Romanian Reports in Physics, Vol. 67, No. 2, P. 386–395, 2015

ELEMENTARY PARTICLE PHYSICS

THE BUCHAREST ATLAS ANALYSIS FACILITY AT IFIN-HH

M. CIUBANCAN, G. STOICEA, C. ALEXA, A. JINARU, J. MAURER,
M. ROTARU, A. TUDORACHE

"Horia Hulubei" National Institute for Physics and Nuclear Engineering, RO-077125 Bucharest-Măgurele, Romania, Email: Mihai.Ciubancan@nipne.ro; Gabriel.Stoicea@nipne.ro; Calin.Alexa@nipne.ro, Adam.Jinaru@nipne.ro; Julien.Maurer@nipne.ro; Marina.Rotaru@nipne.ro, Alexandra.Tudorache@nipne.ro[1]

Received March 21, 2014

Abstract. We present the structure and the functionality of the Tier-3 ATLAS local analysis facility developed in Particle Physics Department of the IFIN-HH. The infrastructure was designed to fulfil the requirements for intensive Monte Carlo simulation and large data volumes analysis. Tests for parallel data analysis using PROOF framework were performed and several studies have been developed for proton-proton collisions at TeV scale. The obtained results confirm the accomplishment of the requirements of both software and hardware infrastructures.

Key words: particle physics, parallel data analysis, Monte Carlo simulations, computing facility.

1. INTRODUCTION

Given the successful testing and validation of the BAAF (Bucharest ATLAS Analysis Facility) prototype [1] and due to the increasing demand for computing resources determined by the intensive involvement of the IFIN-HH ATLAS group in various physics studies, the initial infrastructure was upgraded, both in terms of hardware, by increasing computing resources, and software update to improve the software tools. You can find a description of the ATLAS Experiment at the CERN Large Hadron Collider in Ref. [2].

Our research activities are mainly devoted to phenomenology, prospects for new phenomena discovery potential and detector performance. We are involved in the study of jet production, search for supersymmetry in single and multi-lepton final states, general topological search for new physics, phenomenological studies on new particles production and decay, and Monte Carlo generators validation using data.

In order to handle data volumes of multiple petabytes per year the ATLAS Computing model adopted the Grid model and developed three levels of computing centres. The core of the ATLAS Distributed Computing activities are based on 3 tiers, Tier-0, Tier-1 and Tier-2, establishing a coherent system to perform data processing and management on a global scale and host (re)processing and simulation activities plus group and user analysis activities. With the formation of small computing centres, the model was expanded to include them as Tier-3 sites. Many ATLAS institutes have developed Tier-3 facilities, which consist of non-pledged resources mostly dedicated to running data analysis tasks produced by the local scientific groups.

2. BAAF INFRASTRUCTURE

Implementing BAAF we took into account the specific requirements for running both Monte Carlo (MC) simulations and large data volumes analysis. BAAF is composed of a Tier2 Grid facility RO-02-NIPNE, used for GRID analysis and MC production, and a local batch and interactive computing cluster for local group simulations and data analysis. The local analysis facility cluster used for simulations and parallel analysis is composed of one head/interactive node and four working nodes, in total 256 cores, 2 GB/core RAM, with a storage area of 32 TB, with 10Gbps/40Gbps bandwidth inside a private network which gives access to RO-02-NIPNE resources (Fig. 1).

