2nd ATLAS – South Caucasus
Software/Computing Workshop & Tutorial

ATLAS Status

and Development towards “Run2” (13 TeV)

Accelerator

Detector

Computing

Physics

Hans von der Schmitt, Tbilisi, October 2012

Jörg Wenninger

The LHC Accelerator Complex

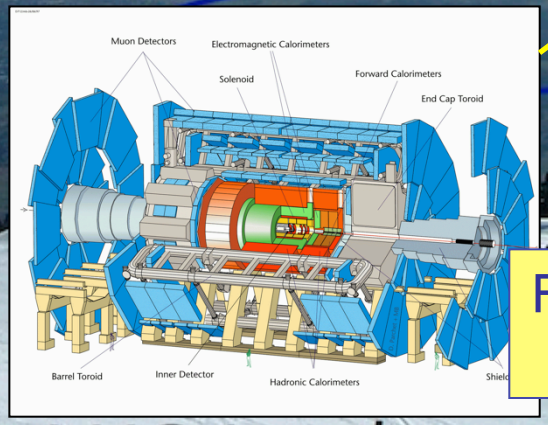

Four large detectors
(here: ATLAS)

The LHC Accelerator Complex

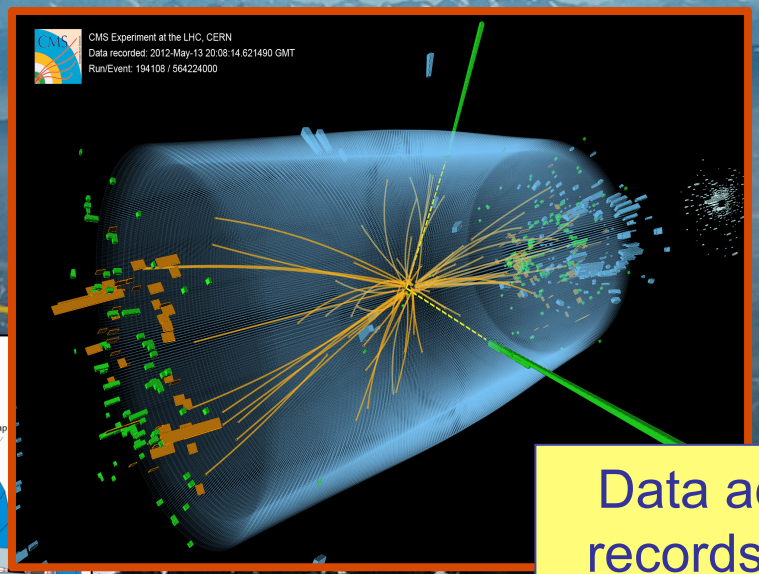

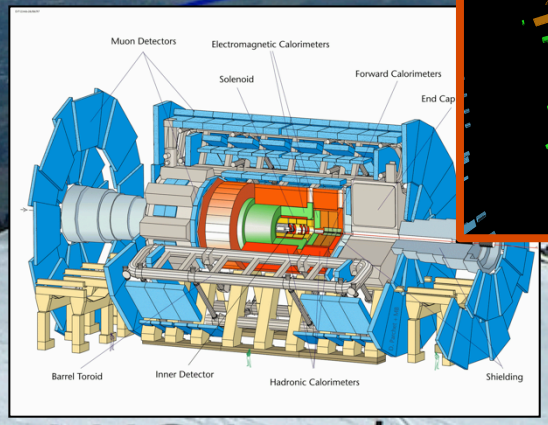

Data acquisition system records individual events (here: Higgs→γγ candidate measured by CMS)

The LHC Accelerator Complex

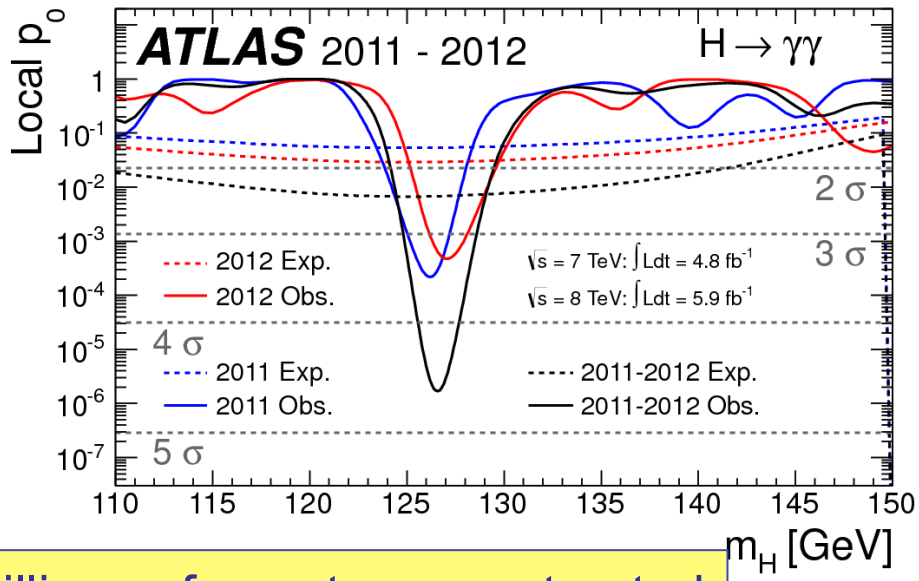

Billions of events reconstructed and analysed statistically for a published physics result

