

# Status of the CMS Experiment

*Robert Cousins on behalf of the CMS Collaboration*

Dept. of Physics and Astronomy, Univ. of California, Los Angeles, California 90095, USA

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

## 1 Introduction

The Compact Muon Solenoid (CMS) Collaboration consists of over 2500 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 became a working experiment in September 2008. In this talk I review some of the highlights of 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; the first maintenance cycle during the 2008-2009 shutdown; and a second cosmics run underway at the time of this conference. I 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 reminder of the exciting possibilities in the months and years ahead of us.

For more information, the CMS main public page [1] 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 [2], 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




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 muon system, leaves deposits in the HCAL and ECAL, and crosses the silicon strip and pixel tracking systems.

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 ( $\text{PbWO}_4$ ) crystals ( $25.8 X_0$  long in the barrel,  $24.7 X_0$  long in the endcaps), 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  $\Delta E/E \approx 100\%/\sqrt{E} \oplus 5\%$ . The CASTOR ( $5.3 < |\eta| < 6.6$ ) and Zero Degree ( $|\eta| > 8.3$ ) calorimeters (made of quartz fibers/plates embedded in tungsten absorbers) cover 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 (DT), 450 Cathode Strip Chambers (CSC), and 900 Resistive Plate Chambers (RPC). 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. 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 the most interesting events in about  $1 \mu\text{s}$ , 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 100 Hz (initial DAQ system is 50 kHz), before data storage. On the Worldwide LHC Computing GRID (WLCG), some 50k cores dedicated to CMS run more than 2M lines of source code.

### 3 Installation of CMS components 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 holds 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 200 km of cables and optical fibres (about 6000 cables). Additionally, more than 20 km 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 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  $ix$  and  $iy$  are indices of the crystals in the horizontal ( $x$ ) and vertical ( $y$ ) coordinates, respectively. (b) Occupancy of the barrel, where  $i\eta$  and  $i\phi$  are indices of the crystals in the  $\eta$ - $\phi$  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 operations 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 150 m upstream of CMS, which yielded sprays (so-called “beam splashes”) containing  $\mathcal{O}(10^5)$  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 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. For the endcap CSC’s, beam halo muons were also valuable for making up for the lack of horizontal cosmics; by using them, local alignment precision of  $270\ \mu\text{m}$  was achieved within each ring of CSCs.

## 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 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.

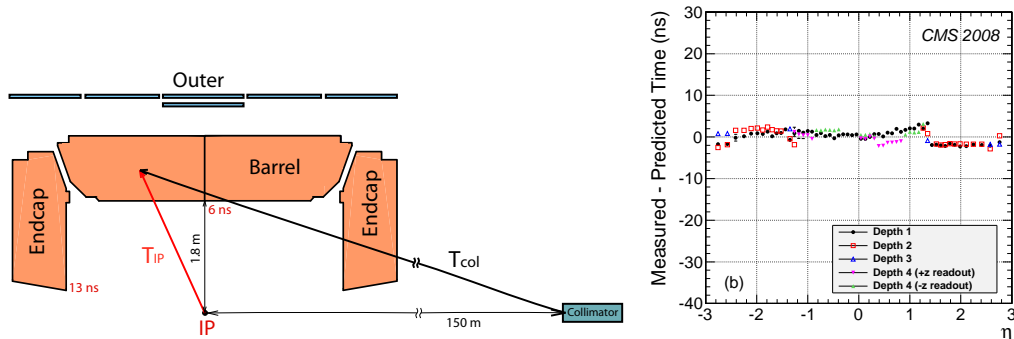


Figure 3: Left: A schematic view of the geometry of beam splash events. This geometry was used to predict the timing of energy deposits and thereby highlight channels in the HCAL that required synchronization. Right: Difference between measured and predicted time as a function of  $\eta$  from LHC beam 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-triggered events with the solenoid at its operating central field of 3.8 T and with at least the silicon strip tracker and drift tube muon system operating at nominal conditions. These data were processed offline 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  $10 \text{ pb}^{-1}$  of LHC beam!), detector calibration, noise cleanup algorithms, as well as 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 submitted to 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; the 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, necessitating a major revision in the ensuing shutdown. 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.

The data from cosmic muons was 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 synchronising 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




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.

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 plot), this was possible only to a limited extent. One method for fine-tuning of the timing delays 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 Level-1 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-tune 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 240 MB/s, and the total volume transferred was about 600 TB.




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

## 7 Data analysis results from CRAFT 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  $\mu\text{m}$  in the barrel and 3–14  $\mu\text{m}$  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.

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 overestimated 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; it is 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.

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  $\mu\text{m}$  was achieved along the higher precision coordinate direction (approximately  $\phi$ ) for the first three DT stations as estimated by a cross-check of extrapolating muon segments from one detector to the next. A few lower-level resolution plots for CSC and DT position measurements were in my talk, and many more are in the CRAFT papers [3].

Figure 8 shows a higher-level result made possible by all this work: muon momentum




Figure 6: CMS magnetic field map.




Figure 7: Left: Muon stopping power measured in  $\text{PbWO}_4$  (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 at  $\eta = 0$ ; (bottom) ratio of data and Monte Carlo simulation predictions (arbitrary scale) for the bottom half of the calorimeter.

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.

## 8 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), at the time of this talk in August 2009, 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; maintenance and (small) repairs involving many sub-systems; and the major revision of the tracker cooling plant to re-



duce leaks and improve reliability. 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 re-reconstructions of the CRAFT data, large Monte Carlo simulation productions were produced and reconstructed (e.g., about 250M fully simulated events at 10 GeV begun at the end of 2008 were completed in February 2009). At the time of the talk, another large MC production was about to begin with CMSSW 3 (since completed).

