

CMS inner tracking

M. Mannelli

CERN/EP Division, CH-1211, Geneva 23, Switzerland, e-mail: marcello.mannelli@cern.ch

Received: 19 August 2003 / Accepted: 19 August 2003 /
Published Online: 13 July 2004 – © Springer-Verlag / Società Italiana di Fisica 2004

Abstract. The CMS Tracker consists of three layers of silicon pixel detectors, surrounded by some 220m² of silicon strip sensors. It is designed to operate reliably in a harsh radiation environment for at least a decade, and to provide accurate and robust tracking in a high occupancy situation. We briefly discuss the layout of the CMS Tracker, and the basic design considerations from which it was derived. We then review the status of the project, with a focus on module assembly. Finally, we highlight a few aspects of the expected tracking performance.

1 The CMS Tracker design

Charged particle tracking at the LHC is a complex task, and provides a versatile and powerful tool for event reconstruction. In the design of the CMS tracker, we have focused on a few simple considerations and key performance benchmarks, which nevertheless are sufficient to ensure adequate performance for the broad range of envisageable applications of tracking in CMS.

A first basic consideration is that the CMS tracker must provide efficient and robust pattern recognition even in a crowded environment. In Fig. 1, for example, the production and decay of a Higgs into four muons is shown, with the minimum bias event pile-up expected at high luminosity.

Also shown is the same event, when only tracks with P_t above 2 GeV are considered. It can be seen that minimum-bias event pile-up greatly complicates the reconstruction of the relevant event topology. Very high P_t jets also are characterized by regions of very high local track density, which must be examined in detail in order, for example, to verify the presence of secondary vertices from b or tau decays. It is evident that fine granularity is required in order to resolve hits from nearby tracks, as well as a fast response time in order to resolve consecutive bunch crossings, and limit the extent of pile up at high luminosity.

A second consideration is that a resolution of 1 to 2% for 100GeV tracks is required, in order to reconstruct narrow heavy objects. With such a P_t resolution, each of the two Z bosons from the Higgs decay in Fig. 1 can be reconstructed with a resolution significantly better than the intrinsic width of the Z.

Finally, a good track impact parameter resolution is required in order to reliably tag b jets.

The requirements of high granularity, fast response, and good position resolution are all well matched to the use of silicon detectors. In fact, the robustness and preci-

sion of each track hit provided by silicon sensors, allows the requirements discussed above to be met with a relatively small number of measurement layers.

The basic layout of the CMS Tracker [1] is shown in Fig. 2a and b. It consists of 3 pixel and 10 micro-strips Barrel layers in the central region, and two pixel and nine micro-strip disc layers in each of the two End-Caps. The silicon micro-strip Barrel is separated into an Inner and an Outer Barrel. In order to avoid excessively shallow track crossing angles, the Inner Barrel is shorter than the Outer Barrel, and there are an additional three Inner Disks in the transition region between Barrel and End-Cap's, on each side of the Inner Barrel.

The read-out strips are oriented either parallel to the nominal beam line, in the Barrels, or in pointing toward it, in the End-Cap discs.

As illustrated in Fig. 2, there are several measurement layers, in both Barrels and End-Caps, which include small angle stereo-modules, mounted back-to-back onto modules with their strips oriented along the primary measurement axis defined above. In total, there are approximately 67 million pixel channels, and almost 10 million micro-strip channels, distributed over a sensitive surface of about 220m² of silicon sensors.

Figure 2 shows the expected average track occupancy, as function of radius, for high luminosity operation of the LHC. Given that efficient and clean track reconstruction is ensured provided the strip occupancy is below a few %, this implies that the typical strip surface in the inner most layers must be kept well below 1cm². This condition is relaxed at larger radii. Nb. This plot is made assuming a 25ns shaping time; it is evident that, at high luminosity, this is an essential requirement in order to achieve sufficiently low strip occupancy.

Equation 1 shows how the P_t resolution is determined by the basic Tracker design parameters, assuming 12 measurement layers, each with a resolution equal to the strip




Fig. 1. The production and decay of a Higgs boson to four muons, shown with minimum-bias event pile-up expected at high luminosity. Shown below is the same event, taking into consideration only tracks with a P_t above 2GeV

pitch divided by the square-root of twelve, distributed at regular intervals over a radius of 1.1m.

It can be seen that, with a solenoidal field of 4Tesla, a typical pitch of order $100\mu\text{m}$ is required, in order to obtain a P_t resolution of about 1% for 100GeV tracks. This can be compared to the expected Sagitta of about 2mm for a 100GeV track.

$$\frac{\Delta p}{p} \approx 0.12 \left(\frac{\text{pitch}}{100\mu\text{m}} \right)^1 \left(\frac{1.1\text{m}}{L} \right)^2 \left(\frac{4T}{B} \right)^1 \left(\frac{p}{1\text{TeV}} \right) \quad (1)$$

Equation 1 P_t resolution for an idealized tracker, with 12 equally spaced measurement points, distributed over a 1.1 meter radius within a solenoidal magnetic field, and assuming a hit resolution equal to the strip pitch over the square root of twelve.

In light of the above, the strip pitch and length in the CMS Silicon Tracker range respectively from about $80\mu\text{m}$ and 10cm in the inner most layers, to about $200\mu\text{m}$ and 20cm in the outer most layers. Note that, in order to allow efficient common mode noise subtraction, the CMS Tracker has an analogue read-out. Given the typical level of charge sharing among neighbouring strips, this allows a hit position resolution with $200\mu\text{m}$ strip pitch approaching that expected for a $100\mu\text{m}$ pitch, digital device.

The increase in noise for the longer strip modules is (more than) offset by the use of thicker silicon sensors ($500\mu\text{m}$) compared to the inner layers ($300\mu\text{m}$), and the resulting increase in charge collected. As a result, and given a typical strip capacitance of about 1.2pF/cm and the excellent noise performance of the APV25 read-out chip, a signal to noise of better than 13 and 15 is expected throughout the full Tracker, even after 10 years of LHC operation, for the short thin sensor modules and the long thick sensor modules respectively [1].