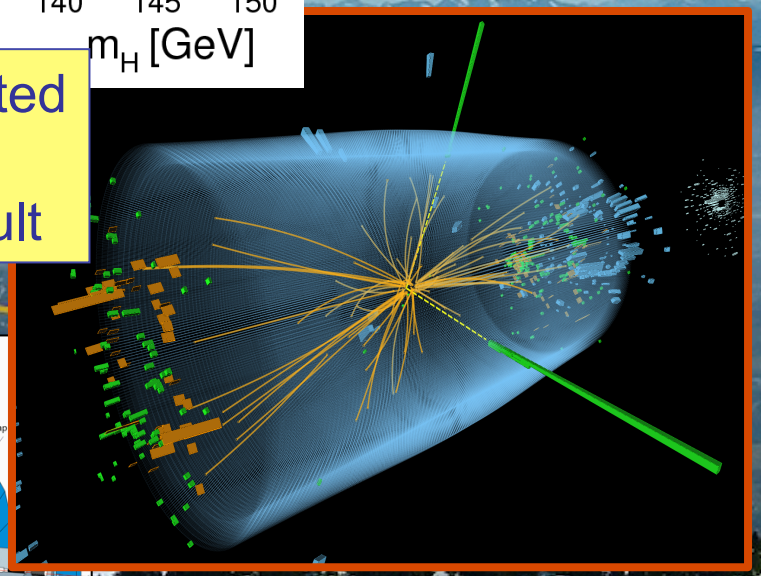

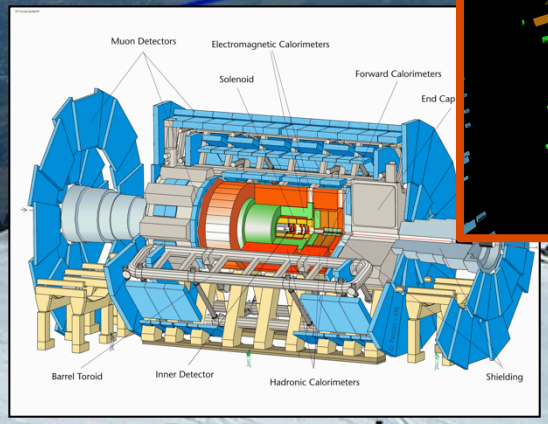

The LHC Accelerator Complex

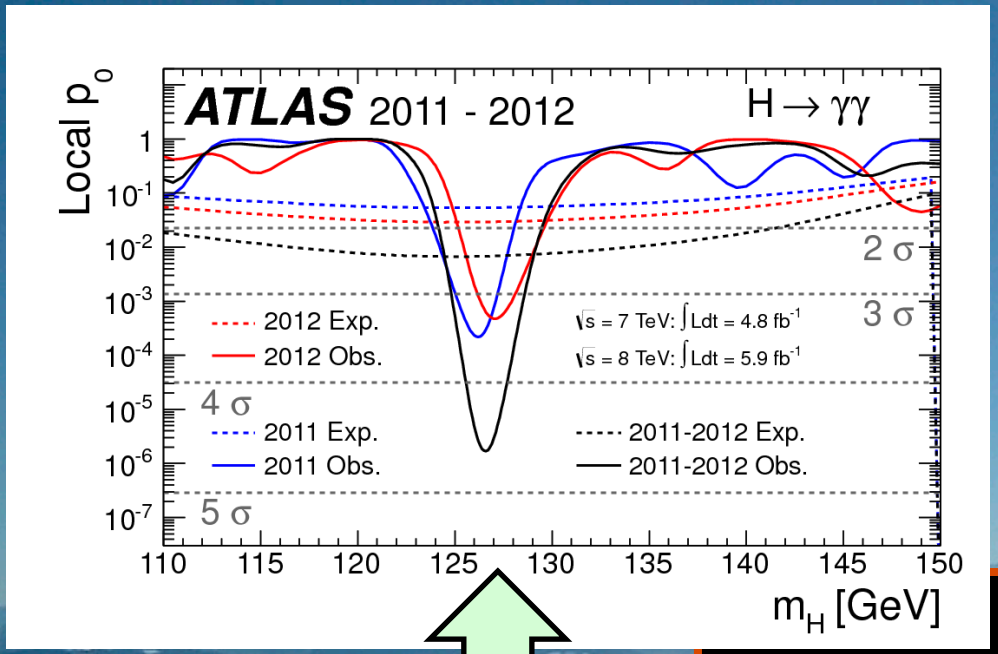

Fast: just days from last data taken to full statistics analysed

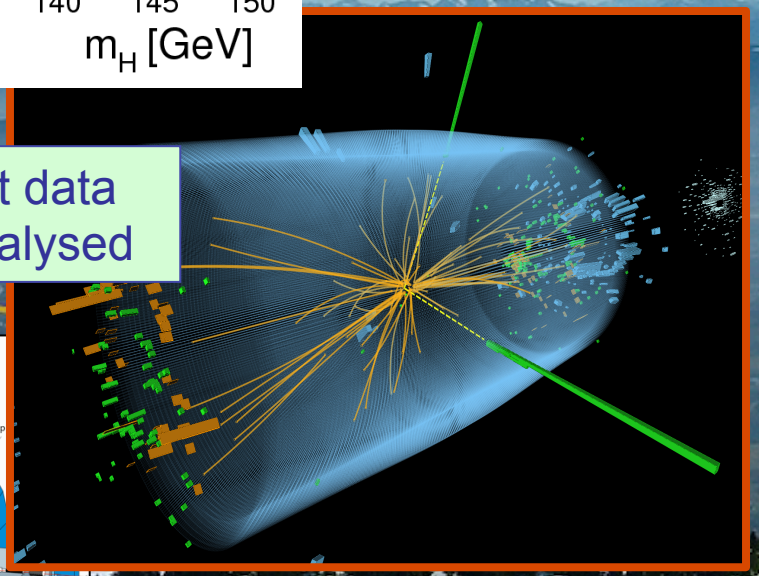

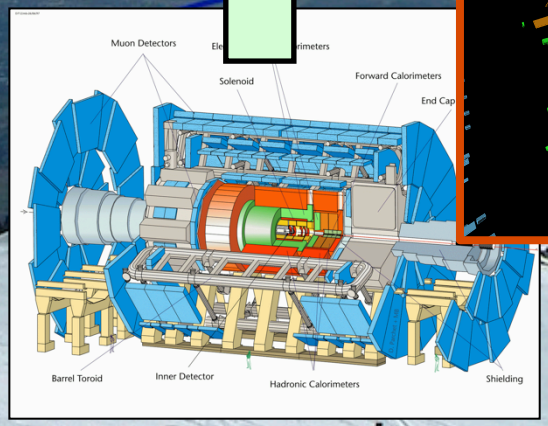

The LHC Accelerator Complex

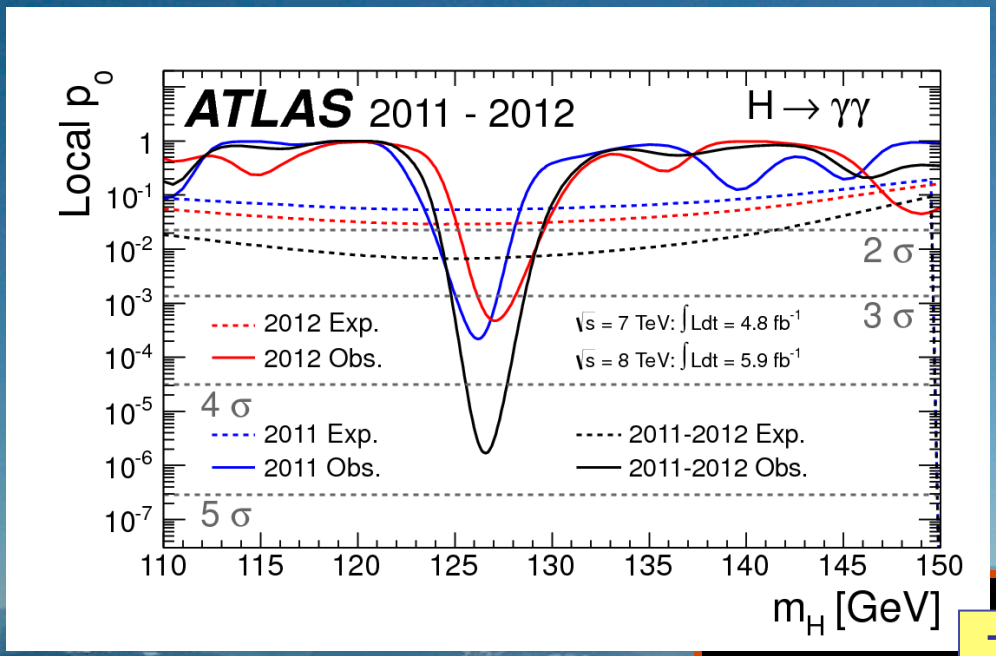

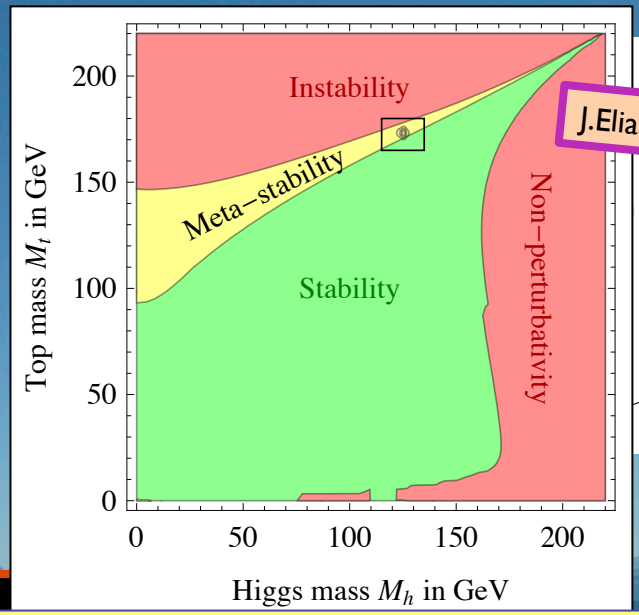

J.Elias-Miró / Barcelona

$+ \; Y^{ij}\Psi_L^i\Psi_R^j\Phi$

Theorists eagerly await the results to compare with their favourite ideas

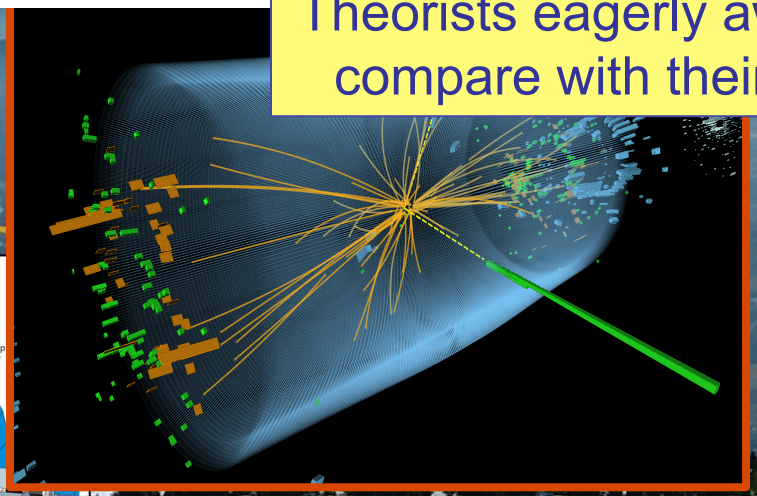

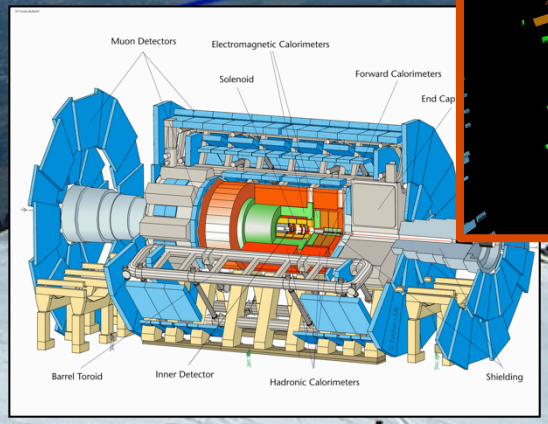

The LHC Accelerator Complex

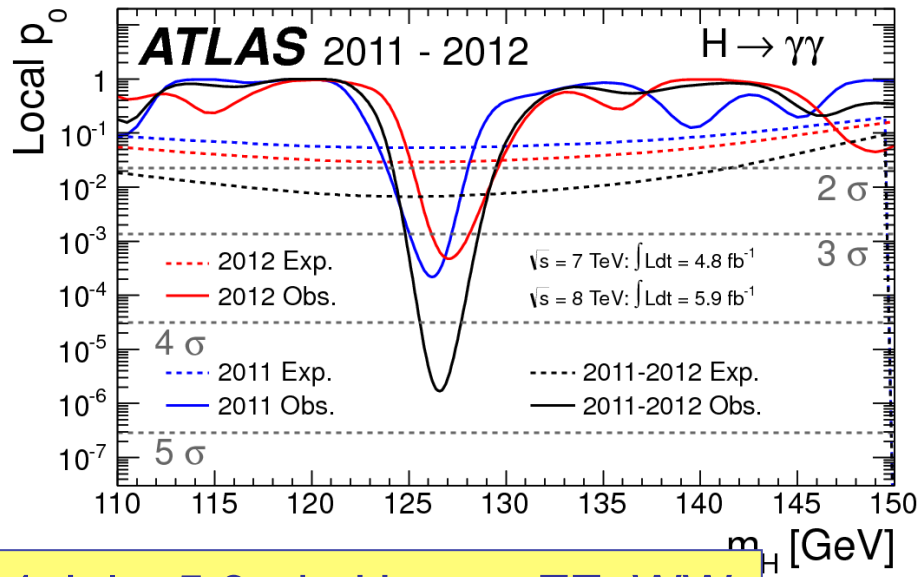

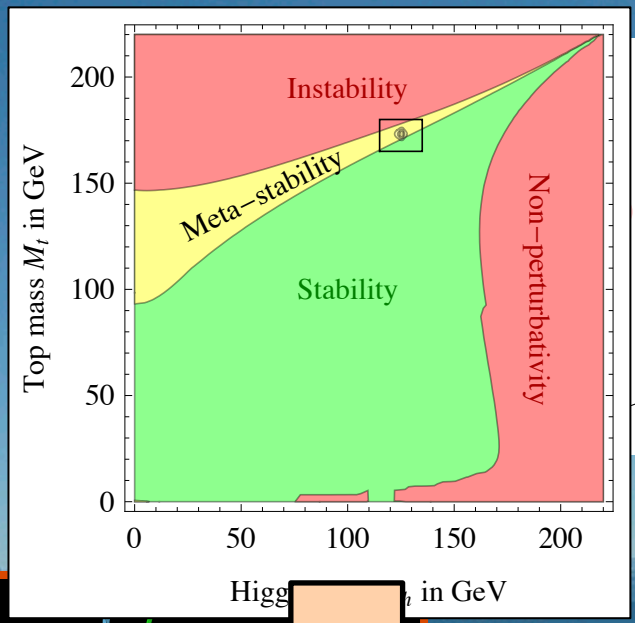

$+ Y^{ij}\Psi_L^i\Psi_R^j\Phi$

4 July: 5.9σ in H → γγ, ZZ, WW
Now: work to include bb, ττ, ...
Assess more boson properties

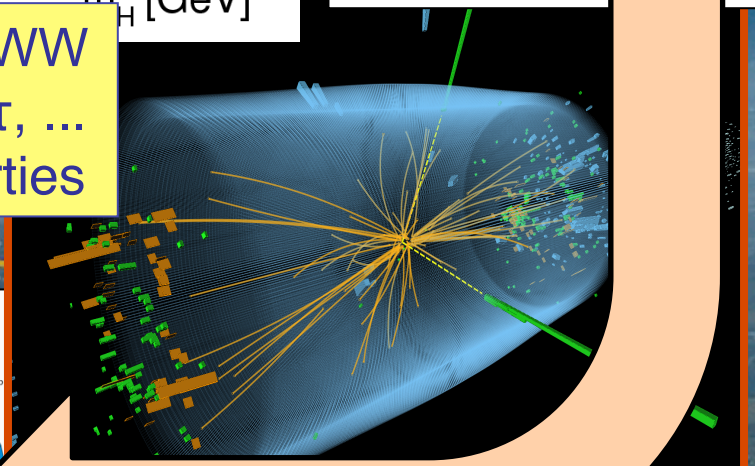

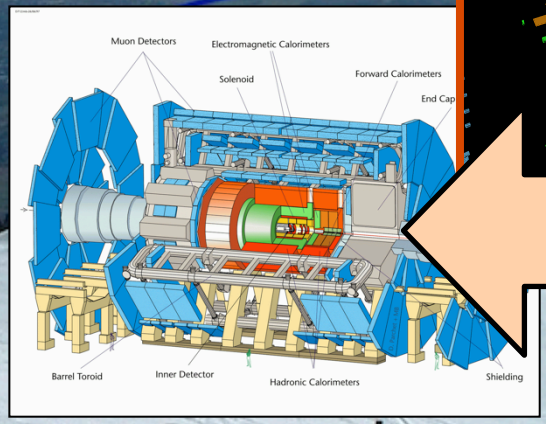

New / refined measurements
to exclude hypotheses

The LHC Accelerator Complex

We need three things…

Glen Crawford / DOE

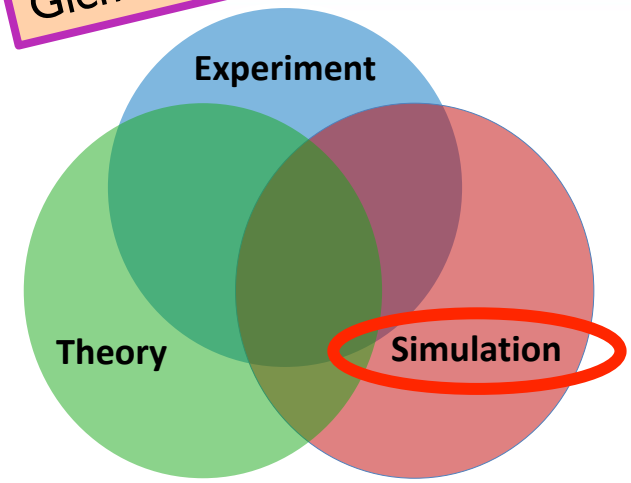

Along Three Paths

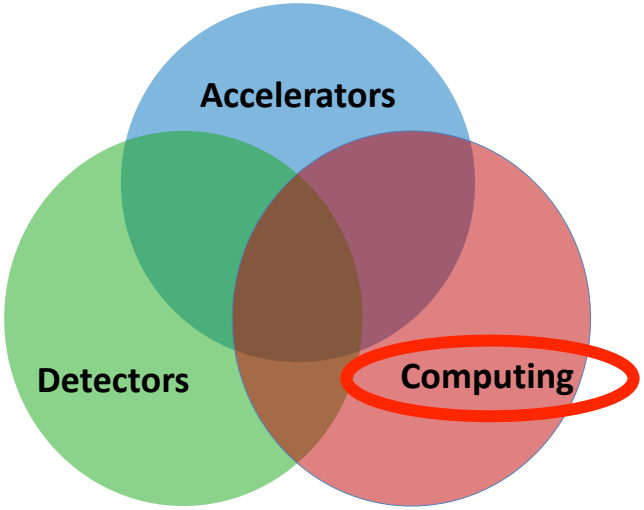

Enabled by Advanced Technologies in:

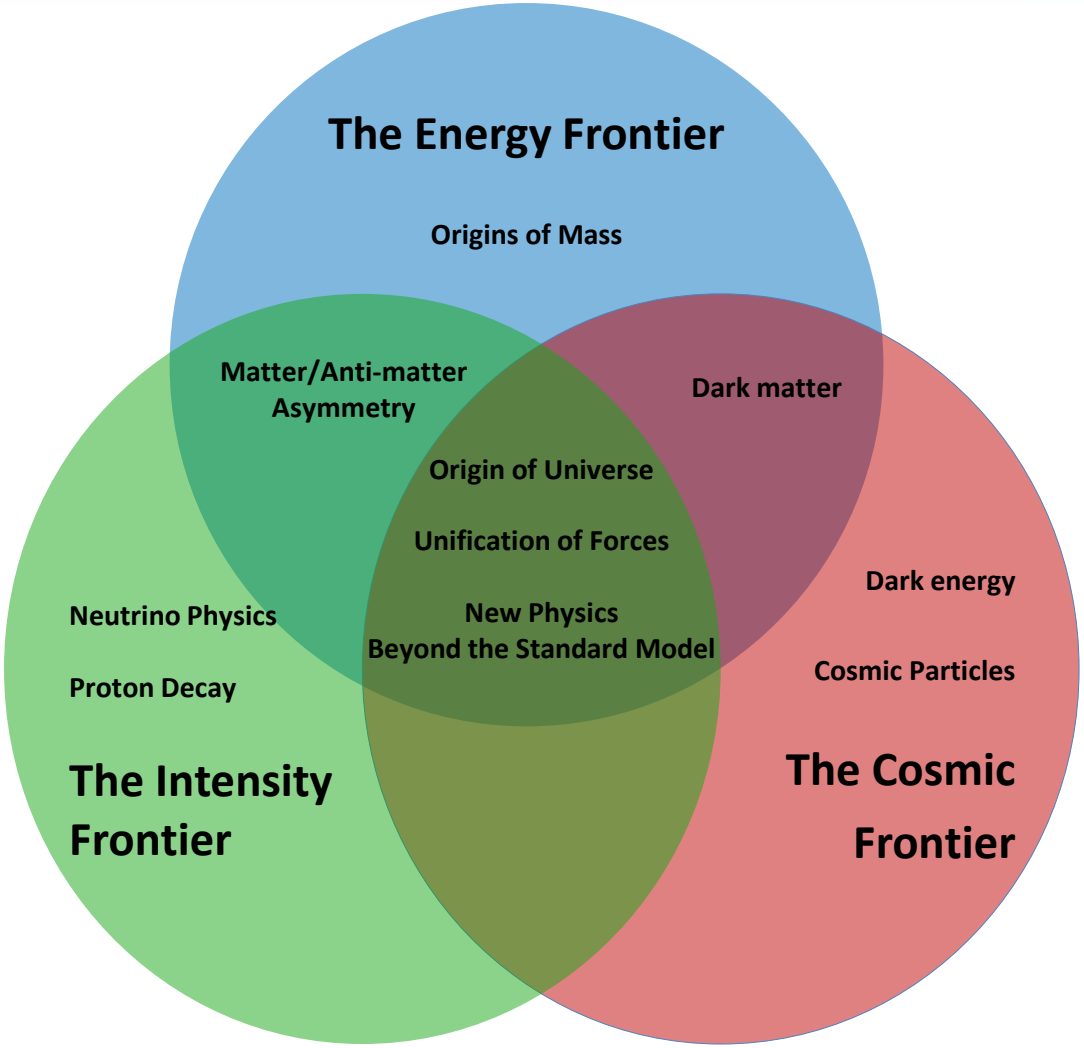

Physics Frontiers

Proton-Proton collisions at LHC – parameters 2012

LHC parameters 2012:

4 TeV	Proton energy
$7*10^{33}cm^{-2}s^{-1}$	Luminosity
1380	Bunches per beam
$1.6*10^{11}$	Protons per bunch
50 ns	Bunch spacing
15 µm	Beam spot size
25	Interactions/bunch crossing

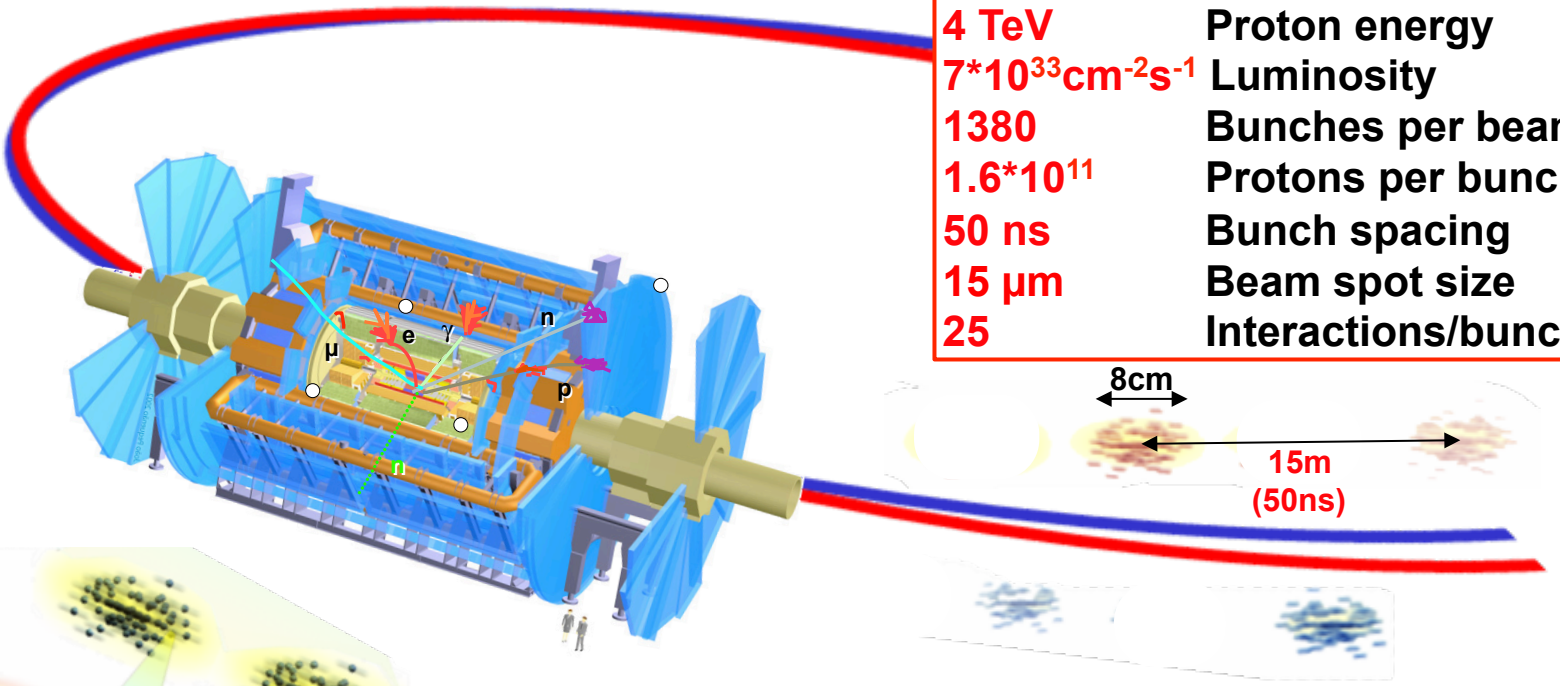

0.4 A
per beam

140 MJoule
per beam

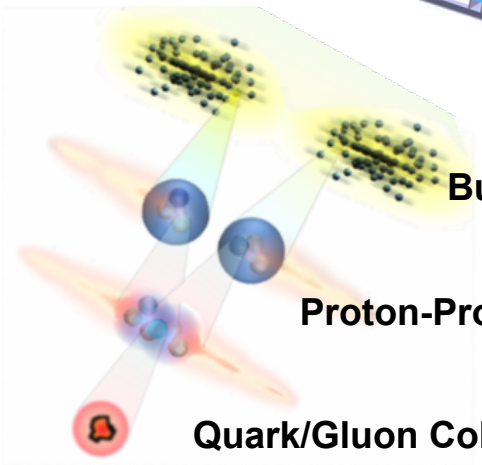

Bunch Crossings $2x10^7$ Hz

Proton-Proton Collisions $0.5*10^9$ Hz

Quark/Gluon Collisions

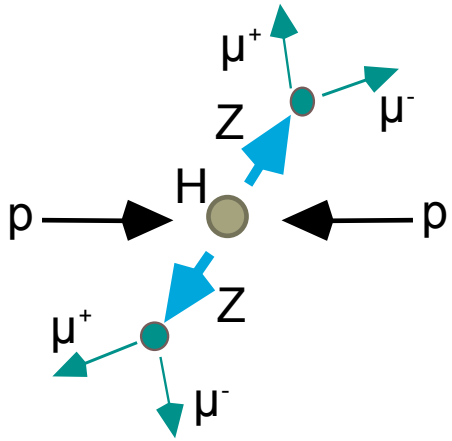

Production of heavy particles $10^{+3...-7}$ Hz
(W, Z, t, Higgs, SUSY,…)