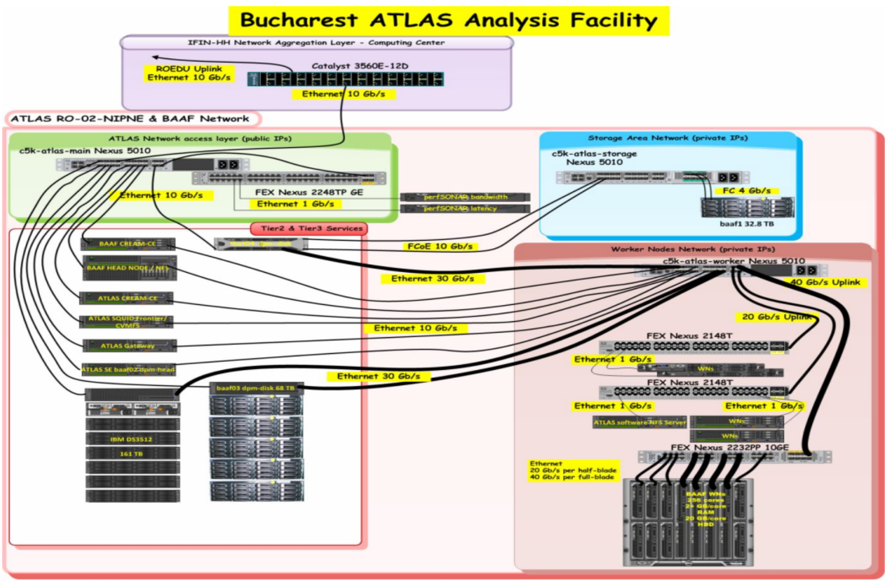

Fig. 1 – BAAF infrastructure schematic view.

The BAAF is structured as a TIER3g [3], having a distributed computing part controlled by PBS (Portable Batch System) Torque/MAUI [4] server and a parallel computing part using PROOF (Parallel ROOT Facility). PROOF is an extension of ROOT [5] enabling interactive analysis of large sets of ROOT files in parallel on clusters of computers or many-core machines. More generally PROOF [6] can parallelize tasks that can be formulated as a set of independent sub-tasks. For parallel analysis we implemented PoD (PROOF-on-Demand) [7], a tool-set which sets up a PROOF cluster on any resource management system. PoD is a specially designed solution to provide a PROOF cluster on the fly, with a very simple and fully automated installation. Its distribution contains preconfigured modules and everything users need to immediately start working with it right after the installation. PBS TORQUE Resource Manager provides control over batch jobs and distributed computing resource, while Maui Scheduler is a policy engine which allows sites control over when, where, and how resources such as processors, memory, and disk are allocated to jobs. In addition to this control, it also provides mechanisms which help to intelligently optimize the use of these resources, monitor system performance, help diagnose problems, and generally manage the system.

Simulation jobs are executed using the local batch system while datasets which are ROOT NTUPLES files are analysed using PoD.

BAAF was designed to perform large data volumes analysis and intensive Monte Carlo simulations. In Fig. 2 we show a schematic view of the data flow.

The main branch is using the ATLAS software framework, Athena [8]. A generator produces events in standard HepMC format, these events are then read into the ATLAS detector simulation [9]. Analysis Object Data (AOD) events are physically clustered on output by trigger or physics channel or other criteria that reflect analysis access patterns. This last output dataset might be used to create an ROOT n-tuple for further analysis. A detailed description of the ATLAS Simulation Infrastructure is presented in Ref. [9].

The second branch presented in Fig. 2, chains various open source software packages together, general-purpose particle physics event generators, detector simulation packages and data analysis tools. This setup is mainly used for students training and fast phenomenological prospects for various physics studies. Eventually, PoD is used for parallel data analysis produced either through the Athena framework chain or over the open source software setup.

On BAAF head node machine it is installed the PBS Torque/MAUI server and a NFS server [10] for storing the datasets and which can be accessed by the worker nodes (WN). The users can access the data stored in the Grid and are able to process and store them locally on NFS server. All the versions of ATLAS software are available *via* CERN Virtual Machine File System (CVMFS) [11] on head/interactive node and on worker nodes too. CVMFS is a network file system based on HTTP and optimized to deliver experiment software in a fast, scalable,

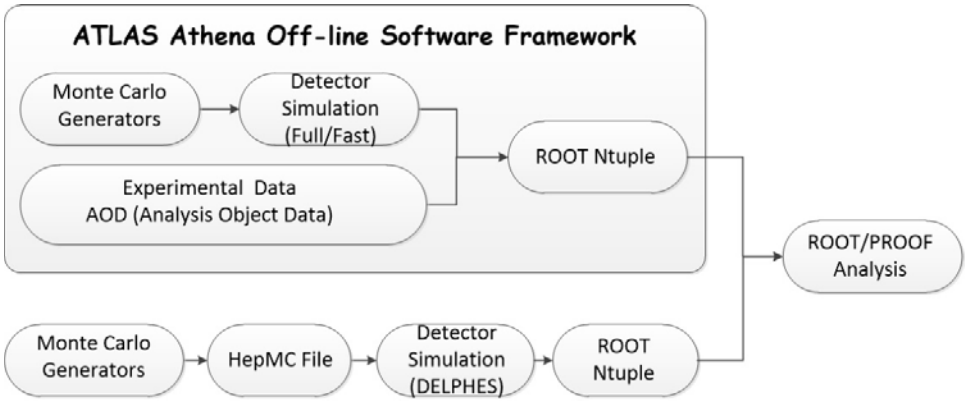

Fig. 2 – Overview of the data flow in BAAF.

