

# Introduction

ATLAS Software and Computing Week  
@CERN Nov. 29 – Dec. 3 2010

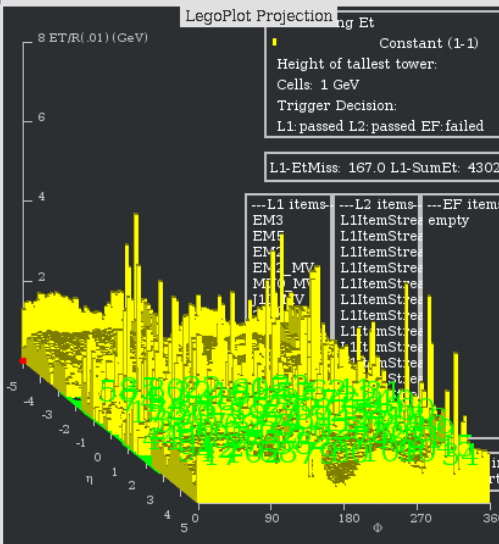

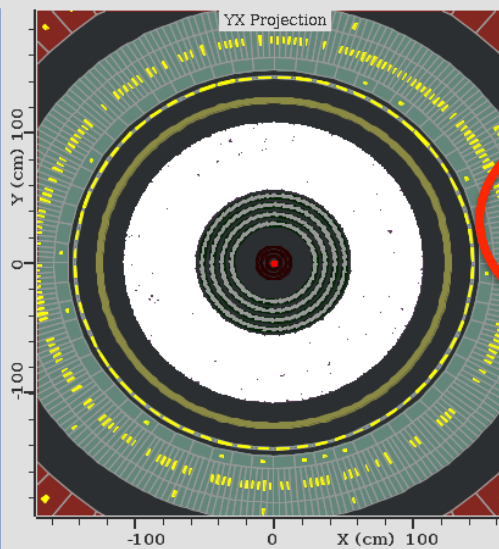
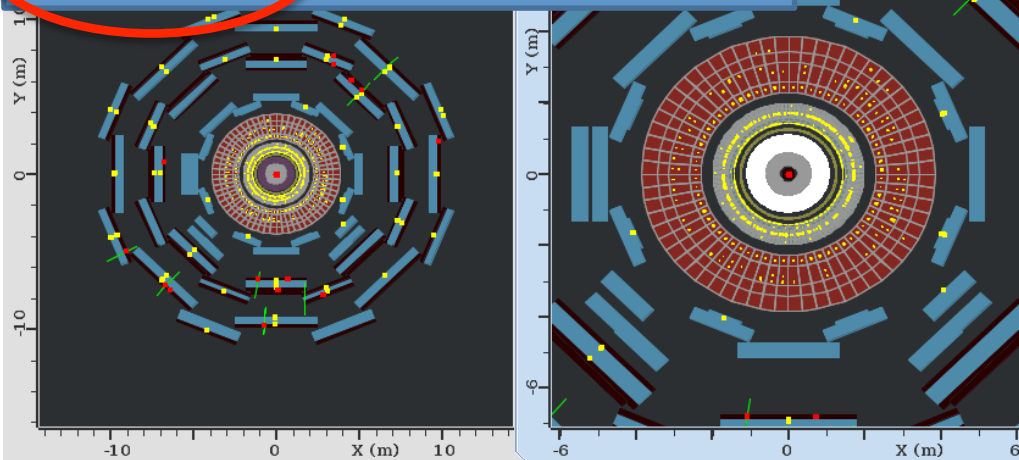
Kors Bos

# First Heavy Ion Collisions

Setting up of the machine started on November 1

Run:168595 ev:930540 lumiBlock:456

Atlantis




First collisions on Sunday, Nov. 7, at 00:30  
No stable beams – Pixel detector was switched off, SCT in standby voltage

# First Heavy Ion Stable Beam Collisions

1<sup>st</sup> stable beams  
(1 bunch colliding)  
Monday, Nov. 8,  
~11:20  
4 colliding bunches:  
Tuesday, 01:01

