https://doi.org/10.1142/9789811280184_0012

Chapter 12

The CMS HL-LHC Phase II upgrade program: Overview and selected highlights

Marcello Mannelli

CERN

1. Introduction

With the unprecedented instantaneous luminosity of the HL-LHC, the Phase II program aims to provide about one order of magnitude more integrated luminosity compared to Phase I (∼3'000–4'000 fb^{-1} vs ∼300–450 fb^{-1}) and offers the potential for an extensive program of precision measurements of Standard Model processes and possible deviations from the predictions. This includes a detailed characterization of the Higgs sector, together with direct searches for rare processes and subtle and/or exotic signatures in the search for physics beyond the Standard Model. Realizing that potential in the HL-LHC environment, with up to 200 almost-simultaneous interactions (in-time pile up) per 40 MHz bunch crossing rate, together with unprecedented levels of radiation exposure to the detector, presents a series of formidable experimental challenges, which define the requirements and drive the scope of the CMS Phase II upgrade program.

Briefly, these can be summarized as follows:

- The High Radiation Environment necessitates:
  - The complete replacement of the Tracker and Endcap Calorimeter systems, which cannot continue to operate at integrated luminosities much higher than the 300∼450 fb^{-1} foreseen for the Phase I program. By then, the performance will have already been substantially degraded due to radiation effects;

  - Cold operation (around 9°C compared to 18°C during Phase I) of the Barrel ECAL to mitigate radiation induced dark currents in the Avalanche Photo Diodes (APDs), which would otherwise result in unacceptable performance degradation through increased noise.

- The High Pile Up motivates:
  - Improved granularity wherever possible, and extended Tracking capability to cover higher pseudorapidity;
  - Novel approaches to in-time pile up mitigation, and in particular the use of Precision Timing (∼30 ps) to discriminate between particles, both charged and neutral, originating from the collision of interest and those resulting from pileup collisions within the same single bunch crossing (∼180 ps RMS).

- The High Luminosity requires:
  - Substantially improved L1 Trigger primitives for better selectiveness, despite the high pileup, which in turn necessitates a longer L1 trigger latency, and motivates the inclusion of L1 Tracking Trigger capability as an integral part of the CMS Tracker upgrade, together with high granularity calorimeter information read-out at 40 MHz, and the introduction of additional End-Cap Muon Stations for use in the L1 Trigger;
  - Adapting the Front-End read-out electronics, L1 Trigger and Data Acquisition (DAQ) systems, to accommodate longer L1 trigger latency, 12 μs compared to the current 4 μs, accommodate an increased event size and allow higher read-out bandwidth, 750 kHz compared to the present 100 kHz, to provide high efficiency across the broader set of physics signatures of interest;
  - Dimensioning the High-Level Trigger (HLT) computer farm to run online event reconstruction on the 750 kHz input stream and select up to 7.5 kHz for permanent storage and further off-line analysis.

Taken together, these considerations lead to a set of CMS Phase II upgrade projects[1] which include:

- Replacement of the existing Silicon Tracker with a new Silicon

Tracker[2] featuring, in addition to improved radiation hardness, adequate for 4'000 fb^{-1};

  - A fourth Inner Pixel layer and higher Outer Tracker granularity to maintain sufficiently low cell occupancy to provide reliable pattern recognition in the high pile up HL-LHC environment;
  - An on-detector front-end system capable of locally identifying, selecting, and transmitting Track vectors for charged particles with $p_T > 2$ GeV/c, together with a back-end system able to process these and reconstruct tracks at 40 MHZ for use in the L1 Trigger;
  - Extended tracking coverage, up to $|\eta| \sim 4$, mostly achieved through an extended Inner Pixel Tracker;
  - Substantially reduced material within the Tracker volume, to improve momentum resolution for charged tracks, as well as improving the calorimeter performance and quality of electron and photon reconstruction, while also reducing tracking inefficiencies and/or reconstruction errors resulting from hadrons undergoing nuclear interactions.

