

The Status of the CMS Experiment at the LHC

Tejinder S. Virdee¹

Imperial College, London and CERN

PH Dept., CERN, CH-1211 Geneva 23, Switzerland

E-mail: tejinder.virdee@cern.ch

Robert Cousins

University of California, Los Angeles

Address, U. S. A.

E-mail: cousins@physics.ucla.edu

The Compact Muon Solenoid (CMS) experiment was commissioned and ready for first LHC beams in September 2008. Muons from beam splash, beam halo, and cosmic rays have provided valuable data for alignment and calibration, detector performance studies and for testing numerous aspects of offline software and computing. The first maintenance cycle is complete, and while eagerly anticipating the first LHC collisions, we also recall the CMS collaboration's physics studies based on simulated data as a reminder of the exciting possibilities in the near future.

¹ Speaker (on behalf of the CMS Collaboration)

*European Physical Society Europhysics Conference on High Energy Physics EPS-HEP 2009,
Krakow , Poland
July 16-22, 2009*

POS (EPS-HEP 2009) 006

1. Introduction

The Compact Muon Solenoid (CMS) Collaboration consists of over 3000 scientists and engineers from 38 countries. After almost 20 years of work by collaborators and associated institutional and industrial partners, from the conception and design, through the construction, installation, and commissioning, CMS [1] became a working experiment in September 2008. Here we review some of the highlights of the installation and commissioning underground; the useful data taken during the brief period of single beams; the very productive month-long cosmics data-taking run after the LHC incident of September 19, 2008; the first maintenance cycle during the 2008-2009 shutdown; and a second cosmics run. We also describe some of the many preparations and tests of software and computing, and end with a quick overview of some of the many physics analyses which have been prepared and extensively exercised on simulated data. These serve as a reminder of the exciting possibilities in the months and years ahead of us.

The CMS main public page [2] has links to extensive public web pages including outreach information, continuously updated physics results for physicists, and an e-commentary during LHC beam periods.

2. The CMS Detector

The central feature of the Compact Muon Solenoid (CMS) detector, illustrated in Fig. 1 alongside a transverse view of a cosmic muon traversing CMS, is a superconducting solenoid, of 6 m internal diameter, providing a field of 3.8 T. Within the field volume are the silicon pixel and strip tracker, the lead-tungstate crystal electromagnetic calorimeter (ECAL), and the brass/scintillator hadron calorimeter (HCAL). Muons are measured in gas-ionization detectors embedded in the steel return yoke. In addition to the barrel and endcap detectors, CMS has extensive forward calorimetry. CMS has an overall length of 22 m, a diameter of 15 m, and weighs 12 500 tonnes.

With 75 848 lead tungstate (PbWO_4) crystals that are $25.8X_0$ long in the barrel and $24.7X_0$ long in the end-caps, the ECAL has an energy resolution of better than 0.5% above 100 GeV. The 15k-channel HCAL, when combined with the ECAL, measures jets with a resolution $\Delta E/E \sim 100\%/\sqrt{E} \oplus 5\%$. The calorimeter cells are grouped in projective towers of granularity $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$ at central rapidities and 0.175×0.175 at forward rapidities. The CASTOR ($5.3 < |\eta| < 6.6$) and Zero Degree ($|\eta| > 8.3$) calorimeters (made of quartz fibers/plates embedded in tungsten absorbers) cover the very forward angles.

Muons with pseudorapidity in the range $|\eta| < 2.4$ are measured with detection planes made of three technologies: about 250 Drift Tube chambers, 450 Cathode Strip Chambers, and 900 Resistive Plate Chambers. The readout has nearly 1 million electronic channels. Matching the muons to the tracks measured in the silicon tracker should result in a transverse momentum resolution between 1 and 5 %, for p_T values up to 1 TeV/c.

The inner tracker measures charged particles within the $|\eta| < 2.5$ pseudorapidity range. It consists of 1440 silicon pixel and 15 148 silicon strip detector modules and is located in the field of the superconducting solenoid. It provides an impact parameter resolution of $\sim 15 \mu\text{m}$ and a transverse momentum (p_T) resolution of about 1.5% for 100 GeV/c particles.

The first level (Level-1) of the CMS trigger system, composed of custom hardware processors, is designed to select, in about 1 μs of processing time, the most interesting events, using information from the calorimeters and muon detectors. The High Level Trigger (HLT)

processor farm further decreases the event rate from up to 100 kHz to a few hundreds of Hz (initial DAQ system is 50 kHz), before data storage. Offline GRID computing of some 50k cores runs more than 2M lines of source code.

