CERN-RRB-2011-067
10 October 2011

ATLAS Progress Report for the October 2011 RRB

1. Introduction and Collaboration Composition

The year 2011 has continued to witness outstanding performance of the LHC and of ATLAS. Peak luminosity has grown consistently throughout the year, now exceeding $3x10^{33}$ $cm^{-2}s^{-1}$. The integrated luminosity recorded by ATLAS had already exceeded the 1 fb^{-1} goal for the year by the time of the EPS conference in July, and is now well in excess of 4 fb^{-1}. The ATLAS detector continues to perform extremely well. The average data-taking efficiency for 2011 is 94%, demonstrating the excellent status and performance of all detector components, as well as of the trigger and data-acquisition systems. The fraction of non-operational detector channels ranges from a few per mil to ~4% depending upon sub-detector, and the fraction of good-quality data used for analysis is 90-98% depending upon the analysis.

The last six months have seen a hundred-fold increase in the pp data statistics. A wide spectrum of updated search results were presented at the summer conferences, including, with CMS, the first – wide – exclusion limits on the Standard Model Higgs boson mass from the LHC. The 2011 and 2012 data together look very likely to reveal whether a Higgs boson exists or not in its pure Standard Model guise. Looking beyond searches, a large number of detailed precision measurements have been finalised and published with 2010 data, and signals for lower cross-section processes such as diboson and single-top production have been established in the 2011 data. These results are documented in a total to date of 76 papers published or submitted for publication, and 246 physics notes presented at international conferences. These achievements also demonstrate that the Collaboration is able to deliver results very quickly, thanks also to the smooth and effective operation of the computing grid infrastructure based on the WLCG backbone.

Excellent LHC performance, including high peak luminosity, large integrated luminosity, and high numbers of interactions per crossing, all exceeding expectations for 2011, creates challenges for the ATLAS detector performance and computing resources.

ATLAS has a full and well-established programme for the 2013-2014 shutdown, including installation of a new inner-detector cooling plant and a redundant magnet refrigerator compressor, replacement of the liquid-argon and tile calorimeter low voltage power supplies, and much additional consolidation work. Furthermore, the so-called “Insertable B-layer” (IBL) will be installed in the Pixel detector. Substantial continued progress has also been made on R&D activities and planning for LHC luminosity improvements, within the overall CERN schedule. Two dedicated task forces have been working to finalize the plans for the two following shutdowns (foreseen approximately in 2018 and 2022). An additional task force has been established to study the physics capabilities of detector improvements. ATLAS phase I plans will be summarized in a Letter of Intent (LoI) prior to the April 2012 RRB meeting, and an LoI for Phase II will be presented about one year later.

The ATLAS Collaboration consists today of 173 institutions from 38 countries with approximately 3000 active scientists, of which 1800 with Ph.D. (hence considered for M&O cost-sharing) and about 1000 students. Since the April 2011 RRB, University of Regina, Canada has withdrawn from the collaboration. At its October 2011 meeting, the ATLAS Collaboration Board (CB) will be asked to admit Kyushu University, Japan and University of Warwick, UK as new member institutions, which would bring the total number of institutions to 175. In addition, Federal University of São João del Rei (UFSJ) and Federal University of Juiz de Fora (UFJF) have become member institutes of the Brazilian cluster. The expansion of this cluster does not change the composition of the Collaboration Board, as each cluster remains one member institution and continues to have one vote in the CB.

2. Maintenance and Operation of the Magnet System

The ATLAS superconducting magnet system comprises the Central Solenoid (CS), the Barrel Toroid (BT), two End-Cap Toroids (ECT), and their common services.

The magnet system has operated with nearly 100% availability for data-taking with stable beams from 2009 to 2011. Some difficulties in operation have been encountered, including power converter unit failures, clogged cold circuit filters, and current lead problems. Failures of power converter units arising from problems in the units' current meter caused a few slow dumps of the toroid magnets in the last year. The power converter manufacturer is expected to provide a reliable solution. A hidden defect in the heat exchanger of the shield refrigerator cold circuit allowed water into the compressor oil, eventually leading to clogging of the filters in the cold box. Installation in the circuit of a dryer in July 2010 eliminated further problems. A permanent dryer will be installed in this circuit and the main refrigeration circuit during the 2013-2014 shutdown. Hot spots in the aluminum-copper bus bars were repaired. Further repairs will be made during the end-of-year technical stop. Additional major consolidation work on the magnet services is planned for the 2013-2014 LHC shutdown. It includes installation of a redundant main refrigerator compressor and adoption of booster compressors, which will preclude a possible year-long ATLAS downtime in case of main compressor failure. Consolidation work is also foreseen for the electrical, vacuum, and controls services.

3. Maintenance and Operation of the Inner Detector

The Inner Detector (ID) combines three concentric sub-systems, the Pixel detector, the silicon strip detectors (SCT), and the transition radiation straw tracker (TRT), listed from inside to outside.