High data rates and high event complexity

20 M bunch crossings / second

25 reactions on average per bunch crossing, overlaid with the interesting interaction (“pileup”)

500 interesting interactions recorded per s: 1-2 B per year

Each takes on average 25 s CPU to reconstruct

Raw data size ~1 MB / event

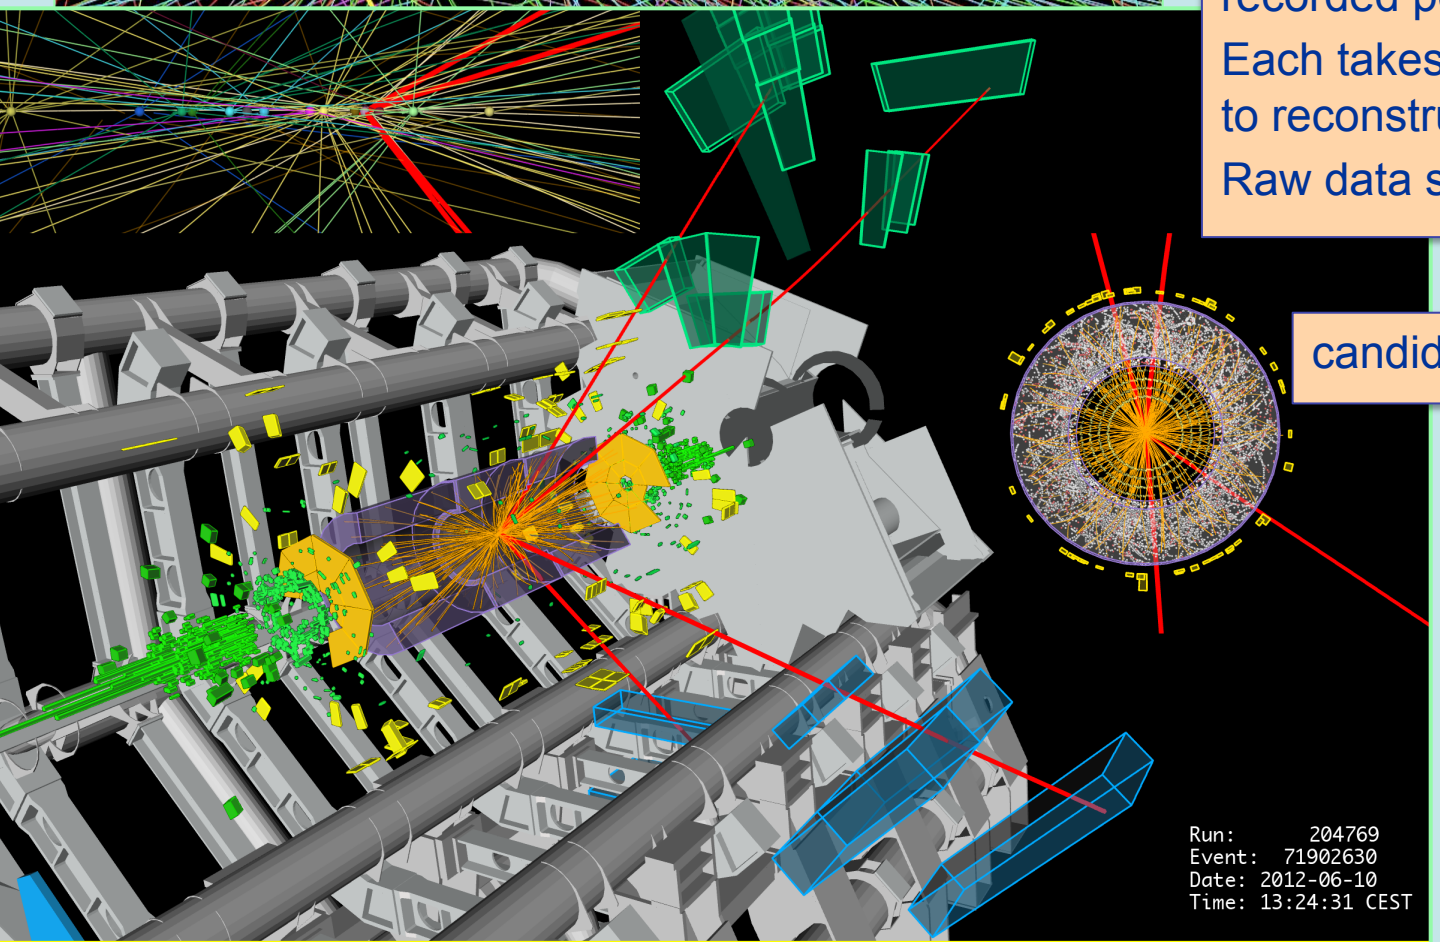

candidate Higgs → 4μ event

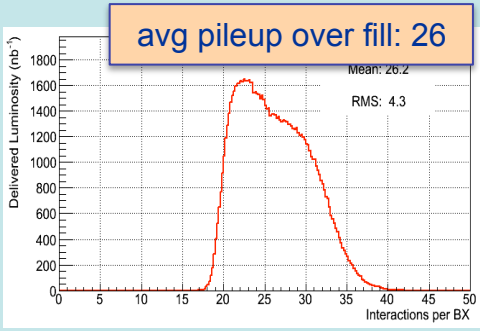

Detector status

ATLAS p-p run: April-June 2012										
Inner Tracker			Calorimeters		Muon Spectrometer				Magnets	
Pixel	SCT	TRT	LAr	Tile	MDT	RPC	CSC	TGC	Solenoid	Toroid
100	99.6	100	96.2	99.1	100	99.6	100	100	99.4	100
All good for physics: 93.6%										

Luminosity weighted relative detector uptime and good quality data delivery during 2012 stable beams in pp collisions at $\sqrt{s}$=8 TeV between April 4th and June 18th (in %) – corresponding to 6.3 fb^{-1} of recorded data. The inefficiencies in the LAr calorimeter will partially be recovered in the future.

Trigger this year

Low-threshold dimuon triggers disabled at highest L due to Level 1 constraint of ~75kHz

Lower (60 GeV) threshold ETMISS trigger below ~ 6 10^{33}/cm^2/s

Trigger optimized to make full use of current ressources

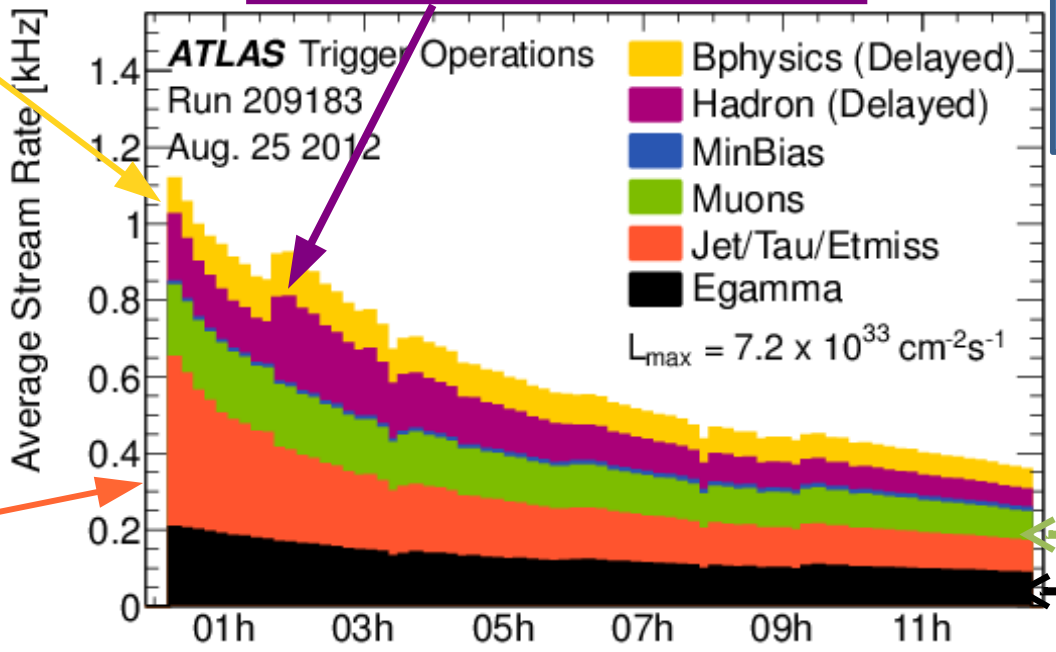

Jet/ETMiss baseline triggers expected to hold till end of the year

Single lepton threshold correspond to offline cuts pT >25 GeV (unchanged during 2012)

(Main) streams sizes up to 3rd technical stop:

Stream	Egamma	Muons	JetTauEtmiss	Total prompt	Hadron delayed	Bphysics delayed	Total Delayed
Events (10^9)	0.47	0.48	0.54	**1.62**	0.22	0.23	**0.47**
Average Rate [Hz]	110	110	120	**370**	50	50	**110**

Distribution of work on the Grid

Tier-0
- RAW data recording into Tape
- FIrst-pass calibrations
- FIrst-pass data processing

Tier-0

CAF

Registration

Transfer

Tier-1

Tier-1 centers
- Reprocessing
- MC Production
- Tape + Disk

Distributed Data Management System (DDM/DQ2)

Transfer

Registration & Transfer Requests

Job

Tier-2

Tier-2 centers
- MC Production
- User Analysis
- Disk storage

Task / Job Definition

Production and Distributed Analysis System (PanDA)

Job

Job Submission

Download

Upload

End-user Clients (pathena/ganga/dq2)

off-grid

Off-Grid facilities
- User Analysis

RAW → ESD → AOD → TAG

ESD → dESD

AOD → dAOD → NTUP

Raw data from detector and data types during reconstruction

WLCG - worldwide LHC computing grid
CPU and storage capacity distributed around the globe

