LHC: An example of a Global Scientific Community

Sergio Bertolucci

CERN

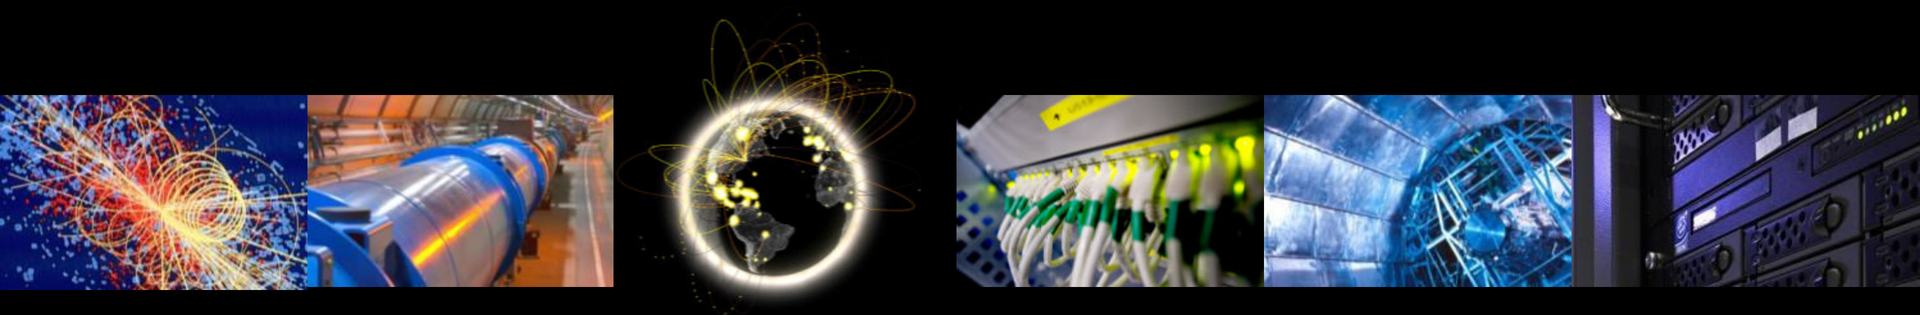

5th EGEE User Forum

Uppsala, 14th April 2010

WLCG
Worldwide LHC Computing Grid

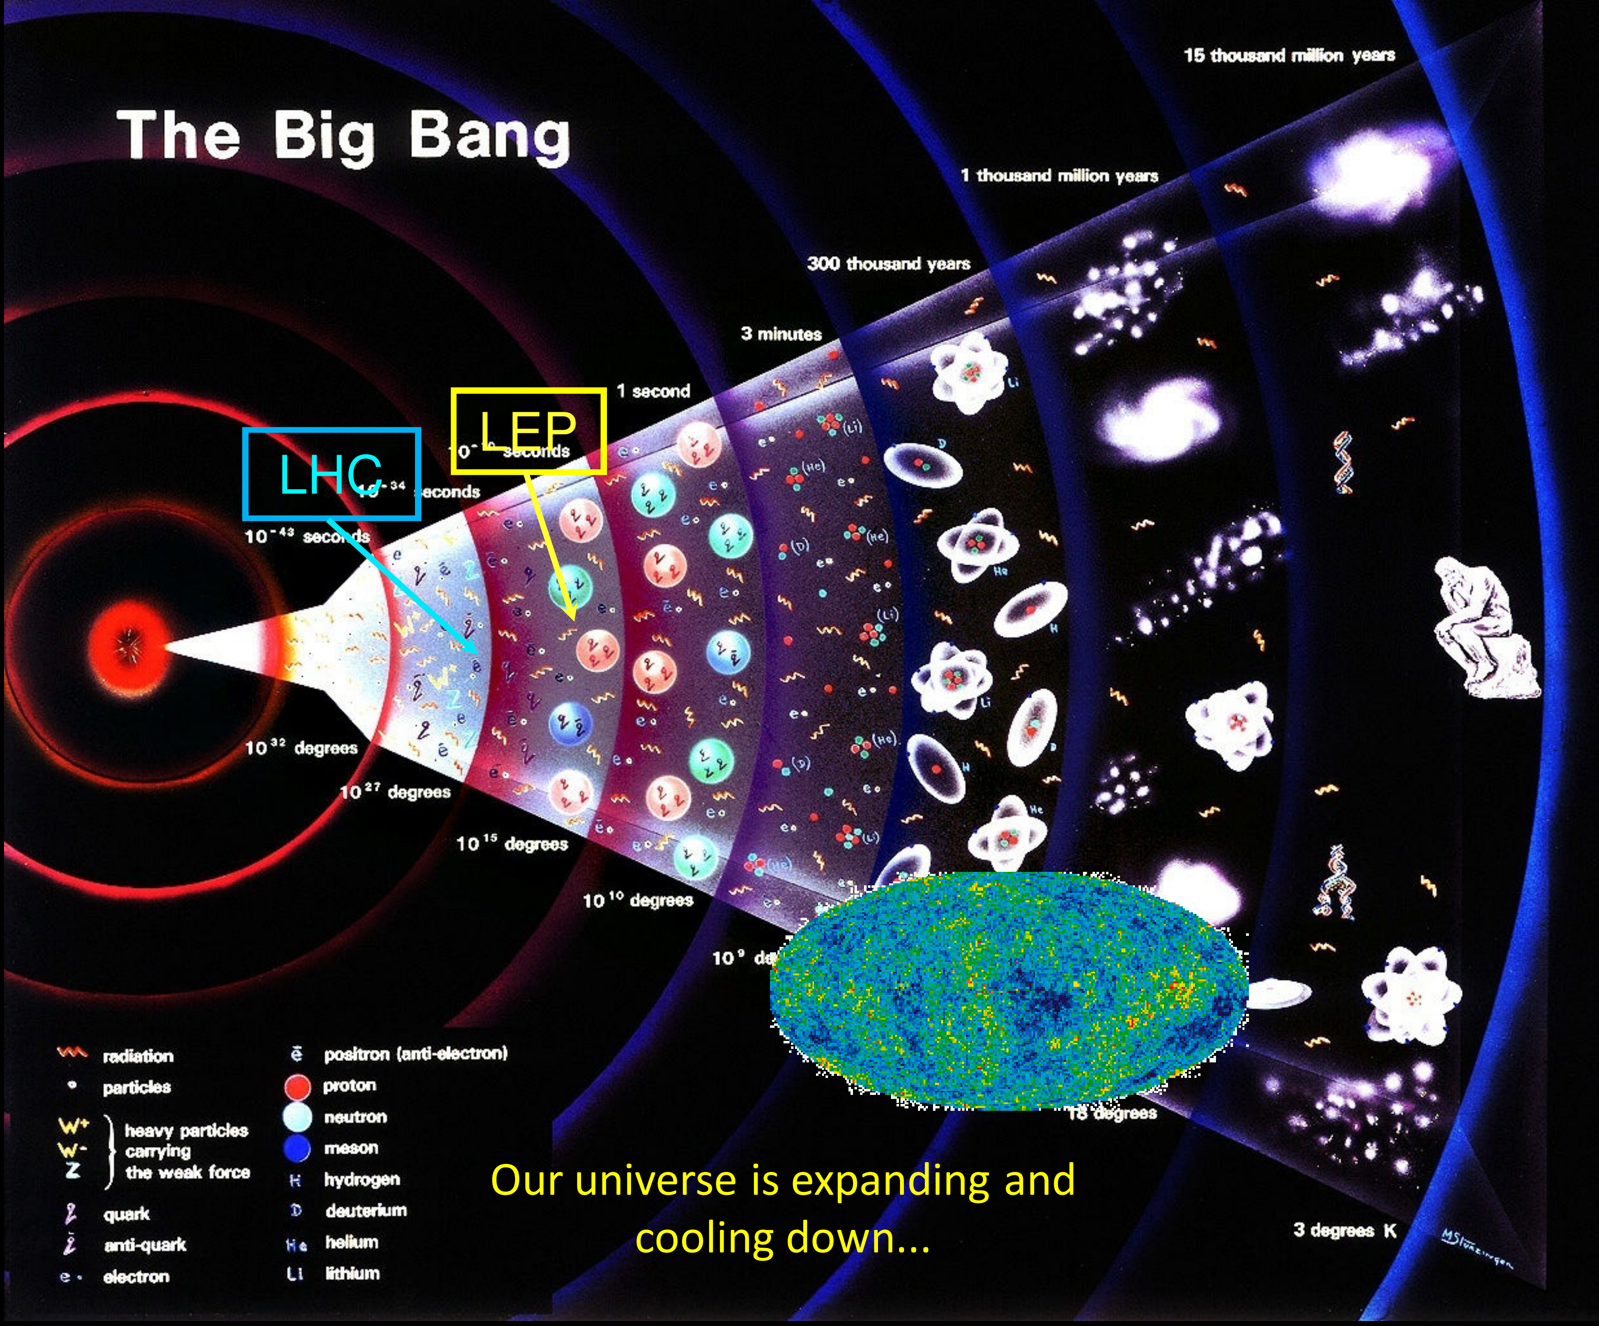

Our universe is expanding and cooling down...