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 (such as track-corrected jets and missing  $E_T$ , and 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 WLCG infrastructure and the software 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 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




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.

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.

Already two years ago, Tag and Probe (T&P) methods (such as those used in Tevatron experiments) were demonstrated in CMS simulations assuming 14 TeV LHC energy, to identify a physics object in an unbiased way in order to study efficiencies. E.g., in  $Z^0 \rightarrow e^+e^-$  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  $94.36 \pm 0.24$ , compared to MC truth of  $94.63 \pm 0.24$  for  $10 \text{ pb}^{-1}$  at 14 TeV.

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




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




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

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 ( $\pi^0$ ) and tracking ( $K_S^0$ ,  $\Lambda^0$ ). (This in fact occurred [1].)

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  $t\bar{t}$  events with at least one  $e$  or  $\mu$  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  $W \rightarrow jj$ , 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  $Z'$  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  $1000 \text{ pb}^{-1}$ , CMS can enter the Higgs discovery era.

Fig. 11 can serve to remind the audience of some discovery possibilities which have been simulated. My talk 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 ( $Z'$ ) if they exist just beyond current limits.




Figure 11: Left: Comparison of 5-sigma reach curves for search for SUSY in  $\geq 3\text{jets} + \cancel{E}_T$  channel for  $100 \text{ pb}^{-1}$  of data collected at 14 and 10 TeV; the color maps gives the significance in each point at 10 TeV. Right: Projected exclusion limits for the SM Higgs boson at 14 and 10 TeV center of mass energies, using only  $H \rightarrow \text{WW}^{(*)} \rightarrow 2\ell 2\nu$  and  $H \rightarrow \text{ZZ}^{(*)} \rightarrow 4\ell$  channels.

## 10 Summary

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 maintenance activities and 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].

At the time of the talk, 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 the software release intended for data taking. The experiment was being closed and another long magnet-on cosmic run (CRAFT 2009) was underway to put CMS once again into “beam-ready” state. CMS was (again) be ready, and eager, for LHC beam. Indeed, after the talk, 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].

## Acknowledgments

I thank the U.S. Dept. of Energy and the National Science Foundation for their support of my work on CMS. Acknowledgments of support of all CMS are given in Refs. [2, 3].

## References

- [1] CMS Public and Outreach web pages, <http://cms.cern.ch>. This includes a link to the real-time CMS “e-commentary” during the runs following this talk.
- [2] CMS Collaboration, “The CMS experiment at the CERN LHC”, JINST 0803:S08004, 2008.
- [3] CMS Collaboration, “Commissioning of the CMS Experiment and the Cosmic Run at Four Tesla,” arXiv:0911.4845v1, and 22 other papers posted at arxiv.org, submitted to JINST.

## Discussion

**Joel Feltesse (IRFU, CEA):** Now that your detector has become reality, and that you better know the performances, have you identified a few physics processes where you are in a better or worse place than ATLAS? Of course the previous speaker (Kerstin Jon-And) can also comment about it.

**Answer:** I think that the detectors are fairly well matched in capabilities, each being well-designed for the the broad physics program. Now the key for each collaboration going forward is to understand well the calibration, alignment, and performance of the detector that they have built. That will determine which discoveries are seen or missed (and when) more than anything else. Both collaborations have a very good start on this by using test beams, beam splash events, and cosmics, as we have described.

**Dimitri Denisov (FNAL):** With the 2010 data set consisting of 7 TeV and 10 TeV samples: What is the strategy of combining these two samples for physics studies? Or is the 7 TeV sample expected to be mainly used for physics commissioning?

**Answer:** Of course, the first thing will be to analyze each set separately, so that a consumer can judge the reasonableness of any answer coming out of a combination. When I think about combining, I conclude that transparency of the combination may put a constraint on how fancy one should get in combinations in our first data sets. But if, as we hope, most of the running is at the higher energy (10 TeV is an upper bound, but it could be lower), then one might see hints at 7 TeV but the more compelling evidence at the higher energy without needing to combine them formally.

**Bennie Ward (Baylor University):** The previous speaker quoted ATLAS efficiency (data taking eff.) at 83% You showed CMS's at 70%. What is the main reason for the difference in the two?

**Answer:** I would first mention that in quoting an efficiency there are choices to be made regarding which time interval to average over, which types of one-off failures (such as CERN site-wide power cut) to include, etc., if you want to make a relevant inference about future running efficiency. So I do not know if the numbers are directly comparable. In our case, the number is a rather inclusive number for several weeks. We have days (when we are not interrupting for tests) in which the efficiency is over 90%.

**Sakue Yamada (KEK):** You mentioned about PFA and said MC gives much better result. Did you try to find the reason?

**Answer:** A lot of it goes back to design choices: the CMS design has a high magnetic field and emphasizes excellent tracking and high-precision extremely-fined-grained electromagnetic calorimetry, which in turn forced some compromises (most obviously non-compensation) in CMS jet calorimetry. So, as you saw, the ATLAS calorimeter has better intrinsic jet resolution than that of CMS. The particle flow algorithm takes advantage of the superb CMS resolution in the tracker and ECAL, and substantially reduces the dependence on the HCAL for measuring hadrons (as neutral hadrons ac-

count for only about 10% of the energy of a jet), except at higher energy where the HCAL resolution is better and tracking resolution worsens. The result is a much smaller correction factor from raw (observed) jet energy to estimated (corrected) jet energy.