Now routinely with  
121 colliding bunches

## Heavy Ion Collision Event



Run 168665, Event 83797

Time 2010-11-08 11:37:15 CET




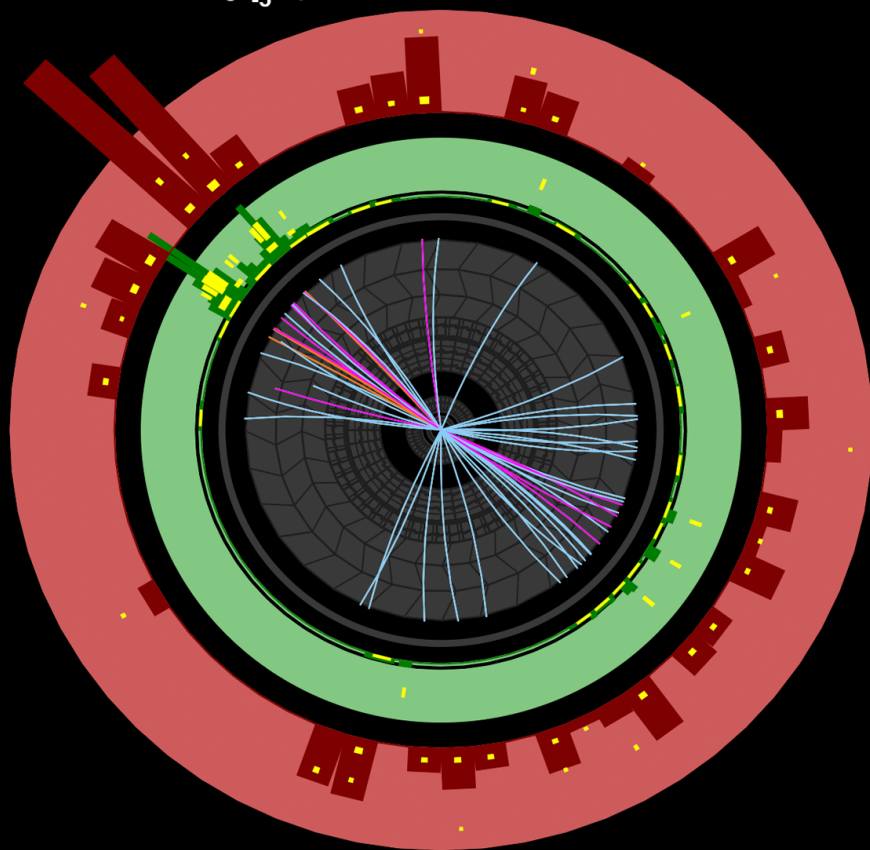


**ATLAS**  
**EXPERIMENT**



# ATLAS EXPERIMENT

Run Number: 169045, Event Number: 1914004

Date: 2010-11-12 04:11:44 CET





Submitted to Phys.Rev.Letters last Thursday  
Accepted on Friday, 10 hours later

## Observation of a Centrality-Dependent Dijet Asymmetry in Lead-Lead Collisions at $\sqrt{s_{NN}} = 2.76$ TeV with the ATLAS Detector at the LHC

G. Aad *et al.* (The ATLAS Collaboration)\*

Using the ATLAS detector, observations have been made of a centrality-dependent dijet asymmetry in the collisions of lead ions at the Large Hadron Collider. In a sample of lead-lead events with a per-nucleon center of mass energy of 2.76 TeV, selected with a minimum bias trigger, jets are reconstructed in fine-grained, longitudinally-segmented electromagnetic and hadronic calorimeters. The underlying event is measured and subtracted event-by-event, giving estimates of jet transverse energy above the ambient background. The transverse energies of dijets in opposite hemispheres is observed to become systematically more unbalanced with increasing event centrality leading to a large number of events which contain highly asymmetric dijets. This is the first observation of an enhancement of events with such large dijet asymmetries, not observed in proton-proton collisions, which may point to an interpretation in terms of strong jet energy loss in a hot, dense medium.

PACS numbers:

Collisions of heavy ions at ultra-relativistic energies are expected to produce an evanescent hot, dense state, with temperatures exceeding two trillion kelvins, in which the relevant degrees of freedom are not hadrons, but quarks and gluons. In this medium, high-energy quarks and gluons are expected to transfer energy to the medium by multiple interactions with the ambient plasma. There is a rich theoretical literature on in-medium QCD energy

The data in this paper were obtained by ATLAS during the 2010 lead-lead run at the LHC and correspond to an integrated luminosity of approximately  $1.7 \mu b^{-1}$ .

For this study, the focus is on the balance between the highest transverse energy pair of jets in events where those jets have an azimuthal angle separation,  $\Delta\phi = |\phi_1 - \phi_2| > \pi/2$  to reduce contributions from multi-jet final states. In this letter, jets with  $\Delta\phi > \pi/2$  are la-

<https://twiki.cern.ch/twiki/pub/AtlasProtected/TempDraftRepository/HeavyIonJets.v3.pdf>

# What happened since the last S&C WS


- Finished the pp run on Oct.21
- Started the HI run on Nov.1, now still ongoing
- Re-processed pp data (Oct. runs separately)
- Did more PD2P and got into some trouble
- Learned about 2011 and 2012
- Learned what the Trigger rate will be
- ..