... and has been doing so for approximately 13.7 billion years

> 95%
UNKNOWN
STUFF
OUT THERE

Black hole

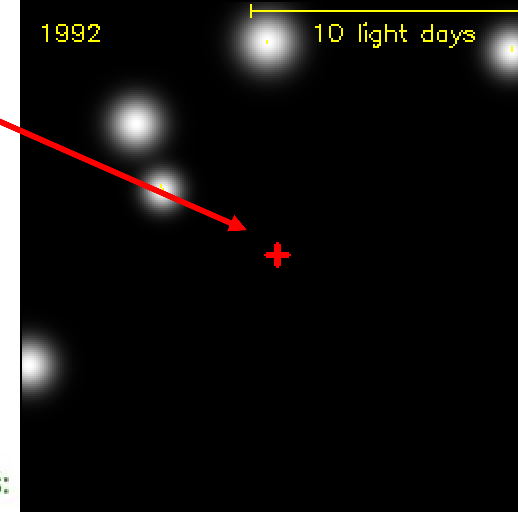

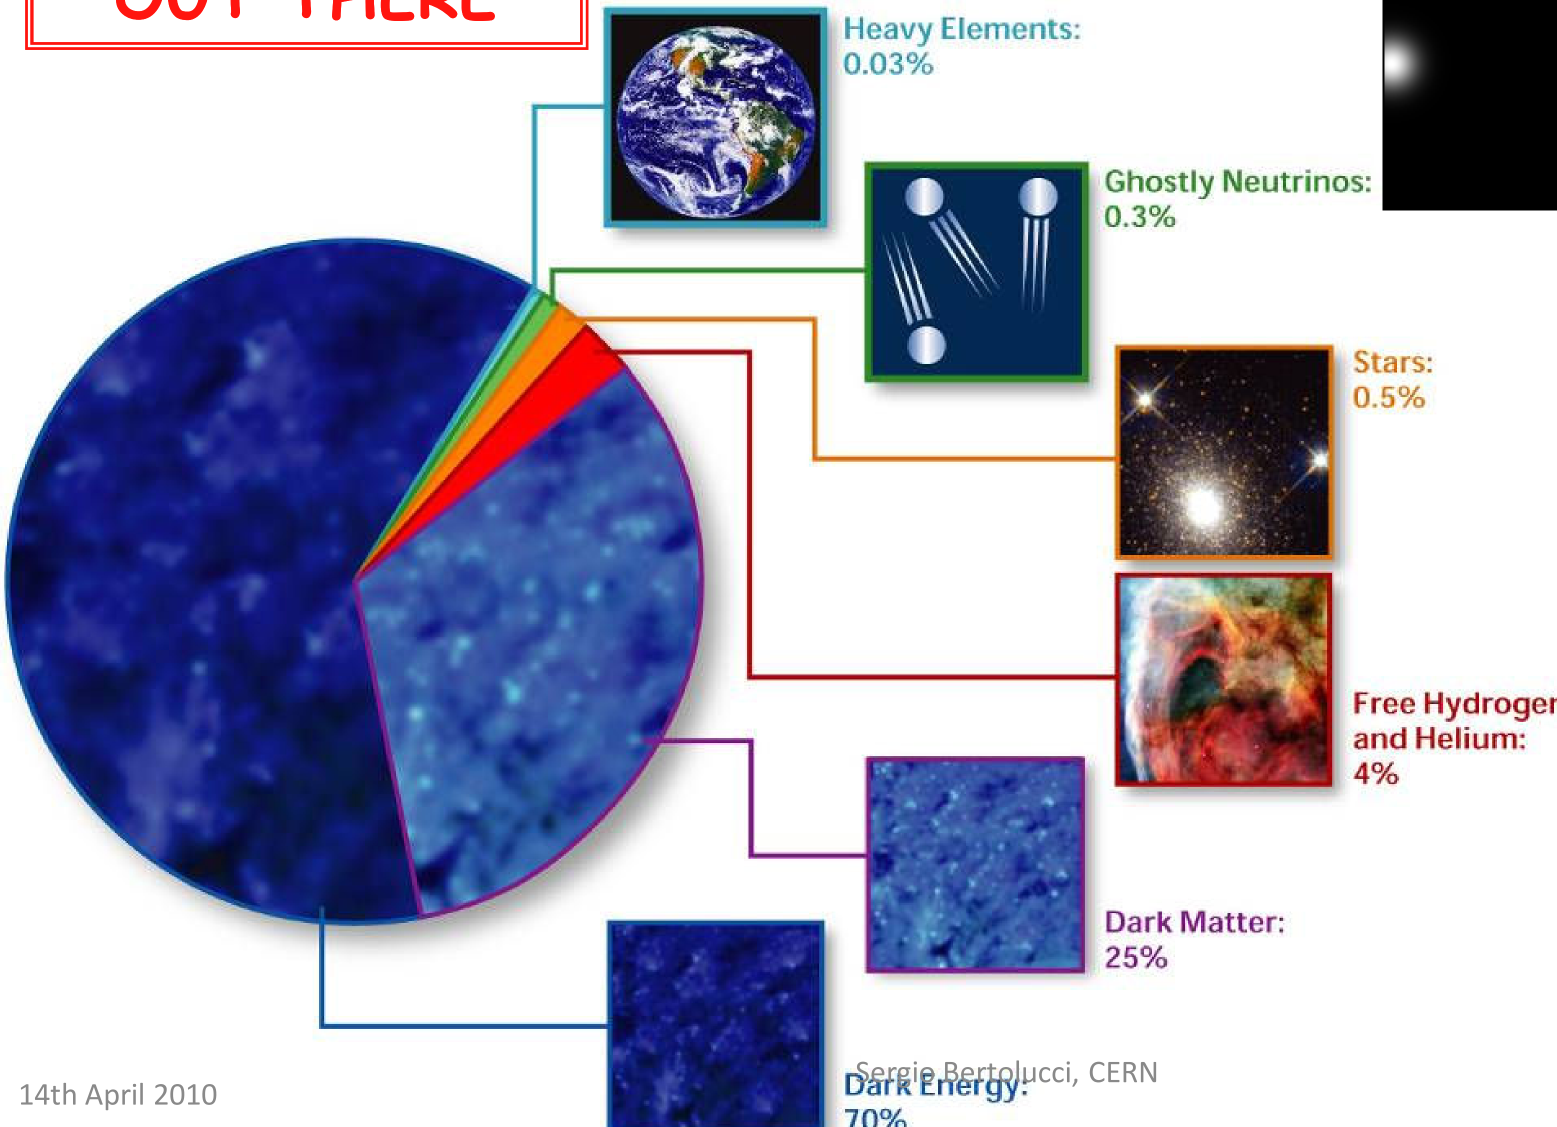

Fundamental Physics Questions

- Why do particles have mass?
  - Newton could not explain it - and neither can we…
- What is 96% of the Universe made of?
  - We only observe 4% of it!
- Why is there no antimatter left in the Universe?
  - Nature should be symmetrical
- What was matter like during the first second of the Universe, right after the "Big Bang"?
  - A journey towards the beginning of the Universe gives us deeper insight

The Large Hadron Collider (LHC), allows us to look at microscopic big bangs to understand the fundamental laws of nature

CERN stands for over 50 years of…

- **fundamental research and discoveries**
- **technological innovation**
- **training and education**
- **bringing the world together**

1954 Rebuilding Europe
First meeting of the CERN Council

1980 East meets West
Visit of a delegation from Beijing

2010 Global Collaboration
The Large Hadron Collider involves over 100 countries

CERN's Tools

- The world's most powerful accelerator: LHC
  - A 27 km long tunnel filled with high-tech instruments
  - Equipped with thousands of superconducting magnets
  - Accelerates particles to energies never before obtained
  - Produces particle collisions creating microscopic "big bangs"
- Very large sophisticated detectors
  - Four experiments each the size of a cathedral
  - Hundred million measurement channels each
  - Data acquisition systems treating Petabytes per second
- Significant computing to distribute and analyse the data
  - A Computing Grid linking ~200 computer centres around the globe
  - Sufficient computing power and storage to handle 15 Petabytes per year, making them available to thousands of physicists for analysis