Worldwide LHC Computing Grid (WLCG): ~ 150 computing centres in ~ 35 countries

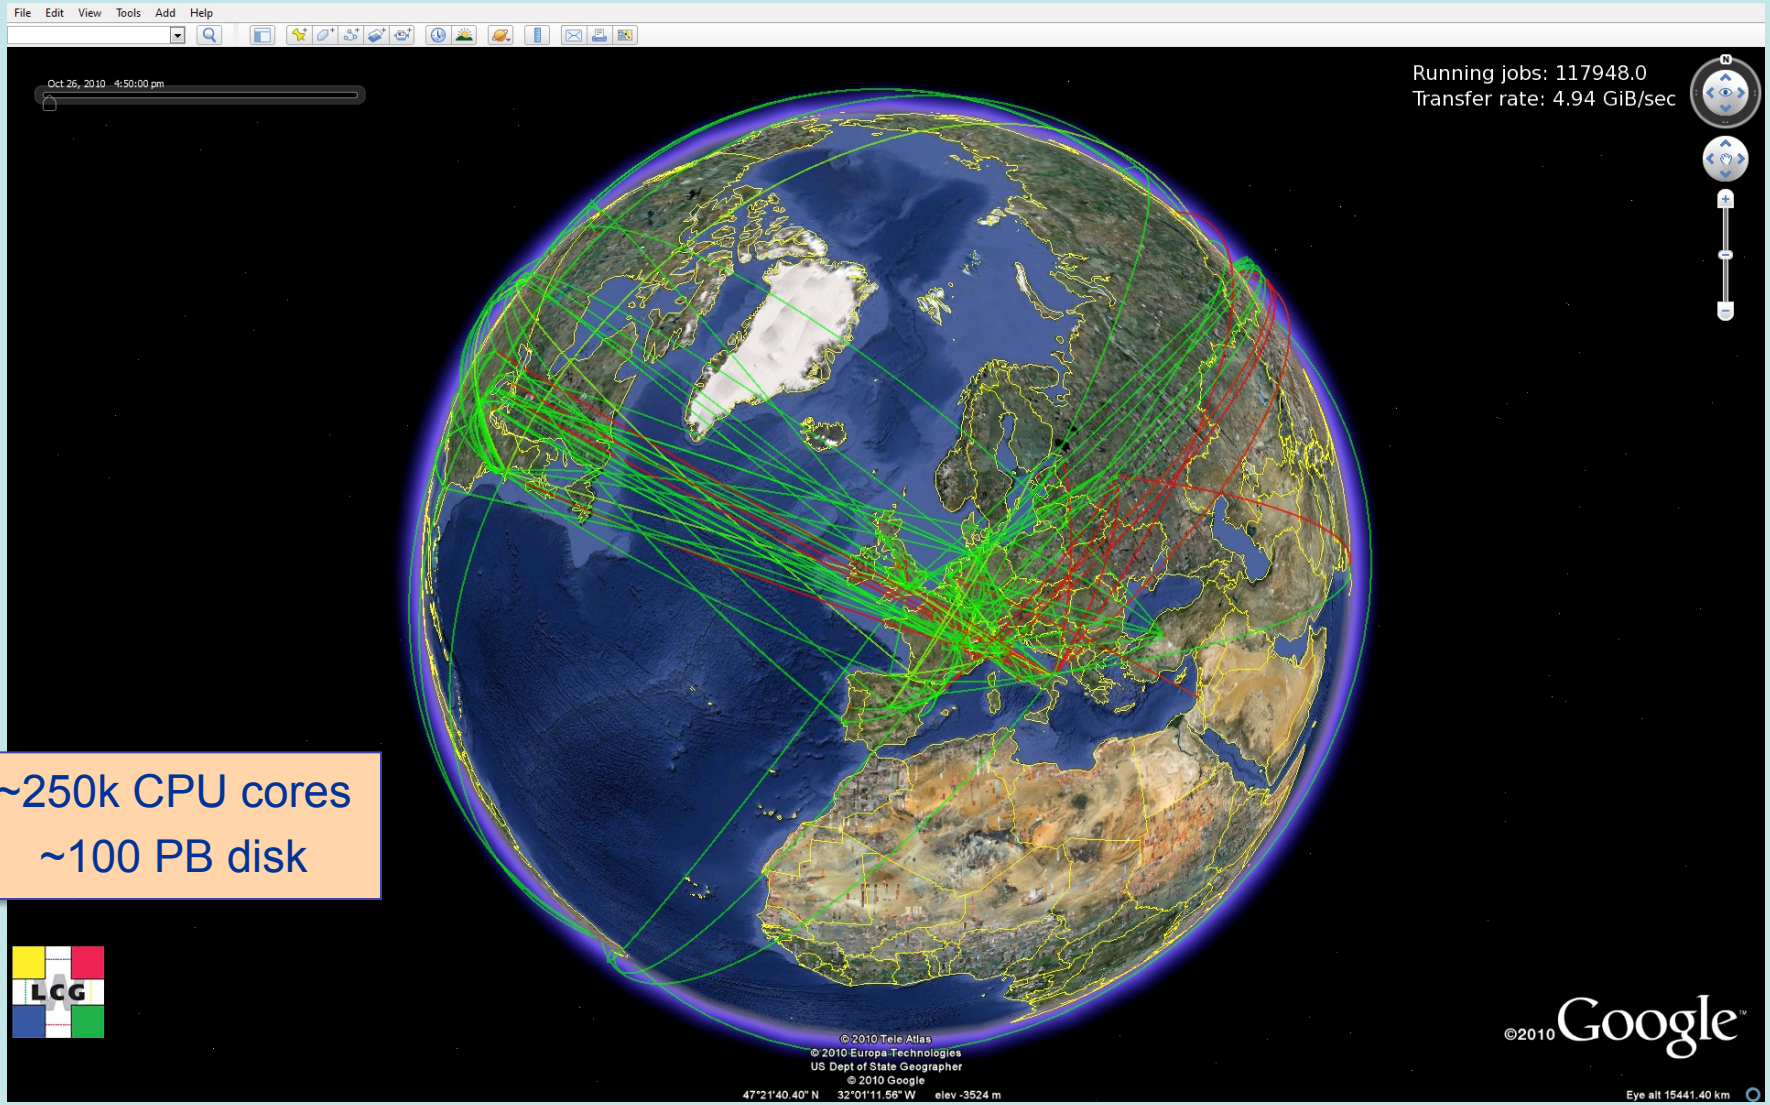

~250k CPU cores
~100 PB disk

Total data volume used on disk and tape
accumulated by ATLAS to date

100 PB

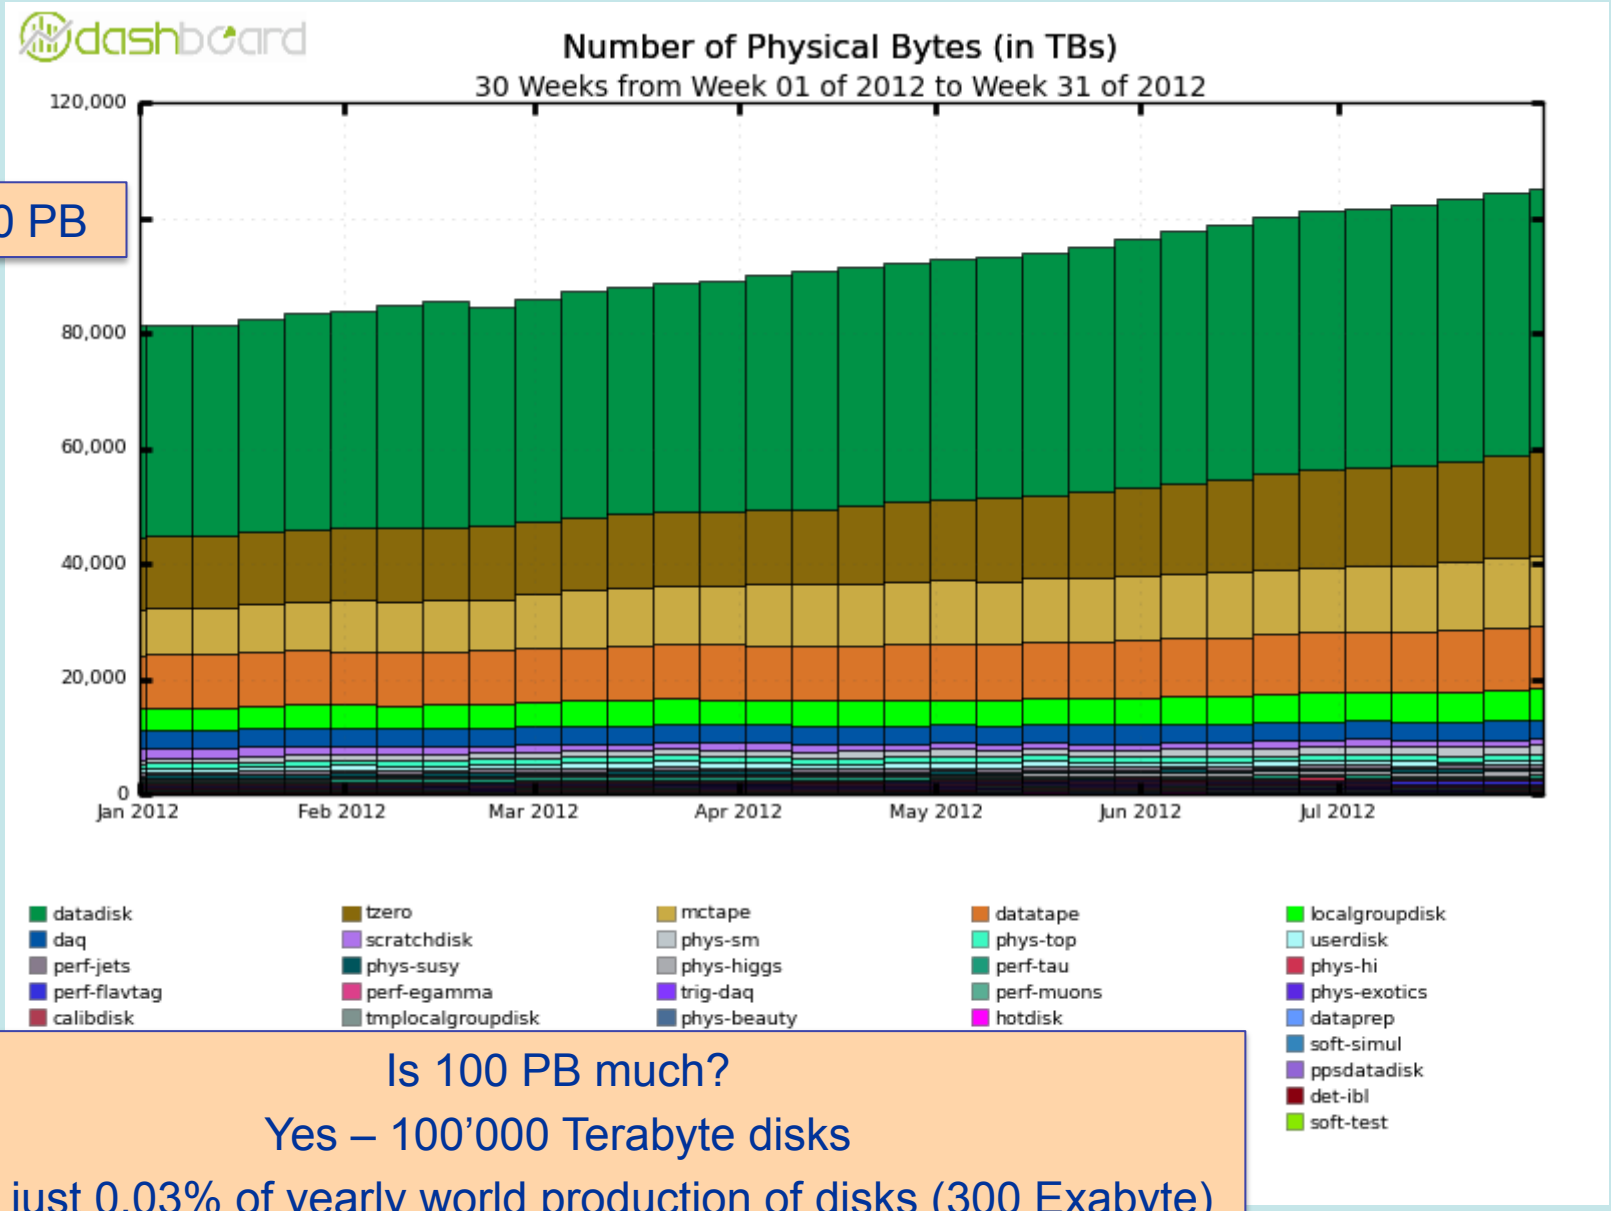

Is 100 PB much?

Yes – 100’000 Terabyte disks

No – just 0.03% of yearly world production of disks (300 Exabyte)

Current status of computing

- **Tier-0: keeping up well, using 6-7000 CPU cores, also done TS1 fast repro**
- **Tier-1 and 2:** We can use much more **CPU** and some more **disk** than pledged - ***thanks!***
- **Extended p-p run gives about 2* planned data volume - covered by extra pledges and early deployment of 2013 resources - *thanks!***

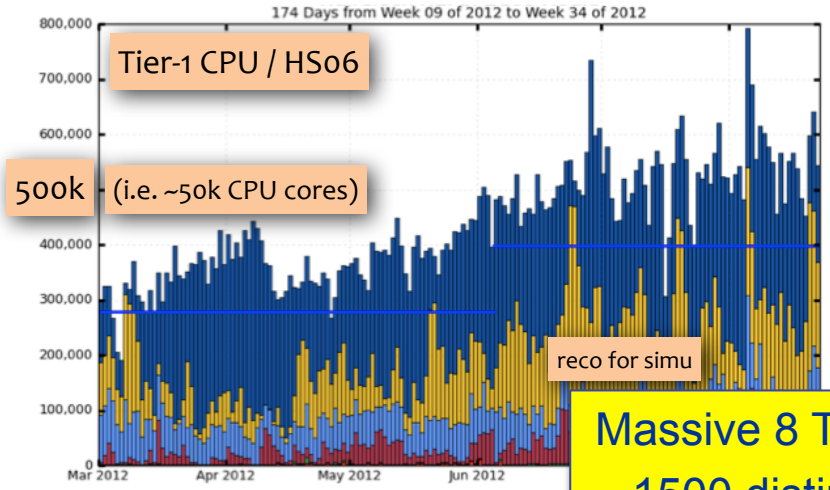

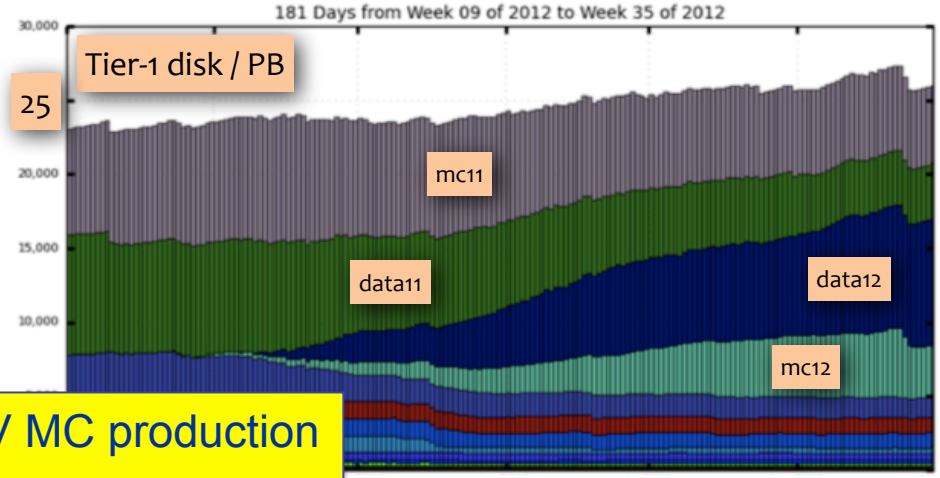

Massive 8 TeV MC production
1500 distinct ATLAS users

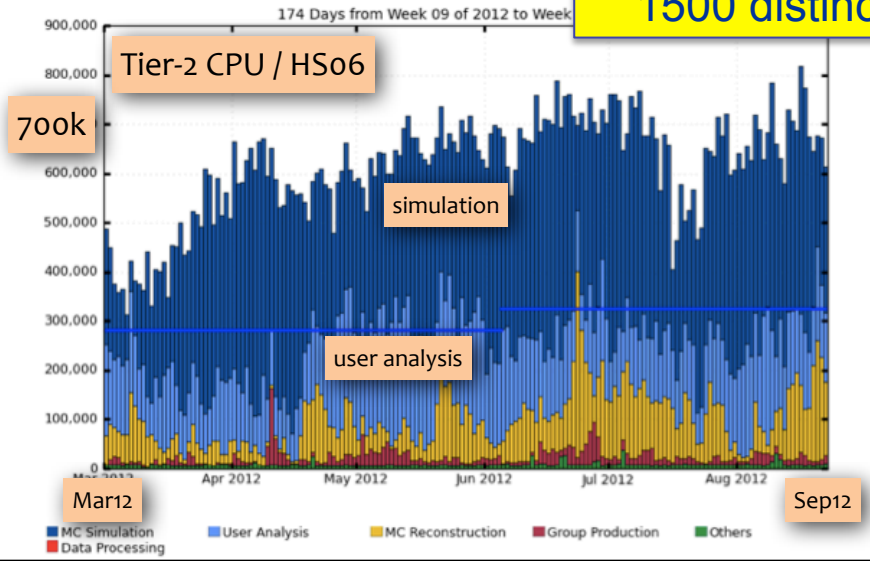

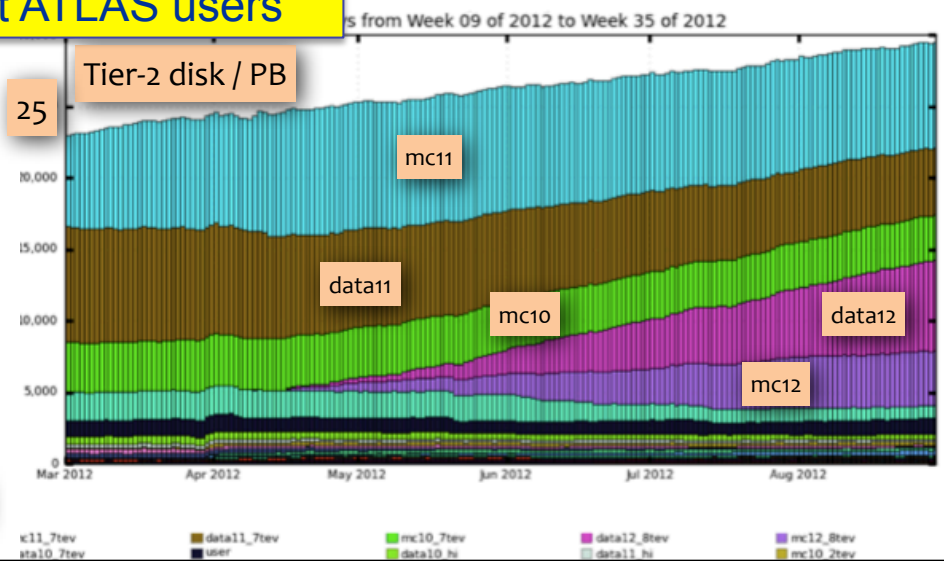

Status and plans for Long Shutdown 1 (LS1)

- **Running very well in 2012, but CPU limitations in MC production**
- **Extra resources for 2012 p-p run extension** are under control: thanks to early deployment! *(plus, had to reduce data distribution)*
- Brief outline of our resource planning guidelines for 2013-2015:
  - **In 2013:**
    - There will be one full reprocessing of 2010-2012 data and MC to further improve the quality of our reconstruction and simulation.
    - (More) new MC for analysis will be produced.
    - Very active group/user analysis.
  - **In 2014:**
    - Largish MC samples for high energy running will be produced and related physics group/user analysis.
    - The final full reprocessing of 2010-2012 data and MC, foreseen to use the evolved event formatting/data model/data distribution prepared for 2015 high-energy data taking.
  - **In 2015:**
    - Processing and reprocessing of new high energy data.
    - Related production of MC samples matching the data.
    - Increased group/user activity.

Data placement

- In order to keep our disk space needs under control we have had to 'tune' down our disk usage (replication policies) and we are making plans to further optimize our Computing Model for 2015.
- A table summarizing the updated replication policy:

Tier-1 disk policy (sum Tier-1s)	**2012**	**2013**	**2014**
Sim RDO disk copies	0	0	0
Sim ESD disk copies (current)	0	0	0
Sim ESD disk copies (previous)	0	0	0
Sim AOD disk copies (current)	2	2	2
Sim AOD disk copies (previous)	0	0	0
Real RAW disk copies	1	1	1
Real ESD disk copies (current)	~~0.2~~ **0.13**	~~0.2~~ **0.05**	~~0.2~~ **0.05**
Real ESD disk copies (previous)	0	0	0
Real AOD disk copies (current)	2	2	2
Real AOD disk copies (previous)	0	0	0
Real DESD disk copies	~~2~~ 1	~~2~~ 1	~~2~~ 1
Tier-1 processing			
Number of reprocessing/year	~~1.5~~ 1	1	1
Tier-2 disk policy (sum of T2's)			
Sim RDO disk copies	0.05	0.05	0.05
Sim AOD disk copies (current)	2	2	2
Sim AOD disk copies (previous)	2	2	2
Sim ESD disk copies (current)	0.2	0.1	0.1
Sim ESD disk copies (previous)	0.2	0.1	0.1
Real RAW disk copies	0	0	0
Real AOD disk copies (current)	2	2	2
Real AOD disk copies (previous)	2	2	2
Real DESD disk copies	~~2~~ 1	~~2~~ 1	~~2~~ 1
Real DESD disk copies (previous)	~~2~~ 1	~~2~~ 1	~~2~~ 1

Extra copies of course!