# to be discussed this week to be decided on Friday

- cvmfs for software distribution
- DB simplifications
- PD2P improvements
- Parameters and data management for 2011
- The same for 2012
- ADCLab launch
- Steps towards a new CM
  - More processing from T0 → T1/2s
  - Less data distribution, more PD2P
  - Event size reduction: RAW, ESD, ...

Needs to be discussed at SIT and ADC  
 Can we decide on Friday ? *Rod?*



# cvmfs



- Currently: application software and conditions databases are installed on a shared (nfs, afs) area on each site
- Problems with:
  - nfs (over) load,
  - space,
  - too many (unused) releases

- Cvmfs is client-server file system using http to access 1 central repository with Frontier/Squid for caching
- Description and tests at RAL and PIC:
  - <http://indico.cern.ch/aetFile.ov/access?contribId=36&sessionId=2&resId=1&materialId=slides&confId=92498>
- ATLAS issue: relocatable code
- ATLAS issue: multiple build:
  - Pacman kits
  - RPMs (used @P1)
  - Pacballs (distributed to sites by DDM)
  - cvmfs (master installation using 1 of the above)
- Can we reduce this ?
  - Drop Pacman
  - Use RPMs to build cvmfs master copy
- ATLAS issue: site specific configurations ?
- ATLAS issue: how to deal with the db files ?

## Tests at PIC – setup time





# Database simplifications 1

## *Conditions Data*

- We used to access Oracle db in each T1 for conditions data
- Then we started testing Frontier servers at the sites with and Squids at the sites without a db
- This was successful and reduced the load on Oracle significantly
- We now have
  - 1 Frontier server at CERN and 6 at the T1s
  - Squids NDGF, SARA and Taipei
- The load is still low and we could further reduce the number of Oracle db + Frontier installations
- The minimum is 2 for redundancy reasons.
- We could better use the freed capacity for TAGs

*Needs to be discussed at DB and ADC  
Can we decide on how many on Friday ?  
Can we decide where to create more TAG DBs*

*Dario?*

# Database simplifications 2

## *LFCs*

- We have 1 LFC per cloud (more in the US)
- The SARA incident showed that we are still very fragile if one breaks
  - The whole cloud was down for 3 weeks
  - We decided to make a backup at CERN
- We now propose to
  - have a backup at CERN for each T1 LFC
  - Switch the “back-up” to the “primary” one by one
  - Merge these primary LFCs into 1 master LFC (can also be done step-wise to study load)
  - Make a separate back-up of this single master LFC
  - Drop all other LFC and work from one central one (and a life backup)
- Need to discuss backup technology (streams, Data Guard, ..)
- Need to test LFC merge script

*Needs to be discussed at DB and DDM and ADC*

*Can we decide on the plan (timing) on Friday ? **Fernando?***

# ADC Lab

- First year of ATLAS computing has been a success
  - Now we will face big challenges: this was only the first year's data
    - Especially in scalability, but also in sustainability
- We see a need for an environment to explore new ideas for the future of our computing
  - Like NoSQL DB, caching, map-reduce, simulators
- So we propose the *ADC Lab*, as an environment to
  - Incubate new projects
  - Foster collaboration
- See CDS document <http://cdsweb.cern.ch/record/1309932>

*Needs to be discussed at many sessions and at coffee  
Can we decide on Friday ? Graeme ?*


