

CMS SLHC Tracker Upgrade: Selected Thoughts, Challenges and Strategies

Marcello Mannelli

CERN

Geneva, Switzerland

E-mail: marcello.mannelli@cern.ch

Upgrading the CMS Tracker for the SLHC presents many challenges, of which the much harsher radiation environment and higher occupancy are two central aspects, as well as the scope for qualitatively new performance goals.

These issues have been the subjects of broad ranging discussions, aimed at first exploring alternative ideas and approaches, before focusing on the most promising ones.

In what follows, rather than attempting to summarize these discussions, we will first review the Desired Functionality, and then highlight two key considerations, namely Granularity and Power Consumption, and Local Occupancy Reduction / Hit Discrimination, and illustrate how they may drive the design for an upgraded Tracker.

*The 16th International Workshop on Vertex Detectors
Lake Placid, NY, USA
September 23rd – 28th 2007*

1. Desired Functionality

The present CMS Silicon Tracker (Pixel and Strips) provides fast, robust, efficient and high resolution Tracking and Impact Parameter determination.

An upgraded Tracker for the SLHC should at least maintain this performance, in the much more difficult higher luminosity environment. Improved impact parameter determination, if reasonably achievable, would clearly be desirable, as this would directly translate into useful improvements to b-jet tagging.

A dominant performance-limiting factor for the present CMS Tracker is the higher than initially hoped for amount of material within the tracking volume. The Tracker material budget ranges from $\sim 0.4X0$ to $0.8X0$ in the central region spanning $|\eta| < 0.8$, rises abruptly to close $1.8X0$ at $\eta \sim 1.5$, before falling again to about $0.8X0$ at $\eta \sim 2.5$ which is the limit of the Tracker acceptance, and down to 0 by $\eta > 3.5$. In addition to substantially degrading the Tracking performance, the material inside the Tracker also affects the performance of the Electro-magnetic calorimeter.

An important goal for the upgraded Tracker is a substantial reduction in the amount of material within the tracking volume.

The very large power dissipated by the present CMS Tracker is an important driving factor for the material inside the Tracker. With this in mind, we are presently pursuing Power Distribution schemes aimed at reducing the amount of Copper required, as well as more efficient cooling schemes, in a concerted effort to minimize the amount of material in the Tracking volume, for a given power dissipation. At the same time, maintaining the present level of power dissipation or, preferably, even reducing it will certainly help to achieve a substantial reduction the amount of material.

Maintaining excellent tracking performance in the face of the much higher track multiplicities at the SLHC will require a substantial increase in granularity, and the interplay between Granularity and Power Consumption will therefore be a very important consideration. Granularity reflects not only on cell occupancy but also, in the case of Pixels, provides a 3D reconstruction capability with great performance potential. Indeed, despite being subject to the highest occupancy in the Tracker, the Inner Pixel layers are used to provide high quality seeds for tracking, thanks to their 3D reconstruction capability.

In the CMS Trigger and DAQ architecture, a Level 1 Trigger is used to reduce the rate from 40MHz to a maximum of 100KHz. Events accepted by the Level 1 Trigger are fully read-

out, and distributed to a High Level Trigger, consisting of a farm of PC's, which in turn discard all but 100~200Hz of events, that are written to mass storage for further analysis.

Only in 2002, when preparing the CMS DAQ TDR, was it realized that the speed with which the pattern recognition could be solved with the CMS Tracker, with Pixel seeding, provides the potential for extensive use of Tracking at the very early stages of the High Level Trigger, to identify Physics Objects and Event Topologies with much better resolution, efficiency and background rejection than would otherwise be possible.

Any increase in the CMS Level 1 Trigger rate would require extensive replacement of the Front-End electronics in many if not all of the CMS sub-detectors, and is not considered to be reasonably feasible. At the SLHC, therefore, we must maintain the same maximum 100KHz Level 1 Trigger rate, despite a tenfold increase in luminosity and some ten to twenty times higher occupancy (depending on the rate of collisions), without sacrificing efficiency and background rejection.