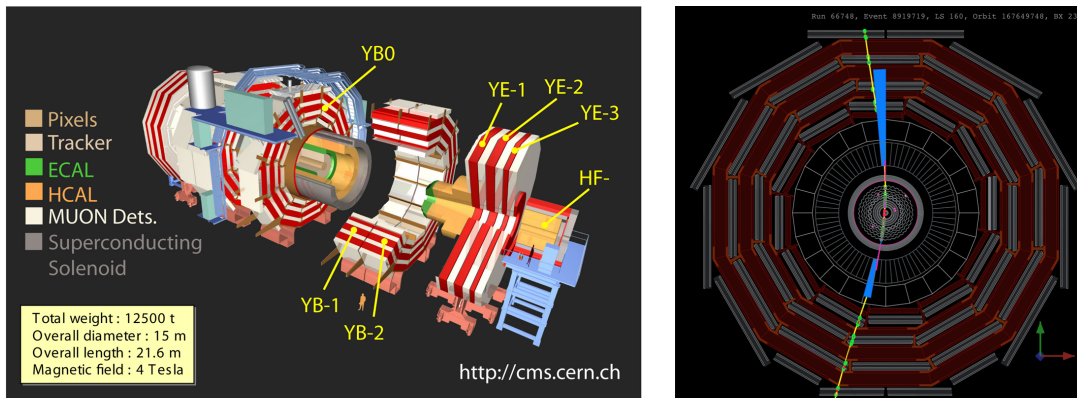


Figure 1: Left: The CMS detector. The z-axis, defined along the nominal LHC beams, has +z direction pointing counter-clockwise. Right: Display of a cosmic muon (recorded during the CRAFT run, Sec. 6) that enters and exits through the DT muon system, leaves deposits in the HCAL and ECAL, and crosses the silicon strip and pixel tracking systems.

3. Installation of CMS Underground

The heavy elements of CMS began to be lowered into the experimental cavern in November 2006, starting with the forward calorimeters and continuing shortly thereafter with the +z endcap disks and barrel wheels, complete with muon detectors and services. The central yoke wheel (YB0), which houses the cryostat, was lowered in February 2007, and by January 2008 the last heavy elements of the -z endcap were successfully lowered into the cavern.

The campaign to connect services for the detectors within the central portion of CMS included the installation of more than 200km of cables and optical fibres (about 6000 cables). Additionally, more than 20km of cooling pipes (about 1000 pipes) were installed. The whole enterprise took place over a 5 month period and required more than 50 000 man-hours of effort. The silicon strip tracker was inserted in December 2007, its cabling was completed in March 2008, and its cooling was operational by June 2008. In the same month, the beryllium central beam pipe was installed and baked out. The silicon pixel tracking system and the endcaps of the ECAL were the last components to be installed, in August 2008. The mechanics and the cabling of the pixel system have been designed to allow relatively easy access or replacement if needed. The preshower detector for the endcap electromagnetic calorimeter was the only major subsystem not installed prior to the 2008 data-taking with beam and cosmics. It was installed in March 2009.

4. Global Run Commissioning and Final Closing of CMS

A series of global commissioning exercises using the final detectors and electronics installed in the underground caverns, each lasting 3–10 days and occurring monthly or bimonthly, commenced in May 2007 and lasted until the experiment was prepared for LHC beams, by the end of August 2008. The scale of these “global runs” steadily increased to include all of CMS, while balancing the need to continue installation and perform extensive detector subsystem commissioning with the need for global system tests.

Many detector subsystems were available in their entirety for global commissioning by May 2008, and thus a series of four week-long exercises, each known as a Cosmics RUn at ZERO Tesla (CRUZET), were conducted to accumulate sizable samples of cosmic muon events from which to study the overall detector performance. By July (for the third CRUZET exercise) the silicon strip tracker was in the data-taking with about 75% of the front-end modules. In the fourth CRUZET exercise, in August, the complete silicon pixel tracker was introduced, along with the both endcaps of the ECAL, thus constituting the first global run with the CMS detector to be used for the September LHC beam. The total accumulated cosmic triggers at zero field exceeded 300 million, including the triggers recorded in September 2008 when the experiment was live for the first LHC beams. These global runs regularly exercised the full data flow from the data acquisition system at the experimental site to the reconstruction facility at the the CERN Tier-0 computing facility, followed by the subsequent transfer of the reconstructed data to all seven of the CMS Tier-1 centres and to some selected Tier-2 centres.

The final sequence of closing the steel yoke was completed on August 25, 2008, and the last pieces of shielding were in place on September 3. After almost 20 years, from conception, design, construction and commissioning, CMS became a working experiment.

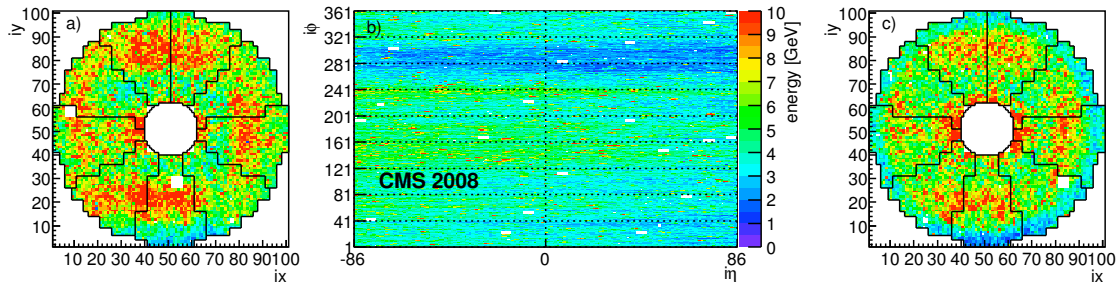


Figure 2: ECAL average energy deposit per crystal for a typical “beam splash” event with muons coming from the $-z$ side. (a) Occupancy of the $-z$ endcap, where i_x and i_y are indices of the crystals in the horizontal (x) and vertical (y) coordinates, respectively. (b) Occupancy of the barrel, where i_η and i_ϕ are indices of the crystals in the η - ϕ coordinates. (c) Occupancy of the $+z$ endcap. The white regions correspond to channels masked in the readout, less than 1% of the total channels in that run. Many of these channels have been recovered subsequently.

