- The process of down-selecting technology options should be better explained in the text. Statements such as "excessive power consumption" should be supported with quantitative arguments for clarity and transparency.

Recommendations

- Develop a detailed plan to validate the choice of GS-HCAL technology in a timely manner. This plan should include the development of glass samples with reproducible and controlled quality, along with a detailed understanding of single-particle and jet energy resolution.
- Prioritize the construction of a full-scale prototype. This prototype should incorporate the preliminary selection of glass tiles and ideally include a first version of both the readout ASIC and the PCB.
- Decide early on the final configuration regarding the number of SiPMs per tile and implement this choice in the prototype.
- Organize the group's work such that the prototype is simultaneously implemented into the simulation framework, including a complete digitization chain, to enable rapid feedback from test beam campaigns.

Muon Detector

Findings

The TDR presents a baseline design for the Muon Detector based on extruded plastic scintillator (PS) bars, coupled with WLS fibers and SiPMs. This approach is cost-effective, and the primary objective is to demonstrate that the system can meet stringent performance requirements across the detector's large active area.

The Muon Detector will consist of six superlayers in both the barrel and endcap regions. It uses long scintillator strips, 4 cm wide, with lengths exceeding 4 meters. Due to WLS fiber routing, effective optical lengths can be even longer, making attenuation length a critical parameter, directly impacting timing resolution and detection efficiency. Enhanced light collection strategies are thus essential.

The readout system adapts designs from the ECAL and HCAL and includes uSiPMs, front-end boards (uFEBs), and management boards (uMBs), tailored specifically for the Muon Detector.

Background studies show a maximum hit rate of ~ 2 Hz/cm², comfortably below the system's rate capability.

The Muon Detector supports trigger functionalities, including tagging long-lived particles (LLPs), although integration into the full CEPC TDAQ system remains preliminary. RPC technology, using environmentally friendly gases, is also under consideration via the DRD1 collaboration.

Comments

The TDR chapter and the review presentation show significant technical progress with clear physics motivation and encouraging prototype results. However, several areas need further development:

- Testing of full-length bars (4.2–5 m) must be completed, fully implementing optimized light collection methods.
- Standardized procedures for optical glue application (thickness, curing time) and validation of coating quality at scale are still required.
- A detailed system-level comparison between NDL and MPPC SiPMs is lacking. Given the Muon Detector's lower light yield and tighter timing requirements compared to calorimeters, selecting the optimal SiPM is critical.
- Consolidated detector specifications (spatial and time resolutions) must be validated through simulations and trigger performance studies.

Recommendations

- Complete comprehensive tests on full-length bars using final SiPM configurations and 2.0 mm WLS fibers, measuring timing, light yield, and efficiency.
- Define standard application procedures and conduct studies addressing mechanical stability, aging, radiation resistance, and refractive index.
- Systematically assess MPPC and NDL SiPMs (including new 20 μ m pixel versions), focusing on gain, DCR, cross-talk, and operating thresholds under realistic conditions.
- Expand simulation studies to include full event samples, especially displaced vertex signatures, to validate and refine time resolution requirements.
- Clearly state detector performance requirements linked to physics goals. Move Muon ID algorithm discussions to the global TDR Performance chapter and focus this chapter on local and standalone performance. Ensure consistent technical terminology and improve figure referencing.

Superconducting Solenoid

Findings

The CEPC Ref-TDR superconducting solenoid design has made significant progress. The solenoid provides a central magnetic field of 3T (for $t\bar{t}$, Higgs, and W runs) and 2–3T (for Z runs), with an inner bore diameter of 7.07 m and a cryostat/yoke half-length of 4.53 m. It is enclosed within an iron yoke that also houses the muon detector. The variation of the solenoid magnetic field in the central tracker (TPC) region is currently calculated to be approximately 10%.

The solenoid coil design is based on four layers of Al-stabilized superconductor, supported by an aluminum-alloy outer support cylinder and cooled using a two-phase helium thermal-siphon system. This general design concept is well established, drawing on the reliable design experience of the CERN-LHC CMS detector.

A key technology development effort focuses on Al-stabilized superconductors, with two approaches to mechanical reinforcement:

- **Option A:** Double-layered, two-step co-extrusion process.
- **Option B:** Single extrusion with micro-alloying for reinforcement, followed by a cold-work process.

Further R&D is planned for 2025.

The maximum von Mises stress on the Al-stabilizer during magnet excitation is evaluated at 96 MPa, which is close to the current yield strength of the Al-stabilizer material (105 MPa). Further mechanical evaluation is necessary to ensure an adequate safety margin, including the mechanical integrity of the NbTi/Cu superconductor itself.

Remaining technical challenges primarily concern the production of large-scale, high-strength Al-stabilized superconductors, as well as ensuring the reliability of the numerous aluminum-pipe welds required in the thermal-siphon cooling system.