During the 2011 data-taking period the Inner Detector continued to take data with very high efficiency. 96.4% of Pixel modules, 99.2% of SCT modules, and 97.5% of TRT channels are operational. Data acquisition efficiencies while LHC beams are stable are 99.9%, 99.2%, and 100% for the Pixels, SCT, and TRT, respectively.
The ID tracking performance has proven robust against the increased pileup rate coming with high luminosity operation of the LHC. An optimization of the reconstruction software has been performed to maintain computing time at an acceptable level

without any significant compromise in performance. An updated alignment resulted in a significant improvement of momentum resolution, especially in the forward region. A new clustering algorithm for the pixel detector shows a significant improvement of impact parameter resolution and an enhancement of track reconstruction efficiency in dense jets. These improvements are now in operation at Tier-0 and were used for the 2011 data reprocessing and simulation campaigns.

The evaporative cooling plant has been operating in 2011 with close to 100% efficiency, and has always been available during collisions. A design for a thermo-siphon system to replace the compressors in 2013-2014 is well advanced and the purchase of the large items is underway.

There have been no additional failures of the thermal enclosure heater pads since three of eight barrel pads failed during June/July 2009. These pads are large-area copper-kapton heaters that should ensure thermal neutrality of the silicon volumes compared to the TRT, and of the ID volume compared to the outside. Although a thermal configuration that does not lead to over-cooling of the TRT has been established, there are concerns for the longer term, after the SCT modules have been exposed to significant radiation and need to operate at reduced temperature. Investigations of C_2F_6 and C_3F_8 as more effective coolant mixtures are progressing well and are yielding encouraging results. A new control card design for the heater pad power supplies, to reduce the chance of further failures, has been manufactured and tested. During the end-of-year technical stop, three (of five) new crates of control cards for the large area heater pads will be installed.

Both the SCT and Pixel detector are now showing evidence of the effects of radiation. The increase in leakage currents and decrease in depletion voltages are in line with expectations based on previous measurements and calculations that include the influence of integrated luminosity and operational temperature. To mitigate the long term effects of radiation damage, the evaporative cooling system will now be in constant operation, including shutdown periods.

The off-detector optical transmitter plug-ins (TXs) used in both the Pixel and SCT detectors that were replaced in 2009 ran smoothly for a number of months; however, they began to fail in April 2010, with a maximum rate of over twenty per week. Resulting inefficiencies in data-taking have remained small because failed Pixel plug-ins are replaced every few days, while most of SCT modules can be recovered by using built-in redundancy. The cause of the failures has been identified as humidity around the VCSEL. A new production of additional “old style” replacement plug-ins (enough to replace every TX in the system) is well under way, and a rolling replacement program of failed units is in place. These arrays have been tested to be more robust against the ingress of humidity. So far, nearly 50% of all of the units have been replaced, and complete replacement will finish by the end of the year. In addition, work has started with a new company to make optical sub-assemblies that will be part of a new backwards-compatible TX plug-in. The Pixel on-detector arrays are of the same type as the off-detector units, leading to possible future lifetime concerns despite the fact that they operate at a lower rate. A task force has studied this problem coupled with the increasing number of failures in the detector due to other sources.

A Pixel task force, convened to study the possibility and implications of on-detector VCSEL failures, has recommended a rebuild of the Pixel services and to move the on-detector transmitters to a more accessible location. This decision is a pragmatic response to the existing situation while investigations of VCSEL failure modes continue. The design and preparation has started and the decision on whether to use these new services will be taken in mid-2012. A large amount of activity is ongoing in the Pixel community to prepare the integration of the fourth, innermost, layer (IBL) with a re-organization of the resources to match the needs in terms of the consolidation of the existing system and the integration of the IBL.

In 2011 the SCT has, in response to 50 ns bunch spacing in the LHC, changed its readout mode to demand a hit in the middle of the three time bins that are read out (so called X1X mode). In this mode of operation, contamination by out-of-time hits is eliminated. A further change to demand the absence of a hit in the first time bin (so-called 01X mode) is expected when the LHC starts 25ns operation. In both the SCT and Pixels, the hit occupancies remain at a manageable level.

With the data recorded in 2011, the TRT group is continuing to improve the detector performance and the contribution of the TRT to the track reconstruction and particle identification capabilities. The detector is now delivering hits to the track reconstruction machinery with a single hit resolution of 120 μm (in some regions even better than 100 μm), which is better than the expectation of about 130 μm. The particle identification capabilities of the TRT are also now well understood, achieving a hadron rejection factor of 20 at 90% electron efficiency, as expected from the specifications.

A problem with the crates of the TRT low-voltage power supply system has required interventions and repairs in the ATLAS cavern. This problem is suspected to be caused by a problematic control connection. The company that produces the crates suggests replacement of all these connections by gold-plated connections, which can only be accomplished during the long 2013/2014 shutdown. During the coming end-of-year technical stop, a few connections will be modified in order to study the functioning and reliability of the fix.

Failures of single TRT high-voltage power supplies have occurred on rare occasions. To avoid possible further problems, a test bench was established to investigate the reliability of the supplies using spare modules. To be prepared in case failures occur at a more frequent rate, additional spare power supplies were ordered and will be available by the end of the year 2011.