Popular data (group production, AODs..) replicated dynamically by PD2P

Summary of resource usage and needs

2012 extended to 2015, based on present sw release

LHC and ATLAS parameters		*Actual p-p (Mar-Aug 2012)*	*Revised model 2012 p-p*	*Model 2012 HI*	*Model 2013 p-p*	*Revised model 2014 p-p*	*Model 2015 p-p (25 ns, mu=19)*
Trigger Rate	events/sec	350 (prompt) + 150 (delayed)	400 (prompt) + 200 (delayed)	200	0	~~1000~~ 0	1000
Live Time	Msec total	3.8	~~3.8~~ 7.3	0.7	0	~~1.6~~ 0	3.8
Average pileup over fill	events	20	25	0	23	23	19
Real data (prompt only)	Bevents total	1.3	~~1.5~~ 2.9	0.14	0	~~1.6~~ 0	3.8
Full Simulation	Bevents total	2.4	1.9	0.01	4	2.5	2
Fast Simulation	Bevents total	1.4	2	0	4.6	4	5
Real RAW	MB/event	0.79	0.8	5	0.8	0.8	0.8
Real ESD	MB/event	2.4	2.5	3	2.5	2.5	2.7
Real AOD	MB/event	0.24	0.35	1	0.35	0.35	0.23
Simulated HITS	MB/event	0.9	1	5	1	1.6	1.6
Simulated ESD	MB/event	3.3	3.5	3	3.5	3.5	3.9
Simulated AOD	MB/event	0.4	0.5	1	0.5	0.5	0.5
Full Simulation	HS06sec/event	3100	3300	48000	3300	5400	5400
Fast Simulation	HS06sec/event	260	310	-	310	500	500
Real Reconstruction	HS06sec/event	210	230	480	230	230	230
Simulation Reconstruction	HS06sec/event	770	830	1200	830	830	830
Group analysis	HS06sec/event		20	20	20	20	20
User analysis	HS06sec/event		0.4	0.2	0.4	0.4	0.4

Luminosity of a hadron collider

$$L = \frac{N^2 k_c f}{4\pi\sigma_x\sigma_y} F = \frac{N^2 k_c f_0 \gamma}{4\pi\,\varepsilon_n \beta^*} F(\theta_c)$$

Hour glass factor: $F = 1 / \sqrt{1 + \left(\frac{\theta_c \sigma_z}{2\sigma^*}\right)^2}$

- Parameters in luminosity
  - No. of particles per bunch N
  - No. of bunches per beam k_b
  - No. of bunches collliding at IP k_c ($k_c < k_b$)
  - Relativistic factor γ
  - Normalised emittance ε_n
  - Beta function at the IP β^*
  - Crossing angle factor F
    - Full crossing angle θ_c
    - Bunch length σ_z
    - Transverse beam size at the IP σ^*

Equal amplitude functions:

$\beta_x^* = \beta_y^* = \beta^*$,

Geometric and normalised emittance:

$$\varepsilon_x^* = \varepsilon_y^* = \varepsilon^* = \frac{\varepsilon_n}{\sqrt{\gamma^2 - 1}}$$

$\Rightarrow$ Round beams at IP:

$$\sigma_x^* = \sigma_y^* = \sigma^* \quad \sqrt{\frac{\beta^* \varepsilon_n}{\gamma}}$$

(N.B. LHC uses RMS emittances.)

ATLAS pushes for 25ns right from beginning of Run2

What if 50ns…

	a): 25 ns	b): 50 ns	c): 50 ns, levelled	d): 50 ns, levelled, extended fill
Initial <μ>	24	2.86*24 = 69	2*24 = 48	2*24 = 48
Average <μ> over fill	0.7*24 = **17**	2.86*17 = 48	1.76*24 = 42	1.57*24 = **37**
Relative fill duration	1	1	1	1.43
Relative avgerage luminosity during fill	1	1.43	1.26	1.11

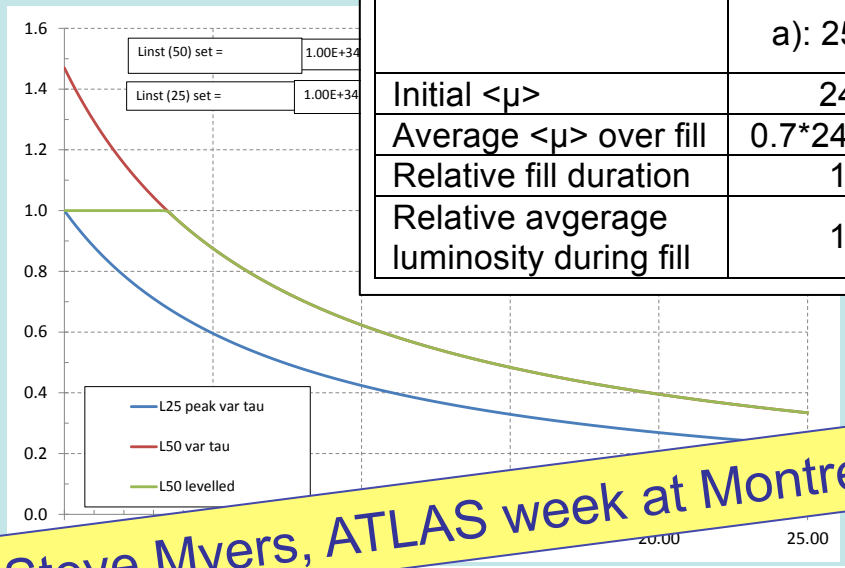

Steve Myers, ATLAS week at Montreux

Standard scenarios used in ATLAS 25ns / 50ns comparisons:

	s1): 25 ns	s2): 50 ns
Initial <μ>	27	2*27 = 54
Average <μ> over fill	0.7*27 = **19**	1.4*27 = **38**
Relative fill duration	1	1
Relative avgerage luminosity during fill	1	1

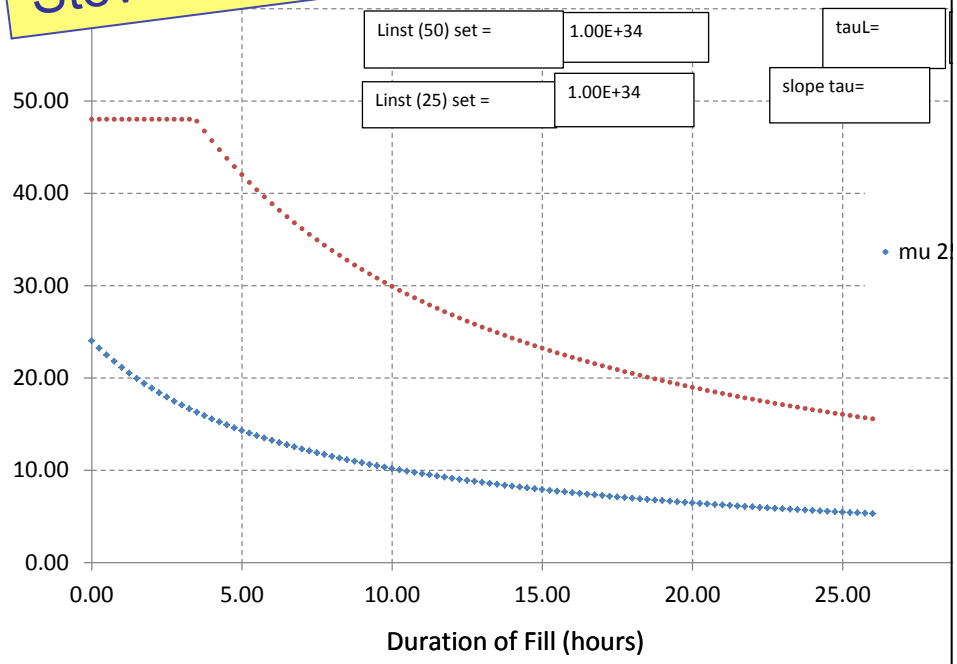

LHC and ATLAS parameters		Model 2015 25 ns, $\mu = 19$	Model 2015 50 ns, $\mu = 38$
Trigger rate	event/sec	1000	1000
Livetime	10^6 sec total	3.8	3.8
Fill-averaged pileup	events	19	38
Real data (prompt only)	10^9 events	3.8	3.8
Full Geant4 simulation	10^9 events	2	2
Fast simulation	10^9 events	5	5
Real RAW data	MB/event	0.8	1.2
Real ESD data	MB/event	2.7	4.6
Real AOD data	MB/event	0.23	0.42
Simulated HIT data	MB/event	1.6	1.6
Simulated ESD data	MB/event	3.9	8
Simulated AOD data	MB/event	0.5	0.9
Full simulation CPU	HS06sec/event	5400	5400
Fast simulation CPU	HS06sec/event	500	500
Real reconstruction CPU	HS06sec/event	230	420
Simulation reconstructioin CPU	HS06sec/event	830	1430
Group analysis	HS06sec/event	20	20
User analysis	HS06sec/event	0.4	0.4

Luminosity and pileup today

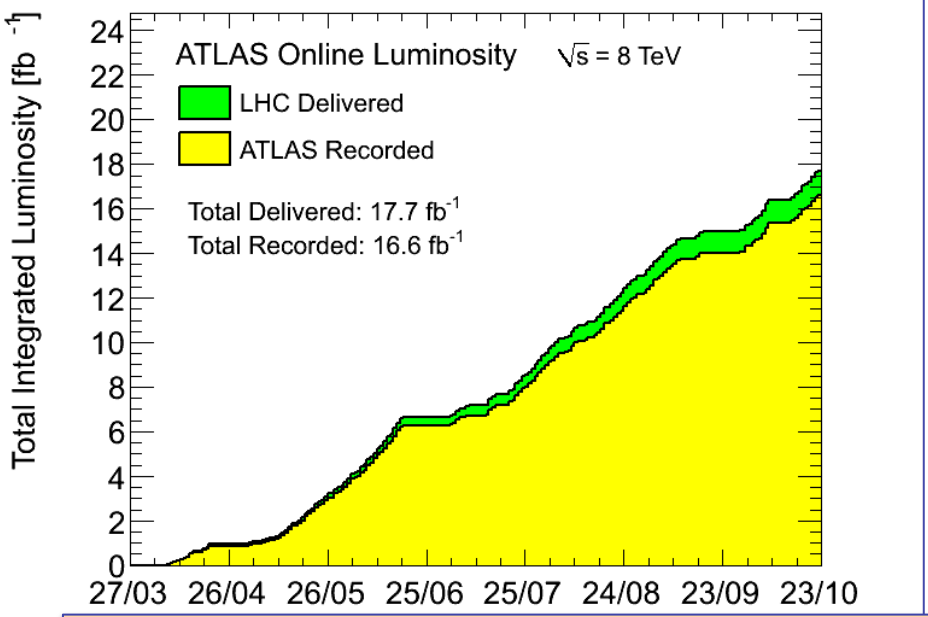

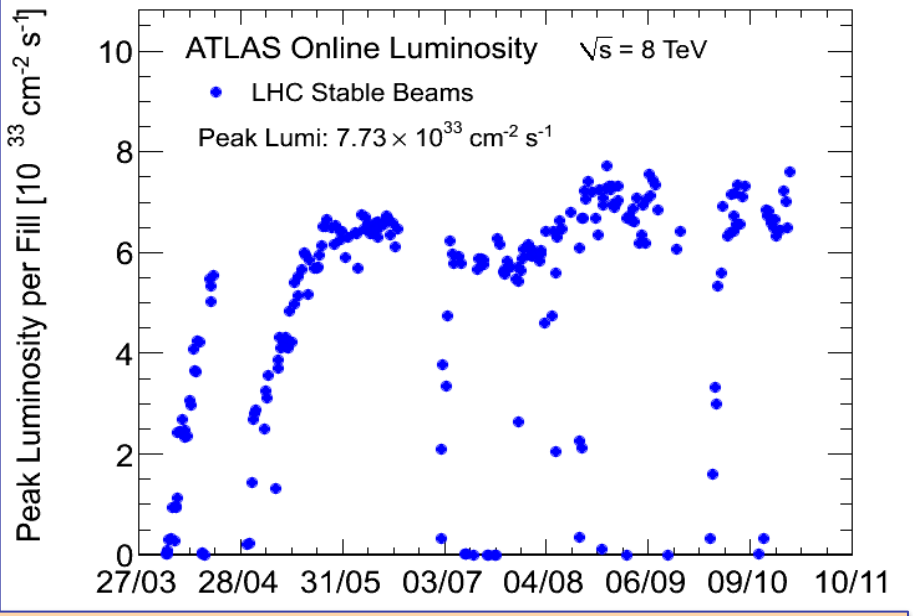

Now ~16/fb recorded at 8TeV – hope for 20-25/fb by mid Dec – plus the 5/fb at 7 TeV

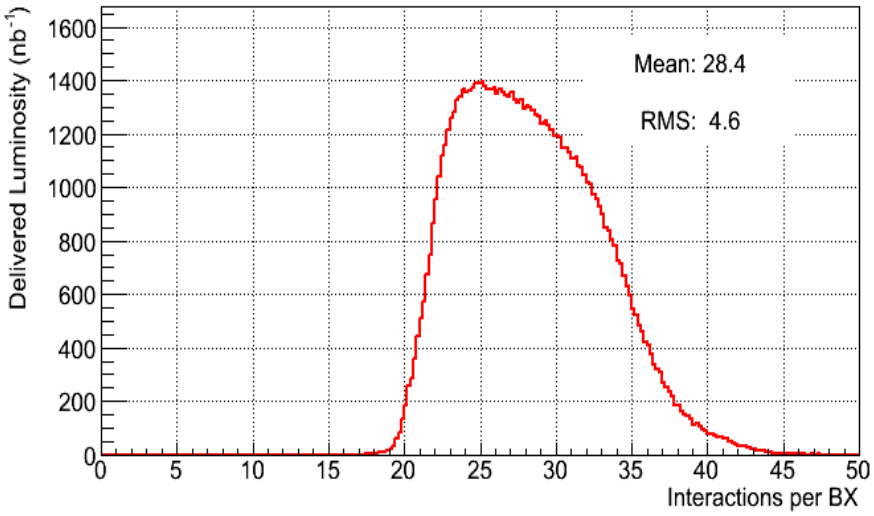

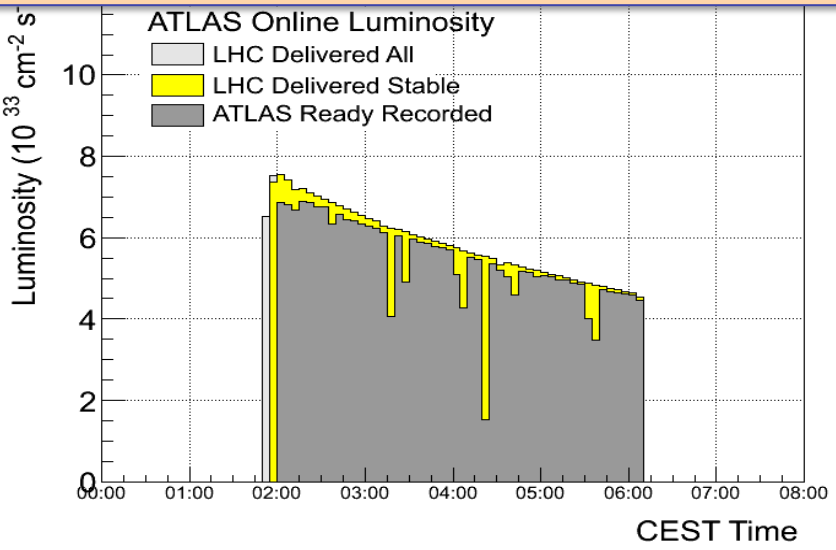

Resource needs for after LS1 (2015):
assume 1kHz, 25ns (but consider higher pileup – 50ns, or early 2-3e34)

- **Assume we will have a flat monetary budget in future for ATLAS computing**
  - this message comes from WLCG and from the scrutiny group
  - increase in Grid resources (disk, CPU, network) by technological progress
    - hope for *extra centres* (all Tiers!)
    - try exploit additional resources, e.g. in *high-performance computing centres, Clouds*
- **Consequently, growing investment in software development required (from reconstruction to distributed computing)**
  - change EDM, make efficient use of CPUs' vector units, caches, many-cores
  - promising development in WAN data access, event-level caching, and application driven network operation to help reduce overall disk volume needs
- **Our credo: computing can't be the limiting factor in LHC physics results**
  - despite growing LHC demands and flat computing budgets
  - need to optimise the usage of the given CPU and storage very substantially, with extra effort, and we ask the collaboration for extra institutional commitments
  - need a long enough lead time for such development - take full advantage of LS1 !
  - the kind of software work involved should be very tempting to a few of you ...