Comments

- Field uniformity in the TPC region is a fundamental boundary condition for detector performance. The current 10% field variation must be carefully assessed to ensure it is acceptable for the required TPC performance. It is important to note that the ALICE detector has a different configuration, and that the ALICE magnet is much larger resulting in significantly better field quality compared to CEPC.
- The committee recognizes the remarkable progress achieved in developing Al-stabilized superconductors, achieving a yield strength of 105 MPa at 4.2 K using micro-alloying (Ni + Be). However, further R&D is required to demonstrate the full mechanical performance

of the conductor and the successful fabrication of full-scale Al-stabilized superconductor segments.

- It is critical to demonstrate that the full conductor (NbTi/Cu + Al-stabilizer) can withstand mechanical stress up to 135 MPa. This threshold is about 50% higher than the maximum von Mises stress (~90 MPa) expected during operation at 3T. If achieving this mechanical strength proves difficult, increasing the thickness of the support cylinder could be considered, although this would be undesirable. As a last resort, if necessary, a reduction in the operating field strength may need to be discussed.
- Developing a scaled model coil, including full cooling and excitation tests, will be an important step. Without such a demonstrator, confidence in the readiness of the solenoid technology for CEPC will be limited.
- Two cooling channel designs are under consideration:
 - Vertical Al-pipe channels, as used in CMS.
 - Sloped or tilted serpentine channels, which may reduce the number of welds and save cryostat space.

The latter option could be advantageous for construction reliability.

Recommendations

- Carefully verify that the magnetic field uniformity of the CEPC Ref-TDR solenoid meets the TPC performance requirements without introducing unacceptable distortions to particle drift paths. If necessary, explore revised designs, such as adding extra windings in the forward and backward regions, similar to solutions studied for ILD.
- Maintain the highest priority for the R&D of the Al-stabilized superconductor, aiming to demonstrate full-scale conductor fabrication, production of sufficient lengths, and mechanical performance that meets the required safety margins.
- Develop a detailed plan for the construction and testing of a superconducting model coil, including demonstrations of cooling and excitation. Achieving this milestone is essential for progressing toward full-scale magnet construction.
- Finalize the cooling system design based on the thermal-siphon concept, with particular attention to the reliability and quality assurance of aluminium-pipe welding procedures during construction.

Readout Electronics

Findings

The reference TDR provides a clear and accurate description of all on-detector and off-detector electronic systems required for the readout of the various detectors, along with their associated requirements—many of which are critical to the overall experiment performance. Since the last

committee meeting, several important decisions have been taken, most notably the adoption of full data transmission (streamed readout), which is now firmly established.

The nine different ASICs, currently at various stages of development, are clearly summarized, with an assessment of their maturity ranging from level 2 (initial design) to level 4 (prototyped) on a 1-to-5 scale. Three of these are common service ASICs used for power distribution and data transmission. The required manpower for the ASIC developments and future steps is well detailed and appears realistic.

The shared back-end architecture and power distribution scheme are also well described, building upon the solid experience gained from large-scale operational experiments such as JUNO.

Comments

The committee is very impressed with the progress made and the thorough implementation of previous recommendations. In particular, the decision to adopt a data streaming readout is a major step forward, significantly impacting the ASIC architecture. The handling of data rates and the application of safety margins are well managed.

The ASICs are well documented, and several key building blocks have already been successfully prototyped, demonstrating the strength and efficiency of the teams involved. The progress is remarkable; however, as the complexity of the chips increases, substantial work still lies ahead. It would be helpful for each ASIC to clearly identify the lead responsible person(s) and the contributors to the various blocks and development tasks.

Verification of the ASICs is critical, especially for large chips that include substantial digital logic (e.g., tracking ASICs). Verification should be conducted by individuals independent of the design teams and should use state-of-the-art tools and methodologies.

Testing of ASICs can never be too extensive: it should involve not only the designers but also multiple user teams wherever possible. Chips should be characterized across a range of temperatures and supply voltages to detect weaknesses early. As larger volumes become available, systematic analysis of failing or underperforming chips (“bad chips”) will be essential.

From a stylistic standpoint, the SIPAC chip (common to three detectors) is extensively described within this chapter, while other ASICs are detailed within each detector’s respective section. For improved readability, it may be preferable to move the SIPAC description to the calorimeter section, retaining in the electronics chapter only a general overview of all ASICs and a detailed view of the shared service chips (e.g., for power and data transmission). For the front-end ASICs, detailed schematics could be reduced in favour of presenting more relevant simulation results, such as pulse shapes or noise versus input capacitance.

The back-end electronics are comprehensively described and benefit from the team's previous experience with large-scale systems. However, the availability of high-performance FPGAs could pose a risk. Maintaining a dual-source strategy for critical components would be a prudent mitigation measure.