Further improvements to track reconstruction algorithms are in active development in order to maintain physics and technical performance with expected higher luminosity and event pileup rate during 2012 data-taking. A seeded Inner Detector reconstruction setup is in preparation for early 2012, which will allow re-initiating the ID reconstruction with different parameter settings seeded by combined objects, such as electromagnetic clusters or forward muons. Use of optimized parameters for specific kinds of tracks will improve the efficiency of the standard reconstruction running for the bulk of the event. Further optimization of track reconstruction minimizing the impact on physics object is already ongoing, in order to find the best working point for next year LHC running, assuming nominal luminosity and developing tunings for both 25 ns and 50 ns bunch spacing.

4. Maintenance and Operation of the Calorimeters

The calorimeter systems include a liquid argon (LAr) electromagnetic calorimeter, a barrel and two extended barrel Tile hadronic calorimeters, end-cap liquid argon hadronic calorimeters (HEC), and liquid argon forward calorimeters (FCal).

All calorimeter systems have performed well throughout data-taking. The precision of calibration is nearing the design values. For instance, the LAr energy calibration precision is ~1% and continues to be improved as more data are analyzed.

The stability of operation and performance of the LAr calorimeter has improved in 2011. After the exchange of some faulty or suspect Front End Boards (FEB) in January 2011, no optical transmitters (OTx) have failed. Recent measurements of the OTx light spectrum show that the width is stable; whereas, failed OTx had abnormal spectra. Decision not to build replacement OTx will probably be taken in December. One controller board failed in April 2011, probably because of a fuse blown shortly after a power cut. This failure led to loss of control and timing signals on eight boards in one barrel front end crate: six FEBs (four connected to the middle layer and two to the back layer), one calibration board, and one tower builder board. A delicate intervention was performed during the July technical stop during which the four middle FEBs and the calibration board were recovered by bringing timing signals through an alternate path. The front-end low-voltage power supplies (LVPS) are still a concern with five units out of 58 being operated without redundancy on one voltage line, but without any data loss. The new set of 70 LVPS is under long term test. A few of these new LVPS will most probably be installed on the detector during the end-of-year technical stop. The high voltage system has stabilized at a low number of trips (typically one or two per fill with an impact of a few minutes) now that luminosity is no longer changing from fill-to-fill, but improvements to the system are still ongoing. Currently a total of 0.02% of LAr cells is dead at the detector level.

In the Tile calorimeter, 3.8% of cells have failed during the 2011 data-taking period due to three front-end drawer and 6 LVPS failures (out of 256 units). The five units of the replacement LVPS design that were installed during the 2010 end-of-year technical stop are operating successfully, showing a significant improvement in noise behavior and an absence of trips, while the old units show a trip rate of ~1 trip/pb^{-1}. Production of replacement LVPS commenced in summer 2011. Forty units will be installed during the end-of-year technical stop, and all 256 units will be replaced during the long shutdown of 2013/2014. The fraction of good quality data from the Tile calorimeter has been 99.7% in 2011.

5. Maintenance and Operation of the Muon Detectors

The Muon Spectrometer is instrumented with precision chambers for momentum measurement (Monitored Drift Tube chambers, MDTs, and for a small high-radiation forward area Cathode Strip Chambers, CSCs) and with fast chambers for triggering (Resistive Plate Chambers, RPCs, in the barrel, and Thin Gap Chambers, TGCs, in the end-caps).

During the whole year 2011 the muon system has been operating with very good efficiency, and the fraction of live channels is at the moment 99.7% for MDT, 97.7% for CSC, greater than 97.0% for RPC, and 98.1% for TGC. Since the start of the 2011 run, the CSC system has lost one plane (out of 128) due to a HV problem, and the TGC system has lost about 40 planes (out of about 3600) with HV problems.

The CSC readout was operated throughout the run and did not cause relevant dead time or loss of data even at the highest trigger rates. The limitations of this system are very close to the specified 75 KHz LVL1 rate.

The RPC system has deployed an automatic procedure to adjust the HV settings according to the variation of ambient temperature and pressure. This procedure enabled an increase in the average efficiency of the chambers and an increase in the voltage on the top sectors that are still affected by the higher temperature of the cavern at that location.

The performance of the Muon Spectrometer has been enhanced mainly by improving the alignment in the CSC region using data taken with the toroidal magnets turned off. In particular, the internal alignment of the four planes of each CSC chamber has been measured together with the overall chamber alignment. These quantities have been included in the global alignment enabling an alignment precision that ranges between 50 and 100 μm throughout almost all the muon spectrometer. High reconstruction efficiency and very good efficiency determination have been achieved with high statistics Z samples.

A small number of maintenance issues continue to be monitored. CAEN high voltage and low voltage power supplies require regular replacement, and Wiener VME power supplies are also showing failures. The power connectors on the 48V power supplies of the RPC system must be replaced during the end-of-year technical stop because they tend to develop high resistance that causes overheating of the cables and connectors. RPC gas inlets are fragile, and gas leaks continue to develop. Broken inlets are being replaced as needed during maintenance periods. Cracking on a few gas jumpers of some MDT chambers has been observed. Some TGCs showed problems related to HV behaviour at high luminosity.