Software & Computing plans for LS1

- **A focus of the S&C workshop last week**
- **CPU speedup, software configuration and build - Reconstruction-type software**
- **Data volume, resources management - ADC-type software**
  - New Distributed Data Management system being implemented
  - WAN data access and data caching (file level, event level)
  - Applications-driven usage of networks
  - New MC production system
  - Tier-0 filtering/streaming
  - Database cleaning-up
- **Simulation: Integrated Simulation Framework**
  - Selection of full/fast/... mode per subevent
  - Speedup of G4 itself tackled by G4 collaboration (Geant4MT, ...)
  - Concurrency within event
  - Plus benefit from reconstruction speedup: offline sw improvements, truth-seeding.

A slide on CPU speedup
(Rolf Seuster)

Instruction Level Parallelism

a.k.a. hyper threading

Single Intruction
Multiple Data

what state of
art CPU offers

well written and
optimized
programs utilize

typical HEP
software utlises

Where are we now?

	SIMD	ILP	HW THREADS	CORES	SOCKETS
MAX	4	4	1.35	8	4
TYPICAL	2.5	1.43	1.25	8	2
HEP	1	0.80	1	6	2

product sum of above

	SIMD	ILP	HW THREADS	CORES	SOCKETS
MAX	4	16	21.6	172.8	691.2
TYPICAL	2.5	3.57	4.46	35.71	71.43
HEP	1	0.80	0.80	4.80	9.60

HEP utilizes only small percentage of what CPU provides

Andrzej Nowak - The growth of commodity computing and HEP software – do they mix?

10

IMHO: it's not THAT bad, but you get the picture ...

A slide on the new production system
(Andrej Filipcic)

- A major re-design of a production system agreed upon
- Core components flexibility → a generic workflow engine, which could hopefully handle “any” future requirements
- Gradual implementation
- Much more flexible and automated system to be in place for the 14TeV operation
- The details will be discussed in the SW&C week in October. All people are welcome to participate. Input from Physics Community is essential.

A slide on the new Distributed Data Management
(Vincent Garonne)

- Current system: DQ2
  - 120 PB
  - 350 million files
  - 800 users
  - 130 sites

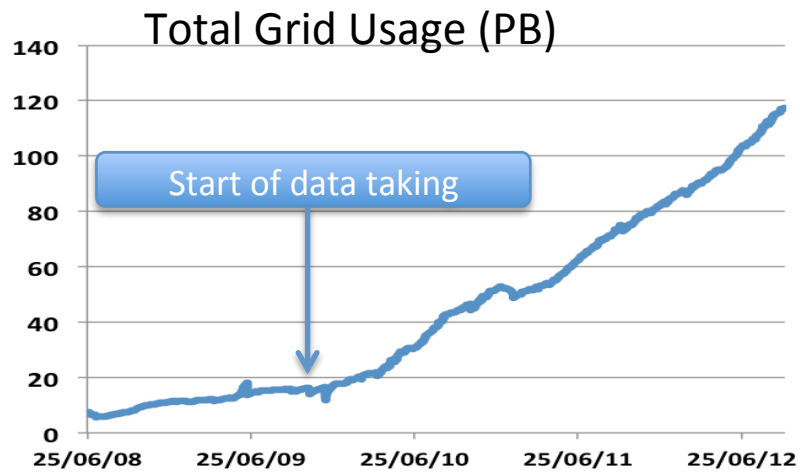

- New major DDM version Rucio to ensure system scalability, reduce operational overhead and support new ATLAS use cases
  - Examples: Better support for group activities; use new technologies
- Timeline

2011. Technical meetings with other LHC experiments
User surveys, Collection of use cases, Rucio Conceptual model

2012. Parallel and incremental development track
Nov. 2012: First prototype with limited functionality

2013. Preparation of migration from DQ2 to Rucio
Preparatory steps, Functional tests
Gradual migration of the external applications (e.g., PanDA)

2014. Rucio in production after LS1

Integrated Simulation Framework (ISF)
(Andi Salzburger)

Simulation Setups

- **Full Geant4** (MC12, ISF)
  G4 in all subdetectors

- **ATLFASTII** (MC12, ISF)
  ID: **Geant4**
  Calo: **Geant4** for muons, **FastCaloSim** for everything else
  MS: **Geant4** (only muons can reach MS, everything else gets *absorbed* by FCS)

- **ATLFASTIIF** (under development, ISF only)
  ID: **Fatras**
  Calo: **Fatras** for muons, **FastCaloSim** for everything else (option to run **parameterized punch-through** simulation)
  MS: **Fatras**

- **FastGamma** (under development, ISF only)
  ID: **Fatras** for particles in cones around EvGen photons
  Calo: **FastCaloSim**

Integrated Simulation Framework (ISF)

	ISF_Kernel::execute/evt	speed-up
○ **Full Geant4** measured with 10 events	560 s	1
○ **ATLFASTII** measured with 100 events	25 s	~25
○ **ATLFASTIIF** measured with 1000 events	0.75 s	~750
measured with 1000 events ○ **FastGamma**	0.18 s	~3000

Higgs couplings

More physics updates: see Open LHCC presentation, H.Bachacou, 26.9.2012

ATLAS-CONF-2012-127

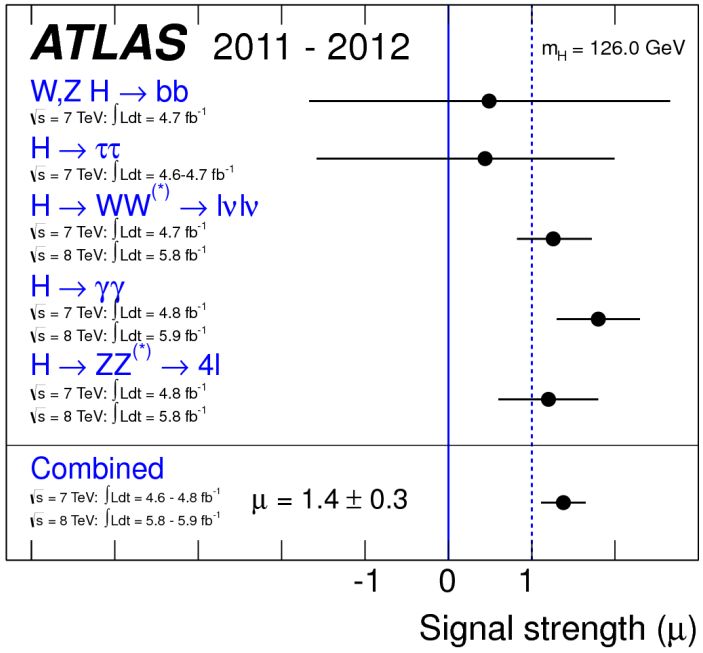

- Using the Higgs LHC cross-section working group formalism: hep-ph: 1209.0040
- Projected sensitivity for 300 and 3000 fb^{-1} studied in the context of ESPP: ATL-PHYS-PUB-2012-001

2D fit of coupling to Fermions (kF) and Vector bosons (kV)

In SM: kF = kV = 1

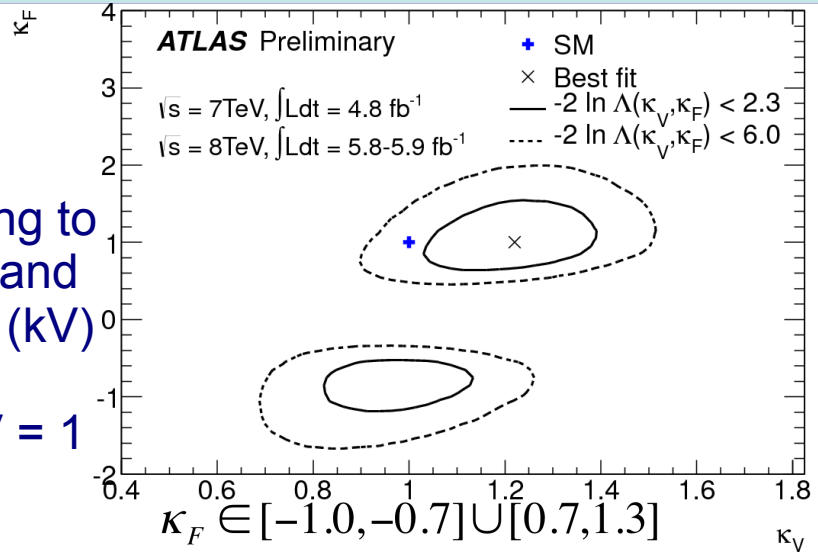

$\kappa_F \in [-1.0,-0.7] \cup [0.7,1.3]$

$\kappa_V \in [0.9,1.0] \cup [1.1,1.3]$

Probe non-SM contributions to gg → H and H→ γγ loops

In SM: kg = kγ = 1

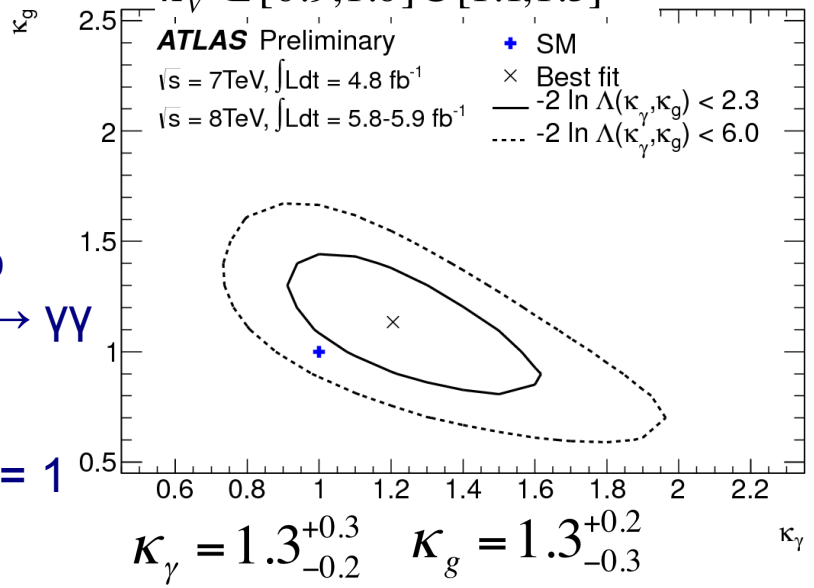

$\kappa_\gamma = 1.3^{+0.3}_{-0.2} \quad \kappa_g = 1.3^{+0.2}_{-0.3}$

Plan for the rest of this year

p-Pb next year – LS1 starting 11 February 2013

	Oct				Nov					Dec			
Wk	40	41	42	43	44	45	46	47	48	49	50	51	52
Mo	1		15	22	29	5	12	Scrubbing				17	24
Tu				Floating MD [24 h]				25 ns set-up	25 ns physics				Xmas
We		MD 3		500+ m [24 h]									
Th												STANDBY	
Fr							Scrubbing run (date tbc)	MD 4					
Sa													
Su													

…but plans change quickly. At present, physics data taking instead of machine development:

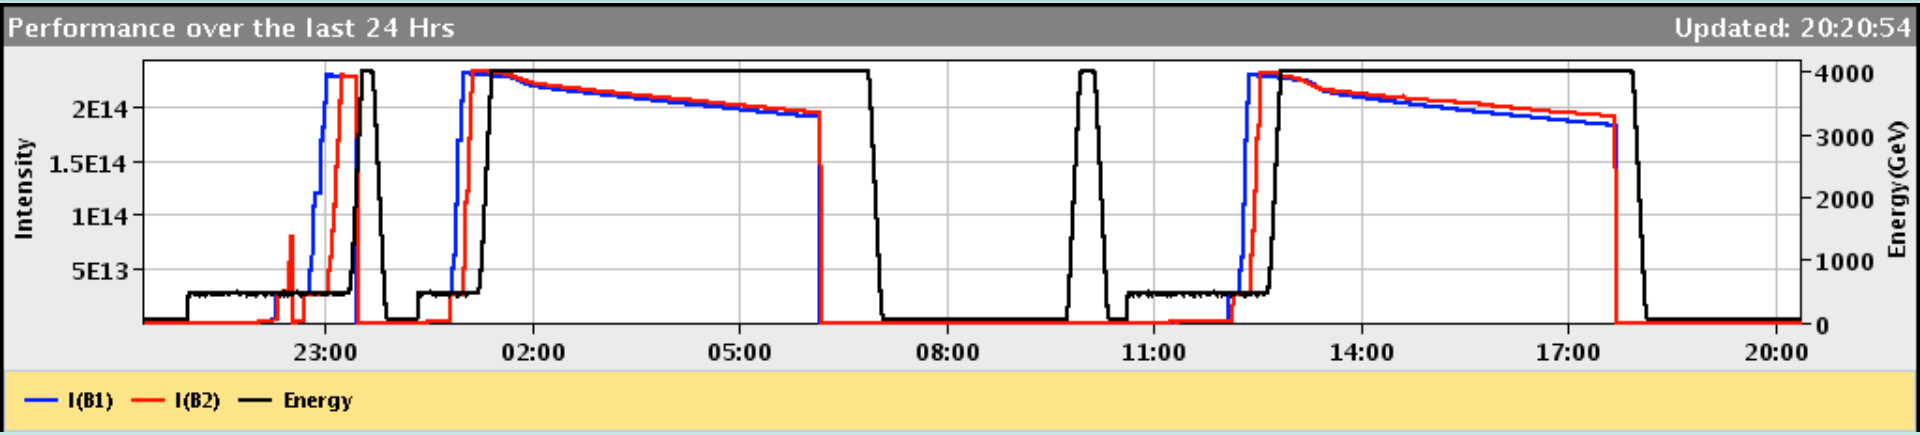

Thanks to many people for the material !

including

Steve Myers, Rolf Seuster, Andi Salzburger, Andrej Filipcic, Borut Kersevan, Ikuo Ueda & the entire ADC team

Thank you for listening!

გმადლობთ!