The pixel size is essentially driven by the surface necessary to house the required number of transistors for each cell of the read-out chip. Recently, the decision has been taken to move to $0.25\mu\text{m}$ technology also for the Pixel read-chips. This has allowed a reduction in the pixel size, to $100\mu\text{m}$ by $150\mu\text{m}$, from the previous $150\mu\text{m}$ by $150\mu\text{m}$. With this pixel size, and making deliberate use of the Lorentz angle in the pixel Barrel, and by inclining the pixel sensors in the End-Caps, in order to produce charge sharing, hit resolutions of about $10\mu\text{m}$ and $20\mu\text{m}$ in the transverse and longitudinal (radial for End-Caps) dimensions can be achieved.

Of crucial importance, as will be seen later, is the fact that such small cell sizes lead to an average channel occupancy of order 10^{-4} even at high luminosity. This implies that, as seen from the eye of the pixels, even the most complex events at the LHC are in fact remarkably clean.

2 CMS Silicon Tracker module components and assembly

Figure 2 illustrates the main components of the CMS Silicon Tracker modules. Beyond meeting the performance and radiation hardness requirements for each module component, the very large numbers of silicon sensors, front-end hybrids, read-out chips, pitch adapters, and carbon fibre frames, required a far more extensive and systematic component industrialisation than in previous projects based on similar technologies.

Much was learnt about the characteristics of silicon micro-strip sensors during the extensive R&D program, which finally led to the definition of a suitable set of specifications for use in the CMS Tracker [1, 2].

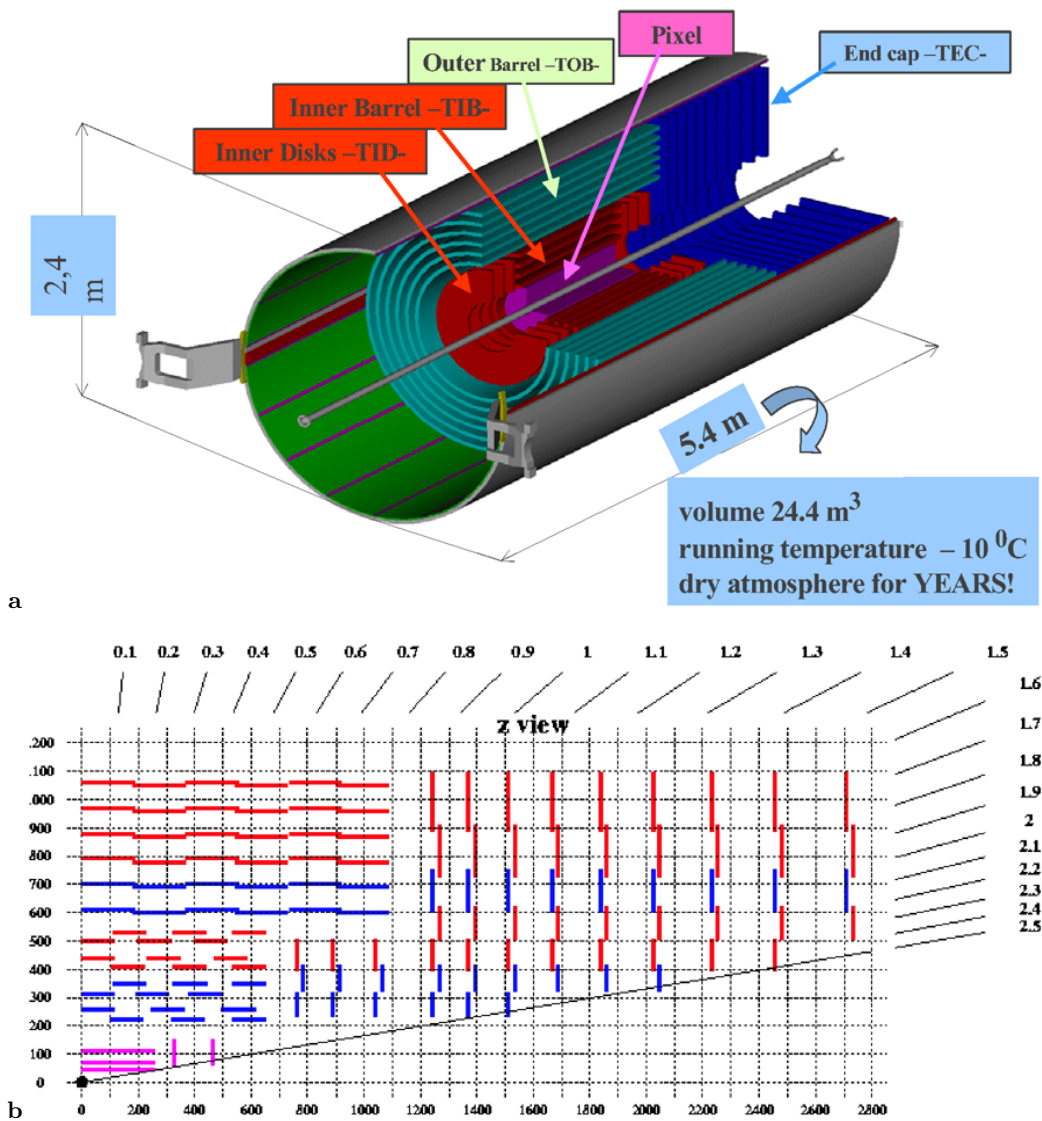


Fig. 2. **a** A schematic view of the CMS Tracker, illustrating the different sub-components of the Tracker. **b** A transverse quarter view of the CMS Silicon Tracker. In red are layers with a simple r - ϕ (r - Z) measurement module configuration. Shown in *blue* are layers with a back-to-back module arrangement, Inner Barrel layers 1 and 2, Outer Barrel Layers 1 and 2, as well as End-Cap rings 1, 2 and 5, which also provides a small angle stereo measurement. In purple, closest to the interaction *point*, are shown the three pixel Barrels and two pixel End-Caps

It was of fundamental importance to the feasibility of the CMS Silicon Tracker project, that radiation hard silicon micro-strip sensors could be produced within the set of processes normally available in large volume industrial 6" wafer fabrication lines. Of equal importance was the systematic program of technology transfer, which allowed industries such as ST-Microelectronics to become a major new producer of silicon micro-strip sensors, along with other companies already well established in this field such as, in particular, Hamamatsu Photonics.