- Deployment of dedicated novel precision timing detectors, in front of both Barrel and End-Cap calorimeters, capable of providing 30 $\sim$ 40 ps precision for Minimum Ionizing Particles (MIPs).[3,4]
- Major overhaul of the existing Barrel ECAL (EB) and HCAL (HB) detector systems.[5] The overhaul:

  - Is compatible with the revised L1 latency time and rate and enables independent read-out of every EB crystal at 40 MHz for the L1 trigger;
  - Implements a lower APD operating temperature, 9°C compared to the current 18°C, which mitigates the effect of radiation induced dark current and the corresponding noise contribution to preserve the intrinsic Barrel ECAL energy resolution;
  - Leads to 30 ps timing precision for electromagnetic (EM) showers E > 30 GeV, such as photons from $H \to \gamma\gamma$ decays, to help mitigate the effects on in-time pile up;
  - Replaces HCAL Hybrid Photodiodes (HPDs) with higher efficiency and more radiation tolerant SiPMs — already done as part of the CMS Phase I upgrade program.

- Replacement of the Endcap calorimeters (Electromagnetic and Hadronic) with a radiation tolerant, integrated high granularity sampling calorimeter,[6] known as HGCAL. The salient features include:
  - Radiation tolerance, which drives the choice of silicon sensors in the regions of the calorimeter with the highest radiation exposure — front sections and highest pseudorapidity regions of the rear section — with SiPM on plastic scintillator tiles in the region less exposed to radiation;
  - High granularity, with the longitudinal segmentation being driven by energy resolution requirements, which the need to maintain the ability to track the effects of radiation damage and calibrate the calorimeter with MIPs translates to small read-out cell transverse sizes to reduce cell capacitance and electronic noise, resulting in a calorimeter with good imaging capability and the ability to resolve nearby showers in a dense environment, and in the presence of extreme pile up;
  - A novel front-end read out ASIC design which, in addition to satisfying the very large dynamic range requirements, also exploits the intrinsic timing characteristics of the silicon signal response to provide timing information with a precision of 20 ps or better for EM showers with $p_T > 2$ GeV, and 30–40 ps precision with an efficiency better than 90% for hadrons of $p_T > 5$ GeV.
- Major overhaul of the Barrel and End-Cap muon read-out systems, to comply with the increased L1 Trigger latency and read-out bandwidth and ensure adequate radiation tolerance for the HL-LHC operation.[7]
- Installation of an additional GEM based End-Cap Muon station (ME0) in the space liberated downstream of the HGCAL, which is denser and thus shorter than the existing Endcap calorimeter system:[7] together with the extended Tracker coverage, this extends the CMS muon acceptance from the present $\eta \approx 2.4$ up to $\eta \approx 2.8$, covering most of the area shadowed by the HGCAL (up to $\eta \approx 3$), and improves the muon trigger performance in the high η region.
- Replacement of the L1 Trigger,[8] and of the DAQ and HLT systems[9] to meet the more demanding requirements of the CMS Phase II operation.

A full overview and in-depth discussion of the CMS upgrade program is well beyond the scope of this article: instead in what follows we briefly summarize some of the salient features of the upgrades, with a focus on the novel techniques they introduce, and their enabling technologies.

2. Tracker upgrade, and L1 Tracking Trigger

The upgraded Outer Tacker coverage will remain limited to $\eta < 2.4$, while the Inner (Pixel) Tracker extends the tracking acceptance out to $\eta < 4$. In addition to substantially improving (VBF/VBS) forward jet reconstruction and MET resolution, the extended tracking allows lepton identification and reconstruction over the full acceptance of the new CMS HGCAL endcap calorimeter, which covers the region up to $\eta < 3$.

Material within the tracking volume adversely affects performance in several ways: multiple scattering may dominate charged particle momentum resolution up to very high p_T, a fraction (proportional to the radiation length of the material present) of photons will convert, and electrons will shower, nuclear interactions may confuse track reconstruction and create displaced vertices which will show up in the background when searching for long lived neutral particles, etc.

The Phase II CMS Silicon Tracker employs thousands of modules, which generate large data flows and dissipate in the order of 100 kW, which in turn requires substantial optical and electrical cable plants and necessitate distributed cooling, all of which must be supported by rigid mechanical support structures. Minimizing the material within the Tracking volume is thus a key challenge for such a detector. Following the experience from the Phase I CMS Tracker, the first all-silicon Tracker ever deployed at a collider experiment, the CMS Phase II Tracker achieves a very large relative decrease in the material both between the Inner (Pixel) Tracker (IT) and Outer Tracker (OT) and in the Outer Tracker volume itself, reducing the total amount of material in the tracking volume by about a factor of two over the range of $0.5 \sim \eta \sim 1.5$. Notably, this is achieved despite the extension of the tracking acceptance, and the inclusion of L1 Track Trigger capability, discussed below.