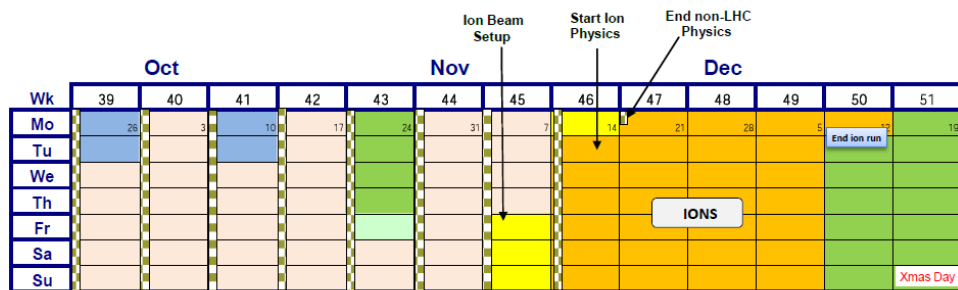


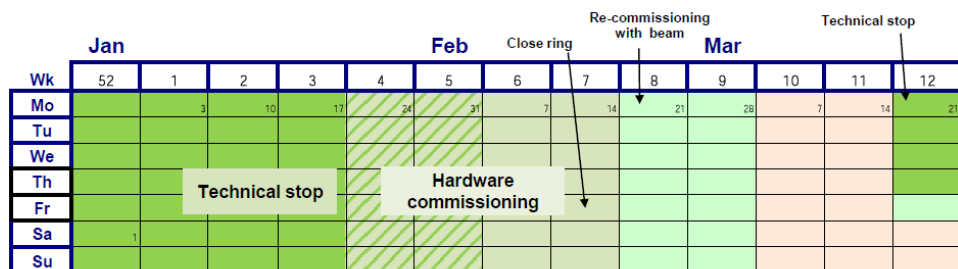
# We must also discuss: Changes to the Computing Model

- Our current resources are nearly full with 2010 data
- We don't know how much increase we get for 2011
- More beam and luminosity in 2011
  - 200 days of uninterrupted running
- Likely even more so in 2012
  - To be decided before the April RRB
- New energy (8 (?) TeV) in 2011
  - To be decided in Chamonix in January
- Request for higher trigger rates
  - 400 Hz and 600 Hz need to be considered
- We need revolution, evolution is not enough

*Needs to be discussed at many sessions and at coffee  
Can we decide which options to pursue on Friday ? Ikuo?*

2011 LHC Schedule  
DRAFT

# 2011



- Beam back around 21<sup>st</sup> February
- 2 weeks re-commissioning with beam (at least)
- 4 day technical stop every 6 weeks
- Count 1 day to recover from TS (optimistic)
- 2 days machine development every 2 weeks or so
- 4 days ions set-up
- 4 weeks ion run
- End of run – 12<sup>th</sup> December

**~200 days proton physics**

Slides by Roger Bailey  
LHCC, November 17

# 2011: “reasonable” numbers

- 4 TeV (to be discussed at Chamonix)
- 936 bunches (75 ns)
- 3 micron emittance
- $1.2 \times 10^{11}$  protons/bunch
- $\beta^* = 2.5$  m, nominal crossing angle

Peak luminosity	$6.4 \times 10^{32}$
Integrated per day	11 pb <sup>-1</sup>
200 days	2.2 fb <sup>-1</sup>
Stored energy	72 MJ

*Usual warnings apply – see problems, problems above*

Slides by Roger Bailey  
LHCC, November 17



# Ultimate reach

- 4 TeV
- 1400 bunches (50 ns)
- 2.5 micron emittance
- $1.5 \times 10^{11}$  protons/bunch
- $\beta^* = 2.0$  m, nominal crossing angle

Peak luminosity	$2.2 \times 10^{33}$
Integrated per day	38 pb <sup>-1</sup>
200 days	7.6 fb <sup>-1</sup>
Stored energy	134 MJ