Recommendations

- Identify the teams responsible for testing each ASIC, ensuring they are different from the design teams. Bugs are often discovered late in development, and broader involvement in chip evaluation improves characterization and helps uncover potential weaknesses.
- Begin detector testing with available “sister chips” where feasible. This parallel approach can advance both detector and chip development and provide valuable benchmarks for comparison.
- Minimize power dissipation at every stage. Each milliwatt of consumption should be justified by a specific performance requirement.
- Proceed with system development in parallel with chip development. System-level issues can often be addressed early at the chip design stage if identified promptly.

Trigger and Data Acquisition

Findings

The TDAQ system is based on the full transmission of data from the front-end electronics to the back-end electronics (BEE), with the trigger operating at the BEE level. Data rates from each subsystem are provided for the different CEPC operational scenarios; however, the underlying assumptions behind these estimates should be discussed more explicitly in some cases.

The TPC represents a special case due to its long drift time, and a more detailed discussion of the implications of this extended integration time is necessary.

The first stage of the trigger is planned to operate on the BEE boards, generating trigger primitives from the calorimeter and muon systems, with the possible addition of primitives for a track trigger. The trigger architecture has evolved to include BEE-resident primitive algorithms alongside a three-layer L1 hardware system. The CEPC team has demonstrated that simple requirements on calorimeter and muon primitives effectively reduce background rates while preserving high efficiency for the primary physics processes of interest, addressing a key recommendation from the previous review.

It was noted that some exploratory work on a fast-track trigger has been conducted, though no results were presented. A silicon-based fast track trigger remains of high interest, as it would complement calorimeter and muon primitives and provide stronger separation between collision and non-collision backgrounds.

This review focuses primarily on the ZH production and low-luminosity Z-pole running planned for the first ~10 years of CEPC operations. Nevertheless, the overall TDAQ design appears appropriate for other planned running modes as well.

The Ref-TDR chapter should first describe the proposed baseline trigger system before presenting proof-of-concept studies for the trigger primitive algorithms.

Comments

The background event rate and event size are critical inputs to defining the TDAQ system dataflow requirements. Assumptions used should be explicitly stated and consistently referenced across all relevant chapters (some inconsistencies were noted in the current draft).

The safety factors applied in system design should be made explicit, clearly motivated, and documented to allow for easy updates and assessment of their impact on different subsystems.

Section 12.4 should be moved after Section 12.6 to improve the logical flow of the chapter.

Recommendations

- Clearly motivate the data bandwidth requirements into the L1 trigger, between trigger layers, and specify the number of links needed. Provide details on the expected trigger primitives from each BEE type and the associated data volumes. Similarly, describe the output products of the local trigger layers and provide data bandwidth estimates into the global trigger layer.
- Update the design of the common trigger board—particularly the bandwidth and input link capacities from the BEEs—to reflect the information presented. Explain the rationale behind the estimated number of trigger boards required at each of the three L1 levels.
- Provide a detailed description of how the TPC data stream will be handled, especially considering its long drift times.
- Continue exploring the development of a fast-track trigger, emphasizing its potential to reduce background rates and ease the load on the HLT.
- Demonstrate the necessity for RDMA technology in transferring data from the BEE to the HLT, including a justification based on system performance needs.
- Describe the system's strategy for handling data flow in the event of full buffers to avoid data loss or system instability.
- Evaluate whether the L1 trigger hierarchy could be simplified, for example by collapsing it into a multiplexer layer feeding directly into boards handling full global trigger primitives.
- Categorize the event rates into physics signals (e.g., ZH production, two- and four-fermion processes, Bhabha scattering), gamma-gamma interactions, beam-related backgrounds. The first category aims at full trigger efficiency, the second one is potentially of some interest but with lower priority and the last one must be reduced as much as possible.
- Analyze the beam background distributions as functions of energy, multiplicity, and polar/azimuthal angle to optimize background suppression algorithms. In particular,

investigate the origin of background peaks such as the one observed at ~ 7 GeV in the ECAL endcap (Figure 12.10).

Offline Software and Computing

Findings

The review team acknowledges the very good progress made in software and computing since the last report. Almost all previous recommendations have been successfully addressed by the collaboration. The current status and advancements are thoroughly documented in the Ref-TDR and were presented and discussed during a dedicated session at this review. Overall, the Offline Software and Computing activities are in a strong position, the outlined approach appears highly feasible, and no significant challenges or showstoppers have been identified.

The offline software, CEPCSW, is based on the international Key4hep project, which is used by all future collider projects. CEPC is making significant and well-aligned contributions to this effort within the global community. In terms of computing, CEPC leverages established tools and methodologies developed for the LHC. The current computing model follows the traditional Tier structure, though a transition to a more efficient model is envisaged, aligned with broader community developments.