A key novel feature of the CMS Phase II Outer Tracker is that it fully integrates a L1 Track Trigger capability. This is achieved through the introduction of so called "p_T" modules. As sketched out in Fig. 1(Left), p_T modules integrate a pair of closely spaced silicon sensors, which can be used to determine not only the location but also the incident angle, in

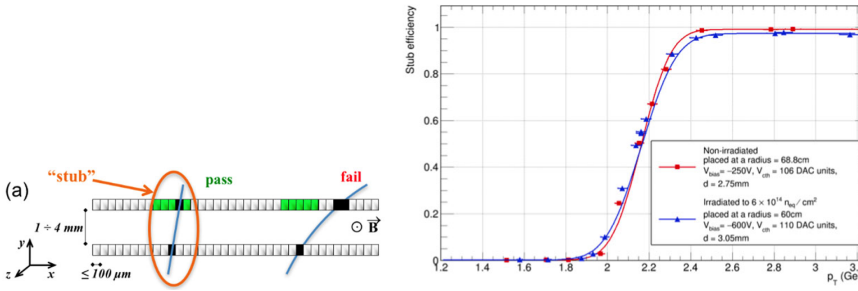

Left Right

Fig. 1. Left: Illustration of the p_T module concept. (a) Correlation of signals in closely-spaced sensors enables rejection of low-p_T particles; the channels shown in green represent the selection window to define an accepted "stub". (c) For the endcap discs, a larger spacing between the sensors is needed to achieve the same discriminating power as in the barrel. Center: Stub reconstruction efficiency for a non-irradiated (red) and an irradiated (blue) 2S mini-module. Right: The 2S module (left) and PS module (right) of the Outer Tracker. Shown are views of the assembled modules (top), details of the module parts (centre) and sketches of the front-end hybrid folded assembly and connectivity (bottom).

the r-φ plane, of charged particles traversing the module. This provides a measurement of the p_T of the incident particle, with sufficient resolution to discriminate efficiently between particles with p_T above or below a ~ 2 GeV threshold. This is shown in Fig. 1(Center), where the performance of modules built with irradiated and non-irradiated sensors are also compared.

Hits from the two sensors are brought together in a single readout ASIC, so that pairs of hits belonging to charged particles with $p_T > 2$ GeV, constituting a small fraction of the total number of hits, can be locally selected on-module, and transmitted to the off-detector L1 Track Trigger system, which in turns reconstructs the corresponding tracks and makes them available to the L1 Global Trigger.

The CMS Phase II Outer Tracker uses two different types of p_T modules. At smaller radii, below about 60 cm, so-called Pixel-Strip or PS p_T Modules are used, in which a strip sensor with 100 μm pitch and 2×24 mm strip length is paired with a second sensor with 100 μm pitch and 1.5 mm macro-pixel length. This granularity is well matched to the higher particle density in the inner region of the tracker, and the use of 1.5 mm long macro-pixels allows longitudinal primary vertex reconstruction with a precision adequate for the needs of the L1 Trigger. At radii above 60 cm, where the particle

density is lower and the distance from the beam too large to allow a useful improvement of the vertex reconstruction, so called Strip-Strip or 2S p_T modules are used, which use a pair of strip sensors with 90 μm pitch and 2×48 mm strip length. These are simpler in construction compared to the PS modules, dissipate less power, and allow both cost reduction and decreased material within the tracking volume.

3. MIP Timing Detectors, MTD

If left unaddressed, particles from 140 to 200 collisions within a single bunch crossing (in-time pile up) will degrade the ability to correctly reconstruct the collision of interest. In addition to having adequate detector granularity to retain accurate reconstruction of individual particle tracks and showers in the resulting crowded environment, an effective means of differentiating reconstructed particles originating from the collision of interest from those due to pile up collisions can greatly aid in maintaining the overall event reconstruction quality.

The LHC luminous region is several millimeters long, so that longitudinal track vertex reconstruction can help distinguish charged particle tracks (from the event vertex) from those originating from pile up vertices. At very high pile up, however, the vertex density is such that these overlap so much in space such that it is no longer possible to resolve them with track vertex reconstruction alone. Figure 2(Left) shows how the addition of precision timing information to charged particle tracks can help resolve such overlapping pile up vertices. It is becoming customary to refer to this use of timing information in vertex reconstruction as "4D reconstruction", as opposed to the usual 3D reconstruction.