Figure 2 shows a typical sensor produced from a 6" wafer, compared to one produced from a 4" wafer. It can be seen that the useful surface available for sensor production on a 6" wafer is more than twice that available on a 4" wafer. This leads to a corresponding reduction in

both cost and number of individual sensors to be manipulated. Also shown is a sensor from the ST-Microelectronics 6" production line, along with the test structures fabricated on the wafer periphery, which are used for systematic Quality Assurance throughout the production [3]

The production of silicon micro-strip sensors for the CMS Tracker is now well underway both in Hamamatsu and ST-Microelectronics. Due to their extensive prior experience, the quality of the Hamamatsu sensors fabricated in their 6" production lines has been excellent since the beginning of the production. ST-Microelectronics also has produced good quality sensors since the beginning of the production run. However, in this case a clear learning curve could be observed, as a number of fabrication and post-fabrication steps have required considerable tu-




Fig. 3. Charged particle occupancy at high luminosity in the CMS Tracker, per square centimetre, as a function of radius

ning in order to arrive at a consistently satisfactory level of quality. In particular, the transfer of all sensor fabrication steps from dedicated R&D lines to large volume production lines required an interruption of a few months in sensor deliveries, after the completion of a pre-series of approximately 5% of the total. Sensor production is now re-starting at ST-Microelectronics, following a dedicated program of process optimization and qualification on the large volume production lines.

Also of fundamental importance for the CMS Tracker project was the ability to produce sufficiently radiation tolerant front-end read-out chips (APV25), using standard $0.25\mu\text{m}$ processes, rather than costly custom radiation hard processes. IBM has already produced most of the ASICS required for the Silicon Tracker. An unexpected drop in the yield of good chips also momentarily slowed the production of APV25 down. This was a peculiar problem, which radically affected the yield of the APV25 and other circuits produced by IBM for the high-energy physics community, but which had no discernible effect on the yield of any of their other commercial ASICS. Eventually, the problem was traced to a process variation that, coupled with a layout feature common, and essentially unique, to the HEP designs resulted in a high failure rate of one of the several dielectric layers within the circuits. Here too, following the identification of the problem, and the stabilization of the process, the full rate of production has since been re-established.

The start-up of module assembly for the CMS Tracker has been considerably delayed by the difficulties encountered in industrializing the pitch adapters and, in particular, the front-end hybrids.

For the pitch adapters, we rely on the Aluminium on glass technology already extensively used in previous experiments. However, identifying producers able to produce the large number of pieces needed, while meeting the very low defect rate required over the full surface of components

much larger than most commercial products manufactured with similar technologies, turned out to be a more difficult challenge than originally anticipated.

Amusingly, of the two companies (RMT and Planar) participating in the production of our pitch adapters, one (Planar) specializes in the production of large area flat fluorescent displays. These consist of glass plates with two layers of fine pitch Aluminium strips, at right angles to each other, separated by fluorescent and dielectric layers. These plates, up to A4 size, are required to have essentially zero defects over the full surface of the plate! When discussing with this company, it soon became apparent that we had succeeded in identifying a large volume industrial producer, which was well adapted to the fabrication of our pitch-adapters.

Similarly, for the front-end hybrids we initially attempted to base the production on a "tried-and-true" technology. In this case, this was mixed thin and thick film gold on ceramic technology similar, for example, to that used by the LEP experiments for their silicon detectors. Here, however, the very small feature size of the circuit excluded all but a couple of potential producers. Furthermore, prototype runs showed that neither of these producers could reliably achieve the required level of quality, nor could they realistically provide the number of pieces required.

Following this abortive attempt, a survey of more modern technologies was carried out, which finally led to the choice of using a four layer Copper on Kapton flexible circuit, subsequently laminated to a ceramic rigidifier to provide the required thermal and mechanical characteristics. One such hybrid, from the ongoing production run, is shown in Fig. 2. In addition, the hybrid ceramic rigidifier is also used to carry the pitch adapter, as can also be seen in Fig. 2.

This choice has several important advantages over other possible ones. In particular, it decoupled the technology used for producing the circuit from the stringent mechanical requirements imposed by the need for reliable automated manipulation of the hybrids and pitch adapters during module assembly, and the thermo-mechanical requirements imposed for safe operation of the hybrid within the assembled module. This decoupling allowed the use of the technology best suited for large volume production of fine-feature circuits, as well as the (literally) seamless integration of a flexible cable into the hybrid circuit. It also allowed the choice of a rigidifier substrate best suited to our thermal, mechanical and cost requirements.

A second challenge for the production of the front-end hybrids was to identify companies able to mount components on the hybrids and, most importantly, precisely place and wire-bond the APV25 ASICS. Several attempts were required before one such company was clearly identified.

The large volume production of the CMS Tracker hybrids is now starting, following a series of successful prototype and qualification runs.

Given the successful industrialisation of the production of module components, the task of assembling some 17'500 modules is itself a formidable challenge. The ap-




Fig. 4. A schematic illustration of a CMS Silicon Tracker module, showing the different main module components




Fig. 5. Shown on the *left* is a typical sensor fabricated on a 6" wafer, compared to one fabricated on a 4" wafer (Hamamatsu). On the *right* is shown a sensor fabricated from a 6" wafer, along with the various test-structures on the wafer periphery, which are used as part of the Quality Assurance throughout the production (ST-Microelectronics)

proach chosen from the outset to meet this challenge is that of automating all the crucial module assembly steps, most of which had previously involved a great of highly skilled manual operator intervention. A survey of available pick-and-place machines used in the semi-conductor industry showed that no commercially available device was well suited to our needs. In particular, whereas most industrial applications required the high speed handling of small dies, with precisions of order tens of microns, our application requires handling of much larger dies (the sensors measure of order 100cm^2), and placing them to a precision of a few microns, with speed being essentially not a relevant issue.