6. Maintenance and Operation of the Forward Detectors

The ATLAS forward detectors consist of a Luminosity Cerenkov Integrating Detector (LUCID) placed around the beam pipe inside the forward shielding at ± 17 m from the Interaction Point (IP), a Zero Degree Calorimeter (ZDC) placed in the TAN absorber structure where the beams enter separate beam pipes at ±140 m from the IP, and an Absolute Luminosity for ATLAS (ALFA) system based on fiber detectors located in Roman Pots at ±240 m from the IP.

Since its final installation in January 2011, the ALFA detector was successfully commissioned and is ready to take data for the total cross section measurement foreseen using special β^*=90 m beam optics in Autumn 2011. During ALFA commissioning, a short data-taking run with β^*=90 m demonstrated the excellent capabilities of the detector to reconstruct inelastic pp interactions. LHC development and qualification of the 90m optics has been successfully completed. Beam-based alignment has been performed with standard collision optics, and will be repeated with 90m optics.

As LHC performance has progressed during 2011, instantaneous luminosity and the number of interactions per bunch crossing have exceeded the design limit for LUCID.

Consequently, detector response started to be sensitive to gain variations from the anode current drawn by its photomultiplier tubes (PMT). These variations are affecting the luminosity measurement at a few-per-mil to one-percent level. Furthermore, at present LHC pileup levels, the event counting algorithms used to determine luminosity saturate. For this reason, the Cerenkov radiator gas has been evacuated from the LUCID, such that light is produced only in the quartz window of the PMTs.

During the end-of-year technical stop, the PMT bases will be upgraded to reduce the dependence of PMT gain upon current. In addition, the readout of the Cerenkov tubes will be replaces by quartz fibers. Meanwhile, LUCID is still able to measure the luminosity with the required precision of a couple percent, although the beam current monitor (BCM) detector now provides an even higher precision.

The ZDC readout, which is optimized for heavy ion runs, is not able to cope with the 50 ns bunch spacing of the LHC proton beams in 2011. In order to maintain the features of ZDC as luminosity monitor and triggering device, the readout chain of the total energy sum from its four-plus-four module was modified to use faster electronics designed for LUCID. This operation was effective, although some further tuning is needed to suppress electronics noise. During the end-of-year technical stop, the eight ZDC total sum PMTs and their bases will be changed to eliminate dependence of PMT gain upon anode current, as in the LUCID.

7. Maintenance and Operation of the Trigger and DAQ System

The major sub-systems of the Trigger and Data Acquisition System (TDAQ) are the Level-1 Trigger (with the sub-systems calorimeter, muon, and central trigger processor (CTP)), the High Level Trigger (HLT – comprising the Level-2 trigger and the Event Filter), the Data Acquisition (DAQ), and the Detector Control System (DCS).

The complete Trigger and Data Acquisition chain continues to perform very well and has proven to be able to adapt efficiently to the changing data-taking needs of the experiment, to the fast ramp-up of the instantaneous luminosity, and to increasing pileup. Through 2011 operations to date, operational efficiency during collision periods has averaged 94%. The system operates typically at a peak Level-2 input rate of ~55 kHz, a peak Level-2 output rate of ~5 kHz, and an average Event Filter output rate of 400 Hz. Following detailed measurements of the HLT computing resource usage and trends using recent data, the HLT computing power was increased during the third quarter of 2011 from 50% of its foreseen capacity to 75% in order to match the demands of increasing pileup on HLT processing. The scope and quality of operation procedures and documentation are the subject of continual review and improvement.

The Level-1 trigger systems are performing efficiently and stably. Additional flexibility has recently been added to the Level-1 calorimeter trigger, implementing electromagnetic and hadronic isolation, and eta-dependent thresholds. Studies are ongoing to further optimize Level-1 trigger calibration parameters.

During the course of 2010 and 2011, a series of trigger selections ("menus") specifically aimed at high-p_T physics have been successively deployed in order to cope with the progressive increase over five orders of magnitude in the LHC instantaneous luminosity up to a maximum of 3 x 10^{33} $cm^{-2}s^{-1}$.The HLT has been used for the online selection of events since May 2010. The current trigger menu is optimized for luminosities up to 3.5 x 10^{33} $cm^{-2}s^{-1}$ and has recently been successfully deployed. A new menu for luminosities up to 5 x 10^{33} $cm^{-2}s^{-1}$ and beyond is in preparation. The many items in the menu have been classified to reflect ATLAS's physics priorities. Their rates are varied during the course of an LHC fill as luminosity decreases, giving full bandwidth to high priority triggers at the start of the fill and, as their rate decreases, allocating some bandwidth to lower priority triggers, thereby keeping the mean Event Filter output rate in the range 300-400 Hz and staying within the available offline computing resources.