*Usual warnings particularly apply – see problems, problems above*

Slides by Roger Bailey  
LHCC, November 17

# New parameters to work with

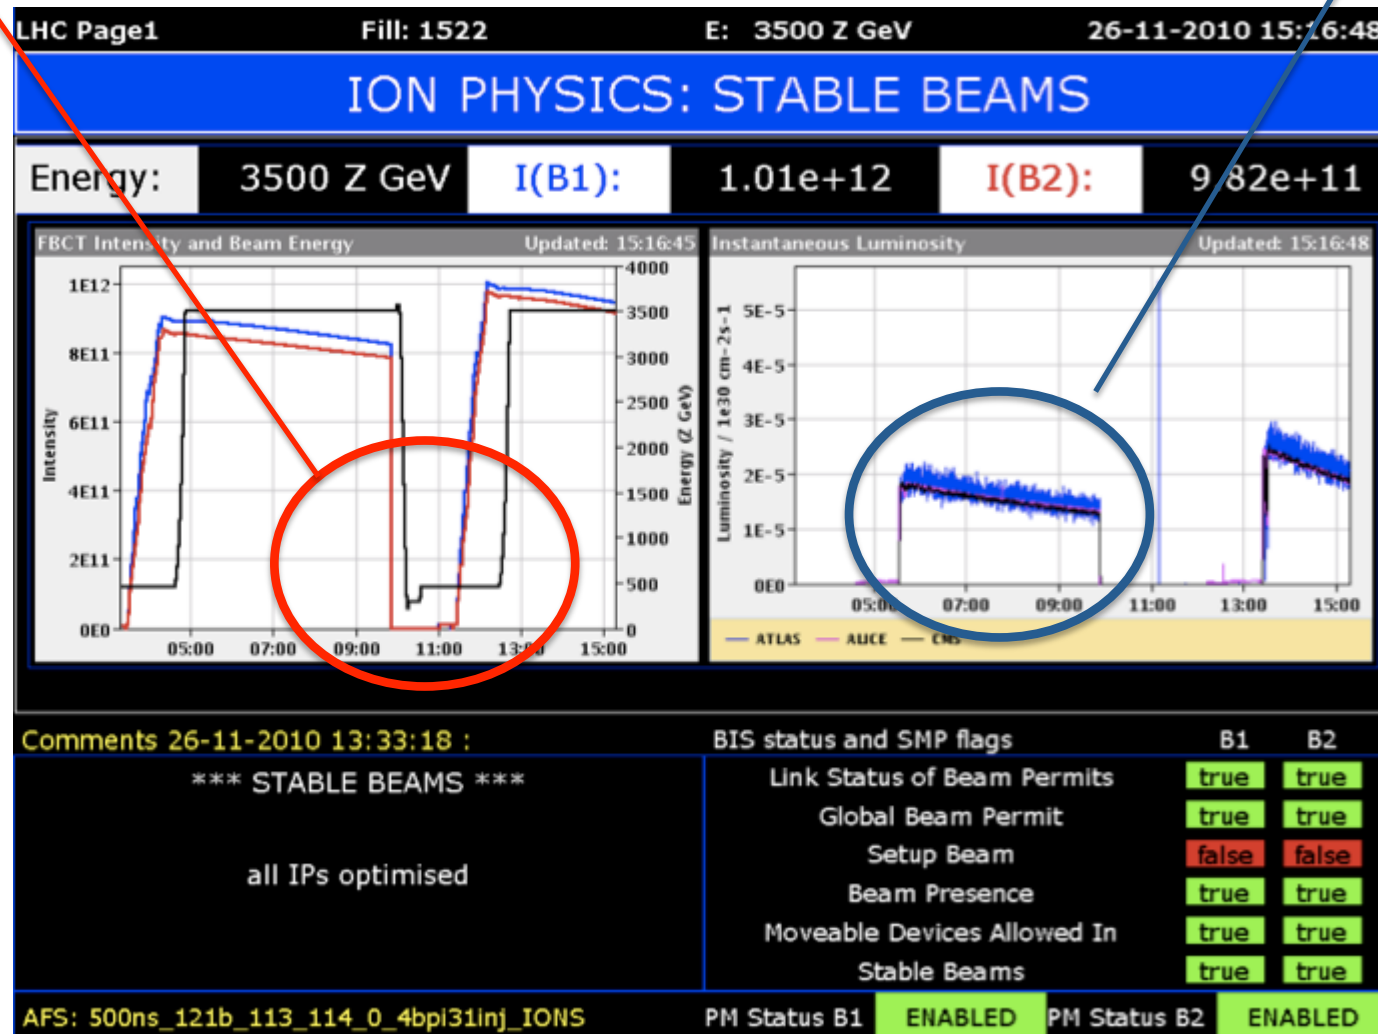
- 200 days of pp and 28 days of HI during 2011
- Most likely the same for 2012
- 50% of the time stable beams
- We will take data at 400 Hz
- Energy will be 4+4 TeV
- 7 TeV MC will still be needed for some time
- HI MC will be needed
- 30% increase of resources on April 1<sup>st</sup> 2011 and 2012

# >40% efficiency now

we assumed 30% for 2010  
but was effectively much less

Less beam loss for protons →  
Longer lifetime of the fill

2 hours for a re-fill



# Revolutionary ideas needed

- Reduce event size from Calo, Indet, Trigg, ..
  - Factor 3 (?) in RAW size, gives also smaller ESD
- Use PD2P everywhere (maybe slower, but ..)
  - And make it more intelligent
- Run bulk processing on the grid
- Get more T1s (stable site with tape) (CERN, ..)
- Custodial copies more distributed:
  - RAW on tape: 1 @CERN and 1 @T1s
  - HITS on tape: 1 @T1s , @T2s
  - ESD on tape: 1 @T1s
- Fewer primary copies on disk:
  - ESD 1 copy @T1s
  - AOD, DESD, etc. on disk: 2 (3?) copies @T1s
  - No a priori data @T2s (rely on PD2P)
- Do away with space tokens (at least in T2s)
  - Just one big cache
  - Will dramatically improve life for Groups

# Workshop dinner

CAFE DU SOLEIL

Place du petit-saconnex 6 / Genève tél : +41 (0) 22 733 34 17

- Wednesday Dec.1 19:30
- Café du Soleil: <http://www.cafedusoleil.ch/>
- Max 35 people, now 32 in Doodle
- If you still would like to join: talk to Ale
- Deadline is tonight