As a consequence of this, a custom pick-and-place machine was developed, tailored to the various operations required for our module construction [4]. Figure 2 shows one of the module assembly gantries.

The CMS Tracker module assembly machine is based on a commercially available large area, precision, x-y "Gantry" table produced by Aerotech, equipped with additional z and phi movement axis (also manufactured by Aerotech). The x-y movement axes of the Gantry are driven by linear magnetic drives, and can be calibrated to achieve positioning precisions of a few microns. In keeping

with the requirements of the semi-conductor industry, they are also capable of extremely high speeds, and better than 3g accelerations!

The Gantry was then equipped with a pneumatically operated custom developed "master-head", capable of picking up a variety of custom tools, ranging from glue dispensers to a series of pick-up tools specially adapted to the variety of components to be manipulated. A series of custom pneumatic base-plates allows the easy loading and retention of components in preparation for the assembly, as well as the retention of assembled parts where necessary.

A pattern recognition program analyzes images from a camera mounted on the master-head, as shown in the inset of Fig. 2, and allows an accurate and automatic determination of the position of the base-plate, and that of all components on it.

The precision achieved for the placement of two sensors, relative to each other within a module, is shown in Fig. 2a. From Fig. 2b, it can be seen that this level of precision will not require the introduction of separate alignment parameters for each of the two sensors within a module, an appreciable simplification for track reconstruction.



Fig. 6. A front-end read-out hybrid and pitch adapter assembly for the CMS Tracker. The hybrid is produced as a four layer Copper on Kapton circuit, which is then laminated to a ceramic rigidifier prior to component assembly and bonding. The ceramic rigidifier also carries the pitch adapter so that it, together with hybrid, can be integrated into the module as a single assembly. Also shown is an adapter card, which avoids the need for multiple connection cycles on the fragile hybrid connector, during the repeated tests which the hybrid will be subject to prior to its definitive installation in the Tracker

Over half a dozen Gantries have been commissioned in as many institutes, each tailored to the assembly of specific module types. An additional Gantry has recently been commissioned, whose task is to assemble the pitch-adapter to the front-end hybrid, together with ceramic spacers under the pitch adapters, as required. By using a number of different inter-changeable base-plates, this single gantry will be able to perform this assembly operation for the full set of different hybrid and pitch-adapter types. This pre-assembly of the hybrid and pitch adapters greatly simplifies the successive module assembly steps.

The 17'500 modules of the CMS Tracker require some 25'000'000 wire bonds. Given the automation of module assembly, as described above, the wire bonding is the most likely bottleneck for the module production.

The quality of the bond surfaces is crucial to successfully make the required number of wire-bonds, as any significant level of re-work could slow the throughput by large factors. Of particular concern in this context, is the quality of the metallization of the pitch-adapters. This has been the focus of much optimization with each of the two manufacturers, and is subject to continuous screening during production.

The other obvious requirement for carrying out 25'000'000 wire bonds in some eighteen months is a park of modern, high-capacity bonding machines and teams of skilled operators. There are some eighteen such machines dedicated to the CMS Tracker module wire-bonding, distributed over fourteen institutes.

The flow of module components to the assembly lines now appears to be assured, and module assembly is finally starting to ramp up to speed.

Modules are operated in electrically independent read-out units of varying multiplicity in the various parts of the Tracker. In the Outer Barrel, for example, a "ROD" is the basic read-out unit. A single ROD integrates 6 modules, which may be either single or back-to-back, according to the measurement layer they are placed in.

Figure 2 shows a schematic of an Outer Barrel ROD, equipped with modules in a back-to-back configuration. In the lower, more detailed illustration, the ROD is shown "unfolded" into Top and Bottom views. The R-phi module configuration is shown on the Top view, and the small angles stereo configuration on the Bottom view.

As well as providing the required mechanical support and cooling for the modules, the ROD includes an inter-connect card which carries all the necessary power and control lines. The analogue electrical signals from the APV25 chips are taken to nearby Opto-hybrids, also serviced by the inter-connect board, where they are converted to analogue optical signals and sent to the Far End Digitizer boards over long optical fibres.

Figure 2 shows a ROD, equipped with six back-to-back module pairs, as used for read-out tests.

The results obtained, which reproduce well the performance previously observed for single modules in an ideal environment, are an important validation of the Tracker read-out, control and power scheme at the level of an electrically complete read-out unit, in this case the ROD.

It may not come as a surprise, that initial results were not so good, and that a number of detailed modifications were required in order to obtain the desired level of performance. This underlies the importance of extensive tests at each successive level in modularity, early enough in the progress of construction of the project so that detailed changes may still be implemented without excessive disruption. The difficulty in achieving this, of course, is the fact that definitive tests of this kind are only possible once final components are in hand, and that pressure on the schedule greatly reduces the time windows for such tests once final components do finally become available.

3 CMS track reconstruction, and tracking performance

In January 2001, the first fully functional track reconstruction prototype implementation in C++ was released. This provided an essential and very successful platform for all subsequent development. In particular, the well thought out factorization of the track reconstruction into basic components, made it possible to very much expand the number of useful developers, and to proceed to a systematic understanding and optimization of the performance of each of these basic components. Furthermore, this provided also the necessary input for other levels of tracking reconstruction and analysis, such as primary and secondary vertex reconstruction, alignment studies etc.