The benefits of this are substantial for all aspects of event reconstruction. As an example, Fig. 2(Right) shows the improvement for the b-tagging performance of CMS.

Precision timing information for neutral particles can also play an important role in limiting the effect of pile up on event reconstruction. The upgrade of the Barrel Electromagnetic Calorimeter (EB) aims at providing ≈ 30 ps time resolution for high p_T photons, such as those from Higgs to $\gamma\gamma$ decays.

Two distinct technologies are adopted for the CMS MTD, as follows.

For the Barrel Timing Layer (BTL), which requires a relatively large surface to be instrumented (38 m^2), and where the maximum fluence remains below $\approx 2 \times 10^{14}$ (1 MeV neq)/cm^2, LYSO crystals doped with small

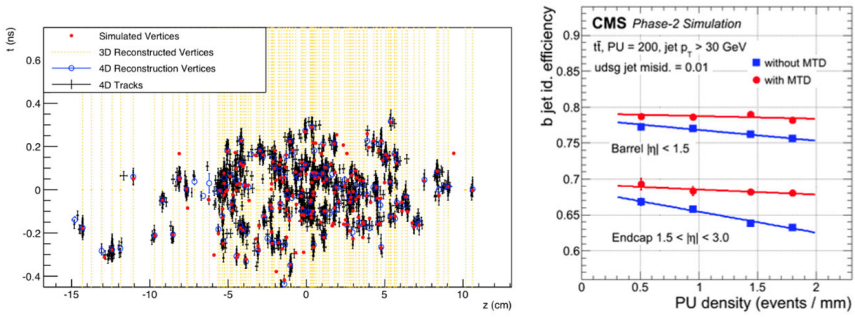

Left Right

Fig. 2. Left: Simulated and reconstructed vertices in a bunch crossing with 200 pile up collisions, assuming a MIP timing detector covering the barrel and endcaps. The vertical lines indicate 3D-reconstructed vertices, with instances of vertex merging visible throughout the event display. Center: Secondary vertex tagging ROC curves for light-quark and charm-quark jets for $|\eta| < 1.5$ (left) and for $1.5 < |\eta| < 3.0$ (centre), and b-tagging efficiency vs. average pile up density, with a constant light-jet efficiency of 0.01 (right). Results with and without precision timing information are compared to the zero pile up case. Right: Space-time diagrams illustrating the concept of hermetic timing for $H \to \gamma\gamma$ events. The reconstructed time for the photons at each vertex (green open dots), with error bars from the uncertainty on the time measurement of photons, can be cross referenced with the time information of the 4D vertices. The green straight lines are drawn to guide the eye. The pile up is reduced to an average of 20 in this case, to improve clarity. For photons with a small rapidity gap, shown here, photon timing alone is not sufficient and the coincidence with a 4D-vertex is necessary to enable accurate vertex location.

amounts of Cerium (LYSO: Ce) with SiPM read-out have been chosen. Early "proof of principle" prototype assemblies of such devices are shown in Fig. 3(Left), which also shows an arrangement used for test-beam studies of the time resolution. Results from such tests are shown in Fig. 3(Right), which demonstrate the target resolution of ≈ 30 ps for un-irradiated devices. It is interesting to note that the R&D for the CMS BTL benefitted from the already ongoing studies of similar devices, as part of an R&D aimed at improving the 3D-imaging capability of PET devices for medical purposes with precision timing (so called "TOF PET"). This is a notable example of a spin-off from High Energy Physics detector development coming around to contribute back into the field.

Maintaining sufficient timing resolution in the face of radiation damage is challenging and has been the focus of extensive optimization studies following the MTD TDR.[4] Among other things, these have led to the adoption of thermo-electrical coolers to reduce the effect of radiation induced