5. LHC Beam Operation in 2008

The CMS experiment recorded data associated with activity from the first LHC beams in September 2008. This activity began with single shots of the beam onto a collimator 150m upstream of CMS, which yielded sprays (so-called “beam splashes”) containing $O(10^5)$ muons crossing the cavern synchronously. With first circulating LHC beams beginning on September 10, CMS recorded beam-halo muons, first from the uncaptured beam, and then (at dramatically reduced beam-halo rate) from the RF-captured beams. The CMS solenoid was off at the LHC’s request, and the silicon pixel and strip tracking systems were powered off for safety reasons.

The first “beam splash” events were used to synchronize the beam triggers from several sources, and to commission the diamond beam condition monitors. The data collected from the “beam splash” events also proved useful for diagnostic occupancy plots and for adjusting the inter-channel timing of the ECAL and HCAL readout channels (Figs. 2 and 3), as the synchronous wave of crossing muons has a characteristic time-of-flight signature. CMS also recorded nearly 1 million beam-halo triggered events during the 2008 beam operations.

6. CRAFT 2008

Following the LHC incident of September 19 and end of LHC beams for 2008, in October and November, CMS conducted a data-taking exercise known as the Cosmic Run At Four Tesla (CRAFT) in late 2008. In addition to commissioning the experiment operationally for an extended period, the cosmic muon dataset collected during CRAFT has proven invaluable for understanding the performance of the CMS experiment as a whole. The objectives of the CRAFT exercise were (1) to test the solenoid magnet at its operating field (3.8 T), with the CMS experiment in its final configuration underground; (2) to gain experience operating CMS continuously for one month; and (3) to collect approximately 300 million cosmic triggers for performance studies of the CMS subdetectors.

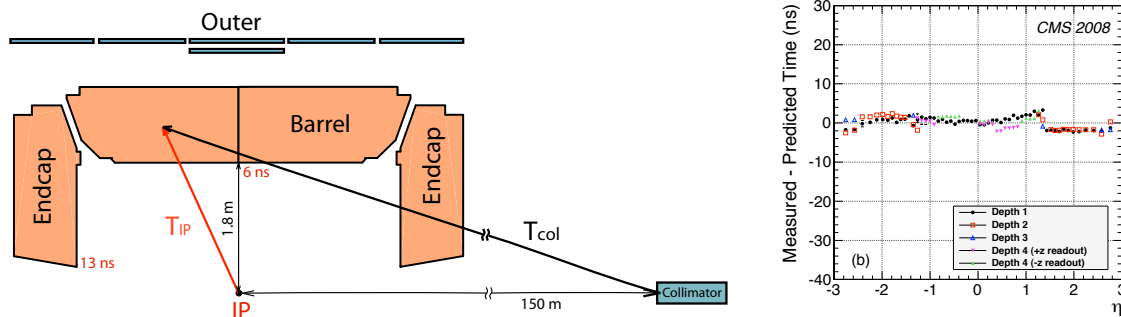


Figure 3: Left: A schematic view of the geometry of “beam splash” events. This geometry was used to predict the timing of energy deposits and thereby highlight channels in the HCAL that required synchronization. Right: Difference between measured and predicted time as a function of η from LHC “beam splash” events, after compensating sample delay settings were loaded for all three calorimeters, barrel, endcap and outer.

These goals were successfully met: over the course of 23 days during the most stable part of CRAFT, the experiment collected 270 million cosmic ray triggered events with the solenoid at its operating central field of 3.8T and with at least the silicon strip tracker and drift tube muon system participating and operating at nominal conditions. These data were processed by the offline data handling system, and then analyzed by teams dedicated to the calibration, alignment, and characterization of the detector subsystems. The precision achieved in detector alignment (equivalent to that foreseen after 10 pb^{-1} of LHC beam!), detector calibration, noise cleanup algorithms, as well as operational experience of running the experiment for an extended period of time, gave confidence that the CMS experiment was ready for LHC beam operations. In the subsequent year 23 papers were written on performance and understanding of the detector, posted at arxiv.org, and published in JINST [3].

The infrastructure and services met the demands of running the experiment continuously for one month, although the exercise indicated areas needing further improvement. Several cooling failures did occur, resulting in the shutdown of some equipment during CRAFT. Leaks were detected on the barrel yoke circuit for wheels YB-2 and YB-1. The leak detection system on one of the cooling circuits fired unnecessarily a few times resulting in three automatic slow dumps of the magnet; threshold has been subsequently adjusted. The cooling plants for the silicon strip tracker also suffered a few trips, The leak rate of the system was also higher than expected. In the shutdown after CRAFT, the cooling plants and piping were significantly refurbished to eliminate these leaks. In total, about 70 hours of downtime were caused by general infrastructure related incidents, about 10% of the duration of CRAFT. This time was dominated by the downtime of the magnet. Aside from this, the typical data collection

efficiency of CMS was about 70% during CRAFT, including periods used to conduct detector calibrations, pedestal runs, and diagnostics to further advance the commissioning of the experiment.