- **Global collaborations essential at all stages**

Distribution of All CERN Users by Nation of Institute on 6 January 2009

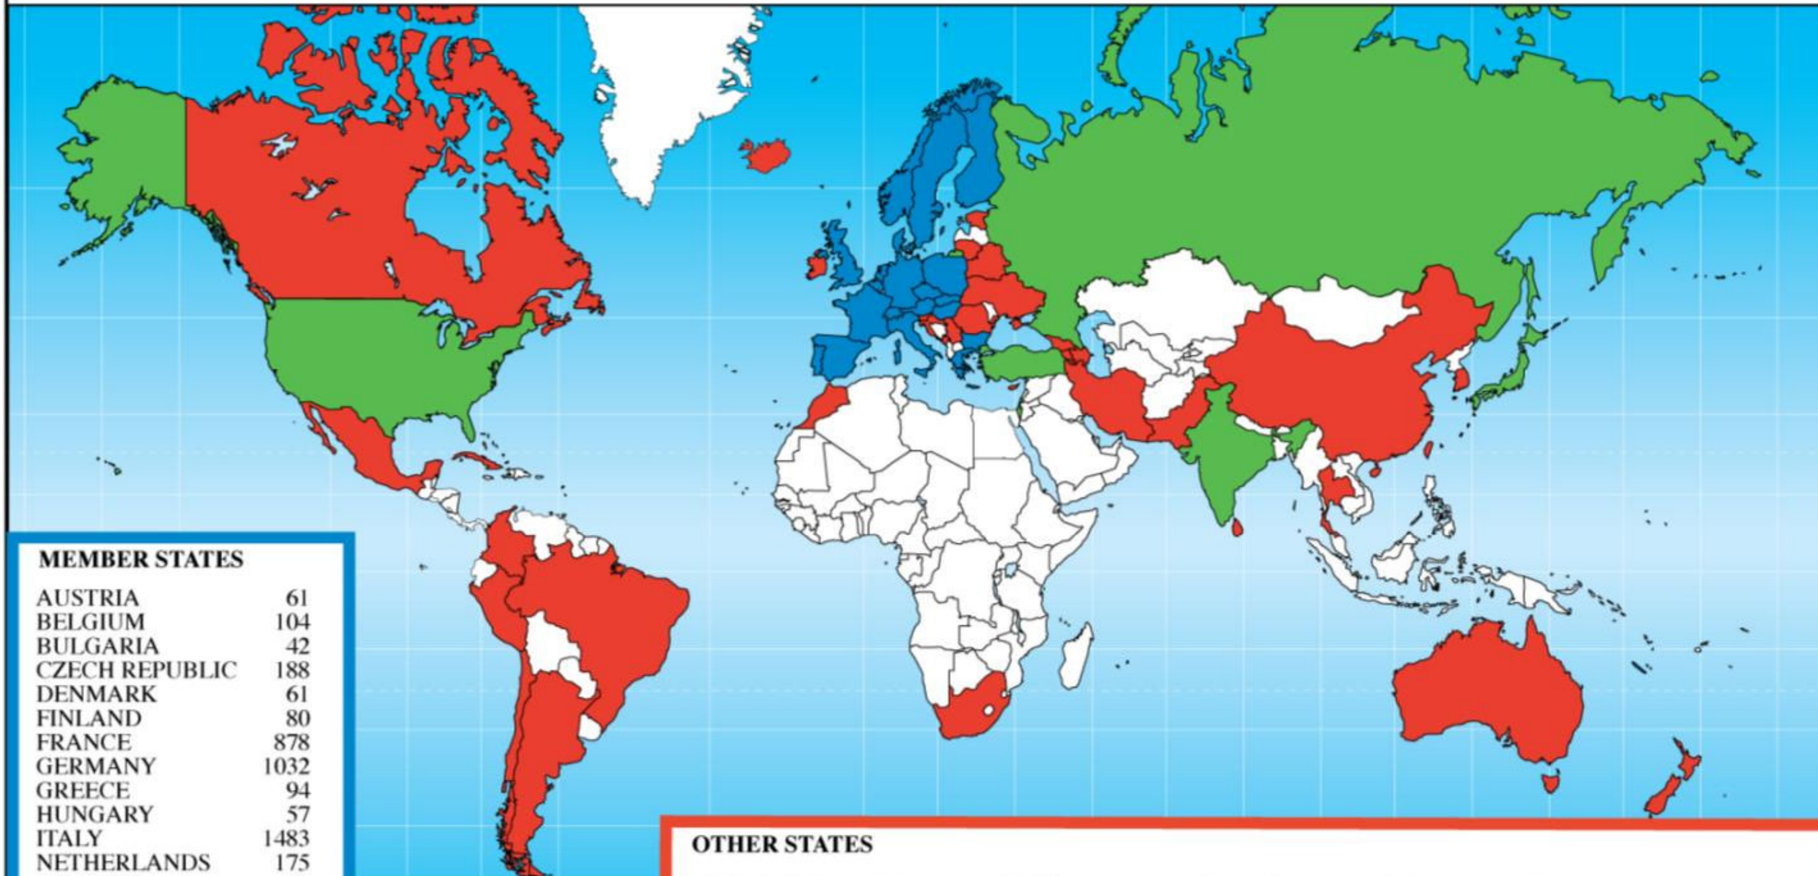

MEMBER STATES

AUSTRIA	61
BELGIUM	104
BULGARIA	42
CZECH REPUBLIC	188
DENMARK	61
FINLAND	80
FRANCE	878
GERMANY	1032
GREECE	94
HUNGARY	57
ITALY	1483
NETHERLANDS	175
NORWAY	78
POLAND	174
PORTUGAL	111
SLOVAKIA	49
SPAIN	286
SWEDEN	73
SWITZERLAND	330
UNITED KINGDOM	715

6071

OBSERVER STATES

INDIA	89
ISRAEL	59
JAPAN	200
RUSSIA	883
TURKEY	52
USA	1485

2768

OTHER STATES

ARGENTINA	10	CUBA	3	MONTENEGRO	1	SRI LANKA	1
ARMENIA	15	CYPRUS	6	MOROCCO	5	TAIWAN	42
AUSTRALIA	14	ESTONIA	11	NEW ZEALAND	6	THAILAND	1
AZERBAIJAN	1	GEORGIA	11	PAKISTAN	24	UKRAINE	18
BELARUS	19	ICELAND	1	PERU	1		
BRAZIL	73	IRAN	12	ROMANIA	49		
CANADA	136	IRELAND	12	SERBIA	17		
CHILE	4	KOREA	51	SLOVENIA	16		
CHINA	64	LITHUANIA	5	SOUTH AFRICA	8		
COLOMBIA	11	MEXICO	28				
CROATIA	20						

696

LHC is in operation!

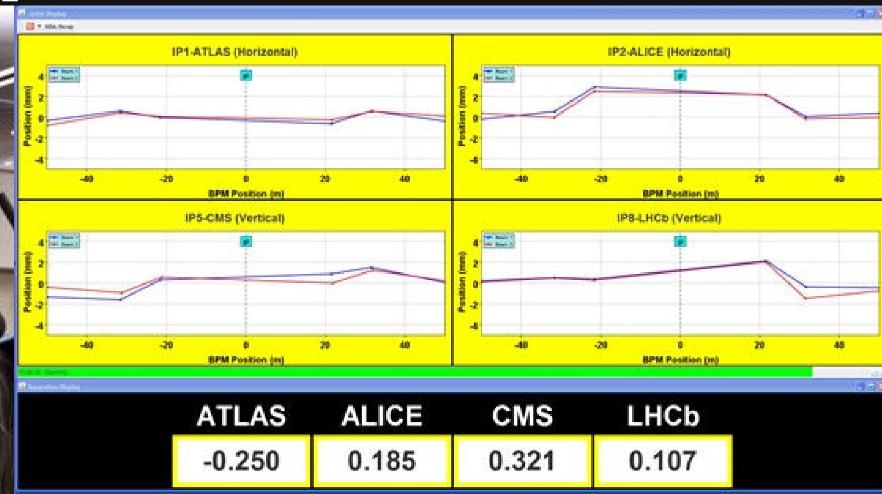

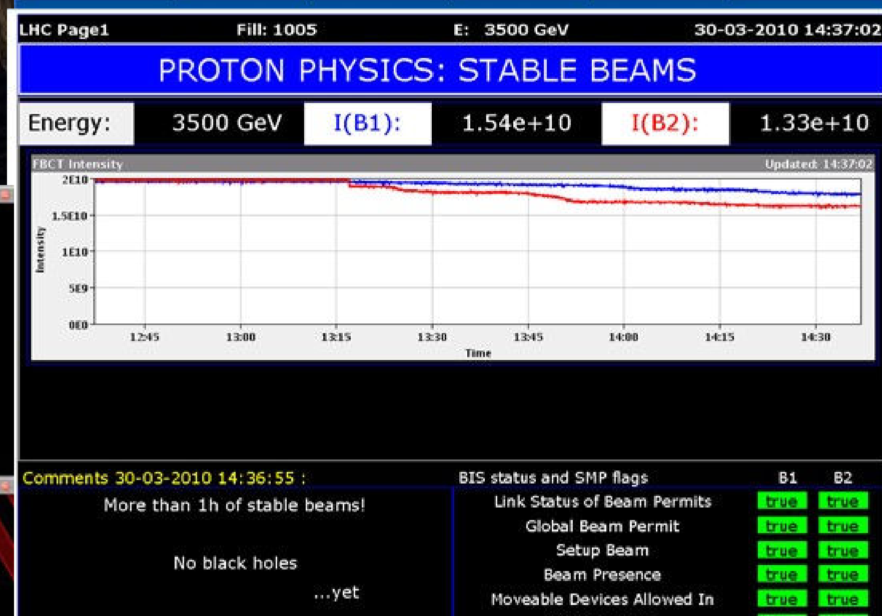

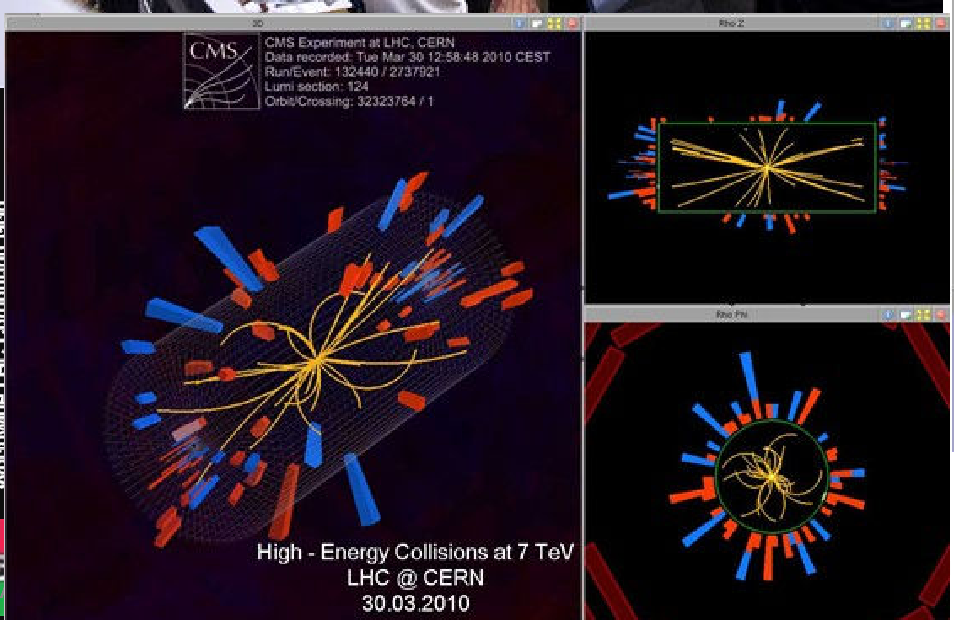

From this (October 2008) ...

Collateral damage: magnet displacements

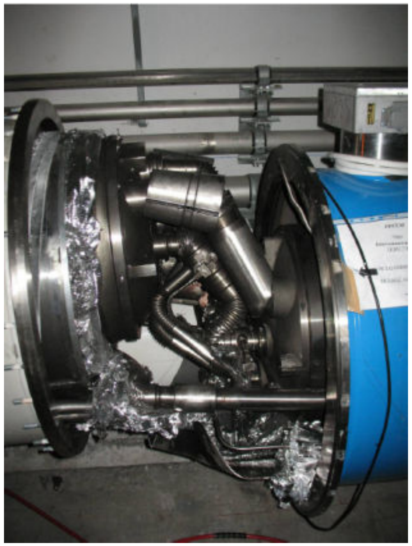

Collateral damage: ground supports

... To this (Nov 2009)