CERN was founded 1954

by 12 European states – today: 20 member states

Distribution of All CERN Users by Nation of Institute on 9 January 2012

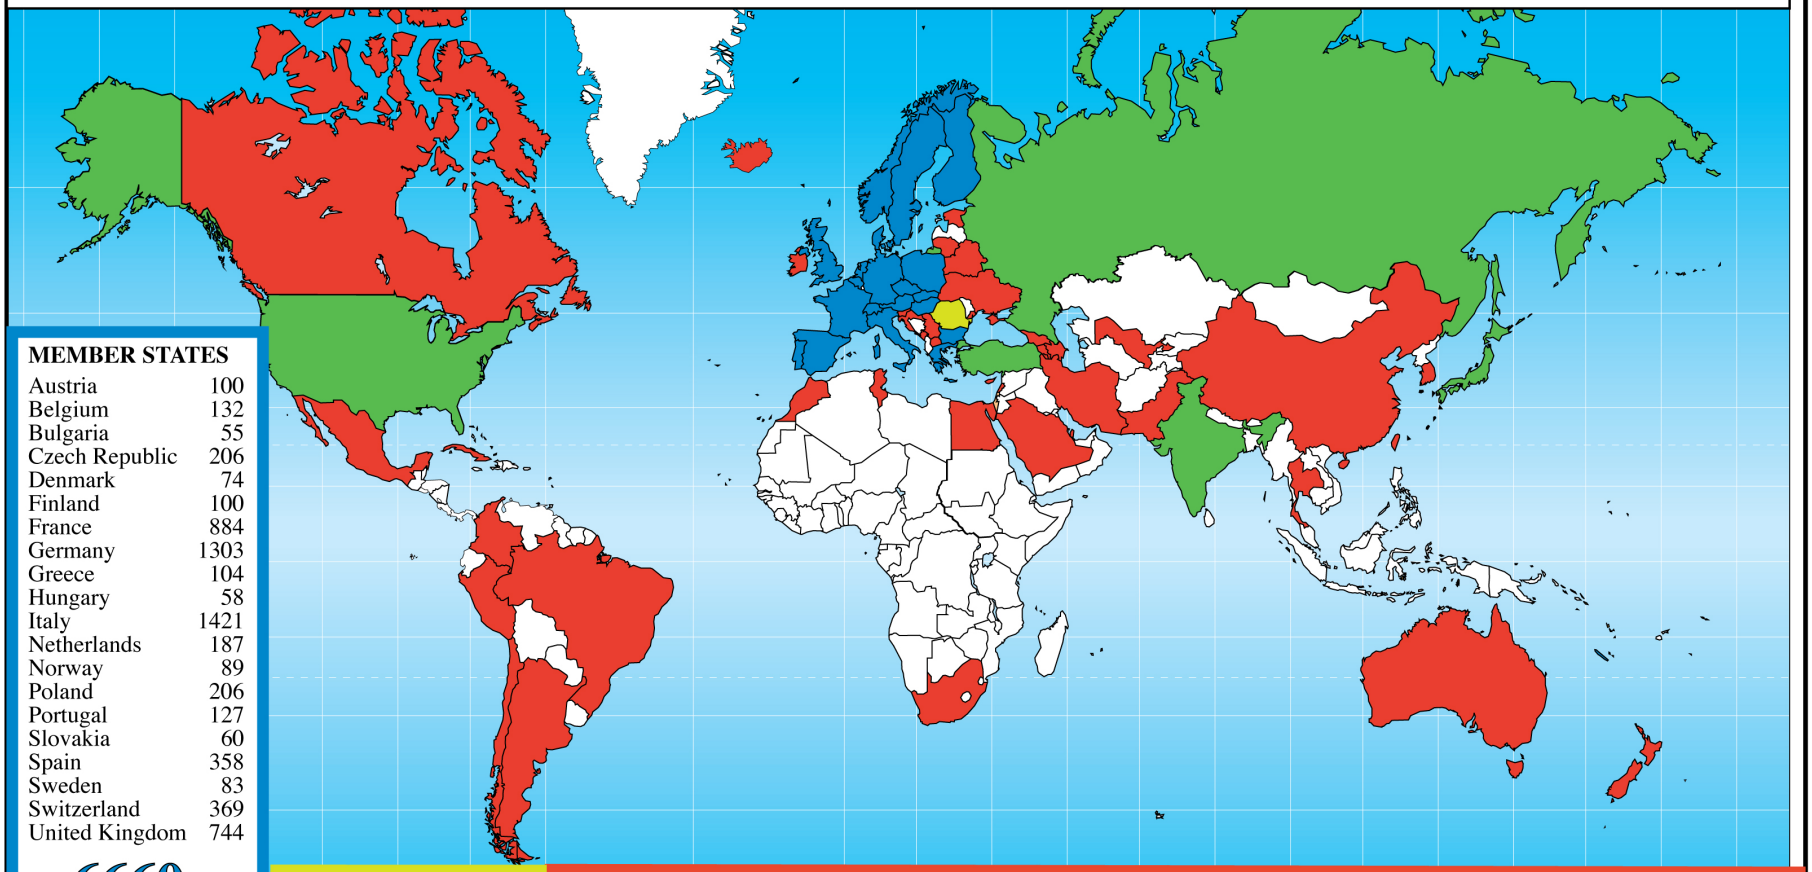

MEMBER STATES

Austria	100
Belgium	132
Bulgaria	55
Czech Republic	206
Denmark	74
Finland	100
France	884
Germany	1303
Greece	104
Hungary	58
Italy	1421
Netherlands	187
Norway	89
Poland	206
Portugal	127
Slovakia	60
Spain	358
Sweden	83
Switzerland	369
United Kingdom	744

6660

OBSERVERS

India	115
Japan	225
Russia	856
Turkey	77
USA	1708

2981

CANDIDATE FOR ACCESSION

Romania	75

ASSOCIATE MEMBER IN THE PRE-STAGE TO MEMBERSHIP

Israel	62

OTHERS

Argentina	18	China	95	Iran	14	Pakistan	19	Ukraine	21
Armenia	12	China (Taipei)	67	Ireland	10	Peru	2	Uzbekistan	1
Australia	24	Colombia	10	Korea	89	Qatar	1		
Azerbaijan	1	Croatia	17	Lebanon	1	Saudi Arabia	3		
Belarus	22	Cuba	4	Lithuania	12	Serbia	26		
Brazil	93	Cyprus	9	Malta	1	Slovenia	37		
Canada	167	Egypt	7	Mexico	43	South Africa	21		
Chile	4	Estonia	18	Montenegro	1	Thailand	5		
		Georgia	10	Morocco	5	T.F.Y.R.O.M.	2		
		Iceland	3	New Zealand	11	Tunisia	1		

907

Higgs combination
Update since presentation at CERN Council of 4 July

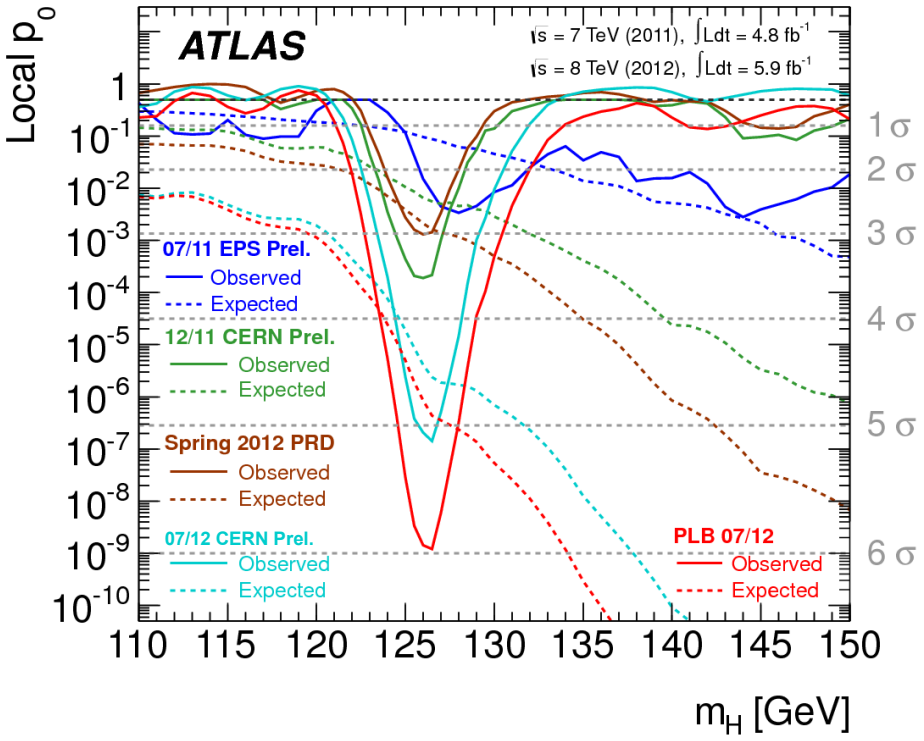

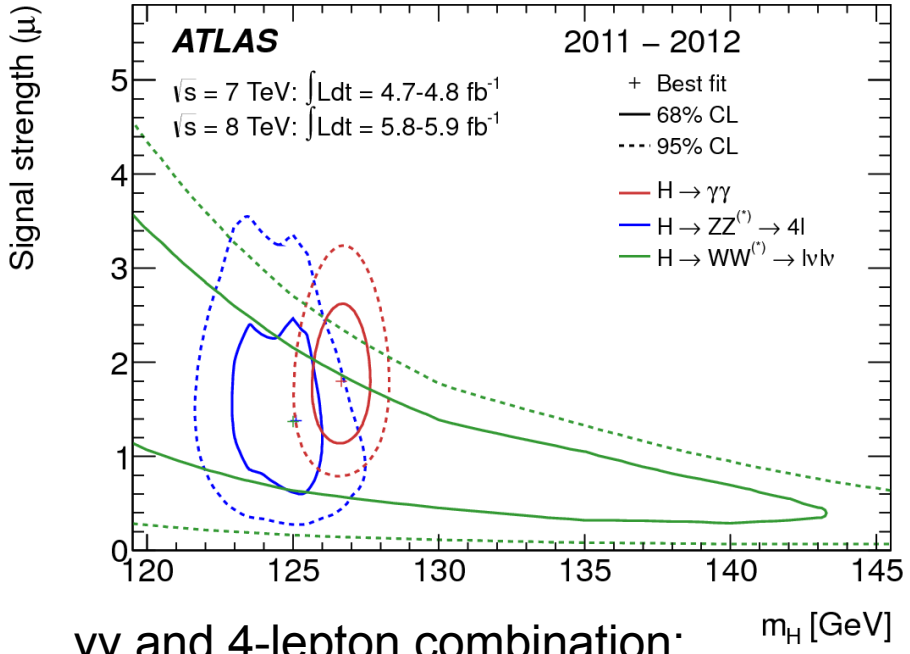

γγ and 4-lepton combination:
m(H) = 126.0 ± 0.4 ± 0.4 GeV

4th July → *Phys. Lett. B 716 (2012) 1-29*

Expected :	4.6 σ	→	4.9 σ
Observed local :	5.0 σ	→	5.9 σ
Observed global* :	4.1 σ	→	5.1 σ

* full range

H → evμv, 8TeV data

- Since July 4th, publication of observation paper including 8 TeV WW → evμv channel
- 2.8 sigma excess in this channel alone (7+8 TeV data)

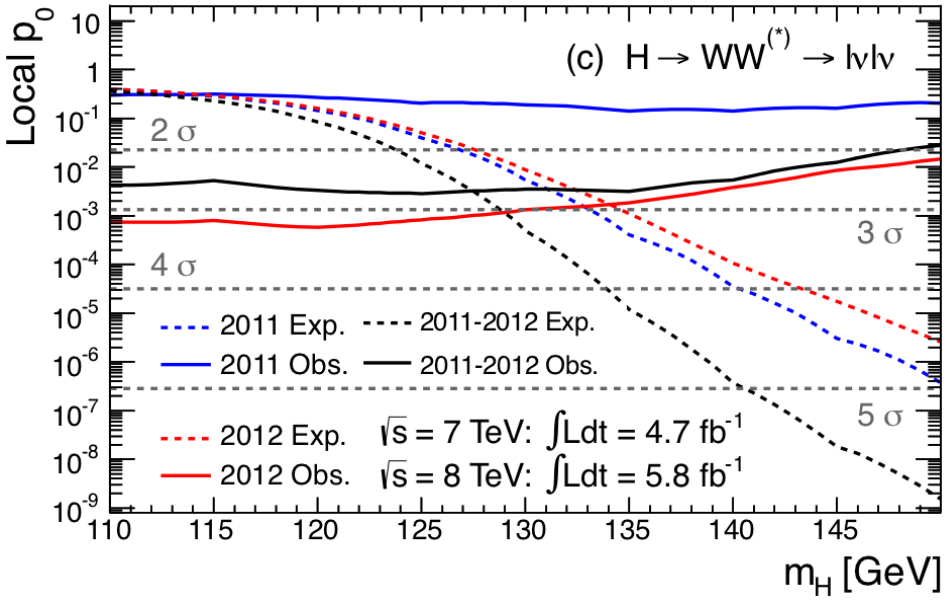

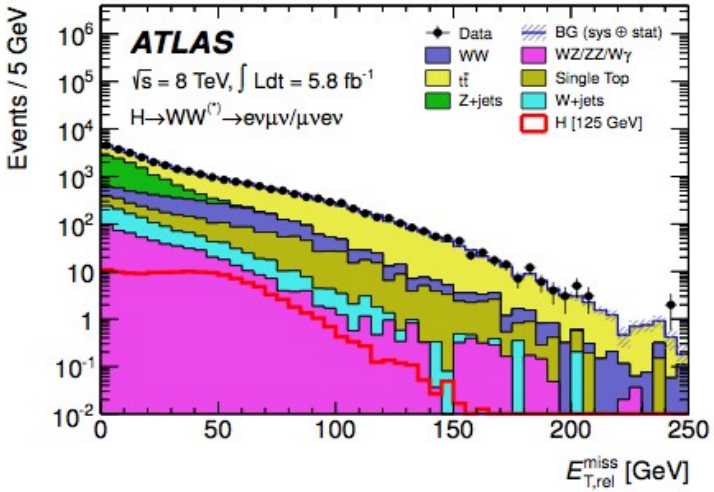

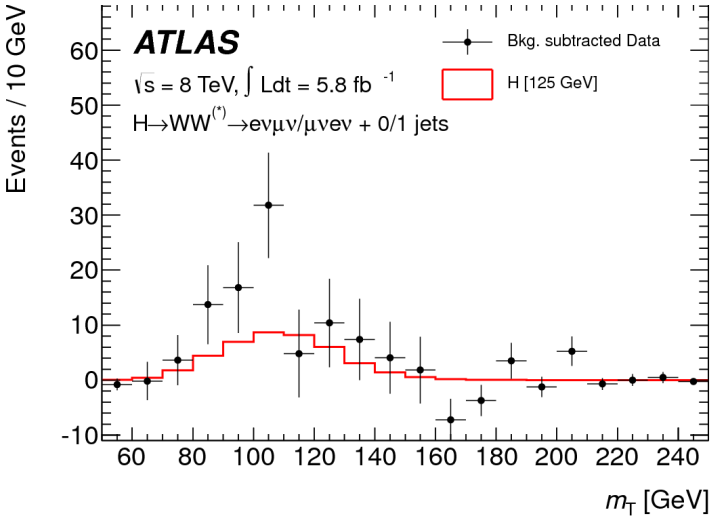

[Phys. Lett. B 716 (2012) 1-29]