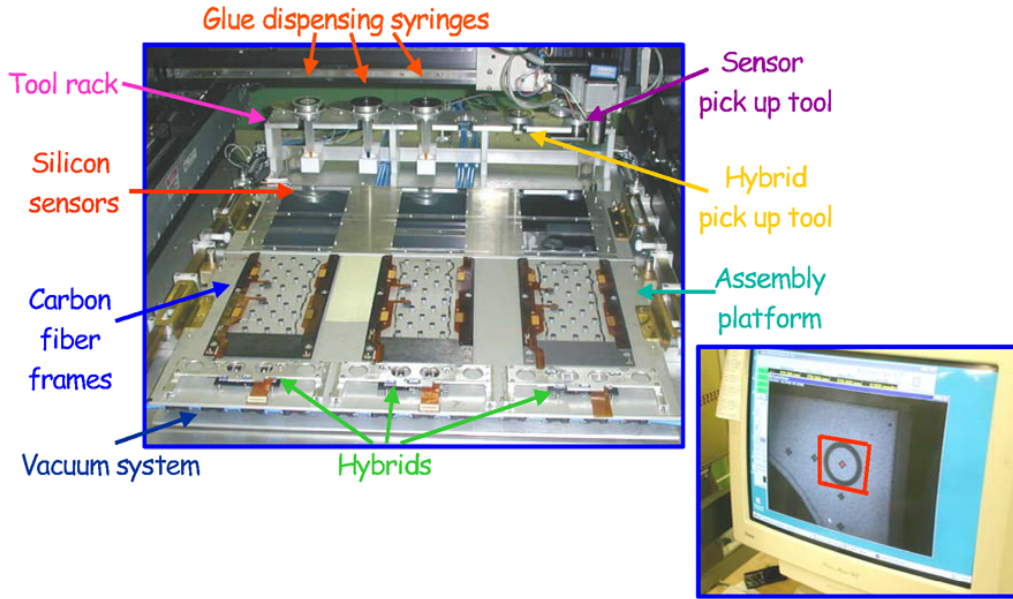


Fig. 7. A CMS Tracker module assembly Gantry. In the inset is shown the camera image, captured by the pattern recognition program, and used to determine the precise position of each object on the gantry table

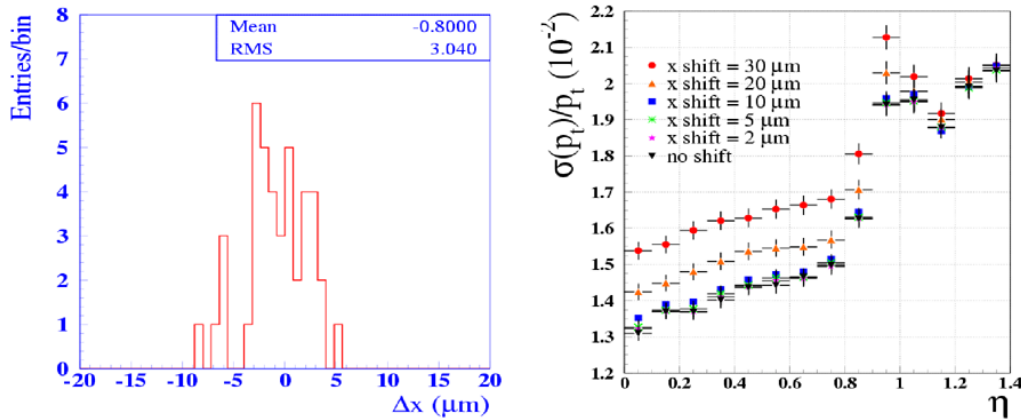


Fig. 8. On the *left* **a** is shown the measured miss-match between strips in the two sensors of a module, perpendicular to the strip direction. On the *right* **b** is shown the expected degradation in P_t resolution as function of this miss-match, if it is left uncorrected in the track reconstruction

A detailed description of the Tracker geometry, in which, for example, all module components that contribute significantly to the material budget are individually included, as well as a detailed simulation of the response of both the micro-strip and pixel detectors have been implemented in GEANT3. A GEANT4 version of these is in the final stages of validation.

Systematic studies of the effects of charged track occupancy in the CMS Tracker have been performed. For example, the fraction of hits along interesting tracks within a 200GeV b jet (those above 1 GeV in P), which are contaminated by hits from other near-by tracks has been measured, according to the detailed simulation.

It is found that, even in presence of high luminosity even pile-up, the fraction of polluted hits is indeed between 2% to 4% in the innermost micro-strip layers, and from 1% to 2% at larger radii. Moreover, it has also been found

that the tracks suffering from one or more polluted hits are typically still reconstructed, and the polluted hit assigned to the reconstructed track. The silicon micro-strip tracker does, therefore, indeed provide robust and clean hits, even in highly collimated jets and at high luminosity.

From Fig. 3, it can be seen that the CMS Tracker will, indeed, provide resolution of about 2% or better, up to a rapidity of around 1.75, corresponding to the length of the CMS Tracker, over which the full radial lever arm is available. Beyond this point, the radial lever arm starts to become reduced and the P_t resolution is accordingly degraded. Also shown, is the P_t resolution for an identical, but mass-less, hypothetical Tracker. From this comparison it can be seen that the difference between the resolution expected from the naïve calculation of 1, and the performance determined from the full simulation, is essentially due to the material within the tracking volume. One con-




Fig. 9. A schematic illustration of a Tracker Outer Barrel ROD




Fig. 10. A Tracker Outer Barrel ROD, as used for read-out tests

sequence of this is that a finer pitch, and higher channel count, would therefore yield only diminishing returns in terms of improving the P_t resolution.

It is important to note, that the single hit resolutions, and the lever arm between consecutive hits, are such that near asymptotic track parameter resolution for tracks between 5 and 10GeV is achieved already with the first 6 reconstructed hits. Typically this means three or four micro-strip hits, in addition to three or two pixel hits respectively. This can be seen from Fig. 3, where the P_t and transverse impact parameter resolutions are shown, as hits are added to the track, starting from the pixel vertex detector and moving out to the micro-strip tracker, for tracks with P_t between 5GeV and 10GeV, in the barrel region.

Also shown is the effect of having all three possible hits in the pixel barrel as opposed to only two. As expected, missing a pixel hit, and in particular the inner most one, leads to an irrecoverable loss in transverse param-

eter resolution, but does not significantly degrade the P_t resolution for the affected track.

This rapid convergence to well measured track parameters with only the first few hits has two important consequences.