Just raising thresholds on such things as high Pt leptons and/or jets is not a viable solution, since this would lead to unacceptable losses of efficiency, for example to W and Z bosons which will remain important for control samples and signatures for much of the likely physics of interest.

In addition, in the case of high Pt muons for example, the measurements provided by the muon chambers for the Level 1 Trigger are such that raising the thresholds will simply not work. Indeed, already at the LHC the high Pt muon Level 1 Trigger rate is completely dominated by miss-measured relatively low Pt muons, and the much improved resolution achievable with Tracking is used to substantially eliminate these in the High Level Trigger. This is illustrated in **Figure [1]**.

The ability to use tracking already in the Level 1 Trigger, much as it is presently used in the High Level Trigger, will be an important part of the CMS Trigger strategy for the SLHC, and is a fundamental new functionality we require of the upgraded Tracker.

The present CMS Tracker produces vast amounts of information for each Level 1 Trigger, at up to 100KHz. For the Strip Tracker, the analogue information for each of the approximately 9 Million strips is read out for every triggered event. For the Pixel Detector, zero suppression is used to read out only the information for clusters above threshold. In the High Level Trigger, only clusters within given regions of interest are considered, at least in the early stages of the selection.




Figure [1]. The rate of accepted muons is shown as function of Pt threshold determined with the Muon chambers, for the L1 Trigger and in the HLT (L2). Also shown is the rate of accepted muons as function of Pt threshold determined in the HLT, but now with the Tracker (L3). It can be seen that the L3 accepted muon rate falls exactly on the generated muon rate, as function of Pt. Once an Isolation cut is also imposed, the accepted muon rate is further reduced and, for thresholds in the range from 20GeV to 40GeV, is typically two orders of magnitude or more below the L1 Trigger accepted rate, using the Muon chambers.

It is implausible to access all this information for the purposes of a Level 1 Trigger at the 20 ~ 40MHz bunch crossing frequency of the SLHC: zero suppression will be required throughout the upgraded Tracker. Even given the plausible advances in data transmission over optical links, however, this will not be sufficient, as it will not be feasible to transmit even the zero suppressed cluster data for the Level 1 Trigger, without unacceptable penalties in terms of both cost and power dissipation. Additional data reduction of at least one to two orders of magnitude will be required, prior to optical transmission over long distances.

This leads us to consider a two-stage read-out scheme, based on Local Occupancy Reduction / Hit Discrimination, by which we mean the ability to discriminate against clusters from uninteresting tracks without having to transmit these over long distances with optical links. In this scheme, only clusters associated to tracks relevant for a L1 trigger decision would be

read-out for every bunch crossing, whereas the full information would be read-out for L1 accepted events.

2. Granularity and Power Consumption

When equipped with a given front-end electronics, an increase in number of channels obviously leads to an increase in power consumption. In the case of Strip Sensors, however, there is an important difference in the impact of changing channel density by varying the strip pitch or the strip length.

On the one hand, channel capacitance plays a fundamental role in defining the power dissipation of the front-end Amplifier: for constant noise performance, the front-end amplifier current varies quadratically with input capacitance, and for constant rise-time response linearly. On the other hand, all else being equal, the strip capacitance is independent of pitch, and linearly proportional to strip length.

Thus, changing the channel count by varying the pitch, which maintains a constant channel capacitance and therefore a constant power per channel requirement, will result in a corresponding variation of the total power dissipation.

However, the effect of changing the channel count by varying the strip length is quite different. Shorter strips have a correspondingly lower channel capacitance, and thus require lower power dissipation per channel. In a rise-time driven regime, the power per channel varies linearly with strip length, so that the total power dissipation is largely independent of strip length and channel count. In a noise driven regime, the power per channel varies quadratically with strip length, so that the total power dissipation would in principle decrease linearly as the strip length is decreased, and the channel count increased accordingly.

In practice, of course, channel count affects power in other ways, and in particular affects the digital power requirements of the Front End readout chip. Nevertheless, the simple-minded analysis above is useful in understanding the more realistic extrapolations discussed below, which are based on the existing Front End readout electronics, for the CMS Strip Tracker and Pixel Vertex detectors.