ATLAS publications

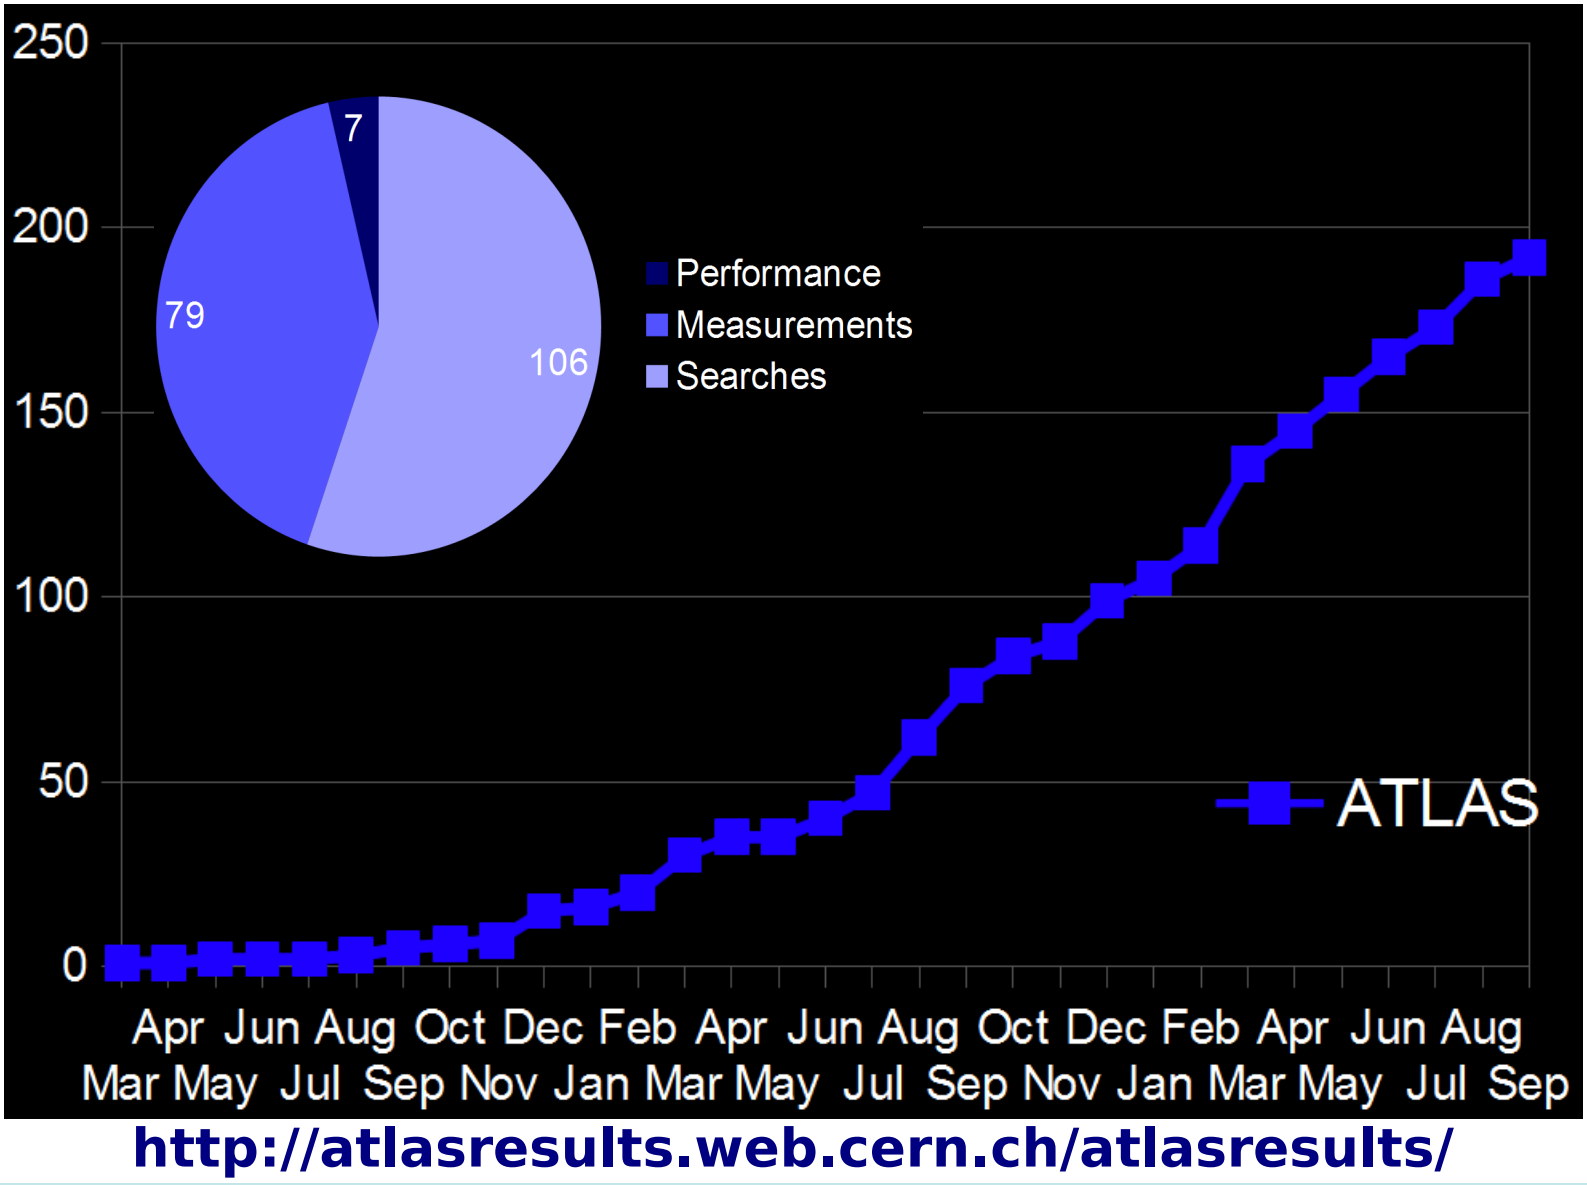

http://atlasresults.web.cern.ch/atlasresults/

ATLAS: a detector for particle physics at LHC

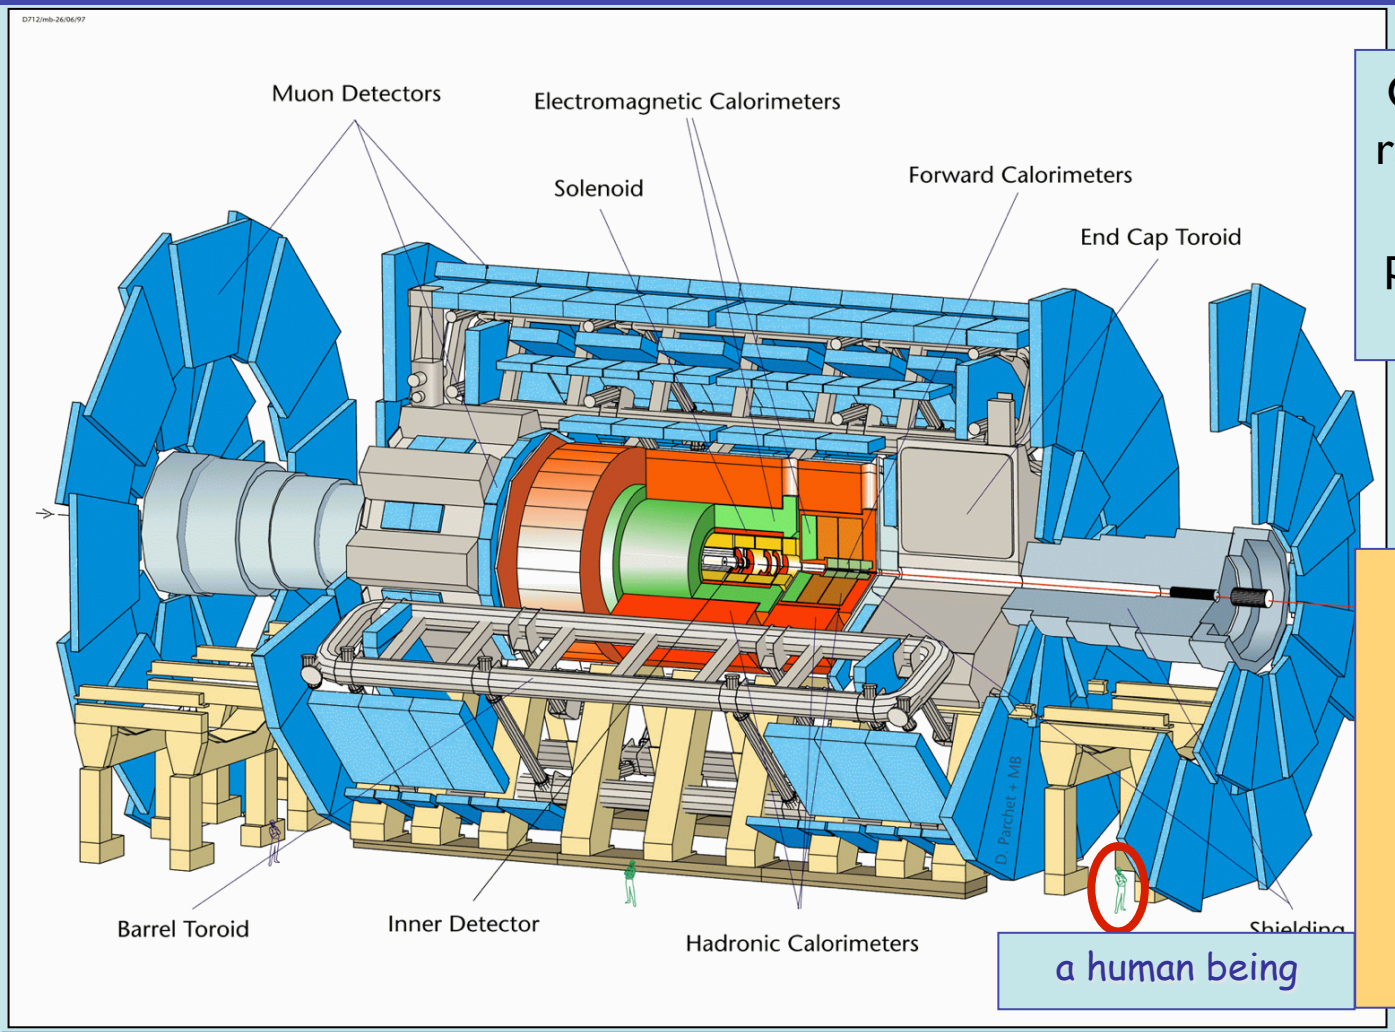

Cover the whole angular range around the collision point to detect as many particles produced in the collision as possible.

LHC detectors are much more complex, performing and challenging than those at previous/present accelerators:
a big jump in concepts and technologies

- Size (length 45m, diameter 25m): to measure and absorb high-energy particles
- Fast response (~50 ns): 20 (2015: 40) million beam-beam collisions per second
- 10^8 electronic channels ("individual signals"): to track ~1000 particles per event and reconstruct their trajectories with ~10 μm precision

Data flow through the Tier-0 at CERN

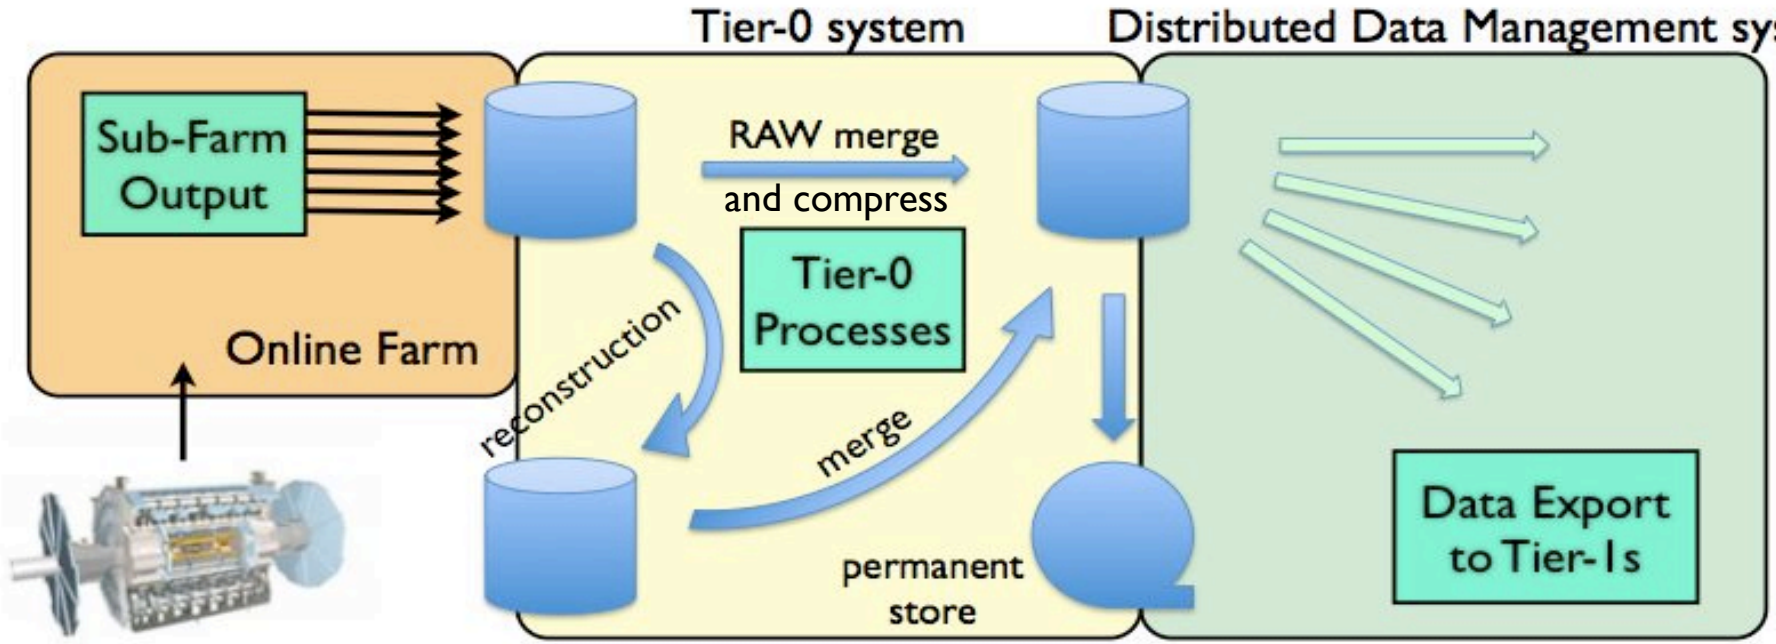

Accepting data from the online system and ensuring it is archived to tape

- Merging small files to adequate size for tape archiving

Processing RAW data (event reconstruction) and archiving the products to tape

- Express stream for prompt calibration and alignment
- First-pass processing of all streams after 36h with calibration and alignment

Registering data to the ATLAS Distributed Data Management system

- Export data to Tier-1 and calibration Tier-2s, as well as CAF

Maximum overall I/O: 6GB/s -- including internal accesses within Tier-0

Integrated Simulation Framework (ISF)

- single-event **simulation** at 1 Hz is possible

  scales linearly with pile-up

- truth bound reconstruction techniques should be explored further

  scales linearly with pile-up (hello, upgrade community)

- this puts quite some **pressure on the digitization**

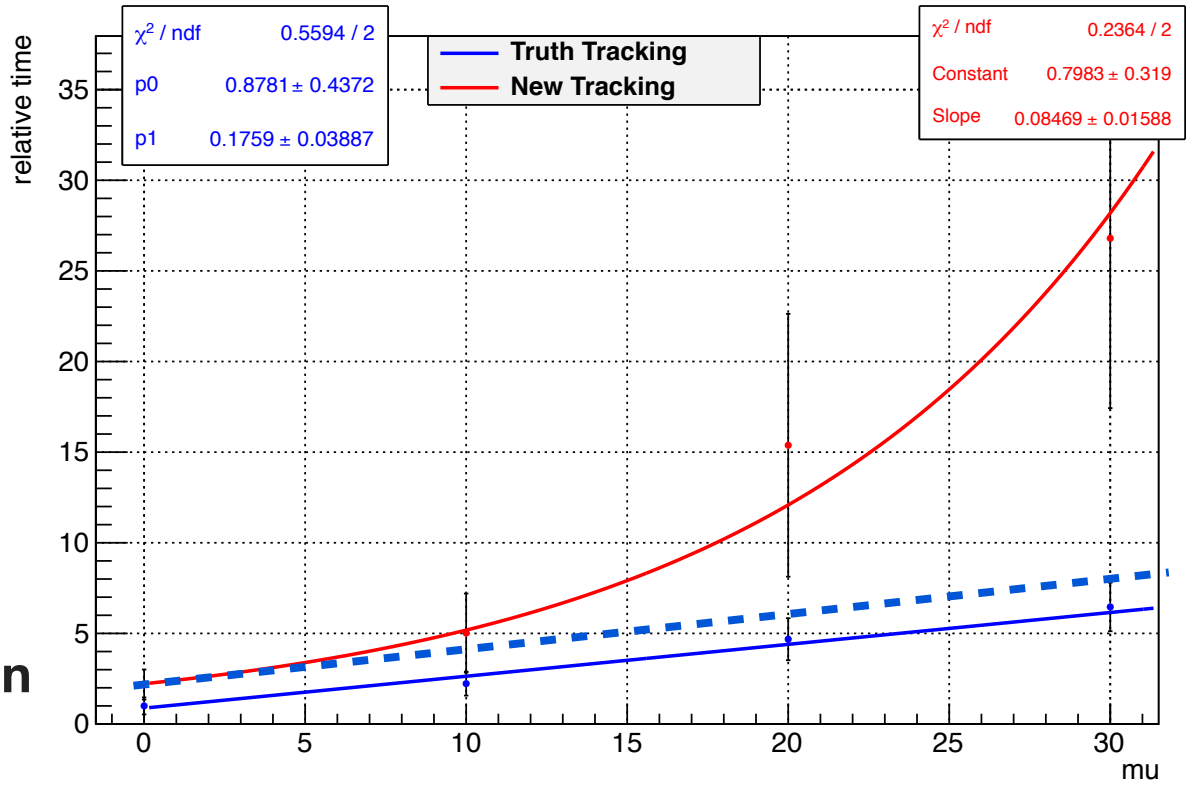

Tier-0 reconstruction this year

- Tier0 reconstruction coping well with luminosity / pile-up
- Tier0 capacity 6k slots. Increased to 7.5k for reprocessing of initial data during TS1 and during the ICHEP ‘rush’.