The trigger system coped well with the LHC heavy-ion run at the end of 2010, using principally Minimum Bias triggers to select events with some limited use of selection at the Event Filter. For the 2011 heavy ion run, the trigger menu will make much more significant use of the Event Filter selection in order to cope with the expected higher luminosity ion beams. This menu is currently under test using minimum bias heavy ion data taken in 2010.

A sophisticated toolkit for predicting trigger rates from data has been developed and is used for menu development. Monte Carlo predictions have been found to be very consistent with the observed trigger rates as luminosity rises. The physics performance of the HLT algorithms is being continually assessed, and they are to date performing robustly in the increasing pileup environment.

Operational efficiency continues to be improved by closely working with the detector systems to better define and subsequently automate operational procedures. Semi-automatic start and end of run sequences have been implemented which automatically initiate the procedures (*e.g.* ramping of voltages) that certain detector systems require during machine operations, for instance during the energy ramp of the beams. More sophisticated procedures for DAQ system performance monitoring and for online and offline trigger data quality assessment monitoring have also been put in place. These procedures contribute significantly to the achieved average data-taking efficiency of 94% stated earlier. The increased use of automatic monitoring, control, and problem correction procedures has allowed TDAQ to reduce the number of shifters in the control room to two, as well as to significantly reduce the number of on-call experts required to run the system.

In addition to regular maintenance and optimisations, notably during LHC technical stops, the following significant activities have taken place since the start of 2011 data taking: (a) the new trigger based on Missing ET significance (XS), which is much more robust than the traditional Missing ET trigger to increasing pileup and which had been previously integrated in the Level-1 calorimeter trigger firmware, is now included in the trigger menu; and (b) the rolling replacement of the ROS system PCs is ongoing and will be completed by the end of the year.

The plan for procurement of the outstanding HLT capacity will continue to be adapted to the evolution of the LHC schedule and performance, as well as to the evolving needs of the experiment. The plan will be reviewed at the end of the 2011 data-taking period. The rolling replacement of computing hardware that has reached the end of warranty continues.

8. Operations

ATLAS continued to operate successfully throughout the first half of 2011, despite the rapid evolution of LHC beam conditions. In particular, the peak luminosity has increased to approximately $3 \times 10^{33}\ cm^{-2}s^{-1}$, and the pileup at the beginning of a fill has increased by a factor more than three from an average number of interactions per crossing $\langle\mu\rangle=5$ to $\langle\mu\rangle=16$. Both, the rising peak luminosity and the rising pileup required frequent redefinition of operating procedures and trigger selections. Efficient operation has also required continual maintenance of general infrastructure, such as electricity, cooling and ventilation, of individual detector elements, and of control, trigger and data acquisition systems. Constant and careful maintenance and development work has enabled ATLAS to achieve a global data-taking efficiency of 94%.

Over the course of 2010 and 2011, steps have been taken to streamline Operation Tasks (OTs) in order to reduce the overall effort required, and to enable more tasks to be performed at remote sites. A first reduction of the shift crew in the control room took place at the beginning of 2011, and further reductions were made in May and at the beginning of October. ATLAS now operates with 10 shifters, compared to 20 initially. Operation Tasks include all activities essential to the operation of ATLAS, from central shifts and on-call tasks at Point-1 to the computing and data preparation tasks, some of which can be executed remotely. Operation, in this broad sense, requires very substantial effort (approximately 650 FTE, plus shifts). Responsibility for Operation Tasks is shared among the Institutions in proportion to their number of authors. Effort on operation of the WLCG is in addition.

9. Computing and Software

The collaboration-wide distributed computing infrastructure is fully embedded into the framework of the WLCG, of which ATLAS is a very active partner.

The entire ATLAS and WLCG computing and software chain has performed very well throughout 2011 data-taking. Data is reconstructed by the Tier-0 centre and distributed worldwide within a couple of days of data-taking. The "prompt calibration loop" operates routinely: intensive calibration work is performed at the CAF (Calibration and Alignment Facility) at CERN within a few hours after the data are collected using small specialized samples. Derived calibration outputs are used in the Tier-0 reconstruction of the physics data shortly afterwards. The global Tier-1 computing infrastructure and the reconstruction software have been used several times in 2011 to reprocess detector and simulated data. Simulation production continues using available CPU resources at Tier-1 and Tier-2 centres, with an average of approximately fifty thousand jobs running concurrently and with approximately 1.5 billion events simulated in 2011. Analysis on the WLCG is in full swing, with several hundreds of people routinely submitting jobs. More than a thousand different ATLAS members have performed analysis on the grid. Large-scale programmes of distributed data analysis at Tier-2 centres have led to the physics results presented at conferences and in publications. Computing and software shifts are operated at the Point-1 control room, and worldwide, to support production operations and user analysis. ATLAS has made efficient use of the computing resources made available at all WLCG sites.