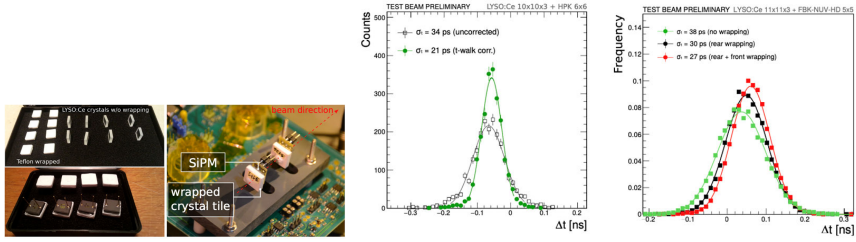

Left Right

Fig. 3. Left: Top left: Set of $11 \times 11 \times 3$ mm^3 LYSO:Ce crystals with depolished lateral faces, before and after Teflon wrapping. Bottom left: 6×6 mm^2 HPK SiPMs glued on LYSO crystals. Right: Crystal+SiPM sensors plugged on the NINO board used for test beam studies. Right: Distribution of the time difference in a pair of LYSO:Ce tiles exposed to a 3 mm wide beam of MIPs hitting the centre of the tiles. Left: Results before and after time walk correction for $10 \times 10 \times 3$ mm^3 crystals read out with 6×6 mm^2 HPK SiPMs. Right: Results for $11 \times 11 \times 3$ mm^3 crystals read out with 5×5 mm^2 FBK SiPMs under different wrapping configurations.

SiPM noise. By lowering their operating temperature several degrees below the $-35°$C minimum temperature that the CO_2 cooling system is able to provide, as well as careful tuning of the LYSO crystal geometry and SiPM characteristics and size.

Even so, the substantially higher fluence (up to $\approx 2 \times 10^{15}$ (1 MeV neq)/cm^2) precludes the use of this technology in the End-Cap region, and the End-Cap Timing Layer (ETL) is based instead on Low Gain Avalanche Diodes (LGADs). This is a relatively recent technology, an extension of the more mature planar silicon technology widely used for strip, pixel, and pad silicon detectors, and is still in a phase of rapid development. The basic novel feature, as illustrated in Fig. 4(Left), is the introduction of an additional implant immediately below the usual surface implant, which generates a sufficiently high local gradient to induce avalanche multiplication of charges, effectively amplifying the original signal.

Barring the constant term, the time resolution is inversely proportional to the signal over noise (S/N) ratio. Achieving time resolutions of order 30 ps for a single MIP signal in traditional silicon sensors would require prohibitively high levels of power dissipation in the front-end read-out electronics for it to be practical over the ≈ 14 m^2 of the CMS endcaps. The signal amplification provided by the LGADs allows the necessary S/N to be achieved with relatively modest electronics power dissipation and thus provides a viable solution. In the context of the ETL design, it has been found that gain of 10 allows a 30–40 ps timing resolution to be achieved.

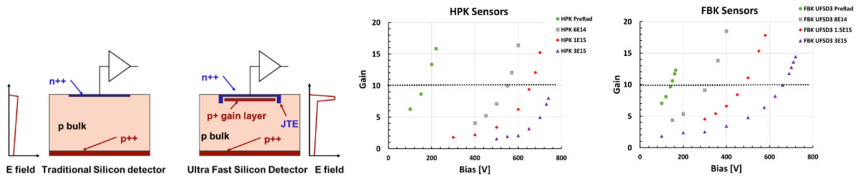

Left Right

Fig. 4. Left: Cross-sectional diagrams comparing a standard silicon detector and a UFSD, with an additional p implant providing the electric field for charge multiplication. Right: Gain as a function of bias voltage for different neutron fluences for LGADs manufactured by HPK, and FBK.

Radiation damage affects both the collection of primary charges (in a similar way as for standard planar silicon sensors) as well as the dopant concentration of the implants and the resulting gain; maintaining the necessary level of performance has been a key goal of the R&D program. Figure 4(Right) shows the gain as function of bias voltage for prototype LGADs from two different producers, after exposure to fluences ranging up to 3×10^{15} (1 MeV neq)/cm^2; these devices can reach a gain of 10, at a bias voltage below 800 V, and have been shown to achieve a timing resolution in the range of 30–50 ps up to the highest fluence.

4. Endcap High Granularity Calorimeter, HGCAL

The existing CMS calorimeter systems were designed for an integrated luminosity up to 500 fb^{-1}. As stated in the Introduction, it is possible to refurbish the Barrel ECAL and HCAL in such a way as to enable their continued exploitation for Phase II operation at the HL-LHC. In the endcap region, however, radiation levels are expected to reach up to 1×10^{16} (1 MeV neq)/cm^2 at 3'000 fb^{-1}, some fifty times higher than in the Barrel. The resulting performance degradation of the existing End Cap calorimeters much beyond ≈ 450 fb^{-1} would be such that these cannot continue to be used, and their replacement is an essential part of the CMS Phase II upgrade.

As the result of several years of dedicated R&D, and having examined several potential alternative designs and technologies, the CMS collaboration settled on HGCAL as the replacement for the existing endcap calorimeters.