The present CMS Strip Tracker has a total of some 9M channels, instrumented with about 73'000 APV25 chips, each reading out 128 channels. The power dissipation per channel of the APV25 chip is about 3mW, which results in a nominal Front End chip power dissipation of about 26kW, compared to a total Power dissipation inside the Tracking volume of about 33kW, including Optical Links, controls etc. Averaged over the more than 200m² of the Strip Tracker, this corresponds to about 120W/m² power dissipated by the Front-End read out chips.

The power dissipation for the CMS Pixel Front-End read out chip, on the other hand, is approximately $30\mu\text{W}$ per channel. With a Pixel area of $100\mu\text{m} \times 150\mu\text{m}$, this corresponds to a power dissipation of about $2\text{kW}/\text{m}^2$. Whereas this is about 16 times higher than for the Strip Tracker, it corresponds to a channel density almost 1'500 times higher.

The analogue power dissipated by the Pixel read out chip is driven by the rise time requirement (for the APV it is driven by the noise requirements), so that it varies only linearly with input capacitance. In addition, the present input capacitance for the Pixel sensor has a substantial contribution from the bump bonding, which is independent of the Pixel size, and the Pixel sensor capacitance itself was not explicitly minimized.

These, and other considerations, make it plausible that the Pixel power consumption per channel could be held constant, even as the Pixel length is increased by some 10 to 20 times. Under this assumption, a Pixel Tracker with, for example, $120\mu\text{m} \times 2\text{mm}$ “Long-Pixels” would have an approximately 100 times higher granularity, but similar power consumption as the present Strip Tracker.

As a matter of fact, there are a number of improvements which can be considered that would even reduce the power dissipation of such a device, including using a slower rise time and compensating as needed by digital signal processing to recover single bunch crossing time discrimination.

If deployed over a similar $\sim 200\text{m}^2$ surface as the present Strip Tracker, such a Tracker would have about a billion channels: this would be a “Giga Channel Tracker”, with full 3D tracking capability.

Of course, such extrapolations need to be backed up with far closer analysis. In addition, power dissipation is only one of a number of major challenges that a Giga Channel Tracker would present, ranging from the need for high-density low cost inter-connect and bump bonding technologies suitable for producing some 200m^2 of long-Pixel sensor modules, to signal processing, calibration and control of a system with an unprecedented count of a billion channels.

On the other hand, whereas detailed studies to quantify this are ongoing, the performance potential of such a Tracker can be expected to be spectacular, and the arguments outlined above show that, at this stage of our thinking for the CMS SLHC Tracker upgrade, it would be premature to exclude this possibility from further consideration.

Indeed, our plan is to pursue the questions and challenges associated with a very large-scale fine granularity “Long-Pixel” Tracker, including realistic assessments of Power Consumption, Cost and Manufacturability, Read-Out Architecture and Systems Implications, while maintaining a more conventional “Short-Strip” option as a backup.

3. Local Occupancy Reduction / Hit Discrimination

Due to the curvature induced by the 4 Tesla field of the CMS solenoid magnet, the tangent to a low Pt track misses the event vertex in the transverse plane, whereas a high Pt track points closely to it. Therefore, a “local” measurement of the track direction vector in the transverse plane can serve to determine the Pt of that track.

The use of “local” track vectors for Pt measurements is well known. The CMS muon Trigger is one example of this. The idea of applying such a scheme to local data reduction for a CMS Tracker Level 1 Trigger, as illustrated in **Figure [2]**, was first presented at a CMS Tracker SLHC Upgrade Workshop in 2005 by J. Jones of the Imperial College group.