A detailed simulation model of the reference detector has been implemented. For track reconstruction, the collaboration is currently adapting tools and algorithms from the linear collider community while also developing more advanced techniques. A new particle flow algorithm, CyberPFA, has been developed, particularly to meet the challenges posed by the novel crystal ECAL design. This builds on prior work from the HCal side and has been successfully integrated into CyberPFA. The full simulation and reconstruction framework has been employed to produce a comprehensive set of technical performance plots, demonstrating that the detector concept, combined with sophisticated algorithms, can meet the required physics performance targets.

The offline software has also been used to estimate the computing needs, based on reasonable—though not highly refined—assumptions regarding data rates, data volumes, processing schedules, and methodologies. However, the estimates also assume optimistic developments in CPU performance and cost trends over the next decade.

Comments

- The corresponding Ref-TDR chapter would benefit from some revisions: reducing the length of more general descriptive sections and expanding technical details where appropriate.

- Some important aspects of the planned computing model remain insufficiently detailed, such as the structure for user analyses, data distribution strategies, open data policies, and long-term data preservation plans.
- The development of CyberPFA is a significant achievement and positions the collaboration as a leader in custom in-house algorithm development. However, the degree of reliance on existing tools like PandoraPFA and ArborPFA needs to be better explained and justified.
- While the event data model (EDM4hep) is well established, more detailed data formats for reduced, compact data levels (e.g., various stages of AODs) are not yet clearly defined.
- The costs associated with the required computing infrastructure—such as whether a new data center will be necessary—are not discussed in the Ref-TDR.
- While future technologies such as quantum computing, machine learning, and AI are mentioned, their potential roles and the motivation for pursuing them should be more clearly articulated, with better-defined staging scenarios.

Recommendations

- Further develop and consolidate the detailed simulation model and reconstruction framework, with particular attention to:
 - Adding realistic estimates for detector imperfections, such as gaps or insensitive material from cables and services, to the simulation model.
 - Performing detailed studies of the CyberPFA performance, including jet energy resolution for light-quark dijet events, single particle responses, neutral hadrons, and neutrino energy measurements in b-decays.
 - Studying the impact of misalignment on tracking performance.
- Develop a systematic performance monitoring framework, organized around a set of benchmark figures and analyses that qualify simulation and reconstruction quality. Each sub-detector should define its own "detector performance" plots, while a standard set of global physics performance plots should be maintained to validate software improvements.
- Once consolidated, establish a clear reference for detector performance within the Ref-TDR, ensuring that future detector and physics performance studies use consistent versions and conditions (e.g., ensuring that jet resolution and b-tagging results are presented with the same software version and setup).
- Develop a more detailed and phased estimate of computing needs during the R&D and preparation phases, for example by organizing a sequence of increasingly challenging "data challenges." These would also serve as benchmarks for assessing the maturity of the offline software and computing infrastructure.

Mechanics and Integration

Findings

Significant progress has been made since the last review in the mechanical design of the CEPC reference detector. The weight and boundary dimensions of each sub-detector are now well defined in the Ref-TDR. Clearances between sub-detectors are typically around 10 mm, which appears adequate to ensure a seamless installation sequence. However, in the critical interface between the integrated beam-pipe/vertex detector assembly and the ITK, the clearance is reduced to just 2 mm.

The issue of sagging of the magnet yoke under its own weight has been effectively addressed. The introduction of end-flanges satisfactorily mitigates the problem, reducing the sagging from about 13 mm to within the required limit of less than 1 mm.

The installation sequence has been carefully developed. The core shaft designed for the installation of the barrel HCAL and ECAL detectors appears adequate, as does the cantilever system proposed for installing the TPC, ITK, and beam-pipe assembly (including the Vertex and LumiCal detectors). The animations presented clearly illustrated the feasibility and logic of the installation process.

Comments

The mechanical connection structures between sub-detectors are well designed, considering the specific requirements of each system. However, the methods by which these structures will achieve the precise alignment of the sub-detectors are only briefly described in the current Ref-TDR. Reference is made to an alignment control network that would provide positional benchmarks, but further details would strengthen the documentation.

Significant progress has also been made in designing the auxiliary facilities, which include the air-conditioning system for the experimental hall, gas and cooling systems for the various sub-detectors, power distribution and electronics systems, the cryogenic system for the superconducting magnet, and the hydraulic pump station used to move the detector to the collision point.

Recommendations

- Further refine the routing plans for cables and services, especially in the critical interface areas between sub-detectors and through the barrel-to-endcap transitions to the outside of the detector and the auxiliary facilities.
- Pay particular attention to the air-cooling and service routing for the vertex detector, where space within the beam-pipe assembly is extremely limited. It is essential to verify that sufficient space is available to ensure the proper functionality of the air-cooling system, without compromising reliability or maintenance access.
- Regarding refrigerant choice for the other sub-detectors:
 - **Supercritical CO₂ (sCO₂)** could be considered as an alternative to water. As a single-phase refrigerant, it is user-friendly for complex, multi-branch systems. Like water, it is non-toxic, non-flammable, and readily available; unlike water, it is

dielectric and harmless to electronics in case of leaks. While sCO₂ brings operational advantages, these must be weighed against some disadvantages (see appendix for details).