The ATLAS data distribution model was changed at the beginning of the 2011 run in order to optimize the usage of our resources, deal with larger than expected pileup, and accommodate higher trigger rates if possible. The ESD (event summary data) format is no longer used directly for physics analysis; consequently, only a fraction of the ESD is made available in rolling buffers at the Tier-1 centres. One copy of the smaller RAW data is distributed, spread over all the Tier-1 disks; however, the RAW is now compressed, saving about 60% of disk space that it would otherwise occupy. Users are satisfied with this evolution away from analysis with ESD. Physics analysis is being performed using AOD (analysis object data) and D3PD (derived physics data; Ntuples). These data types are widely distributed over disks at Tier-1 sites and dynamically over Tier-2 disks.

A new software release (r17) was prepared for data reprocessing. It implements improved calibration derived from the data itself, and it implements many improvements to manage increased pileup. For instance, clustering of hit pixels is completely reworked in order to better separate nearby tracks. The CPU penalty caused by more sophisticated algorithms has been compensated by many improvements in execution time. Nevertheless, the increase in CPU and disk resource requirements due to pileup is still sizeable. Further, improvements of a number of reconstruction algorithms, both purely technical improvements and improvements in physics performance, are still in development. Release 17 is now in routine use at Tier-0, and it has been used at Tier-1 to reprocess all of the 2011 data recorded to date.

While full simulation continues to be the work horse for most physics analyses, ATLAS fast simulation has now matured to the point that its use is increasing and that it has been used to produce results. A completely reworked simulation framework allowing better integration between full and fast simulation is being developed to be in production use next year.

Although several novice newcomers have joined the software project as their technical contribution for authorship qualification, the availability of adequate expertise to maintain ATLAS computing and software remains a concern.

The present LHC machine schedule, with long data-taking period through 2012 with ongoing increases in pileup, will require the full computing resources requested, in particular for 2012. Timely availability of these resources will be critical to prompt ATLAS physics productivity, and to the ability of everyone in the Collaboration to access the data easily and quickly.

The core computing infrastructure and services tasks, defined as M&O Category A, continue to play a crucial role for the smooth operation of the full software and computing chain. They enable ATLAS to exploit the large investments of computing resources made worldwide by the WLCG collaboration partners.

10. Data Collection and Physics Output

The performance of the LHC and ATLAS during the 2011 proton-proton data-taking, now approaching completion, has been outstanding. At the end of July the LHC reached a peak luminosity of $2x10^{33}$ $cm^{-2}s^{-1}$ (see Fig. 1), ten times the 2010 peak, and increased further to $3.3x10^{33}$ $cm^{-2}s^{-1}$ in September. The integrated luminosity recorded by ATLAS is now in excess of 4 fb^{-1}, four times the goal for the year, with excellent prospects of reaching 5 fb^{-1} by the end of the 2011 run. The 2011 data has been taken in conditions rather different from 2010, in that the number of collisions per bunch crossing (pileup) has been high, up to 20 interactions per crossing, close to the maximum expected at the full LHC design luminosity.

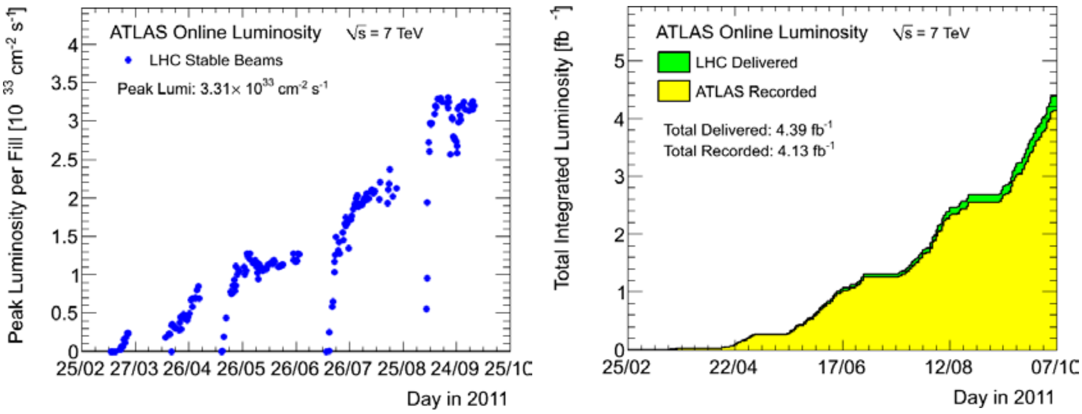

Figure 1: Peak (left) and integrated (right) luminosity vs. time in 2011 proton-proton data taking. (Snapshot of 7 October 2011.)

The "prompt calibration" scheme, whereby the bulk physics data from Point-1 are processed within a couple of days after they are collected, has continued to operate smoothly in 2011, ensuring – together with the stable detector performance – that publication-quality data is available for analysis a few days after it is taken. A major reprocessing campaign has recently been completed on the Grid using the new software release r17, enhancing the data quality in several respects. Intense activity is ongoing at present, to improve further the understanding of physics object performance, to allow more benefit still to be reaped from the reprocessed data.