J. Jones (~2005)
CMS Tracker SLHC Upgrade Workshops

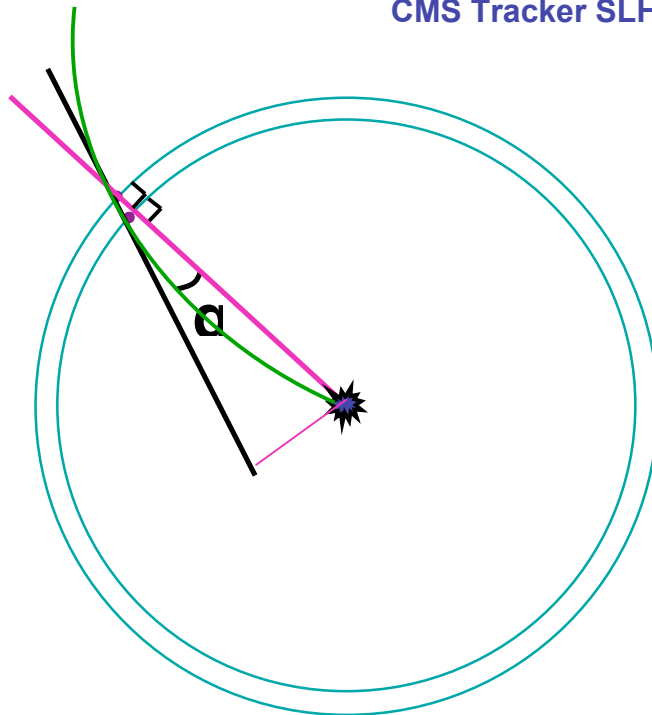


Figure [2]. In this r-phi view illustration it is shown that, whereas the tangent of a high Pt track extrapolates back to the event vertex, the tangent to a low Pt track, curved by the magnetic field, misses the vertex.

In what follows, we develop this idea to discuss a possible scheme for hierarchical Local Pt Discrimination, and Occupancy Reduction.

The Pt spectrum of charged tracks in a representative sample of CMS Level 1 Trigger accepted events falls very steeply with Pt. Tracks below 1GeV Pt are usually not very interesting in the early phases of event selection. Reading out hits only from tracks above a 1GeV Pt threshold would reduce the data rates by an order of magnitude. For Level 1 Trigger purposes, we are typically also interested in identifying tracks above 10 ~ 20GeV Pt. As can be inferred from **Figure [3]**, reading out hits only from tracks above a 10 ~ 20GeV Pt threshold would (nominally) reduce the rate by an additional four to five orders of magnitude.

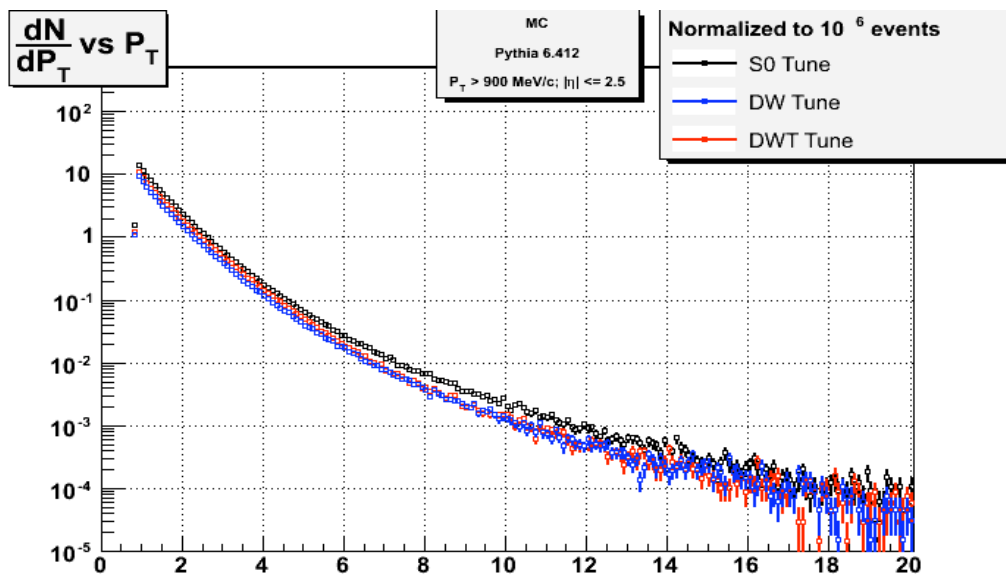


Figure [3]. The Pt spectrum of charged tracks, within the Tracker acceptance of $\eta < 2.5$, for minimum bias events. A threshold of 900MeV has been applied.