The CMS HGCAL uses silicon sensors for the electromagnetic section of the calorimeter, as well as for parts of the hadronic section that are exposed

to the highest radiation levels. Plastic scintillator tiles with direct (on-tile) SiPM readout are used for those sections of the hadronic calorimeter that will be exposed to less than $\approx 5 \times 10^{13}$ (1 MeV neq)/cm^2 after 3'000 fb^{-1}. The CMS HGCAL employs almost 600 m^2 of silicon sensors with ≈ 6 M readout channels, and close to 370 m^2 of plastic scintillator with about 240'000 readout channels.

The choice of HGCAL was motivated by the demonstration that simple planar silicon sensors (single-sided, DC-coupled, n-on-p) can tolerate the required levels of radiation, while retaining adequate signal charge collection efficiency even after exposure to fluences of 1.5×10^{16} (1 MeV neq)/cm^2. On the other hand, the commitment of key industrial partners towards silicon sensor production on 8" lines made possible cost effective sensor production on the very large scale required for the CMS HGCAL, without undue interference with the silicon sensor production for the Phase II ATLAS and CMS Trackers on well-established 6" production lines.

The silicon sensors are of hexagonal shape, the largest tile-able polygon, which allows the most efficient use of the sensor wafer. In combination with the use of 8" wafers this minimizes the number of modules to be assembled and integrated into the system, reducing it by well over a factor two compared to the more typical square sensors produced on 6" wafer sensors.

Even so, with approximately 26'000 silicon modules, there is strong emphasis on a simple, mechanically robust module design well adapted to automated robotic assembly and ease of handling. The HGCAL silicon modules, shown schematically in Fig. 5(Left (a)), include a base plate onto which the sensor is glued. A front-end read out PCB "Hexaboard" is then glued on top of the silicon, which covers the full area of the sensor. The Hexaboard is connected to the sensor by wire bonds via through holes in the PCB. The base plate is made of Cu/W for the Electromagnetic section of the calorimeter, where it forms part of the absorber, and a carbon fiber plate for the hadronic part. A Kapton$^{\text{TM}}$ foil is laminated onto the baseplate to provide both bias (high) voltage DC protection as well as AC de-coupling of the silicon sensor backside.

For the scintillator part of the HGCAL, building on a design developed by the CALICE Collaboration,[10] modules consist of large area PCB boards with surface mounted SiPMs onto which wrapped scintillator tiles, with a central dimple to house the SiPM, are glued (SiPM-on-tile) (Fig. 5(Left (b))).

The requirement to calibrate the detector through the MIP Landau

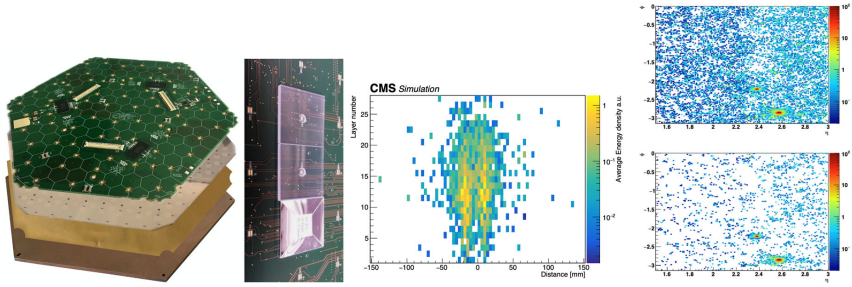

Left (a) Left (b) Center Right

Fig. 5. Left: (a) An HGCAL Silicon Module, from Electromagnetic section of the calorimeter, showing the stacked layers, next to (b) an example of three CALICE 3 × 3 cm^2 scintillator tiles mounted on a PCB that holds one SiPM per tile. The left two scintillators are unwrapped to show the SiPM within the small dome at the centre of the tile, while right-most tile is wrapped with reflective foil. Center: Energy deposited in HGCAL cells by pairs of unconverted photons. The photons have an energy of E = 80 GeV (p_T = 14.4 GeV) at $\eta = 2.4$ in the HGCAL, and are separated by $\Delta R = 0.05$ (in a random orientation), corresponding to a separation distance of about 30 mm. Reconstructed hits are projected onto the plane defined by the axes of the two showers. The colour code represents energy density. Right: Hits with a charge > 12 fC, from a VBF Higgs $\rightarrow \gamma\gamma$ event, projected to the front face of the calorimeter: (upper plot) without a timing requirement, and (lower plot) after removal of hits with $|\Delta t| > 90$ ps.