- **Boiling CO₂** (subcritical, operating around 20–25 °C) is also simpler than two-phase CO₂ systems, which require more complex design, maintenance, and expert operation. Thus, unless two-phase cooling is clearly necessary, a simple and efficient single-phase refrigerant (such as water or supercritical CO₂) is preferable.

Detector and Physics Performance

Findings

Although further improvements are possible and recommended, the collaboration's response to the previous evaluation has been excellent and deserves recognition. The detector software, particularly the simulation component, now enables comprehensive studies of detector performance and physics benchmarks. Several design decisions have been made based on both detector and physics performance criteria. This progress has allowed the definition of a baseline detector concept. While still perfectible and requiring further detailed investigation, this is a significant and commendable step forward. The scope of physics performance studies has broadened considerably since the last assessment. The range of channels explored and the methodologies employed respond well to previous recommendations and collectively address key detector areas and primary physics topics.

Comments

- The Ref-TDR text could be improved and shortened by presenting certain aspects (e.g., particle identification, jet studies) more concisely, while still ensuring that the algorithms are described clearly and transparently.
- Many additional physics analyses are planned. Their presentation and motivation should be aligned with the main purpose of this document: demonstrating the feasibility and physics potential of the reference detector.
- Algorithms for particle identification, jet tagging, etc., are mentioned in multiple places and are still evolving. A systematic approach should be adopted to define and refer to these algorithms consistently across the document.

Recommendations

- The studies encompass both physics benchmarks and detector performance metrics. A clearer distinction between these two aspects would be beneficial. The editorial team is encouraged to organize the content into:
 - *Sub-detector technical performance* — technical performance figures (used for sub-detector configuration decisions) should be placed in the relevant sub-detector chapters.

- *Physics-related performance* — to demonstrate baseline detector capabilities for physics analyses (e.g., particle identification, global variables like ET_{miss}), and to present the physics analyses themselves. This should be the main focus of the performance chapter.
- The physics benchmarks listed in Table 15.3 are intended to demonstrate the performance of the reference detector. Each listed study should be discussed explicitly, explaining the role of the detector performance in achieving the result. Currently, only a subset (e.g., Higgs recoil mass, Higgs branching ratios, weak mixing angle from $Z \rightarrow \mu\mu$) are covered. Other channels (exotic Higgs decays, LLPs, CVP in D-meson decays) should also be summarized, possibly in a summary table.
- The technical improvements planned for more detailed simulation (noise, event overlap, misalignments, calibration effects) should be pursued, and their impact on physics performance carefully demonstrated.
- Several specific technical issues should be clarified and potentially improved:
 - The photon efficiency behaviour around $E = 1$ GeV appears unusual and should be either justified or corrected, as it could impact EM/hadronic separation in PFA and influence missing energy and mass resolution.
 - Jet energy resolution for light quarks should be studied more systematically and used as a benchmark metric.
 - Missing energy reconstruction should be further investigated, particularly in b- and c-jet events with tagged leptons and in BSM channels with large ET_{miss} .
 - (Longer term) Tracking performance should be tested in exotic scenarios, such as long-lived particles.
 - (If possible) Include photon conversions in tracking studies as a material probe and describe their treatment in PFA.
 - Particle identification needs further organization and development:
 - Currently, simple cuts are applied; more sophisticated algorithms (including ML-based methods) should be considered, balancing the ambition against available time and resources. A simpler multivariate approach could serve as an intermediate step.
 - Different working points ("tight" for high purity, "loose" for high efficiency) should be defined and used consistently across analyses (e.g., Figure 15.7, where a 90% WP for muon/electron ID is mentioned). Coherence with PFA must be ensured (avoiding double-counting residual energy, etc.).
 - The description of jet flavor identification needs to be streamlined. Currently, information is dispersed across sections (vertexing, tracking, PFA):
 - A brief overview of standard Jet Flavour Tagging (JFT) is given in Section 15.2.6, while Jet Origin Identification (JOI) is discussed in 15.2.7. However, it is unclear what performance gains are achieved by moving from JFT to JOI. Benchmark comparisons (b/c-tagging efficiency versus misidentification rates at Z-pole and ZH 240 GeV) should be provided to evaluate performance systematically.
 - (Longer term) Consider a dedicated chapter for jet flavour tagging, especially given the comprehensive nature of JOI (which involves many

sub-detectors). Comparative studies between "ideal" and "compromised" performance would also help derive systematic uncertainties in the AI-based approach.