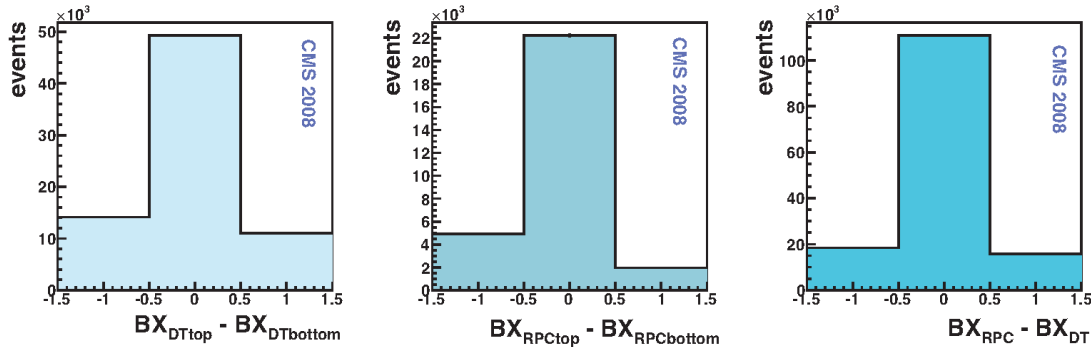


Figure 4: Time differences at the Global Muon Trigger (in units of 25-ns bunch crossing, BX) between Level-1 muon candidates - created in most cases by the same cosmic ray muon - from the top half and the bottom half of the DT system (left) and the RPC (middle) and between the RPC and DT system (right). The majority of the signals are synchronised and the skew at the clock edges is balanced.

The data from cosmic muons were used to refine the timing of the Level-1 trigger, both in further cross-timing the elements within each subdetector beyond that performed with splash events, as well as synchronizing across subsystems. As an example, trigger signals from the three muon systems are processed by the Global Muon Trigger (GMT), so it is necessary to ensure that the signal created by the same muon in different detectors enters the GMT in the same clock cycle. With cosmic rays (which, unlike eventual collision events, arrive at random phases within the LHC clock cycle, and do not originate from the beam collision point), this is possible only to a limited extent. One method for fine-tuning of the timing delays using the cosmic ray data involved measuring the signal arrival time from a particular subdetector with respect to triggers from another. A direct comparison is possible using the readout of the GMT, which records all input muon candidates and reads 3 consecutive clock cycles centered at the trigger. Fig. 4 shows that in most cases L1 muon candidates from different muon systems, induced by the same cosmic ray muon, arrive at the same BX. Occasional difference by 1 BX is unavoidable due to the fact that cosmic rays are asynchronous to the clock of the experiment and because the relative synchronisation between different detector parts obtained with cosmic rays has a finite precision of several ns. Such plots were used to fine-time the CSC timing during CRAFT, and similarly, delays of the various calorimeter trigger inputs to the Global Trigger were adjusted to provide the highest coincidence rate above the noise continuum.

Data were promptly reconstructed at the Tier-0 computing centre at CERN to create high-level physics objects with a job latency of about 8 hours, but with a broad distribution. These data were transferred to the CMS Analysis Facility (CAF) at the CERN Meyrin site and to several Tier-1 and Tier-2 centres worldwide for prompt analysis by teams of physicists. The average export rate from the Tier-0 centre to the Tier-1 centres during CRAFT was 240MB/s, and the total volume transferred was about 600 TB.

7. Results from CRAFT 2008 Data

The data collected during CRAFT facilitated a wide range of analyses on the performance of the CMS subdetectors, the magnitude of the magnetic field in the return yoke, as well as the calibration and alignment of sensors in preparation for physics measurements. Alignment of the silicon strip and pixel sensor modules was improved significantly from initial survey

measurements by applying sophisticated track-based alignment techniques to the data recorded from approximately 3.2 million tracks selected to cross the sensitive tracking region (with 110 000 tracks having at least one pixel hit). The precision achieved for the positions of the detector modules with respect to particle trajectories, derived from the distribution of the median of the cosmic muon track residuals, is 3–4 μm in the barrel and 3–14 μm in the endcaps for the coordinate in the bending plane. Other silicon tracking measurements performed with the CRAFT data include calibration of the absolute energy loss in silicon strip sensors, Lorentz angle measurements, hit efficiencies and position resolutions, track reconstruction efficiencies, and track parameter resolutions.

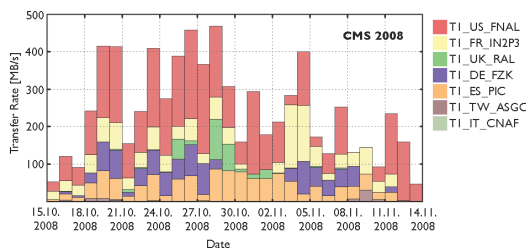


Figure 5: Transfer rates from Tier-0 to Tier-1 centres during CRAFT. The average was about 240 MB/s.