and reliable way; files and file metadata are cached and downloaded on demand. For accessing the ATLAS software it is used the already configured SQUID server [12] for RO-02-NIPNE. The cluster monitoring is performed with Nagios using check_mk [13] interface; we do monitor the load, network bandwidth, etc. and possible errors of CPUs, hard drives, network and essential services and we have the possibility to send alerts to administrators when an error occurs. In this scope we developed different nagios plugins to handle BAAF infrastructure (hardware and software).

3. TESTS AND ANALYSES

To use PoD users has to setup their environment. In order to this it has to edit a bash script file in order to export on the worker nodes the same environment, for example to use a specific version of ROOT software. Once the environment is set the users can easily start the PoD server using specific command and to reserve their working nodes on the cluster via PBS/MAUI. Also they have the possibility to check when the resources are available for their analysis.

In order to use PROOF user need to parallelize the analysis code. This can be done *via* ROOT accessing a Tree n-tuple and generating the software skeleton for it, which is done automatically. This will create two files: Analyzer.h inheriting the methods from the analysis code and Analyzer.C containing the methods that will run on the PROOF nodes. In this file the user will add his new objects definitions (*e*. *g*. Histograms) and specific analysis code.

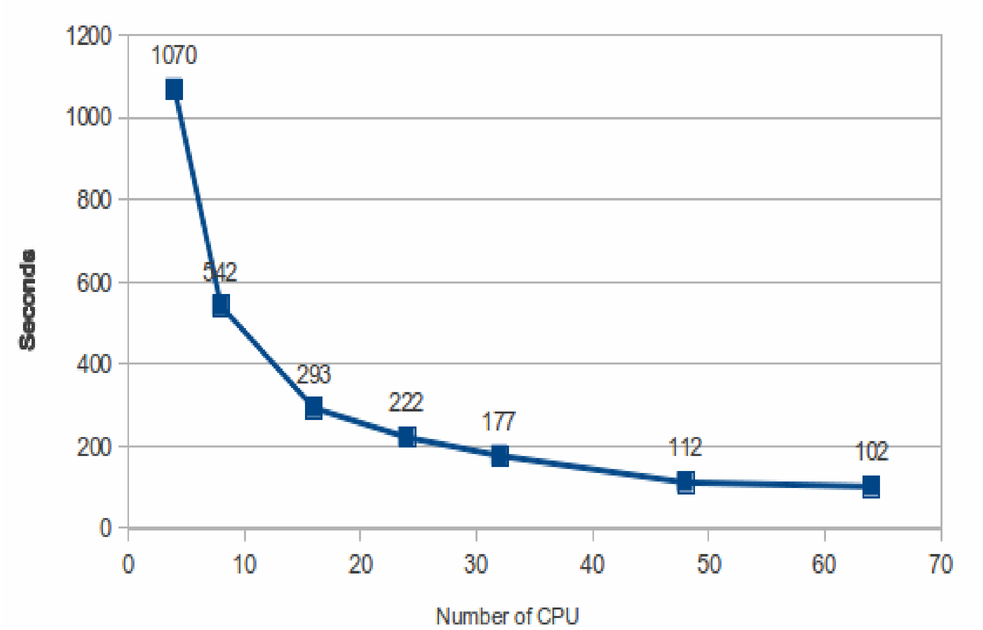

Fig. 3 – Performance distribution of the BAAF cluster.

Using BAAF we have done an analysis test running 50 million events using a specific PROOF benchmark which generates random events on the fly.

One can see in Fig. 3 that the running time decreases very fast with the number of processing cores and gets saturated over about 50 cores.

We repeated this test for a real analysis case using ATLAS data. We used a sample of 188 GB (3.119.725 events) in two scenarios, we placed the data locally on the worker node disk and accessing remotely the data from the storage server using NFS.

In Fig. 4 we can see that the execution time follows the same trend as in the previous test, but the decrease is not so steep because the running time is dominated by the I/O operations. The data are processed much faster when are stored locally, as was expected, but the gain is bigger only when a large number of cores are used.

The same behaviour can be seen in Fig. 5, where we present the event processing rate as a function of the number of processing cores.