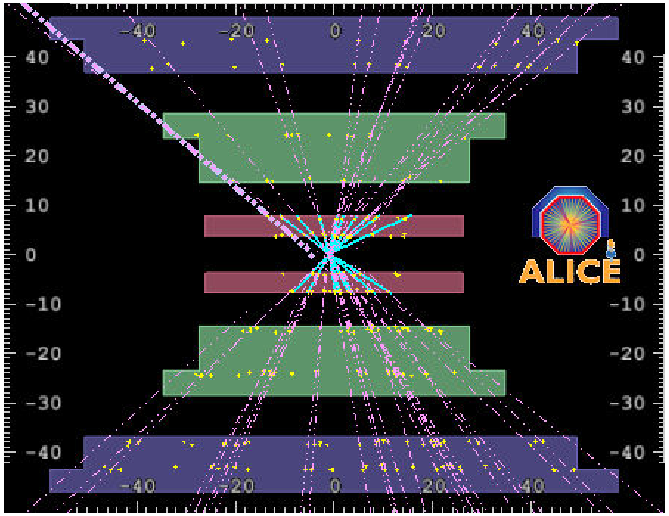

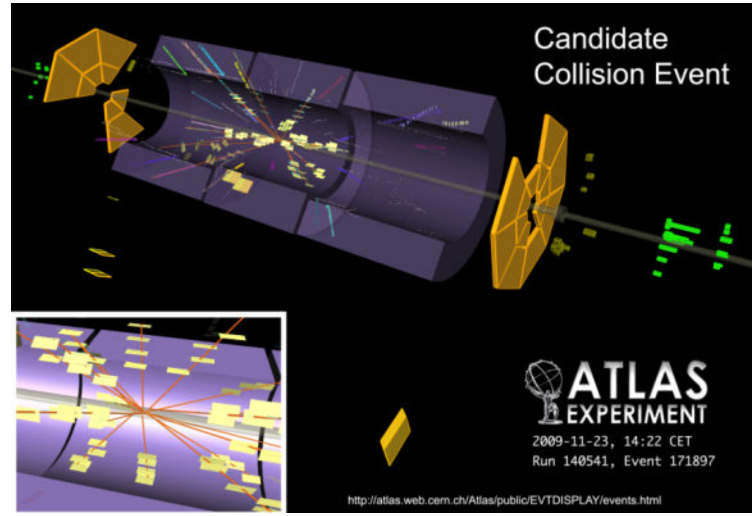

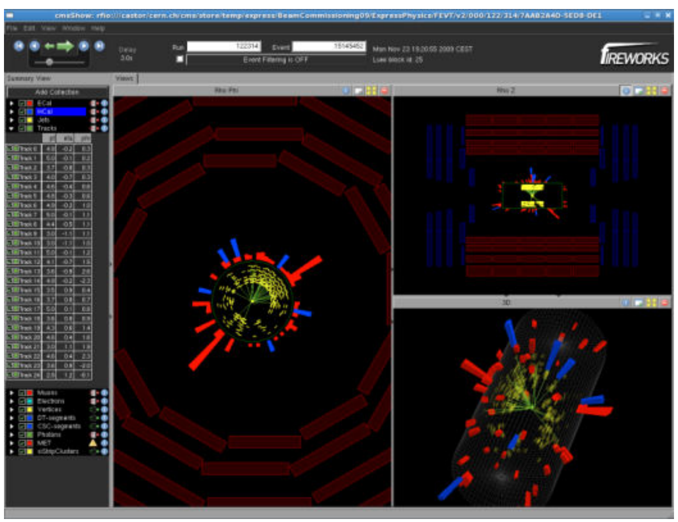

LHCb Event Display

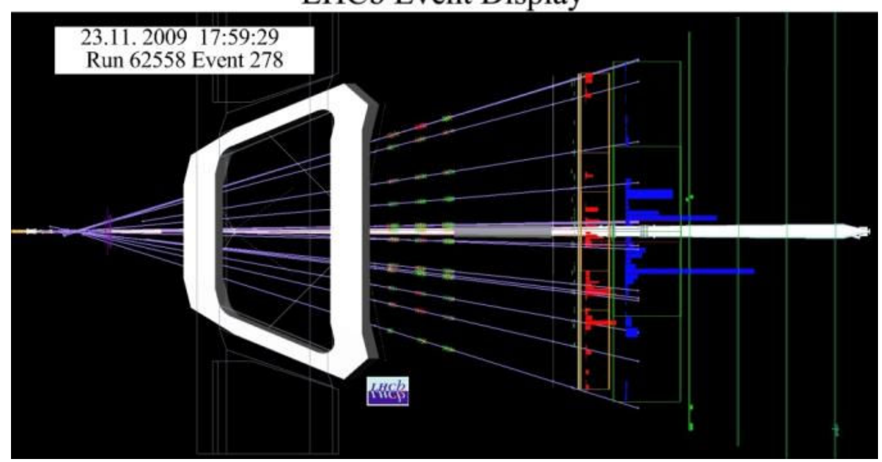

… And now at 7 TeV

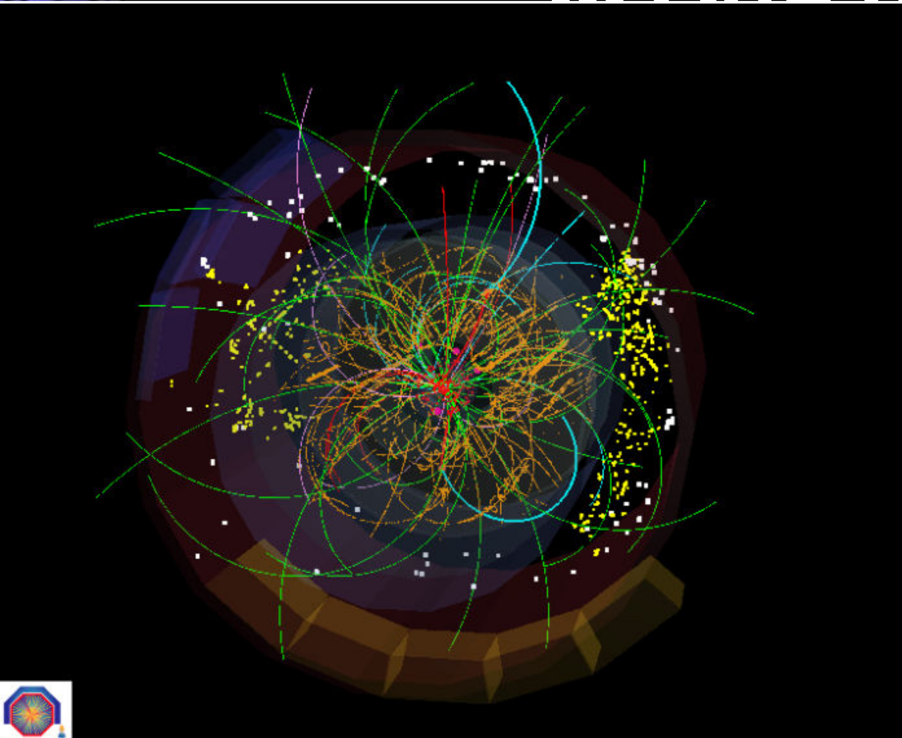

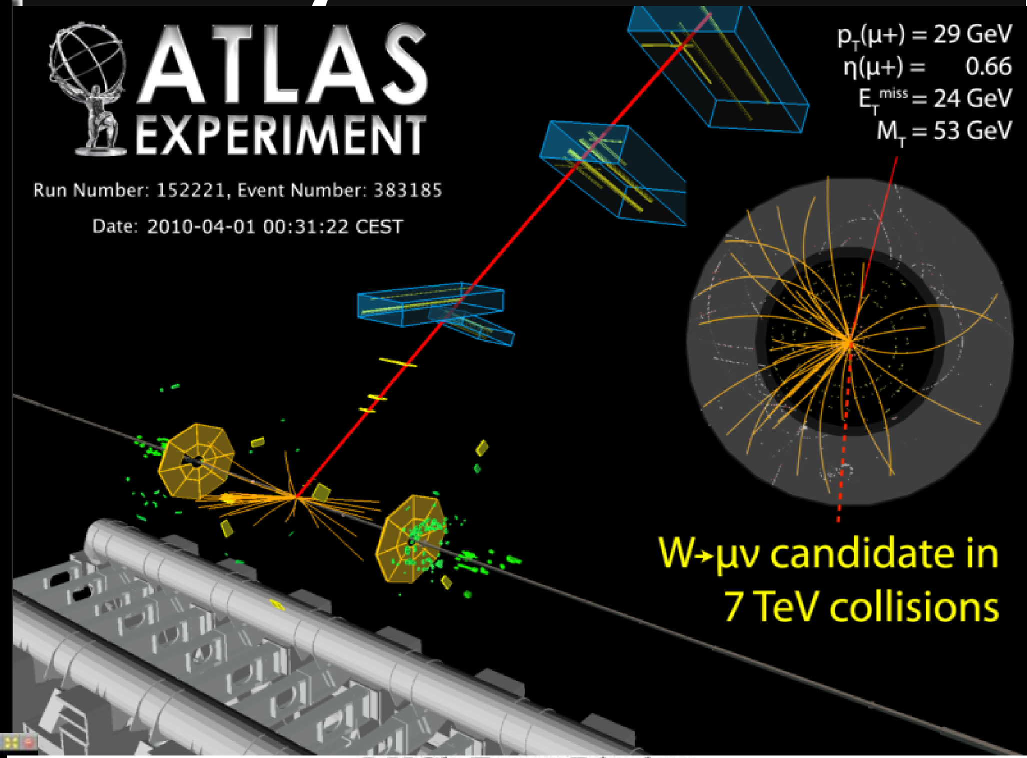

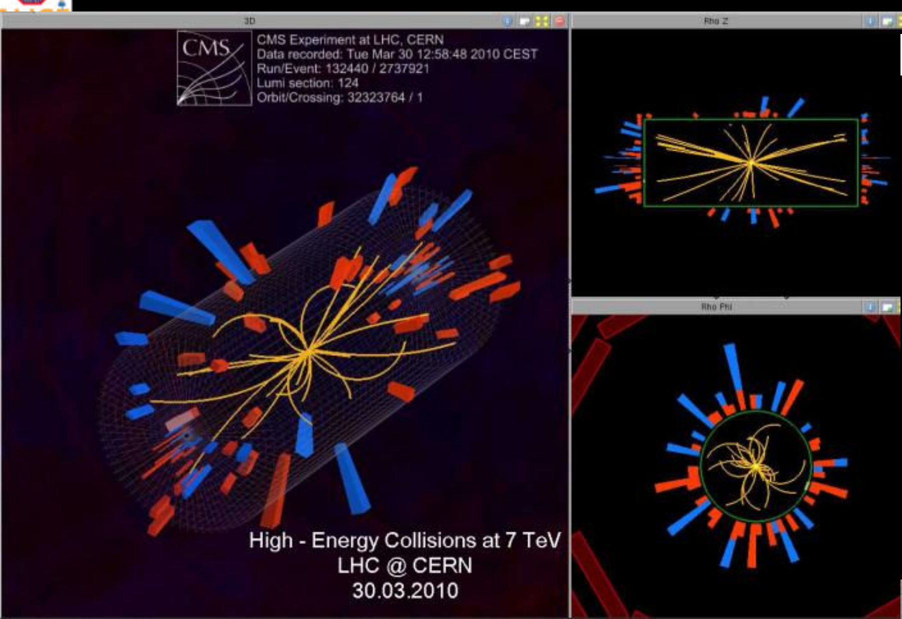

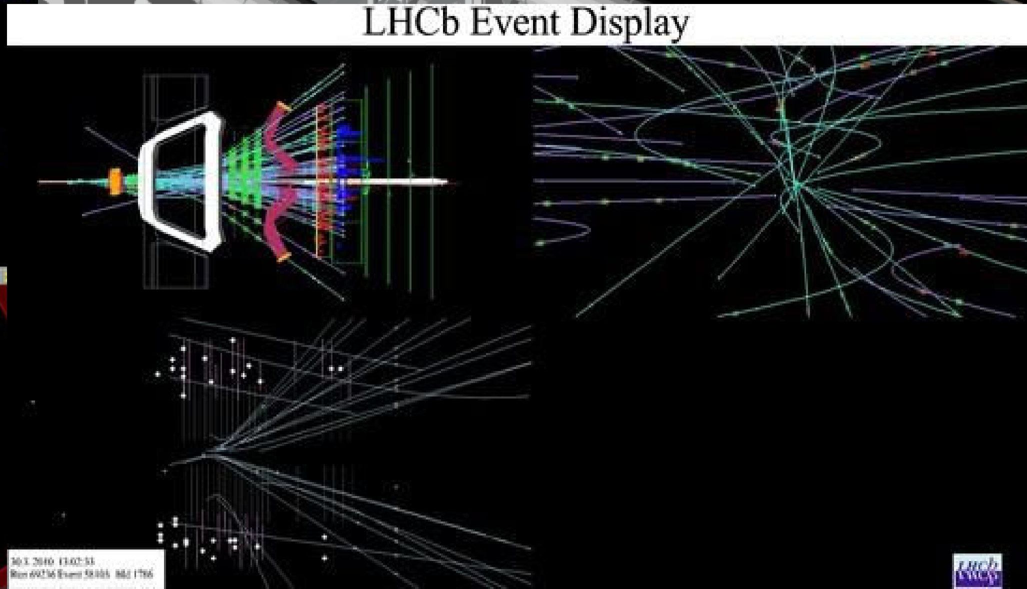

Scale of ATLAS and CMS?

ATLAS superimposed to the 5 floors of building 40

CMS

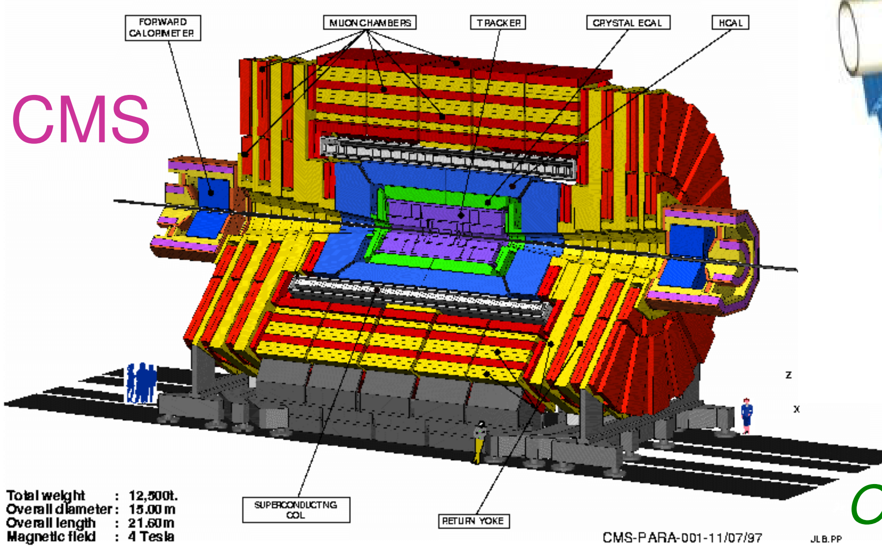