During the year, physics analyses have been a mixture of studies on 2010 data, aimed at detailed precision measurements, and fast studies of the incoming 2011 data, focusing both on established low cross-section processes such as diboson and top production, and, crucially, on searches for new physics. Highlights of the studies on 2010 data submitted or approaching publication are: the measurement of the W and Z cross-sections inclusively and differentially vs. p_T, and in association with inclusive and b-flavoured jets; measurements of the inclusive lepton, photon and diphoton cross-sections; measurements of inclusive jet, dijet and multi-jet cross-sections up to dijet masses of 4 TeV; measurements of the inclusive b-jet cross-section; measurements of jet properties such as profiles and fragmentation; as well as a first measurement of the proton-proton inelastic cross-section at 7 TeV centre-of-mass energy. The 2010 heavy ion data have been used to measure the centrality dependencies of particle flow, charged particle properties, W, Z and J/ψ production, and to probe further the characteristics of jets in lead-lead collisions. Employing 2011 pp data, a number of preliminary results on top production, mass and properties have been presented, including a more than 7 standard deviation observation of electroweak single-top production. Diboson signals have been established and measured for the complete set of WW, Wγ, Zγ, WZ and ZZ production (see Fig. 2).

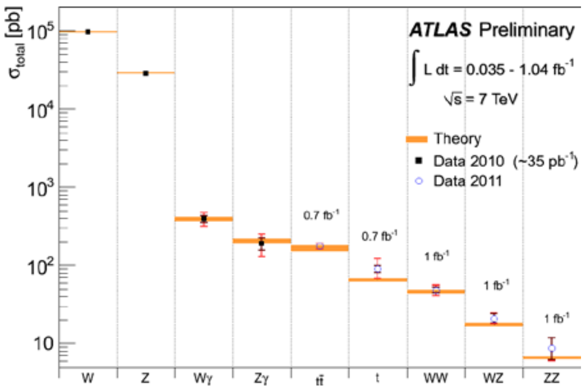

Figure 2: Summary of cross-section measurements of various processes .

Building on the understanding of the physics processes and detector evidenced by the pp measurements, an extensive range of searches for new physics have been carried out, far too many to list. In the beyond-the-Standard Model sector much stricter constraints have been placed on supersymmetry, including on stop and sbottom production.

Much excitement was evident at the main summer conferences, in Grenoble and Mumbai, prompted by the rapidly increasing sensitivity of ATLAS and CMS to the production and decay of the Standard Model Higgs over a wide mass range (see Fig. 3). ATLAS alone excludes Standard Model Higgs production at 95% confidence level (CL) over the mass ranges 146-230, 256-282 and 296-459 GeV.

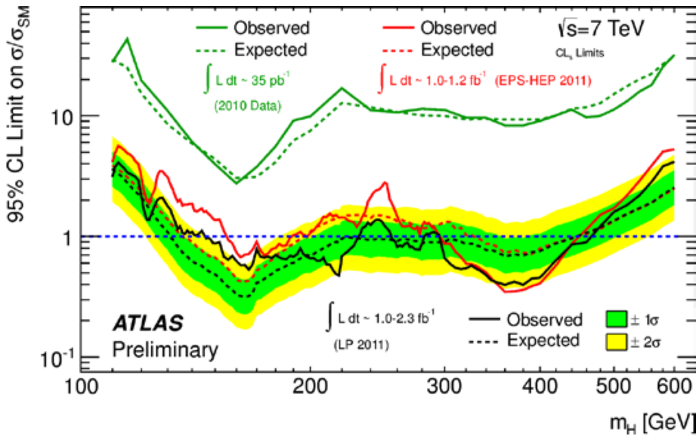

Figure 3: Constraints, at 95% CL, on the cross-section for Standard Model Higgs boson production as a function of its mass. The solid black line shows the current constraint: cross-sections above the line are excluded at 95% CL or more. The green/yellow bands show the range of expected exclusion in the absence of Higgs production. The red and green solid lines show earlier limits, as of the EPS-HEP conference in July, and with only 2010 data, respectively.

At the time of writing, a total of 77 physics papers using collision data have been submitted or published – 45 since the last RRB in April, and 91 conference notes were completed in the same interval.

11. Updates on Due Construction Contributions

ATLAS is grateful to the Funding Agencies for their continuous support over nearly two decades. At present, the Collaboration still faces an income deficit of 2.8 MCHF in the total accepted construction costs (CORE+CtC), including open commitments, mainly due to late payments of baseline Common Fund contributions. Due contributions arriving from Funding Agencies are reported elsewhere (CERN-RRB-2011-070).

The Collaboration most strongly urges all Funding Agencies that have not yet committed to their full calculated share of CtC funding, or have not yet financed their baseline Common Fund contributions, to continue their utmost efforts to secure the missing resources. Only a strong solidarity across all funding partners will allow the Collaboration to complete the full detector and to exploit fully the immense LHC physics potential as early as possible.