We decided to implement in BAAF the NFS access to the data files because the local storage space on the WN is limited and the gain for local data processing time is smaller compared to the advantage of accessing large data volume on a single storage partition for all WN's.

Other accessing protocols to the data like xrootd, rfio, etc., will be tested in the future.

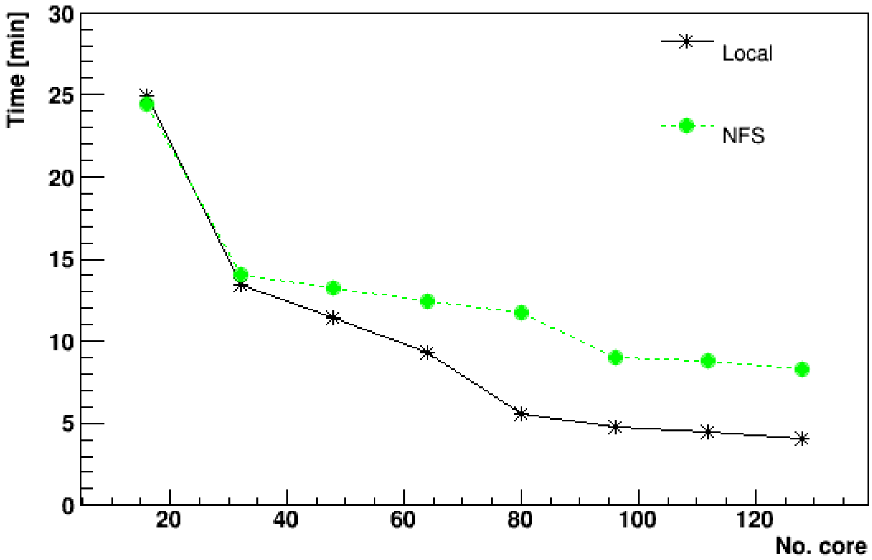

Fig. 4 – Execution time variation as a function of the number of processing cores using ATLAS data.

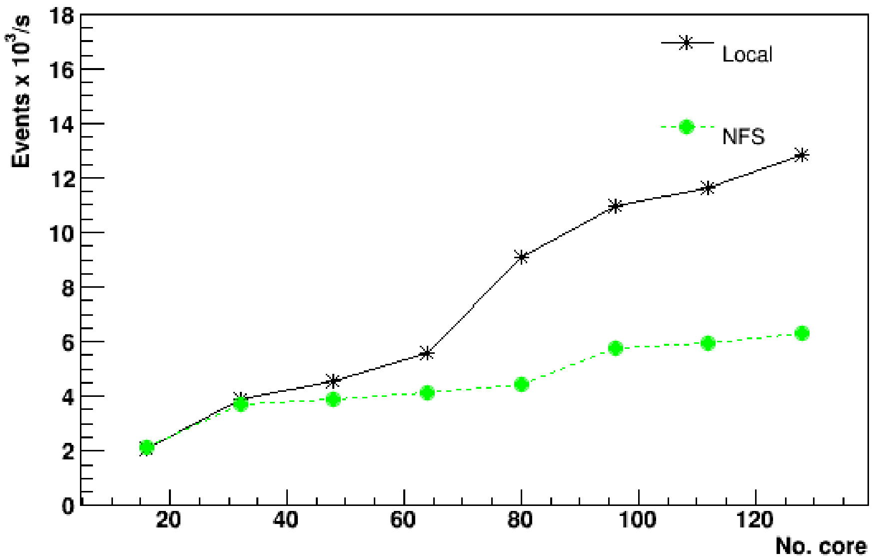

Fig. 5 – Number of events per second as a function of the number of processing cores using ATLAS data.

4. PHYSICS STUDIES

To demonstrate the efficiency and capability of the local analysis facility we will present results obtained running Monte Carlo simulations and data analysis.

4.1. W′ PRODUCTION STUDY

In the context of G(221) models, W' boson production in proton-proton collisions a TeV scale has been studied [14]. W' boson decay into $hH^{\pm}$ was considered, where h stands for a light boson that can be indented with the SM Higgs, and $H^{\pm}$ for a heavy charged Higgs boson yet to be discovered. For this particular decay channel, partial decay widths, branching ratios and production cross-sections were calculated in a recent paper [15].

Using PYTHIA6 and PYTHIA8 [16, 17], W′ heavy gauge boson was produced according to the extended gauge model [18]. The topology was generated using different parton distribution functions (PDF), and different masses for W' and $H^{\pm}$.