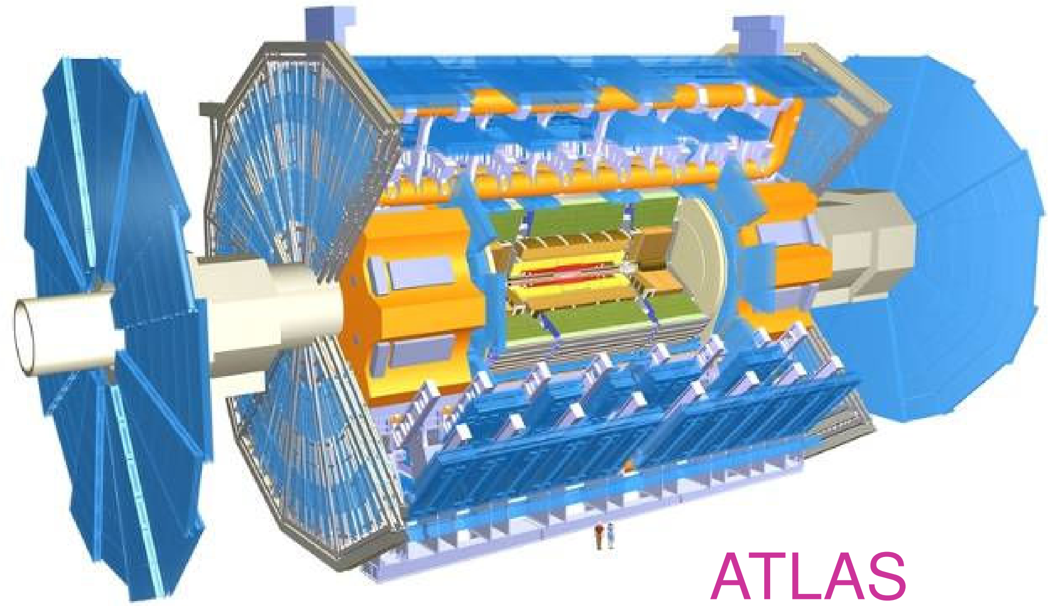

ATLAS

	ATLAS	CMS
Overall weight (tons)	*7000*	*12500*
Diameter	*22 m*	*15 m*
Length	*46 m*	*22 m*
Solenoid field	*2 T*	*4 T*

ATLAS is the product of >20 years sustained activity by a worldwide scientific community.

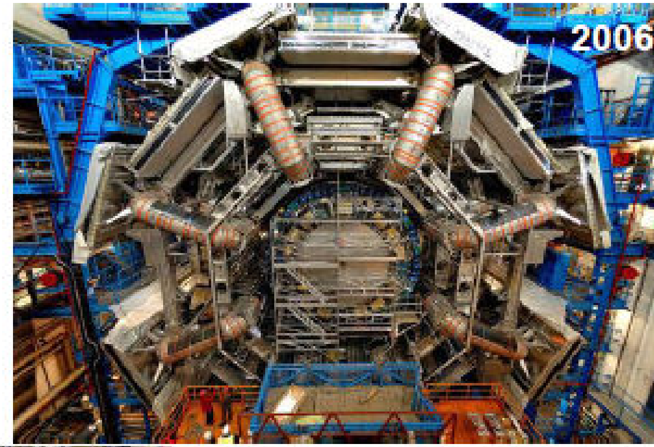

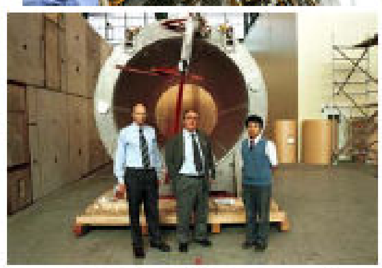

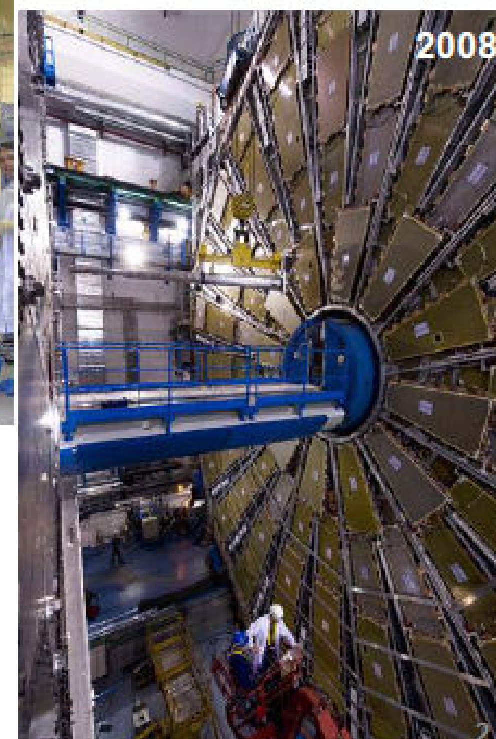

ATLAS is the product of >20 years sustained activity by a worldwide scientific community.