- (Longer term) The confusion matrix (Fig. 15.22) suggests JOI could distinguish quarks from anti-quarks. If validated, this could significantly improve flavour-specific AFB measurements. Physics benchmarks involving b/c-quark AFB (or even strange quarks) should be added if feasible.
- The offline software environment is evolving rapidly. Performance studies should be conducted with synchronized and version-controlled frameworks, especially as CyberPFA depends critically on tracking, particle ID, calibration, and alignment inputs.
- (Longer term) Organizing "data challenges," as mentioned in the "Offline Software and Computing" section, could be valuable. These would serve both as benchmarks for detector performance and stress tests for computing models (through massive production, analysis, and quality checks).
- A centralized database tracking the produced samples and their statistics should be maintained and updated (extending Table 15.4). Technical samples (e.g., single electrons, muons, decaying kaons for PFA studies) should be included and documented similarly for use in detector performance validation.

Overall Construction Cost and Timeline

Findings

The cost estimation is provided as an ancillary table outside the Ref-TDR, following a standard Work Breakdown Structure (WBS). Only summary tables for the baseline detector option are included within the TDR text itself. Costs are primarily based on raw material prices (by volume, weight, or area), projected over the next 5 to 10 years, and multiplied by a fabrication factor.

Currently, only a very rough timeline for detector construction is presented in Chapter 16 of the TDR.

Comments

For future, more informed cost reviews, it would be beneficial to provide the full WBS information, with a breakdown of:

- Raw material costs
- Production costs (including fabrication losses)
- Labour costs
- Integration costs

Additionally, for each item, the following should be clearly specified:

- Unit description (clearly linked to the detector design)
- Unit cost (in original currency)
- Basis of the estimate (e.g., vendor quote, internal prototype, previous projects, catalogue pricing)
- Quantity required for the detector
- Quantity including yield loss
- Total quantity and total cost

It should be explicitly stated whether the current costs include allowances for yield loss.

A structured, consistent table format for each major cost item would greatly improve clarity and traceability.

Concerning computing costs: while the extrapolations appear reasonable, the model could be further refined by considering data size, number of re-processing, number of data copies, and data accumulation schedules. Even a small additional effort to refine or document these factors would be valuable.

The projections for hardware performance and price evolution over 10 years seem somewhat optimistic, as they combine assumptions of improved CPU performance per core with reductions in cost per core.

Regarding the detector construction timeline, the information currently provided is too limited to allow for a meaningful assessment.

Recommendations

At this stage, we highlight areas where additional checks and refinements of the cost evaluation are advisable:

- **ECAL BGO Crystals:** The cost of BGO crystals significantly impacts the overall budget. The table presented dates back to 2019. We recommend engaging with multiple vendors to ensure the required quality, production rate (given the enormous volume), and lowest possible price.
- **SiPM Packaging Costs:** The cost of packaging could be as high as 50% or more. It should be clarified whether packaging costs are included in the quoted SiPM price, and this should be explicitly reflected in the WBS. The current estimate of \$1.25 per channel may be optimistic.
- **HCAL Cost Estimates:** Currently, the lowest informal offer is used for the cost of glass plates. Until it is confirmed that the lowest bidder meets the required specifications, it would be more prudent to use the average of the three quotes. Additionally, the cost estimate assumes one SiPM per glass plate, whereas the detector design currently

requires four SiPMs per plate for viability. Some cost uncertainty reflecting the outcome of ongoing R&D should be included.

- **Muon Detector Electronics:** The cost book reports 43,000 ASICs, 43,000 FEE units, and 43,000 Readout FEE units. However, the readout architecture has recently been updated to a three-stage system, and the number of FEE boards (uFEBs) should be significantly lower than the number of SiPM channels. The cost estimate should be updated to reflect this new architecture. Additionally, the TDR mentions the need for 72 Management Boards for slow control and DAQ interfacing—these boards are currently missing from the cost book and should be added.
- **Magnet Cost Comparison:** The TDR quotes a magnet cost of about 22 MCHF, compared to around 130 MCHF for the more complex ILD system (which includes the yoke). A cross-check should be performed to understand and justify the difference between these two estimates.
- **TPC Cost Comparison:** Similarly, the TPC cost is estimated at 5 MCHF, whereas the ILD TPC estimate was 36 MCHF. This discrepancy should be investigated and explained.

APPENDIX I - Detailed Comments

Introduction

None

Concept of CEPC Reference Detector

None

Machine Detector Interface and Luminosity Measurement

None

Vertex Detector

None

Silicon Trackers

None

Gaseous Trackers

None

Electromagnetic Calorimeter

None

Hadron Calorimeter

- As with other sections, the presentation was significantly better than the corresponding document and should serve as a guideline for revising the written material.