peak drives several of the basic design HGCAL parameters. For the silicon sensors, it limits the input capacitance, thereby setting the cell size, which is ~ 1.2 cm^2 over most of the detector, and 0.5 cm^2 in the regions where the highest radiation level is expected and thinner (120 μm) sensors are used, and sets demanding requirements for the front-end readout chip of a dynamic range from 0.2 fC to 10 pC with a noise of less than 2'500 e^- for a 60 pF input capacitance. In turn, the corresponding S/N level enables a timing resolution of better than 100 (20 ps) for cells with 3 (10) MIPs in 300 μm silicon. Similarly, for the SiPM-on-tile the MIP calibration requirement drives the choice of combination of tile sizes (which range from 4 cm^2 at the inner radii to 30 cm^2 at outer radii), plastic scintillator material (both cast and molded are used), and SiPM sizes (which range from 9 mm^2 at inner radii to 4 mm^2 towards the outer radii).

Figure 5(Center) illustrates the shower imaging capability and resolving power in the HGCAL Electromagnetic section which results from such a granularity, and an example of the impact of precision timing on pile up removal and mitigation is shown in Fig. 5(Right).

5. Summary and Outlook

Following a long period of reflection and focused R&D effort, the CMS Collaboration has engaged in a coherent program of refurbishing and overhauling where possible, and replacing and upgrading where necessary, the infrastructure and detector systems of the CMS experiment to make effective use of the physics potential afforded by HL-LHC. In response to the demands of precision physics at very high luminosity the CMS upgrade will deploy, in addition to many incremental improvements, several novel experimental features which include a full L1 Tracking Trigger, High Granularity Endcap Calorimeters, as well as charged and neutral particle precision timing for pile up mitigation. These will be complemented by the development of novel analysis techniques and the deployment of advanced computing architectures such as GPU's, to make best use of the upgraded detector while maintaining a cost-efficient computing model.

This is a challenging program, for which the R&D phase is by now mostly reaching completion and which is entering into the final qualification phase to prepare for the start of construction. Its success continues to depend crucially on the substantial commitment by the Collaboration with the strong backing of the Funding Agencies involved and of the CERN Laboratory. It will set the ground for at least another decade of scientific exploitation of the LHC complex, building on the discovery of the Higgs Boson in 2012 and the precision measurements made since then, in view of further exploring physics beyond the Standard Model.

References

1. The Compact Muon Solenoid Phase II Upgrade Technical Proposal / CERN-LHCC-2015-010 / CMS-TDR-15-02 / 01-06-2015.
2. The Phase-2 Upgrade of the CMS Tracker Technical Design Report / CERN-LHCC-2017-009 / CMS-TDR-014 / 01-07-2017.
3. Technical Proposal for a MIP Timing Detector in the CMS Experiment Phase 2 Upgrade / LHCC-P-009 / 27-11-2017.
4. A MIP Timing Detector for the CMS Phase-2 Upgrade Technical Design Report / CERN-LHCC-2019-003 / CMS-TDR-020 / 29-03-2019.
5. The Phase-2 Upgrade of the CMS Barrel Calorimeters Technical Design Report / CERN-LHCC-2017-011 / CMS-TDR-015 / 12-09-2017.
6. The Phase-2 Upgrade of the CMS Endcap Calorimeter Technical Design Report / CERN-LHCC-2017-023 / CMS-TDR-019 / 09-05-2018.
7. The Phase-2 Upgrade of the CMS Muon Detectors Technical Design Report / CERN-LHCC-2017-012 / CMS-TDR-016 / 12-09-2017.

8. The Phase-2 Upgrade of the CMS Level-1 Trigger Technical Design Report / CERN-LHCC-2020-004 / CMS-TDR-021 / 10-03-2020.
9. The Phase-2 Upgrade of the CMS Data Acquisition and High-Level Trigger Technical Design Report / CERN-LHCC-2021-007 / CMS-TDR-022 / 17-06-2021.
10. Design, Construction and Commissioning of a Technological Prototype of a Highly Granular SiPM-on-tile Scintillator-Steel Hadronic Calorimeter / e-Print: 2209.15327 [physics.ins-det] / 30-09-2022 / CALICE-PUB-2022-003 Note: to be submitted to JINST.