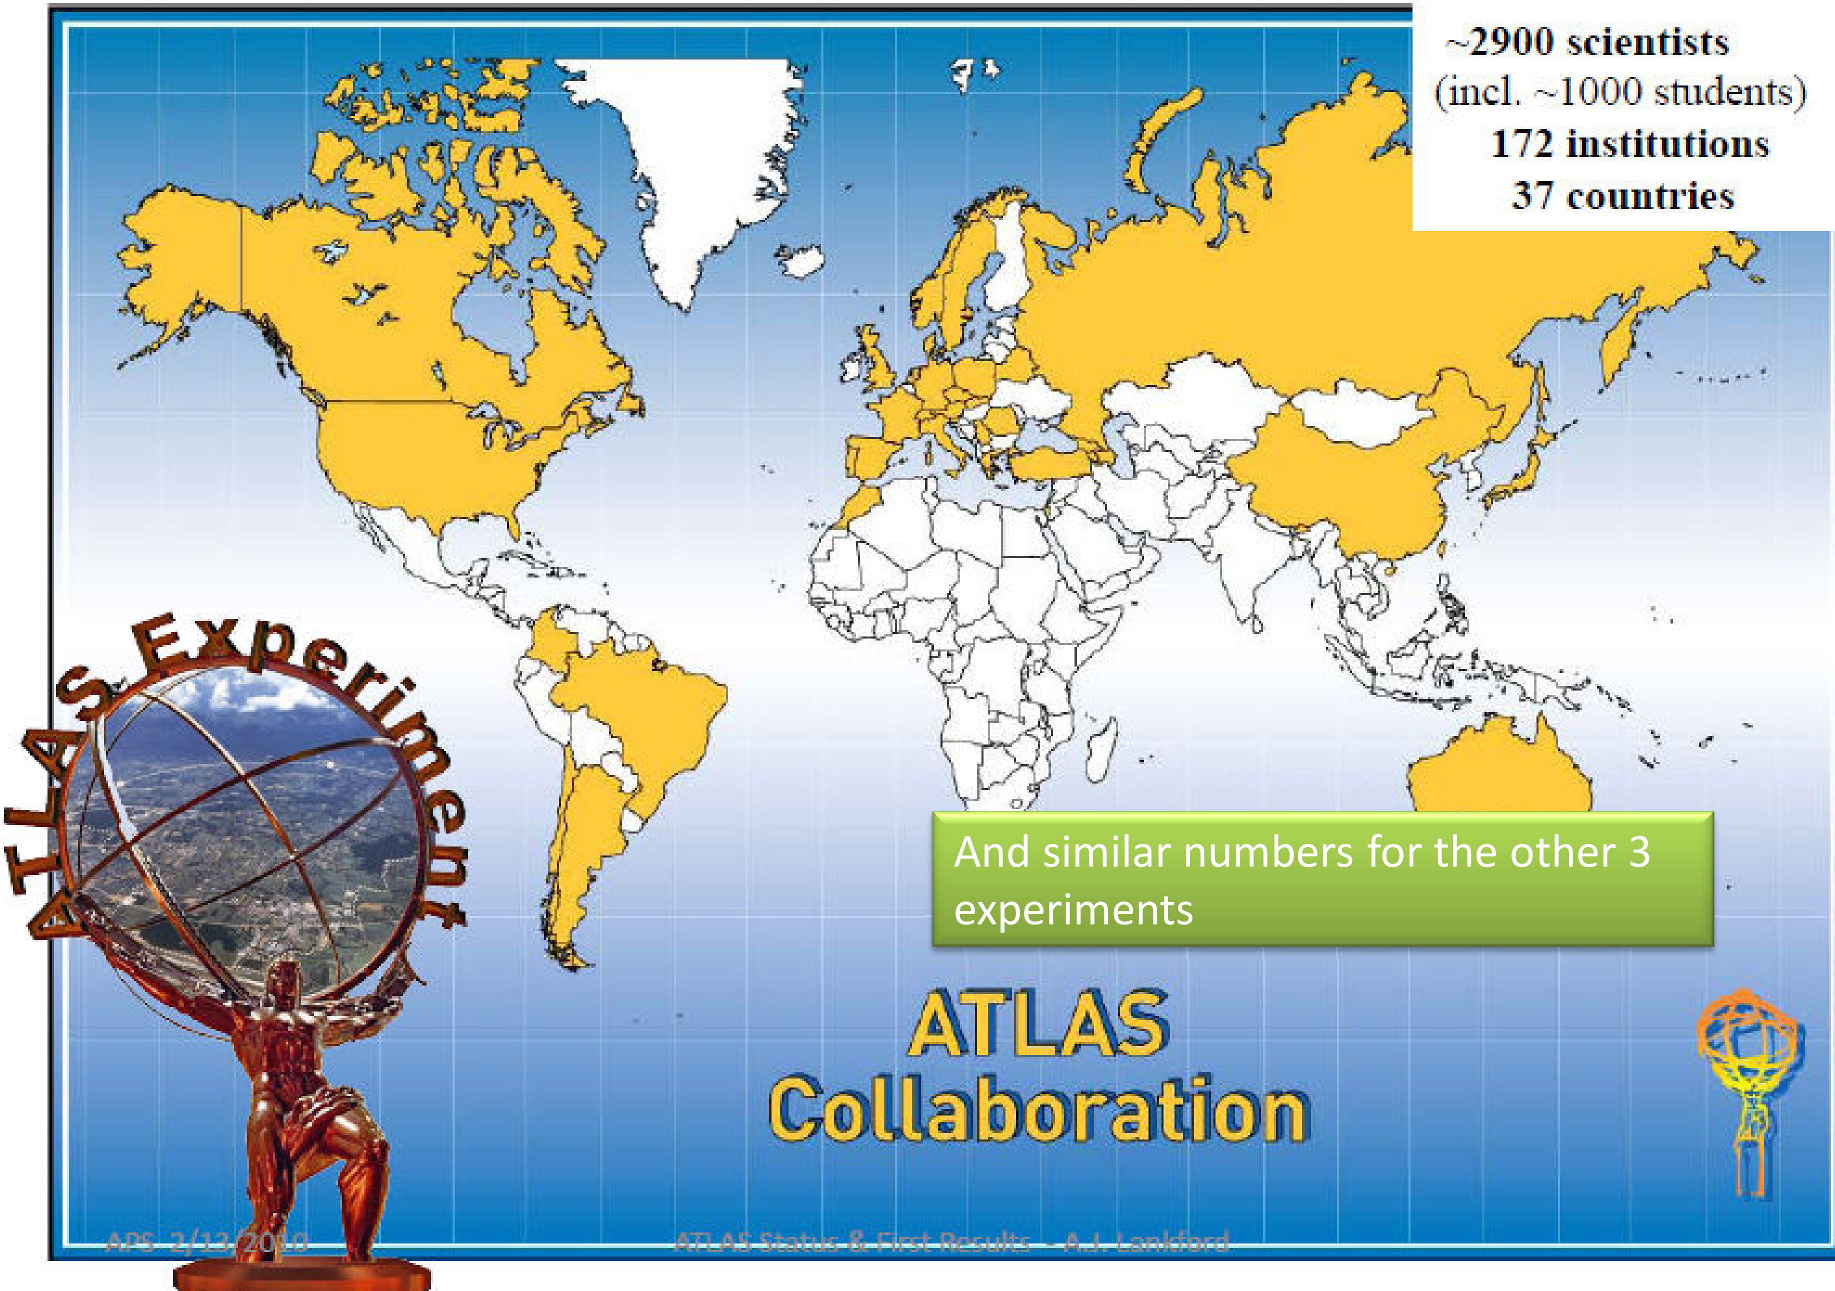

~2900 scientists
(incl. ~1000 students)
172 institutions
37 countries

And similar numbers for the other 3 experiments

Collisions at the LHC: summary

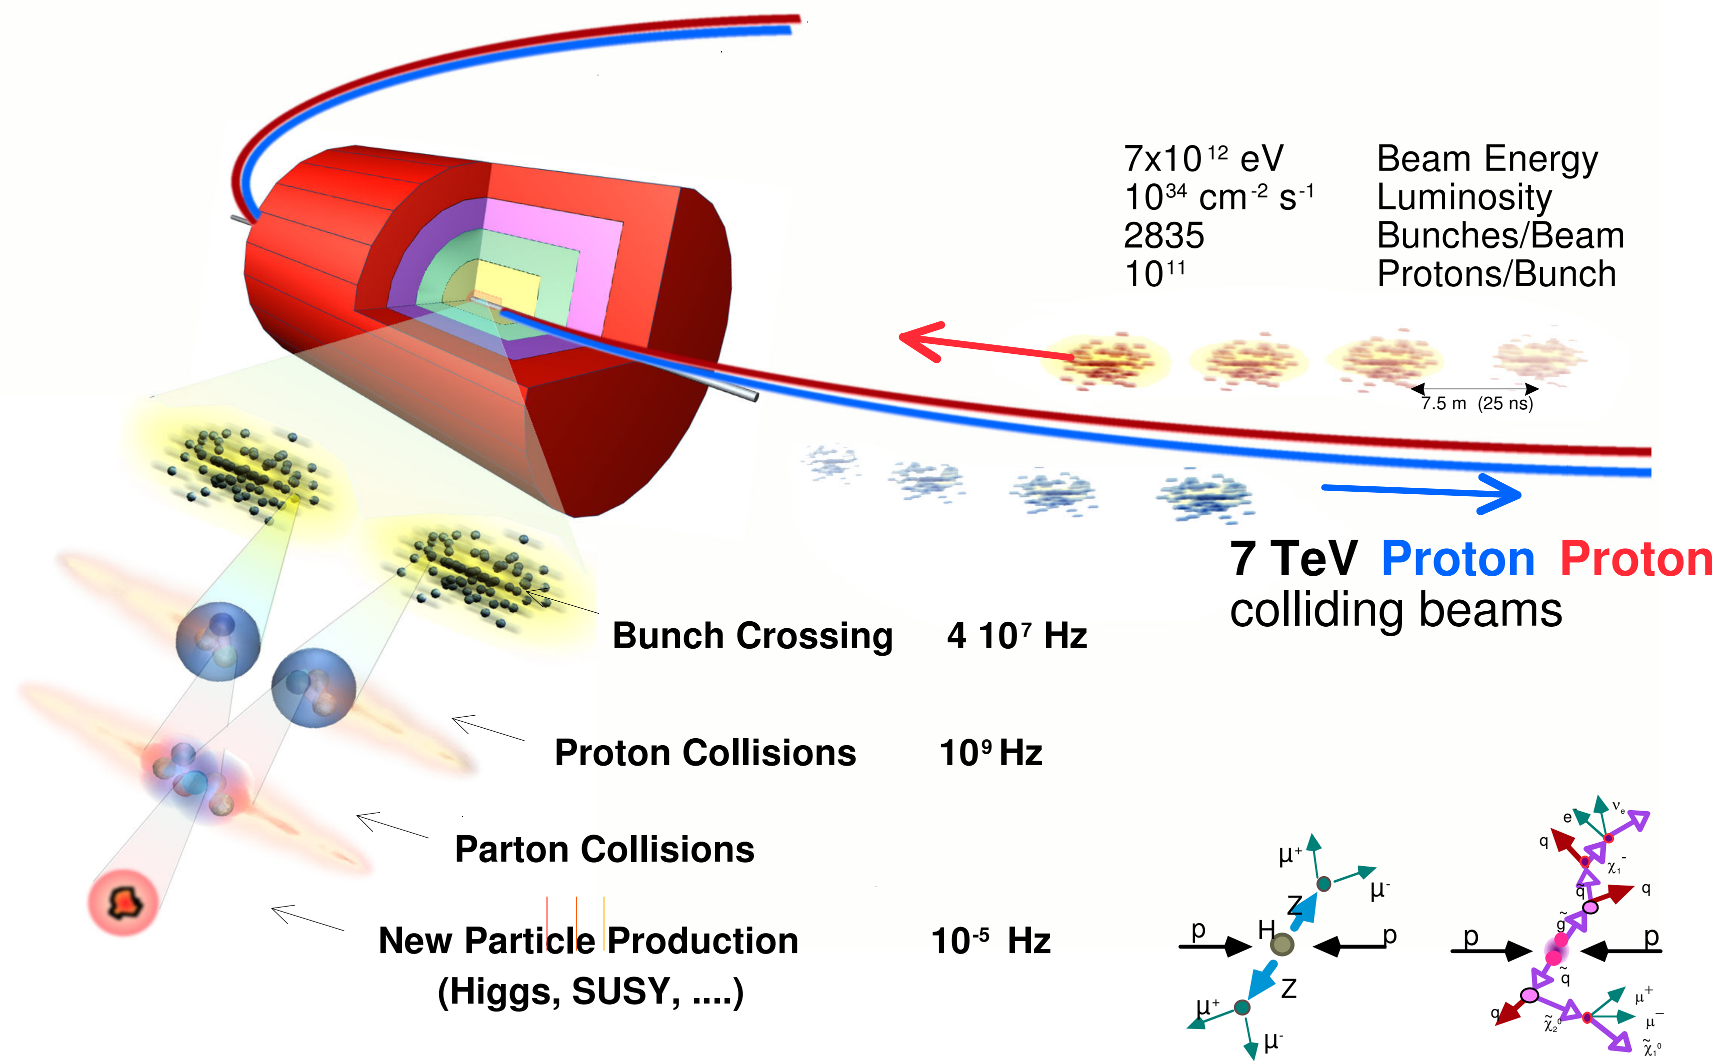

Selection of 1 event in 10,000,000,000,000

pp collisions at 14 TeV at 10^{34} $cm^{-2}s^{-1}$

A very difficult environment …

How to extract this:
(Higgs → 4 muons)

From this:

With:

20 proton-proton collisions overlap

And this repeats every 25 ns…

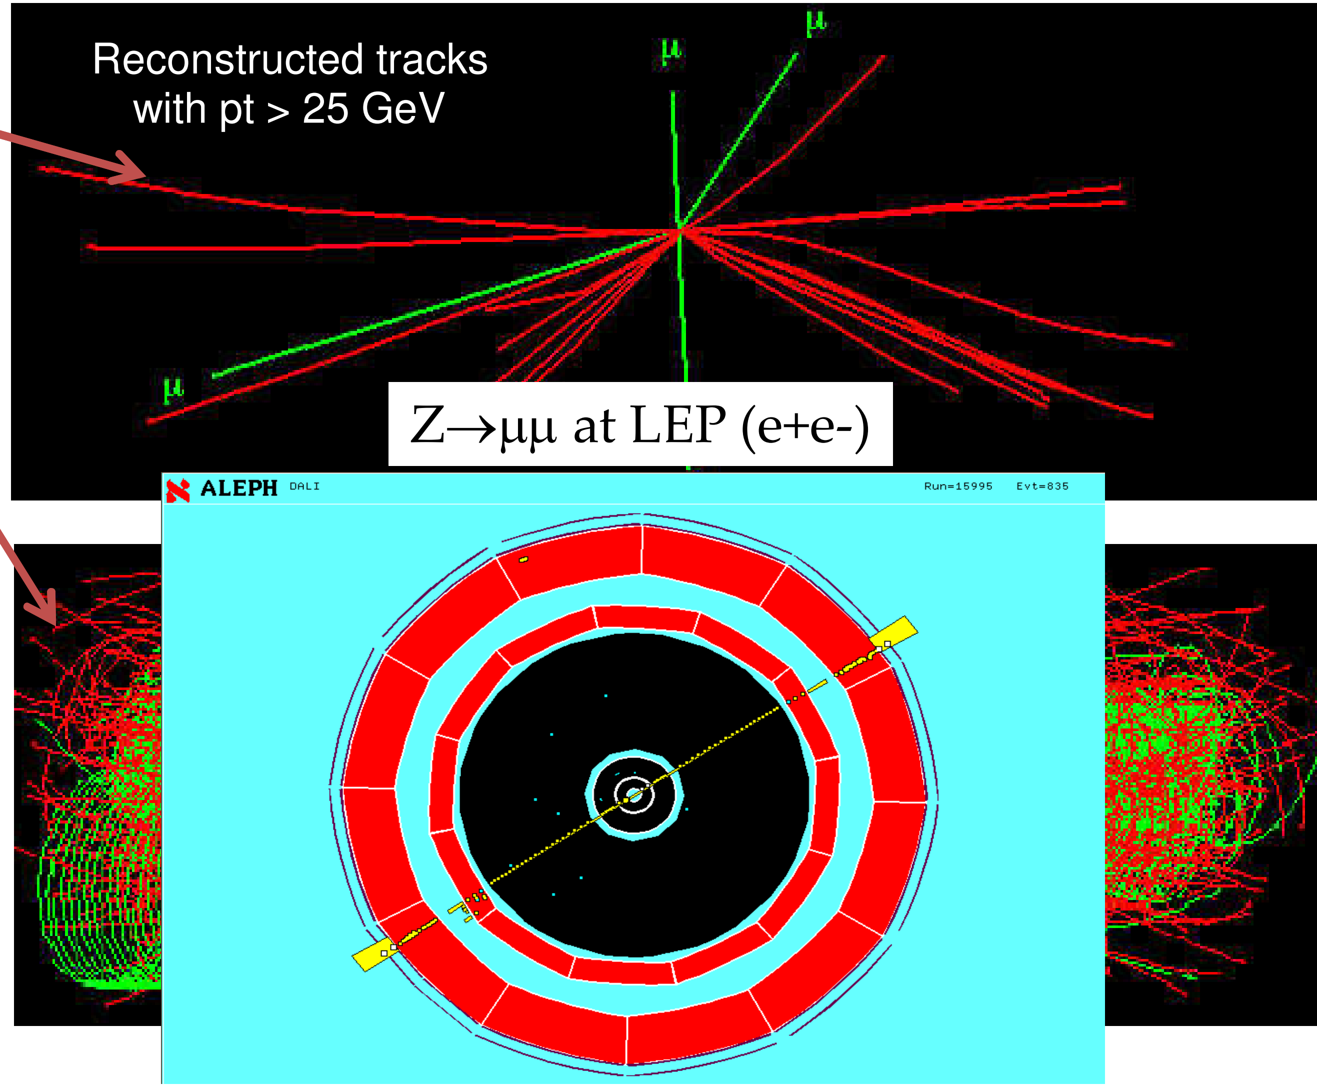

The LHC Computing Challenge