The first, is that search windows when propagating track candidates from one silicon layer to the next become quite small, about $200\mu\text{m}$ in ϕ and $400\mu\text{m}$ in z , as soon as at least one micro-strip hit is added to the two or three pixel hits which constitute the track seed. As a result, the combinatorial pattern recognition problem is essentially solved at that point, which allows for very fast and efficient track reconstruction.

This can clearly be seen from Fig. 3, which shows the fraction of track candidates associated to spurious hits, when propagating from the third Pixel layer to first micro-strip layer (about 15% to 20%), and when propagating from the first micro-strip layer to the second micro-strip layer (about 1%).

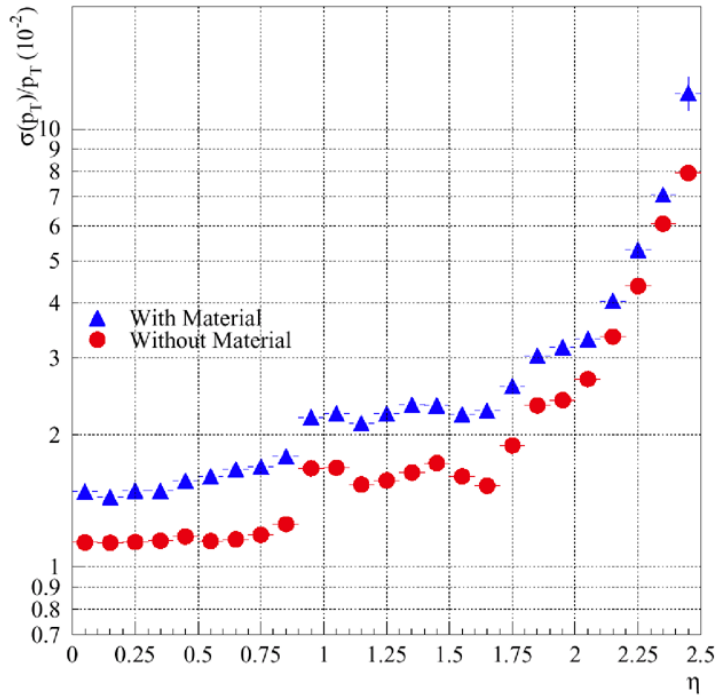


Fig. 11. The P_t resolution of the CMS Tracker, for 100GeV muons, as a function of rapidity

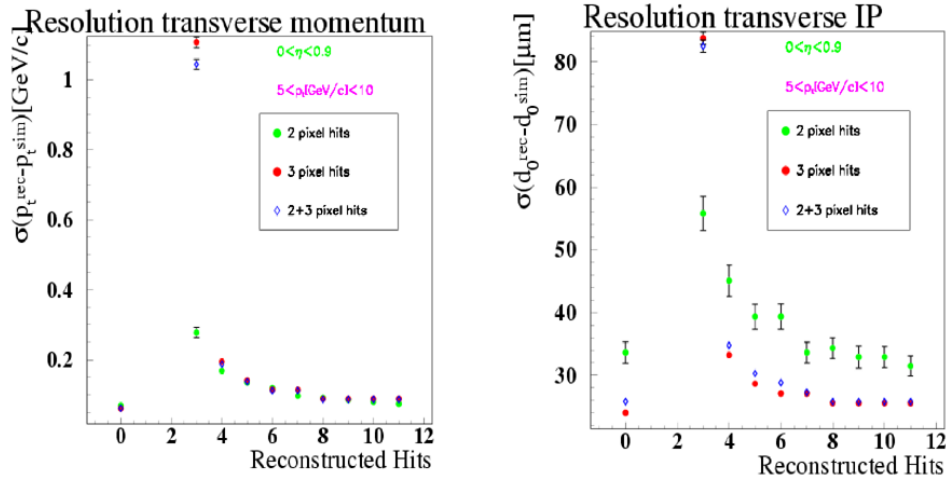


Fig. 12. The improvement in P_t and impact parameter resolutions for tracks with P_t between 5 and 10 GeV, in the central rapidity region, as each successive hit is added to the reconstructed track, from the vertex point outward. The Pixel layers are not perfectly hermetic, and in some cases only two of the nominal three hits are present on a track. Tracks are shown separately according to whether two or three pixel hits are present (*green and red round symbols respectively*). Results for the full set of tracks, with two or three pixel hits, are shown as the empty rhomboidal symbols

The results shown are in the absence of event pile-up, but adding this does not significantly change the results.

The second important consequence is that the vast majority of tracks need not be fully reconstructed in order to obtain sufficiently well measured track parameters for most purposes. This is particularly important for the High Level Trigger of the CMS experiment, as it greatly reduces the reconstruction time required to extract the relevant information.

In the CMS experiment the full L1 trigger rate of 100KHz is passed on to a large multi-CPU High Level

Trigger (HLT) farm, which must reduce the rate by a factor 1000 in order to output to storage about 100Hz of selected events [5]. The time budget for event rejection/selection at HLT, normalized to 1GHz CPU, is projected to be about 500ms.

Together with a systematic optimization of the track reconstruction code, and in association with the use of "regional" and "conditional" track reconstruction, the intrinsic capability of the Tracker to rapidly solve the pattern recognition combinatorial problem, and define precisely the parameters of track candidates at an early stage