In order to make the comparison with the ATLAS experiment results we used the Delphes package [19], a C++ framework for fast simulation of a generic collider experiment. Delphes is including transport of the primary and secondary particles through the detector material accounting for the various detector inefficiencies, the dead material, the imperfections and the geometrical details. Eventually, the package provides reconstructed physics objects, such as leptons, jets, photons and missing energy. The data analysis was performed with ROOT [5].

Final states with leptons, jets and missing transverse energy was studied by comparing them against the same SUSY final states on a basis of kinematical distributions, such as MET, jet and lepton leading and sub-leading transverse momentum distributions. The preferred decay channel for H^{+} was $t\bar{b}$, and for h, WW or ZZ pairs. Another comparison tested the current ATLAS sensitivity to the $W'\rightarrow hH^{\pm}$ decay channel. More detailed studies of this decay were presented in a dedicated article [15] and in a Ph D Thesis [20].

The study performed for testing the platform from generator level to analysis demonstrates the functionality and the efficiency of local analysis facility

4.2. MULTI PARTICLES STATE GENERATION

The BAAF cluster has also proven to perform well on the topic of Monte-Carlo event generation.

To extend the interpretation of the outcome of several ATLAS searches for gluinos pair production [21], it has become interesting to look into regions of the phase space with a small mass gap between gluino and neutralino. In this case, the particle cascade usually characterizing decays of strong superpartners mainly occurs through virtual particles, and gluinos decay directly into a potentially large number of products [22]. When the number of particles in the final state is large, the phase space integration performed prior to the actual event generation can become a very large time-consuming process if one wants to rely on the full matrix elements instead of approximate formulas. The parton-level generator MadGraph [23] allows the leading order computation of these matrix elements in the case mentioned above. Furthermore, this program is very well adapted to parallel computing, and the phase space integration can be splitted in many independent subjobs that can be either ran simultaneously on a multicore machine, or even submitted to computer clusters of different architectures, in particular the BAAF cluster. We have indeed observed, by testing the generation of complex processes such as the one mentioned previously, very important reduction of the total running time compared to a single machine. The results obtained using the platform were presented in several group meetings from ATLAS experiment. The achieved validation of the use of MadGraph distributed computing utilities on BAAF is therefore important for the present and future activities of the group in this domain.

4.3. A GENERAL SEARCH FOR NEW PHENOMENA IN PROTON-PROTON COLLISIONS

The analysis based on general search for new phenomena in proton-proton collisions at a centre-of-mass energy of 8 TeV was processed *via* PROOF event loop, using ROOT version 5.34.10. The general search is a model independent search for deviations from the Standard Model predictions in all final states. No model is assumed to look for new physics. Due to the generality of the analysis this search might find signatures which have been missed by other searches.

The goal of the analysis performed on BAAF was to validate the software setup using a benchmark defined by a specific cut flow used by the ATLAS general search group. Based on a fixed set of requirements, we obtained the same result concerning the event cleaning as the official physics group from CERN. This achievement will allow us to use the PROOF setup for extensive studies dedicated to general search. We could thus highly increase the speed at which analysis algorithms are developed and tested by achieving faster response times compared to the conventional batch or GRID approaches.

5. CONCLUSIONS

Imposed by the increasing demand for computing resources requested/generated by the strong involvement of our ATLAS group in various physics studies, a robust and performant Tier-3 analysis facility infrastructure was implemented, Bucharest ATLAS Analysis Facility (BAAF).

The studies performed for testing this platform from generator level to analysis demonstrates the achieved functionality.

The results obtained for W’ boson production, multi-particle state generation and general search for new phenomena, confirm the accomplishment of the requirements of both software and hardware infrastructures.

Using BAAF, the speed of data processing, developing and testing various algorithms and methods will achieve a fast response being important for the present and future activities of the ATLAS IFIN-HH group.

Acknowledgements. This work was supported by the Sectorial Operational Programme Human Resources Development (SOP HRD), financed from the European Social Fund and by the Romanian Government under the contract number SOP HRD/107/1.5/S/82514, and by the Romanian Ministry of Education through the ATLAS – Capacities Module III project.