- Signal/Noise: 10^{-13} (10^{-9} offline)
- Data volume
  - High rate * large number of channels * 4 experiments
  - ➔ **15 PetaBytes of new data each year**
- Compute power
  - Event complexity * Nb. events * thousands users
  - ➔ **200 k of (today's) fastest CPUs**
  - ➔ **45 PB of disk storage**
- Worldwide analysis & funding
  - Computing funding locally in major regions & countries
  - Efficient analysis everywhere
  - ➔ **GRID technology**

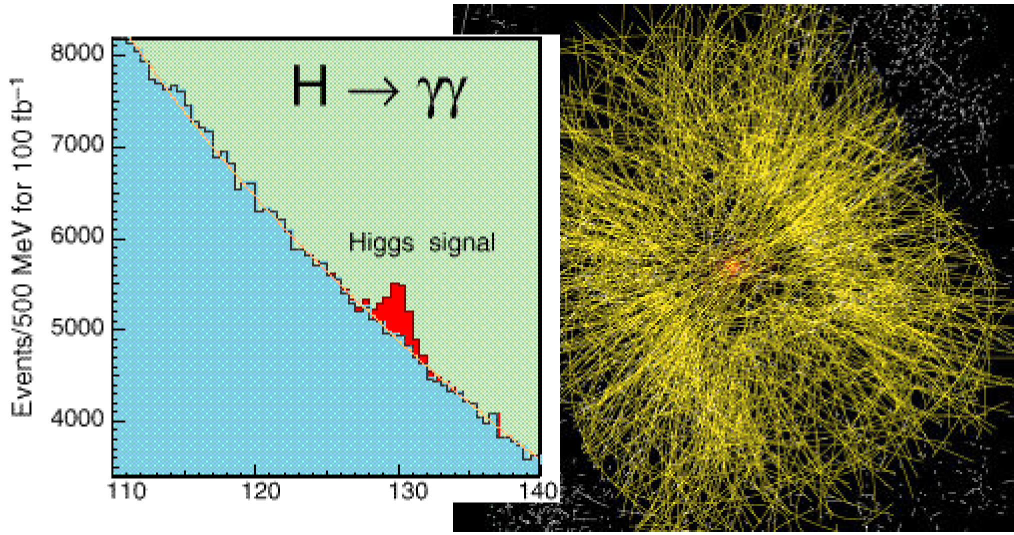

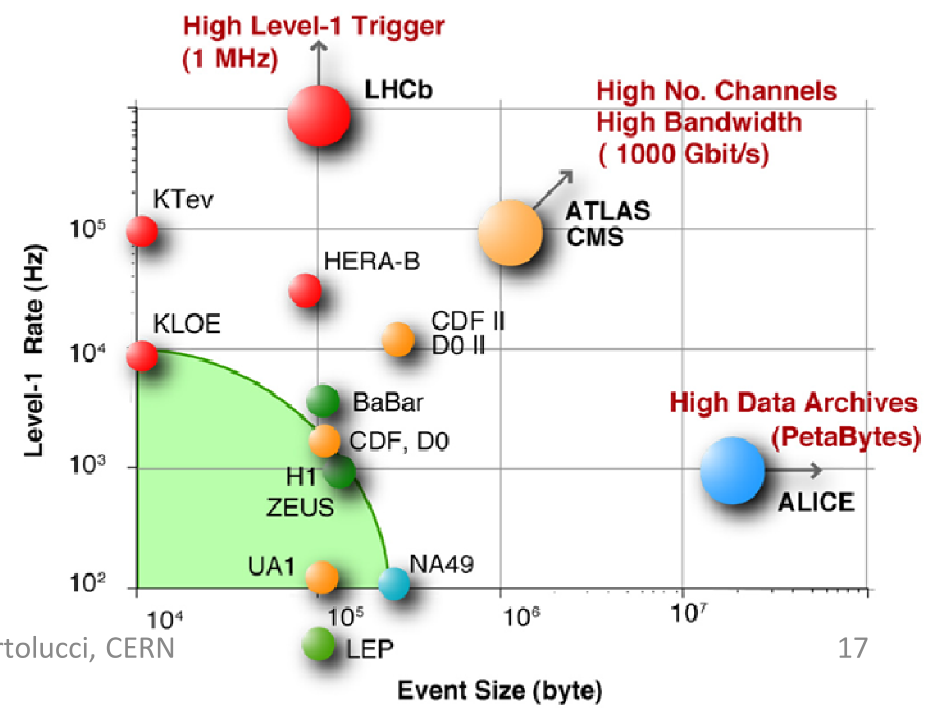

WLCG – what and why?