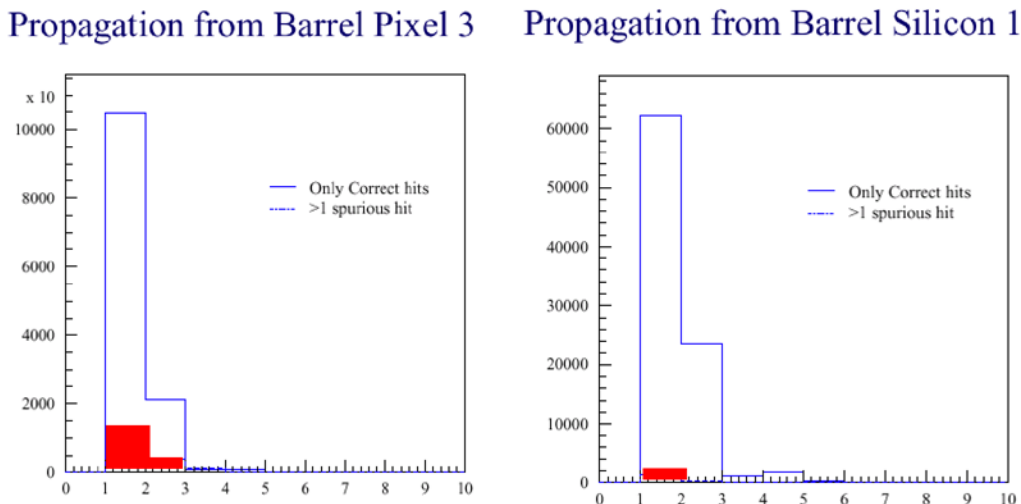


Fig. 13. Shown on the *left* is the number of track candidates formed when propagating a valid track from the third Pixel layer to the first silicon layer. Shown on the *right* is the similar quantity, when propagating from the first to the second silicon layer. Shown in *red* is the number of spurious track candidates in each case. The results are shown for 100GeV b jets, in the absence of pile-up. Adding the effect of pile up does not significantly change the results. Note that the overlap of sensors within a layer occasionally results in a single track leaving two hits within a layer. In the plots above, each of these is shown as a separate track candidate

of their reconstruction, allows the use of detailed tracking information on essentially the full 100KHz L1 trigger rate.

The simplest example of this is the validation of L1 muon triggers. The vast majority of muons passing the L1 trigger originate from b decays, and have a true P_t well below the nominal trigger threshold. By only reconstructing tracks within the "region" of interest defined by the L1 muon trigger, and by discarding tracks as soon as they are no longer compatible with the "condition" of having a P_t higher than the nominal muon trigger threshold, the tracker can efficiently reject these events within about 50ms and 150ms at low and high luminosity respectively, on a 1GHz CPU.

A more complex application of the tracking at the HLT is, for example, the case of tau jet tagging. In this case, a hierarchy of regions of interest are used. The reconstruction of a given track is stopped if it is found to have a P_t below significantly 1GeV, in which case the track is irrelevant to the event selection, or if at least six hits have already been associated to the track, in which case the track parameters are sufficiently well measured for the event selection. This is illustrated in Fig. 3.

In a first step, a search for leading track candidates with a P_t of more 6GeV is performed, within the "signal cone" R_s , defined by the calorimeter based L1 "tau" trigger. The event is rejected if no suitable leading track candidate defining the tau jet is found.

The candidate tau jet should be narrow and isolated. The next step in the reconstruction and analysis of the event therefore is the search for additional tracks outside a jet-track matching cone R_m , within which tracks are assumed to be part of the tau jet, and within a jet isolation cone R_i . The event is discarded as soon as a track with P_t over 1GeV is found in the region thus defined.

A calorimeter cluster analysis is also performed to confirm the narrow and isolated nature of the L1 tau trigger




Fig. 14. An illustration of the different regions of interest used for the identification of tau jets

candidate, and is used in association with the tracking information to achieve the required rejection factor of 1000, while maximizing the efficiency for signal events.

It is found that the optimal combination results from applying a rejection factor of about 3 based on the calorimeter cluster analysis, and of about 330 based on the tracking. The resulting signal efficiency, for example for a Susy Higgs with mass between 200GeV and 500GeV decaying into a pair of tau jets is above 40%, at both low and high luminosity, and the average time required to reject/select L1 triggered events is within the allocated budget in both cases.

Both inclusive and exclusive tagging of b jets may also be performed at the HLT level. Figure 3 shows the efficiency vs. purity curve for an inclusive b -tagging algo-




Fig. 15. The signal vs. purity for b-tagging for a signed impact parameter algorithm. The performance of the b-tag based on regional, conditional tracking, as used for the HLT, is shown with the *empty* symbols; that of the same b-tag using fully reconstructed tracks is shown with the *solid* symbols. In both cases, tracks are used to define the jet direction, as the jet direction from the L1 calorimeter trigger has a relatively poor resolution, which results in randomizing the sign of the impact parameter

rithm, both in the context of the HLT and in the more usual "off-line" scenario, which uses the full information available for all tracks in the jets. It can be seen that there is not a significant difference in the two cases. The time needed for the HLT inclusive b tagging algorithm is about 300ms and 1 second per jet, at low and high luminosity respectively, so that such b tags may be applied very early on in the HLT event filtering. The impact of this, for example in extending the SUSY physics reach of CMS, is currently under study.

Moreover, as mentioned above, events passing the L1 muon triggers are dominated by b decays with muons well below the nominal trigger P_t threshold. Rather than simply rejecting all such events, tracking information may be used in order to also select interesting exclusive b decays. For example, an HLT filtering scheme has been implemented, designed to select B_S to $J-\Psi$ decays, in which the $J-\Psi$ is tagged by its di-muon decay mode. The time required to scan candidate events for a $J-\Psi$ is, on average, about 250ms. The time required to fully reconstruct the B_S , in events with a tagged $J-\Psi$, is about 800ms. This particular algorithm will select about 84'000 b to $J-\Psi$ events for a 10fb^{-1} integrated luminosity, at the cost of less than 1.5Hz of events output to storage. An algorithm similar to that used to tag the $J-\Psi$ may also be used to search for rare B_S to $\mu t\mu$ decays. In this case, about 50 selected events are expected for a 10fb^{-1} integrated luminosity. In light

of this, it appears that the ability to use detailed tracking information in the HLT will allow the CMS physics program to also include some interesting aspects of b physics, in particular for the first physics runs at low luminosity.

An important design parameter for a tracking detector is the amount of material distributed within the tracking volume, and considerable effort has gone into the minimization of the CMS Tracker material. For example, lightweight Carbon composites are used for all mechanical support structures; Aluminium conductor is used for all cables within the tracking volume, etc.