The spatial resolution of Silicon sensors with $\sim 100\mu\text{m}$ pitch in the transverse plane is such that, at least provided the multiple scattering through the material inside the Tracking volume is sufficiently contained, good discrimination against tracks below 1GeV pt can be obtained even with pairs of sensors separated by one to a few millimeters or so in radius. Track vectors measured with pairs of sensors separated one to a few centimeters or so in radius can provide effective identification of tracks above 10~20GeV.

Note that, pairs of sensors separated by a centimeter or more in radius would not be effective in discriminating against low Pt tracks, below a GeV or so, because the window of acceptance for a hit in one layer onto the other layer would be so large that the discrimination

would suffer from a high rate of random coincidences, due to the high occupancy environment of the SLHC.

Pairs of sensors at $\sim 1\text{mm}$ radial separation could be conveniently arranged in a “double-sensor” module, and would provide module level discrimination against hits from very low Pt tracks; pairs of double-sensor modules with a $\sim 2\text{cm}$ radial separation could be arranged in super-layers, and would in addition provide local discrimination against tracks below $10 \sim 20\text{GeV Pt}$.

Figure [4] shows an r-phi cut of an illustrative layout, based on this scheme. In this example, 12 Tracking layers are arranged in 3 pairs of closely spaced hermetic rings of double-sensor module-pairs, or super-layers. At the center of these Tracking layers will be several Vertex layers.

Each super-layer has the ability to filter hits from very low Pt tracks, and to identify sets of hits compatible with high Pt tracks. Having 3 super-layers provides essential redundancy, so that both high efficiency and rejection power can be ensured, while allowing for inevitable detector inefficiencies.

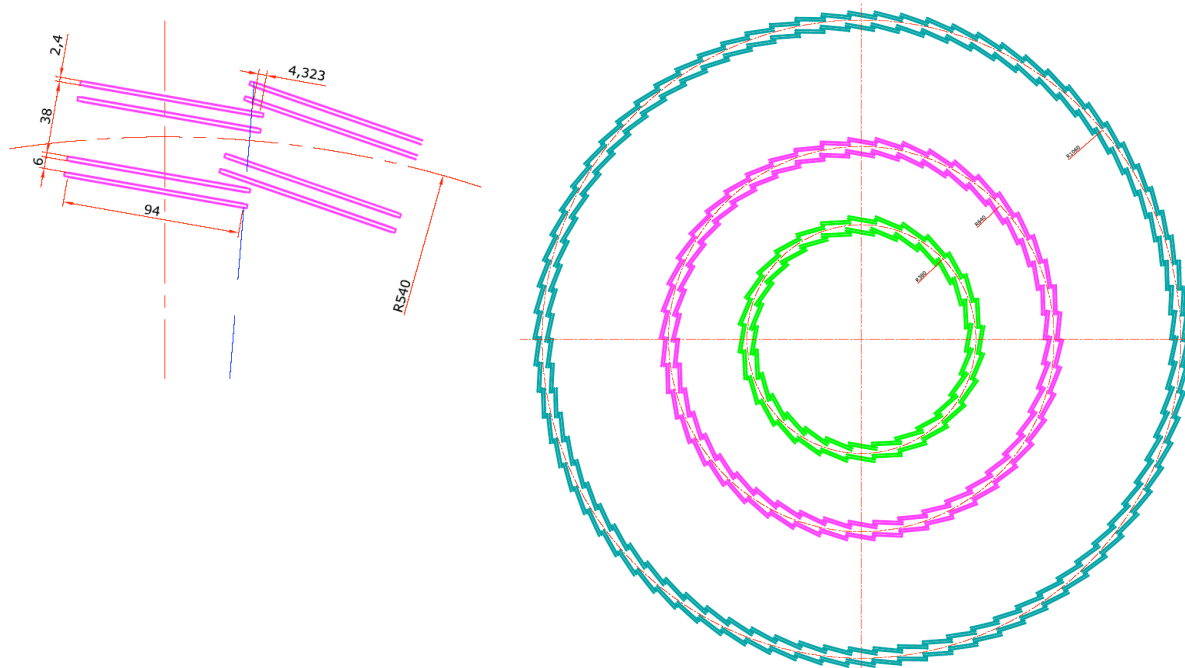


Figure [4]. An r-phi view illustration of a Tracker Layout with 3 super-layers, each consisting of 2 rings of double-sensor modules.