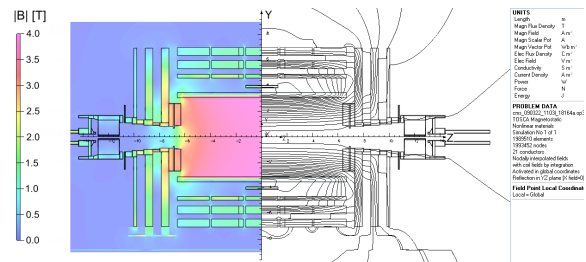


Figure 6: CMS magnetic field map.

Track-based alignment techniques using cosmic muons were also applied to align the DT muon detectors in the barrel region of the experiment. An alignment precision of better than 700 μm was achieved along the higher precision coordinate direction (approximately ϕ) for the first three DT stations as estimated by a cross-check of extrapolating muon segments from one detector to the next. A local alignment precision of 270 μm was achieved within each ring of CSCs using LHC beam-halo muons recorded during beam operations in 2008. While the accuracy of the magnetic field map (calculated using TOSCA) had been previously confirmed to be better than 0.1% in the tracker region (where it is extremely uniform), the accuracy in the barrel steel yoke, used in muon reconstruction, was checked carefully for the first time using CRAFT 2008 DATA. The muon bending measured by DT chambers should agree with that predicted by extrapolating the track parameters measured by the inner tracking system. During CRAFT a discrepancy was noted; the measurement in the yoke was too high by up to 20%. This was later traced to boundary conditions that had been set too restrictively in the TOSCA calculation, causing the magnetic field to be over-estimated in the iron. The analysis also suggested improving the treatment of asymmetric features in the map. A field map based on these results has residual differences between data and the calculation reduced to about 4.5% in the middle station of the barrel yoke and 8.5% in the outer station, and are empirically corrected to better than 2% using the CRAFT measurements.

Measurements of the cosmic muon energy loss, dE/dx , in the ECAL and HCAL barrel compartments validated the initial calibration of individual channels obtained prior to CRAFT. (The endcap studies suffered from small sample sizes.) The measured dE/dx as a function of muon momentum agrees well with first-principles calculations, as shown in Fig. 7.

Lower-level resolution plots for muon chamber position measurements can be found in the CRAFT papers [3]. Figure 8 shows a higher-level result made possible by all this work: muon momentum resolution as determined from muons such as the one in Figure 1, for which one can compare the position and momentum measurements made independently using hits in the two halves of the track.

In the p_T region below approximately 200 GeV/c, where the resolution is dominated by

multiple-scattering effects, the inclusion of muon hits does not improve the resolution beyond that obtained with the tracker-only fits. In the high- p_T region, the resolution obtained with the dedicated high- p_T muon reconstructors (labeled as TPFMS and TMR is better than that of global muons and of tracker-only tracks, as expected.

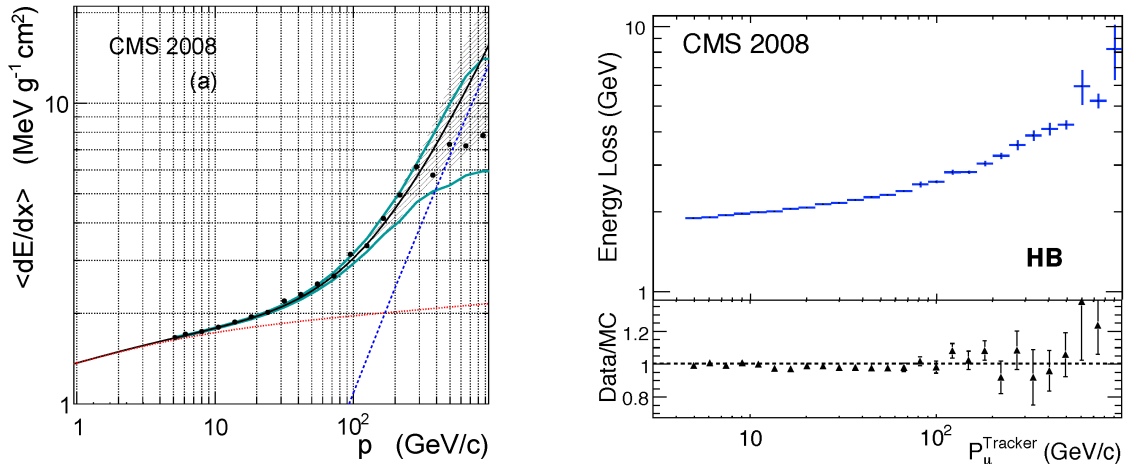


Figure 7: Left: Muon stopping power measured in PbWO₄ (dots) as a function of muon momentum compared to expectations (continuous black line). The expected contributions from collision and radiative processes are plotted as well (red dotted line and blue dashed line respectively). Right: (top) Energy loss of cosmic ray muons in the barrel HCAL after correcting the signal mean value for the muon path length and normalizing to the thickness of HB at $\eta = 0$; (bottom) ratio of data and Monte Carlo simulation predictions (arbitrary scale) for the bottom half of the calorimeter.