REFERENCES

1. M. Cuciuc, M. Ciubancan, V. Tudorache, A. Tudorache, R. Paun, G. Stoicea, C. Alexa, *Building and testing an infrastructure for studying proton-proton collisions at TeV scale*, Rom.Rep.Phys., **65**, 122–132 (2013).
2. ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, JINST, **3**, S08003 (2008).
3. ATL-COM-SOFT-2010-053, *ATLAS Tier-3*, *Doug Benjamin on behalf of the Atlas Collaboration*, Conference on Computing in High Energy and Nuclear Physics 2010, Taipei, Taiwan, 18–22 Oct 2010.
4. PBS Torque http://www.adaptivecomputing.com/products/open-source/torque/; MAUI http://www.adaptivecomputing.com/products/open-source/maui/.
5. R. Brun and F. Rademakers, *ROOT: An object oriented data analysis framework*, Nucl. Inst. & Meth. in Phys. Res., **A 389**, 81–86 (1997).
6. M. Ballintijn, M. Biskup, R. Brun, P. Canal, D. Feichtinger, G. Ganis, G. Kickinger, A. Peters and F. Rademakers, *Parallel interactive data analysis with PROOF*, Nuclear Instruments and Methods in Physics Research, **A** 559, 13–16 (2006);
   G. Ganis, J. Iwaszkiewicz, F. Rademakers, *Data Analysis with PROOF*, Proceedings of ACAT 2008 Conference, 2006.
7. Proof on Demand http://pod.gsi.de/
8. ATLAS Collaboration, ATLAS Computing Technical Design Report, ATLAS-TDR-017, CERN-LHCC-2005-022 (2005);http://atlas-computing.web.cern.ch/atlas-computing/packages/athena-Core/athenaCore.php
9. ATLAS Collaboration, *The ATLAS Simulation Infrastructure*, Eur.Phys.J., **C70**, 823–874, 2010
10. Mike Eisler, Ricardo Labiaga, Hal Stern, *Managing NFS and NIS,* 2nd Edition, O’Reilly Media, July 2001.
11. J. Blomer and T. Fuhrmann, *A Fully Decentralized File System Cache for the CernVM-FS*, Computer Communications and Networks (ICCCN), 2010; Proceedings of 19th International Conference on, pp. 1–6, 2–5 Aug. 2010; doi: 10.1109/ICCCN.2010.5560054;
    J. Blomer *et al*.; *CernVM-FS: delivering scientific software to globally distributed computing resources*; Proceeding NDM '11 – Proceedings of the First International Workshop on Network-Aware Data Management; pp. 49–56, ACM New York, NY, USA, 2011.
12. Duane Wessels, *Squid: The Definitive Guide*, O’Reilly Media, January 2004.
13. Nagios, 2nd Edition, Wolfgang Barth Publisher, *System and Network Monitoring*, No Starch Press,

October, Check_mk: http://mathias-kettner.de/checkmk.html
14. Ken Hsieh, Kai Schmitz, Jiang-Hao Yu, C.-P. Yuan, Phys.Rev., **D 82,** 035011 (2010); arxiv:1003.3482,.
15. A. Jinaru, C. Alexa, I. Caprini and A. Tudorache, http://arxiv.org/abs/1312.4268.
16. T. Sjöstrand, S. Mrenna, and P. Z. Skands, JHEP, **0605**, 026 (2006)
17. T. Sjöstrand, S. Mrenna and P. Skands, Comput. Phys. Comm., **178** (2008).
18. Altarelli G, Mele B and Ruiz-Altaba M, Z. Phys., **C 45**, 109 (1998).
19. *Delphes*, *a framework for fast simulation of a generic collider experiment*, http://arxiv.org/abs/0903.2225.
20. A. Jinaru, *Study of W'boson production in the ATLAS experiment at the LHC*, Ph D Thesis, 2014.
21. ATLAS Collaboration, *Search for squarks and gluinos with the ATLAS detector in final states with jets and missing transverse momentum and 20.3 fb^{-1} of $\sqrt{s} = 8$ TeV proton-proton collision data*, ATLAS-CONF-2013-047.
22. ATLAS Collaboration, *Search for strongly produced superpartners in final states with two same sign leptons with the ATLAS detector using 21 fb^{-1} of proton-proton collisions at $\sqrt{s} = 8$ TeV*, ATLAS-CONF-2013-007.
23. J. Alwall, M. Herquet, F. Maltoni, O. Mattelaer, T. Stelzer, *MadGraph 5: Going Beyond*, JHEP, **1106**, 128 (2011).