POS(Vertex 2007)038

Earlier in the discussion, we pointed out that reducing the material budget is one of the goals for the upgraded Tracker. We have arrived to the layout in **Figure [4]** through a series of considerations driven by the need for local data suppression and generation of L1 Trigger primitives. It is interesting to note that, we could have arrived at a similar layout independently of these considerations, driven by the goal of reducing the material in the Tracker.

Indeed, if we take the example of the present Strip Tracker Outer Barrel, we learn that layers equipped with back-to-back modules provide two hits for only about one and a half times more material than layers with single modules, which provide only one hit. This is the result of being able to use a single mechanical support structure for a pair of hits, rather than two independent sets of supports, and of the ability to reduce other overheads in terms of material, with a shared set of services for the pair of hits rather than two independent ones.

All else being equal, grouping the 12 hits of the upgraded Tracker in six layers of double-sensor modules can therefore result in a large reduction of material compared to a layout with 12 layers of single-sensor modules. Similarly, one should be able to exploit the grouping of the six layers of double-sensor modules into three super-layers, in order to further reduce the material compared to a layout with six independent layers of double-sensor modules. This lends additional motivation to study more closely this somewhat unconventional Tracker layout.

The concentric barrel super-layer geometry discussed so far lends itself naturally to the central region of the Tracker. Several possibilities are under consideration for extending this scheme to the forward part of the Tracker acceptance. One possibility is extending the deployment of barrels also to the forward region.

There are a number of parameters which characterize the proposed layout, and which will affect the performance. These include sensor thickness and cell geometry, the radial spacing of pairs of sensors within a module, the radial spacing of modules within a super-layer and the radial positions of the super-layers, as well as the relative alignment precision of all these elements. These parameters also have obvious implications on crucial aspects of the Tracker design.

Studies are currently underway to understand how the performance of such a Tracker depends on the detailed choice of parameters, which will provide essential input to an eventual optimization of these. As part of these studies, we plan to include comparisons of the “Long-Pixel” layouts with a fallback “Short-Strip” option, as well as with a more conventional layout in which the six layers of double-sensor modules are equally spaced, as a reference benchmark.

4. Summary and Conclusions

The present CMS Tracker builds on the experience gained with previous generations of Strip Vertex detectors, and deploys this technology on an unprecedented scale to realize a full Silicon Strip Tracker. Three Inner Pixels are used to provide high-resolution 3D vertex reconstruction. The 3D reconstruction capability of the inner Pixels layers is also used to provide seeding for track reconstruction.

In thinking about an upgraded CMS Tracker for the SLHC, we are currently studying the possibility of very large-scale deployment of “Long-Pixels”, based on the experience gained with the present Pixel Vertex detectors. This would result in a Giga Tracker, with about a billion channels, and full 3D reconstruction capability. The incremental approach, more closely based on the present Strip Tracker, of deploying “Short-Strips” remains a natural fallback.

The need to use Tracking already at the Level 1 Trigger, and therefore the need for local data suppression, lead us to consider a somewhat unusual layout, based on a small number of super-layers (three), each consisting of two closely spaced layers instrumented with double-sensor modules.

Such a layout also appears to lend itself well to efficient use of mechanical support and distribution of services, with which to pursue the further goal of substantially reducing the material inside the Tracking volume compared to its present level.

Detailed studies are underway, to validate the performance of this layout and optimize its design.

This is an ambitious set of ideas, but one with an attractive performance potential. It is illustrative of the sorts of issues which we want to develop a solid understanding of, before committing ourselves to a definitive baseline layout, in an approach based on making best use of the experience with the present CMS Tracker, and driven by the functionality and performance goals we wish to achieve with the next generation CMS Tracker, for the SLHC.