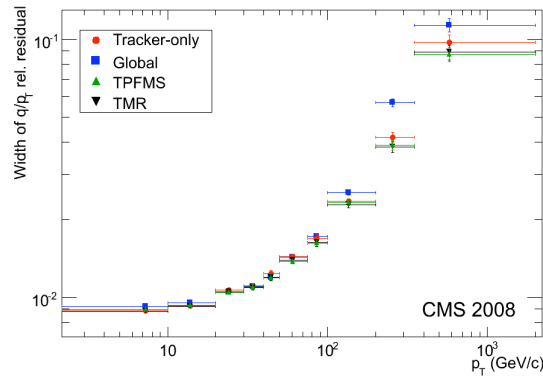


Figure 8: Widths of Gaussian fits to the distributions of the relative differences of measurements of charge over transverse momentum (q/p_T) made on the upper and lower halves of tracks such as that in Fig. 1 (divided by $\sqrt{2}$ to obtain single-track resolution), for various muon reconstruction algorithms, as a function of p_T of the reference track.

8. CMS Activities During the 2008-2009 LHC Shutdown

After the cosmics run ended, the detector was opened for carefully selected maintenance, consolidation, and repair activities, as well as the installation of the preshower subdetector and the CASTOR calorimeter. Work progressed according to the schedule laid down in Nov. 2008. After this seven-month long, successful maintenance period (resulting in only a few/mille channels not functioning), CMS was closed once again and the CRAFT-2009 run was underway, once again taking cosmics data at the operating field of 3.8 T. The plan (since implemented) was to be smoothly taking data in a “beam-ready” state by the time of the first

LHC beams of 2009.

Some highlights of the shutdown work on the hardware included the installation and commissioning of the preshower on both endcaps; the removal, repair, and re-insertion of the forward pixel system; the installation of CASTOR ($5.2 < |\eta| < 6.6$) calorimeter the maintenance and (small) repairs involving many sub-systems; and the revision of the tracker cooling plant. As repaired components came online, re-commissioning of CMS proceeded with resumption of Mid-Week Global Runs and CRUZET-2009, interspersed with final maintenance and consolidation activities. Open, inclusive Level-1 and High-Level Triggers menus for early LHC running have passed tests at high throughput rates (100 kHz for the full DAQ system except the HLT filter farm at 50 kHz).

In parallel to several reconstructions of the CRAFT data, large Monte Carlo simulation productions were produced and reconstructed (e.g., about 250 M fully simulated events at 10 GeV begun at the end of 2008 were completed in February 2009). Meanwhile, there was a large effort to prepare software for 2009 data taking, including improving stability and reliability of the computing infrastructure. Of course all of the results discussed here depend on extensive offline software (SW) (much of which also runs online for monitoring and high-level-triggering). During 2009, the major annual release known as CMSSW version 3 was developed and released in a sequence of updates converging on the version used for the major simulations and data-taking, and the respective analyses thereof. Among the many features added or refined, in many cases based on input from CRAFT08 and analysis exercises, were: GEANT4 and ROOT major updates; mixing of event data from more than one source, including simulated pileup; more realistic detector response, including dead/noisy channels; numerous improvements and innovations in reconstruction (track-corrected jets and missing E_T ; particle flow approach to event analysis); the new magnetic field map and empirical corrections to it; and the new version of Physics Analysis Toolkit used across the collaboration as a common language for analysis objects.

Much work also paid off in improving workflows, converting many manual operations to be more automatic, and reducing elapsed time, for example by pre-staging tapes to disk when re-reconstructing at Tier-1 centers. First analyses of the 2008 data used the Worldwide LHC Computing Grid (WLCG) infrastructure and the SW release (CMSSW 2) that had been destined for 2008 data-taking. Re-reconstruction and further analyses proceeded quickly using updated alignment and calibration and refined CMSSW 2, and then CMSSW 3 after it was available.

All the software usage in turn depends on the enormous computer infrastructure at CERN and on the WLCG, which has been scaled up in previous years and routinely exercised, with each exercise pointing to some needed improvements which had been subsequently implemented. In addition to participating successfully in extensive WLCG-wide tests as the Scale Testing for the Experimental Program (STEP'09) in spring, CMS made a concerted effort to increase the site readiness of all Tier-1 and Tier-2 centers. As first LHC collisions approached, CMS computing was ready.

9. Physics Commissioning

In parallel with installation and commissioning of the CMS detector, numerous physics analysis strategies have been “commissioned”, necessarily using simulated data, but attempting wherever possible to exercise the eventual “data-driven” methods to extract detector resolutions, scale calibrations, and physics backgrounds from the data themselves. With simulated data, such results can be compared with “Monte Carlo truth”, i.e. the assumed true values of resolution and backgrounds underlying the simulation, a comparison known as the “closure”. Already two

years ago, implementations of Tag and Probe (T&P) methods (such as those used in Tevatron experiments) were demonstrated in simulations assuming 14 TeV LHC energy, to identify a physics object in an unbiased way in order to study efficiencies. For example, in Z^0 to e^+e^- events, using one tight electron as the tag, the other electron can be the probe, provided the invariant mass of the pair is M_Z . In a typical example of such an exercise, the measured efficiency from T&P was 94.36 ± 0.24 , compared to MC truth of 94.63 ± 0.24 for 10 pb^{-1} at 14 TeV.