12. Status of FDL Activities and Planning for LHC Luminosity Upgrades

The 2002 Completion Plan reduced the scope of the Full Design Luminosity (FDL) detector as a temporary measure. The staged items included common elements, such as shielding and processors, as well as components of the Inner Detector, calorimeter systems, and muon systems. Some of these items have meanwhile been restored, whereas the fate of other items depends upon the measured performance of the detector. It is clear that infrastructure, particularly shielding and cooling, gas, and cryogenic systems, will require improvement and consolidation during the 2013-2014 LHC shutdown. Improvements to the hermeticity of the muon trigger system are also planned.

The Insertable Pixel B-layer (IBL) is part of the FDL detector activities. The Technical Design Report (TDR) was submitted to the LHCC (CERN-LHCC-2010-013, ATLAS TDR 19) in September 2010, with installation now targeted for the 2013-2014 LHC shutdown. The project is on track in terms of both major decision milestones and overall delivery schedule. Following the selection of sensor technologies for the IBL, the Memorandum of Understanding is being finalized for signature by the Institutes and Funding Agencies concerned. Along with IBL installation, which includes a beryllium beam pipe, aluminium beam pipes in the forward regions will be installed. An option to replace the pixel Service Quarter Panels also exists. A status report on FDL detector activities is provided in CERN-RRB-2011-068.

As previously reported (CERN-RRB-2009-020, April 2009), in accordance with CERN plans to increase the luminosity of the LHC, ATLAS established an R&D programme to study the detector improvements required to operate ATLAS beyond design luminosity conditions. Two phases of improvement are being planned. The first phase is targeted at reaching the “ultimate” LHC luminosity of approximately $2x10^{34}$ $cm^{-2}s^{-1}$, with installation foreseen in 2018. The second phase of detector improvements is targeted at very high integrated luminosity during the High Luminosity LHC (HL-LHC) era, with installation foreseen for 2022. ATLAS plans to proceed by drafting Letters of Intent (LoI) for each upgrade phase. Each LoI will be followed by project-specific Technical Proposals, Technical Design Reports, and Addenda to a Memorandum of Understanding.

Operation of ATLAS at luminosities significantly beyond nominal towards the end of this decade will require a number of detector, electronics, and triggering improvements. The first project (as opposed to the many currently endorsed R&D programmes) to receive internal approval by the Collaboration Board is the “FTK: a hardware track finder for the ATLAS trigger”. The FTK aims to address, between the level 1 trigger and the HLT, charged track reconstruction challenges arising from the high track density environment implied by of order 50 or more collisions per beam-crossing as instantaneous luminosities rise to $2\times10^{34}\,cm^{-2}s^{-1}$ and beyond.

Detector improvements in preparation for the Phase 1 LHC target integration of ~300 fb^{-1} by the start of the next decade, exploiting peak instantaneous luminosities of $2\times10^{34}\,cm^{-2}s^{-1}$ and above. Improvements focus on enhancing ATLAS trigger capabilities in order to maintain good physics selectivity in face of higher luminosity and background rates. Required changes are proposed to include replacement of the “small wheels” of the muon spectrometer, because the $|\eta| \geq 1.2$ regions, which dominate the muon level-1 trigger rate, could overload the available level-1 trigger bandwidth beyond design luminosity. New small wheels could also eliminate saturation effects and loss of hit efficiency in the muon drift tubes arising from cavern background rates. Improvements to the level-1 calorimeter trigger, using finer granularity information and more sophisticated algorithms, are under study in order to cope with pileup levels well beyond those for which the detector was designed. The goal of these improvements is to keep single lepton and missing E_T thresholds at level-1 as low as possible, within constraints of the allowable level-1 rate. Finally, on the time-scale of a 2018 shutdown, ATLAS is also studying enhancing its very far forward detector capability with a view to detectors able to operate in high pileup and to study rare processes in γγ scattering (with W^+W^- final state, for example) and extreme diffractive QCD processes. The LoI detailing Phase I plans, including preliminary cost estimate, is currently being prepared. The LoI will also summarize physics studies that motivate and guide the design of the Phase I detector improvements.

Detector improvements in preparation for the HL-LHC era target integrating 250-300 fb^{-1} per year, exploiting peak luminosity of approximately $5\times10^{34}\,cm^{-2}s^{-1}$, with high radiation and event pileup levels. Such operating conditions will quickly exhaust any remaining radiation budget in the tracker system, requiring replacement of the full Inner Detector. They also will potentially lead to radiation damage to the cold preamplifier electronics of the HEC, and will probably preclude operation of the present FCal. In addition, such operating conditions will require improvements to much of the currently installed readout, trigger, and data acquisition systems. ATLAS is presently conducting a program of urgent R&D on long lead-time improvements. For instance, replacement tracker requirements have received significant attention and prototyping effort, focussing on an all-silicon detector with the very high radiation hardness and granularity suitable for such an environment. Programmes to address the extreme radiation requirements for the pixel layers have also been undertaken, and benefitted greatly from the intensive programme to qualify technologies for the IBL that will survive until the shutdown for HL-LHC installation. The Letter of Intent describing the detector improvements for this phase is presently planned to be prepared approximately one year after the Phase I LoI.