- The authors are encouraged to make the text more concise where possible, while still providing sufficient detail where necessary.
 - Captions such as that for Fig. 8.12 ("The real AHCAL prototype") are inadequate and should be made more informative.
 - The authors should carefully review the text to ensure that all figures are properly motivated, and that the overall argumentation is logical and coherent.
 - For example, it is unclear why the emission spectrum shown in Fig. 8.52 is included—is it representative or used for digitization? This should be explicitly explained.
 - In Fig. 8.45, while the test beam setup is now described, the reason for the visible shoulder (likely due to a double MIP peak) is not explained. Since the test beam is a major highlight of the R&D effort, the authors should focus on presenting it clearly and concisely, possibly selecting a few key highlights for emphasis.
- Section 8.5 contains an extended discussion of SiPMs, mixing general background information with results from the collaboration's own R&D. This section should be streamlined to separate general background from specific experimental achievements.
- Since SiPMs are used extensively across the calorimeter and muon systems, it would be more effective to consolidate the discussion of SiPM R&D into a single section. This would avoid redundancy and present a more coherent overview of the topic.
- Mechanical integration aspects, currently occupying much of Section 8.7, could be better placed within a general "Detector Integration" section. The authors should use their judgment to decide which integration details are most relevant to retain in the specific HCAL section, while moving broader topics to a centralized discussion.

Muon Detector

R&D Progress

- **WLS Fibers:** 2.0 mm diameter Kuraray Y11(200) fibers achieve approximately 2.8× higher photon yield compared to standard 1.2 mm fibers.
- **Optical Coupling and Coating:** Optical glue improves fiber-scintillator coupling by 25–30%. Aluminum foil wrapping enhances far-end light collection by ~38%, with further improvements achieved using TiO₂ or Teflon reflective coatings.
- **Scintillator Composition:** Enhanced scintillator formulations yield an approximate 50% increase in fluorescence output.
- **SiPMs (Room Temperature Operation):**
 - **MPPC SiPMs (50 μm pixels):** High gain with relatively low dark count rate (DCR).
 - **NDL SiPMs (15 μm pixels):** Higher DCR and optical cross-talk; new 20 μm pixel models from NDL show improved gain and a moderate reduction in DCR.

Prototype Results

- **Short Bar Prototypes (1.5–1.6 m):**
 - Detection efficiency >98%
 - Time resolution of ~1.5 ns (NDL) and ~1.2 ns (MPPC)
- **Improved Strips (1.6 m):**
 - Achieved ~50 photoelectrons (p.e.) at the far end using 2.0 mm fibers, optical glue coupling, and aluminum foil reflectors.
- **Long Bar Prototype (4.0 m):**
 - Preliminary results show >90% detection efficiency with a ~27 p.e. threshold at the far end.

Superconducting Solenoid

None

Readout Electronics

Figure 11.15 has the CERN IpGBT numbers on it, not the CEPC ones.

Trigger and Data Acquisition

For estimating storage needs, it would be valuable to provide separate estimates of event sizes for signal and background events passing the trigger. This separation would allow for future refinements as the understanding of background processes improves. While Table 12.5 offers separate readout event size estimates, the storage size is currently assumed to be constant. The motivation for these storage size assumptions, along with the applied safety factors, should be clearly documented.

Figure 12.15 is foundational to understanding the overall trigger architecture. However, it currently includes elements that are only relevant for high-luminosity Z running, which is planned for approximately 10 years in the future (if pursued at all). It would be helpful to provide a diagram that clearly distinguishes baseline design elements from possible future extensions—or to present only the baseline design to avoid confusion.

There is an inconsistency in the use of nomenclature between "BEE" and "ODE" throughout the document. For example, Figure 12.1 refers to "Off-Detector Electronics," while Figure 12.15 uses "BEE." A consistent terminology should be adopted throughout the TDR.

Additionally, the low-luminosity Z mode is inconsistently referred to as both "10 MW" and "12 MW" in different parts of the TDR. This should be corrected for consistency.

Offline Software and Computing

None

Mechanics and Integration

Detailed Considerations on the Advantages and Disadvantages of a Refrigeration System Based on Supercritical CO₂ (sCO₂) Compared to Water.

Advantages of sCO₂ Refrigeration:

- sCO₂ is fully dielectric, making it inherently safer for use near sensitive electronics.
- sCO₂ requires much smaller piping for the same flow rates due to its lower viscosity and higher operating pressure.
- It has a higher heat transfer coefficient, provided that operational conditions (temperature and pressure) are well optimized.
- The high fluid pressure allows for larger acceptable pressure drops in the heat exchanger (e.g., in the detector stage's tubes). With optimized heat exchanger design, it is theoretically possible to induce the necessary pressure drops to maintain nearly isothermal conditions—achieving single-phase flow with the same inlet and outlet temperatures while efficiently removing heat.
- Experience shows that low-pressure circuits tend to develop leaks over time due to lower quality construction standards, while high-pressure systems, requiring stricter quality controls, are more robust. This leads to lower long-term maintenance costs (both financial and operational).