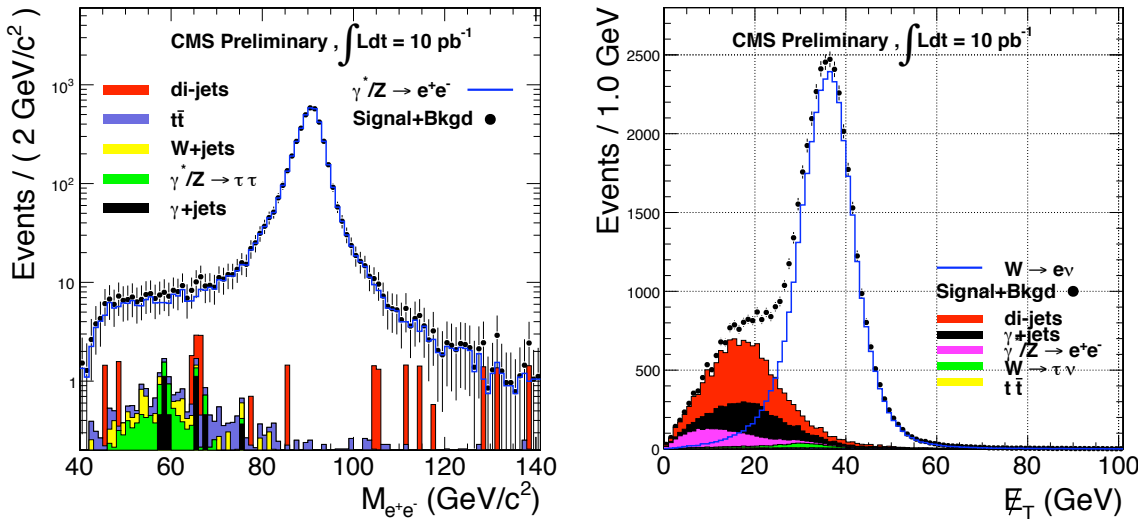


Figure 9: The M_{e_e} distribution (left) for the $\gamma/Z \rightarrow e^+e^-$ signal, and the *missing* E_T distribution (right) for the $W \rightarrow e\nu$ signal, each also showing the considered backgrounds, and each after selection cuts applied for 10 pb^{-1} of integrated luminosity.

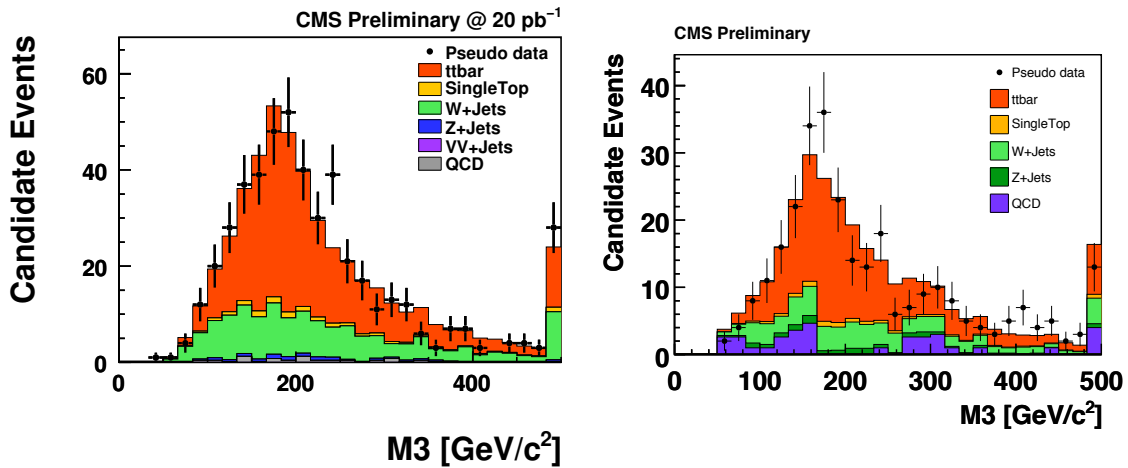


Figure 10: In the simulated search for $t\bar{t} \rightarrow \mu^+ \geq 4 \text{ jets}$ (left) and $t\bar{t} \rightarrow e^+ \geq 4 \text{ jets}$ (right) for 20 pb^{-1} at 10 TeV, the invariant mass of the three jets with the highest vectorially *summed* E_T for the final selection. The pseudodata are a set of randomly selected, Poisson fluctuated, simulated events drawn from each process.

A brief sketch of the (extensively developed) plan for first physics analysis as the LHC increases intensity and energy is as follows. As emphasized above, much detector commissioning has already been achieved in the CRAFT run, supplementing earlier major test beam campaigns. After the LHC restart, splash events will be used for quick checks, and then

first collisions at injection energy will be a long-anticipated next step, providing not only particles emerging from the interaction point for synchronization and alignment refinement, but also the first copious supplies of hadrons, including resonances which can demonstrate the performance of the ECAL (π^0 , η) and tracking (K_s^0 , Λ^0). (This in fact occurred between the time of the talk and this writeup.)