The result is shown in Fig. 3. It can be seen that the material within the tracking volume peaks at just over 1 radiation length, at $\eta \sim 1.6$. Services and mechanical support structures at the outer periphery of the Tracker, account for an additional 0.2 to 0.4 radiation lengths in the End-Cap region, from η greater than about 1.6. As already mentioned, this significantly degrades the P_t resolution of muons even up to 100GeV, and it obviously dominates the resolution for softer tracks. Moreover, hadronic interaction of pions with the material in the tracker is also the largest single source of tracking inefficiency.

However, as may be expected, the most important negative effect of the Tracker material concerns electrons and, to a lesser extent, photons. A specific track reconstruction algorithm has been developed to properly account for the energy loss of electrons in the tracking material. Normally, the effects of hit resolutions, multiple scattering, energy resolution etc., are each modelled with a single Gaussian. A single Gaussian is clearly a rather poor approximation to Bethe-Heitler parameterization for the Brehmstrahlung energy loss of electrons in material. The basic idea, therefore, is to model this as a mixture of different Gaussians, in order to reproduce the varying degrees of hardness of the electron Brehmstrahlung in a layer of material.

The propagation of a single electron through each layer of material within the Tracker, may then be regarded as the propagation of a set of different track hypotheses, each corresponding to one of the Gaussians modelling the Bethe-Heitler energy loss parameterization. This method is called the Gaussian Sum Filter, or GSF [6]. Clearly this method suffers from a possible combinatorial explosion, as the number of track hypotheses would increase by a power of one for each layer traversed, if left unchecked. In practice, it is possible to obtain satisfactory results while reducing to a manageable level the number of different track hypotheses under consideration at each layer in the track propagation.

The results may be seen in Fig. 3, compared with those obtained with the usual Kalman-Filter technique, tuned to ensure the correct mean value for the energy loss of electrons in material, as well as the correct variance. It can be seen that the GSF provides a considerably improved P_t resolution; the FWHM is reduced by close to a factor of two, with significantly reduced tails. Moreover, the GSF also provides a reasonably flat probability distribution, which allows for the correct statistical treatment of the resulting track parameters in a straightforward way.

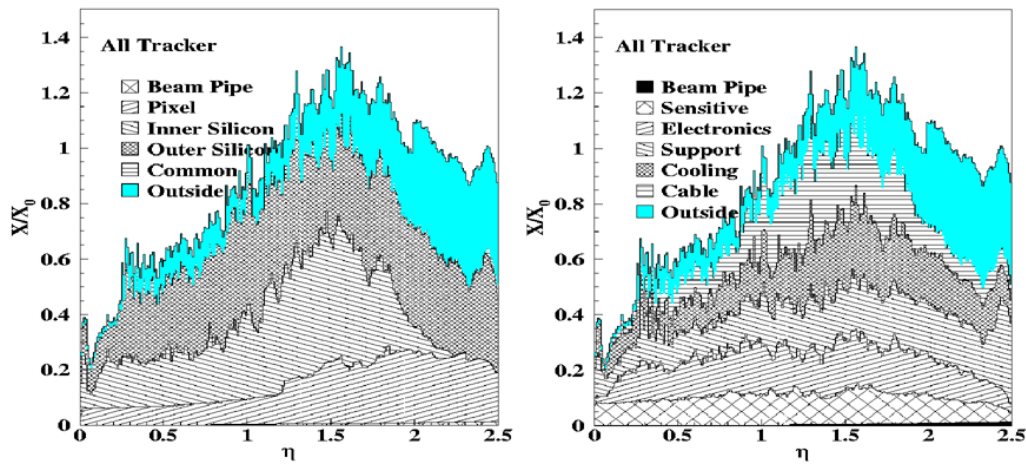


Fig. 16. Distribution of material within the CMS Tracker, in units of radiation lengths, as function of rapidity




Fig. 17. Shown on the *left* is the $1/P$ residual distribution for 10GeV electrons in the Barrel region. On the *right* is shown the probability distribution. Shown in the (*red*) shaded area, and as the continuous line, are the results from the Gaussian Sum Filter, compared with those from the normal, and appropriately tuned, Kalman-Filter

The GSF method is currently being generalized to other cases of interest, such as for example a description of non-Gaussian tails in the inner most Pixel layers, for an improved impact parameter determination.

4 Conclusions

The CMS silicon Tracker calls for the largest scale deployment of silicon strip sensors to date. As such, it has required an unprecedented attention to the successful industrial manufacture of, in particular, all module components. It also has required a novel approach to many aspects of module assembly, with an important emphasis on automation.

The flow of all module components to the assembly lines now appears to be assured, and module assembly is currently starting to ramp up to speed.

The reconstruction software for the CMS Tracker is well advanced, following the first release of a fully functional prototype in C++ in January of 2001.

The CMS Tracker will provide P_t resolution of order $1 \sim 2\%$ for 100GeV muons. It will also allow detailed inspection of jets and, thanks to the Pixel vertex detector, efficient b tagging.

Moreover, the extremely fine pixel segmentation, and the resulting low cell occupancy even at high luminosity,

coupled with the robustness and resolution of each additional hit along a track, will allow use of detailed tracking information at the earliest stages of event analysis and selection in the High Level Filter.

References

1. CMS: The Tracker Project, Technical Design Report, CERN/LHCC 98-6
2. CMS: Addendum to CMS Tracker TDR, CERN/LHCC 2000-016
3. S. Braibant et al.: Investigation of design parameters for radiation hard silicon micro strip detectors, NIMA, 485 (2002) 343-361
4. CMS: Silicon micro-strip sensor quality assurance, CMS Internal Note
5. A. Honma et al.: An automated module assembly system for the CMS Silicon Tracker, CMS Note 2001/005
6. CMS: The Trigger and Data Acquisition project, Volume 2, Data Acquisition and High Level Trigger, Technical Design Report, CERN /LHCC 2002-026
7. R. Fruwirth: Track fitting with non-Gaussian noise, Computer Physics Communications 100 (1997) 1
8. R. Fruwirth and S. Fruwirth-Schnatter: On the treatment of energy loss in track fitting, Computer Physics Communications, 110 (1998) 80