Disadvantages of sCO₂ Refrigeration:

- High-pressure systems are generally estimated to be 15–20% more expensive than equivalent low-pressure systems.
- Operation at 30–35 °C might impact surrounding systems and needs to be considered during integration.
- A dedicated pressure control component ("bladder" or similar device) is required to maintain stable system pressure.
- Depending on the system volume, there are potential safety issues (especially personnel safety) associated with storing large quantities of CO₂ in confined underground spaces.
- The combined effects of all parameters (temperature, pressure, flow rates) on the thermal-fluidic behaviour of supercritical CO₂ are not yet fully understood. However, significant research and development work in this area is currently ongoing.

Detector and Physics Performance

None

Overall Construction Cost and Timeline

None

APPENDIX II – Charge

The Circular Electron Position Collider (CEPC) project, a Higgs factory proposed to be constructed in China, plans to have two experimental interaction regions. After the CEPC project approval, the final two detectors will be developed as International Collaborations. In the meantime, a Reference Detector is being designed to demonstrate the readiness and feasibility of detector technologies.

Following the successful publication of the CEPC Accelerator Technical Design Report (TDR) in December 2023, the CEPC Design Group is devoting its attention to the CEPC Reference Detector TDR. This Reference Detector TDR, together with the accelerator-engineering designs, will complete the preparation of the CEPC proposal to the governments (central and local) for project approval in China.

The Reference Detector described in this TDR could be but does not necessarily need to be, one of the two detectors that will eventually be installed in the CEPC. The technologies used in the Reference Detector are forward-looking, and although not necessarily at hand today, they could be achieved within the timescale of the start of construction. The reference TDR also includes alternative options for each sub-detector, which are often more conservative, based on technologies available today, and for which more advanced prototypes might have already been produced.

The primary focus of the second meeting of the CEPC International Detector Review Committee (IDRC), scheduled for April 2025, is to present the first draft of the Reference Detector TDR and initiate its comprehensive review. The CEPC Design Group will make presentations to cover each chapter of the Reference TDR. The presentations will cover the physics requirements, a technology survey including a justification of the choices made and technical challenges, the progress since the 2024 IDRC meeting, research and development (R&D) including the current status of research, details of the current design, performance from simulation, research team, and the future work plan. In this second meeting, the CEPC International Detector Review Committee (IDRC) is charged with:

1. Reviewing the overall status of the research and development (R&D) efforts for the CEPC detectors, which will be highlighted in the presentations.
2. Assessing the detector design objectives, ensuring they are clearly defined and aligned with the CEPC project's overall goals.
3. Reviewing the detector design and key technologies, examining their effectiveness in achieving the specified performance goals.
4. Scrutinizing the feasibility and readiness of critical technologies and components for the detector system, considering the proposed Reference Detector and alternative technologies, and ensuring that they will be available or achievable by 2028 through ongoing R&D programs
5. Identifying and assessing the primary technical risks associated with the detector design and propose mitigation strategies.
6. Determining the project's readiness for construction upon completion of the outlined R&D and engineering preparations, ensuring any issues identified are addressed in a timely manner.
7. Providing advisory comments on the overall cost evaluation of the detector
8. Providing additional observations or suggestions for improvements across the project's spectrum.
9. Preparing for subsequent reviews by establishing the framework for ongoing evaluation and feedback processes, if necessary and as needed.

The first draft of the Reference Detector TDR was made available early April 2025, and subsequent updates are expected following the review meeting. Future iterations with the IDRC are expected to finalize the document.

APPENDIX III – Committee Contributions

- MDI – **H. Yamamoto**, Kramberger, Titov
- Vertex Detector – **Kramberger**, Vila, H. Yamamoto
- Silicon tracker – **Vila**, Kranberger, Kowalewski
- Gaseous Detector – **Titov**, Colas, Colaeo
- Electromagnetic Calorimetry – **Tabarelli de Fatis**, Poeschl, Brau
- Hadron Calorimeter – **Poeschl**, Tabarelli de Fatis, Brau
- Muon detector – **Colaleo**, Colas, Han
- Superconducting Solenoid – **A. Yamamoto**, H. Yamamoto, Brau
- Readout Electronics – **De la Taille**, Gay, Titov
- Trigger and DAQ – **Gay**, Tenchini, Kowalewski
- Offline Software and computing – **Gaede**, Diaconu, Kowalewski
- Mechanics and integration – **Schmidt**, Kramberger, Tabarelli de Fatis
- Detector and physics Performance – **Diaconu**, Tenchini, Gay, Han
- Cost and timeline – **Tenchini**, Schmidt, De la Taille
- Chair – Bortoletto