In the next step (scheduled for spring 2010), with $10\text{-}20\text{ pb}^{-1}$ at 7 TeV total LHC beam energy, the trigger will continue to be commissioned and refined, and more complex “physics commissioning” can proceed to “rediscover” the heavier objects of the Standard Model, and as luminosity becomes sufficient, measure rates of yet higher- p_T jets, leptons, and W, Z, and top (while of course being watchful for any extraordinary new signatures). At 7 TeV total energy, the approximate yields for detection of particles per inverse pb are 3000 W decays to electron or muon; 300 Z decays to ee or $\mu\mu$, and 5 ttbar events with at least one e or μ from the W decays. Some examples simulated at 10 TeV are in Figs. 9 and 10. Improved understanding of physics objects will result in, e.g., the jet energy scale from W to jj, and extensive use (and understanding) of b-tagging. In parallel there will be extensive efforts to measure and understand backgrounds to SUSY and Higgs searches. As data accumulates, one can look for excesses such as those from SUSY and Z^0 resonances. The search for new physics can then be significantly expanded as the luminosity increases and in particular as the energy increases towards 10 TeV total beam energy. With about 1000 pb^{-1} , CMS can enter the Higgs discovery era.

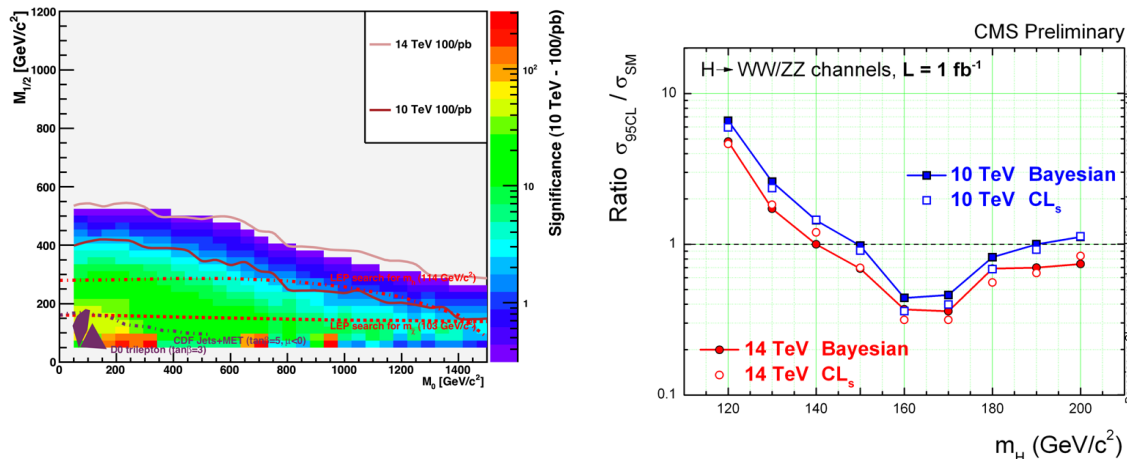


Figure 11: Left: Comparison of 5 sigma reach curves for 100 pb^{-1} of data collected at 14 and 10 TeV. The significance in each point at 10 TeV is shown as a color map. Right: Projected exclusion limits for the SM Higgs boson at 14 and 10 TeV center of mass energies. Only $H \rightarrow WW^* \rightarrow 2l 2\nu$ and $H \rightarrow ZZ^* \rightarrow 4l$ channels are included.

Fig. 11 can serve to remind the audience of some discovery possibilities which have been simulated. The presentation contained a number of other plots, for example those showing the dramatic improvement in jet transverse energy and missing transverse energy resolutions in the reconstruction of simulated high- p_T events if a complete “particle flow” algorithm is applied; it is of course of great interest to see how such sophisticated algorithms perform in real LHC high- p_T data. Other plots pointed to the early potential to see high-mass dilepton resonances (Z^0) if they exist just beyond current limits.

10. Conclusions

During the autumn 2008 LHC beam and cosmic runs, the sub-detectors, online, offline, computing, and analysis systems all performed well. The ensuing shutdown included broad maintenance activities and a programme of carefully selected repairs interleaved with installation of the preshower detector. Much very useful information has been extracted from the 2008 CRAFT data, summarized in 23 papers [3] posted on the arxiv and submitted to JINST.

In autumn 2009, the software, computing systems, and analysis systems were being exercised with newly taken cosmic data and by generating, distributing, and analysing 200M events to update 10 TeV “physics analyses” (and subsequently 7 TeV) using software release intended for data taking. The experiment was being closed and another long long magnet-on cosmic run (CRAFT 2009) was underway to put CMS once again into “beam-ready” state. CMS was (again) ready, and eager, for LHC beam. Indeed, circulating beams were quickly followed by collisions at injection energy and up to 2.36 TeV total beam energy – that story is unfolding on the CMS web site [1] will be told in the proceedings of the winter 2010 conferences.

Acknowledgements

This paper is dedicated to all those who have worked on the making the CMS experiment a reality. Acknowledgments of support of all CMS are given in Refs. [1 and 3].

References

- [1] CMS Collaboration, “The CMS experiment at the CERN LHC”, JINST 0803:S08004, 2008
- [2] CMS Public and Outreach web pages, <http://cms.cern.ch>
- [3] CMS Collaboration, “Commissioning of the CMS Experiment and the Cosmic Run at Four Tesla”, <http://arxiv.org/abs/0911.4845> and references therein. Accepted by JINST, March/April 2010.