- A distributed computing infrastructure to provide the production and analysis environments for the LHC experiments
- Managed and operated by a worldwide collaboration between the experiments and the participating computer centres

- The resources are distributed – for funding and sociological reasons

- Our task is to make use of the resources available to us – no matter where they are located
  - We know it would be simpler to put all the resources in 1 or 2 large centres
  - This is not an option ... today

WLCG Worldwide LHC Computing Grid

LCG

(w)LCG – Project and Collaboration

LCG was set up as a project in 2 phases:

- Phase I – 2002-05 - Development & planning; prototypes
  - End of this phase the computing Technical Design Reports were delivered (1 for LCG and 1 per experiment)
- Phase II – 2006-2008 – Deployment & commissioning of the initial services
  - Program of data and service challenges

- During Phase II, the WLCG Collaboration was set up as the mechanism for the longer term:
  - Via an MoU – signatories are CERN and the funding agencies
  - Sets out conditions and requirements for Tier 0, Tier 1, Tier 2 services, reliabilities etc ("SLA")
  - Specifies resource contributions – 3 year outlook

WLCG Today

Tier 0; 11 Tier 1s; 61 Tier 2 federations (121 Tier 2 sites)

Today we have 49 MoU signatories, representing 34 countries:

Australia, Austria, Belgium, Brazil, Canada, China, Czech Rep, Denmark, Estonia, Finland, France, Germany, Hungary, Italy, India, Israel, Japan, Rep. Korea, Netherlands, Norway, Pakistan, Poland, Portugal, Romania, Russia, Slovenia, Spain, Sweden, Switzerland, Taipei, Turkey, UK, Ukraine, USA.

Today WLCG is:

- Running increasingly high workloads:
  - Jobs in excess of 650k / day; Anticipate millions / day soon
  - CPU equiv. ~100k cores
- Workloads are:
  - Real data processing
  - Simulations
  - Analysis – more and more (new) users
- Data transfers at unprecedented rates ➔ next slide

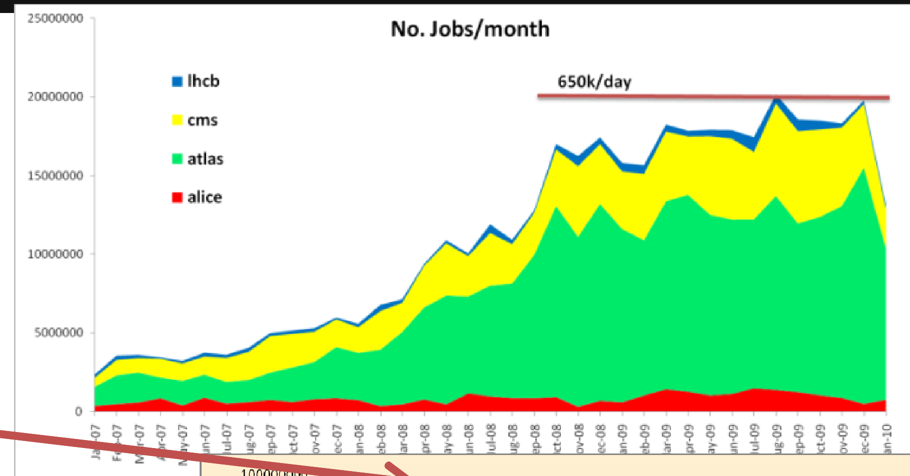

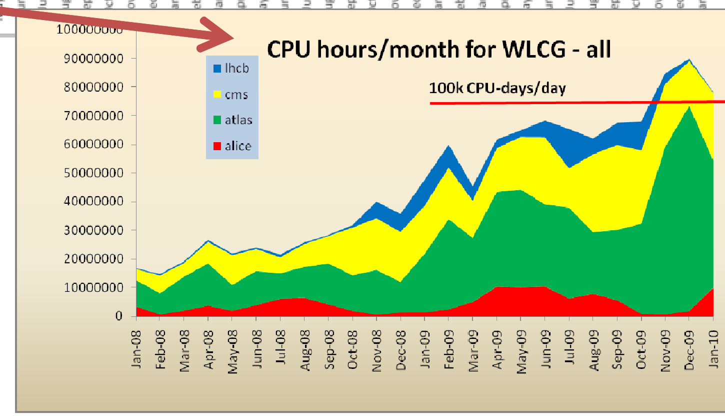

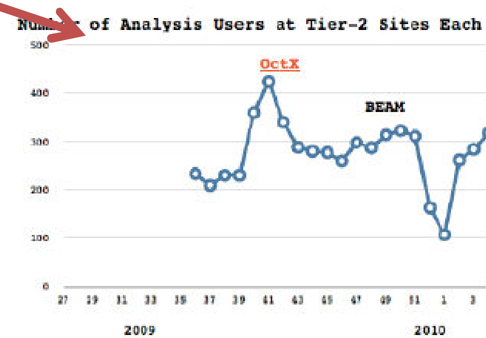

e.g. CMS: no. users doing analysis

Data transfers

Final readiness test (STEP'09)

Preparation for LHC startup

LHC physics data

2009: STEP09 + preparation for data

Nearly 1 petabyte/week

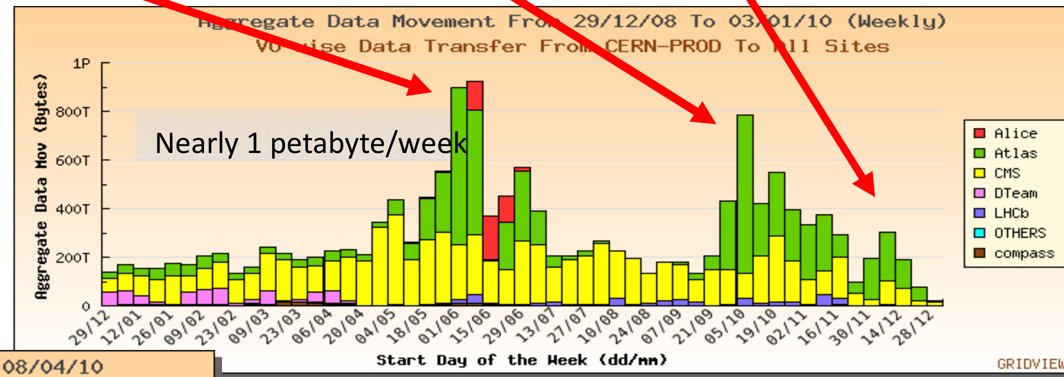

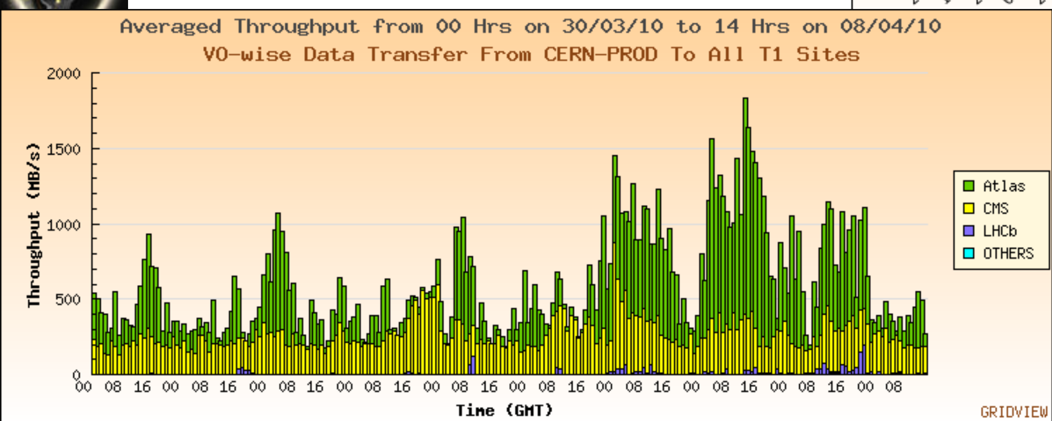

Real data – from 30/3

Castor traffic last week:
- > 4 GB/s input
- > 13 GB/s served

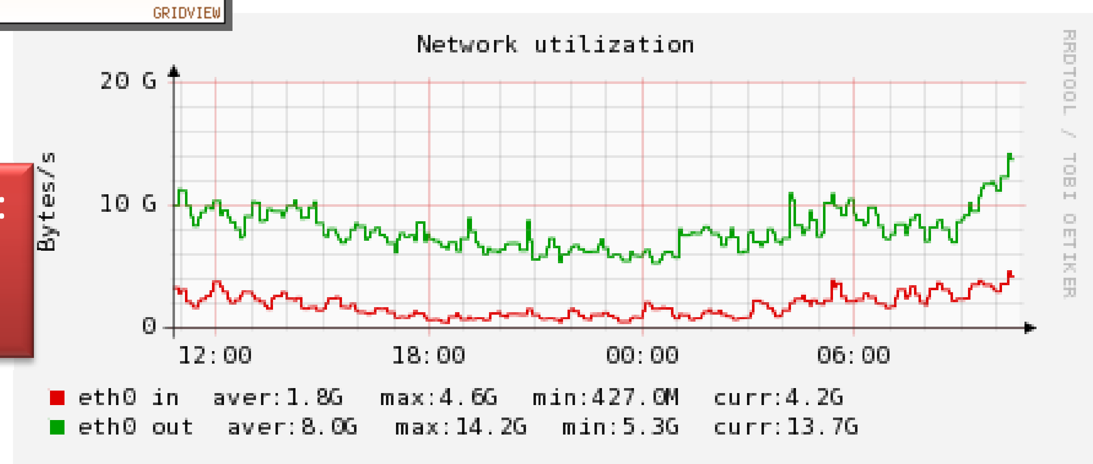

WLCG Worldwide LHC Computing Grid

LCG

LHCOPN – current status

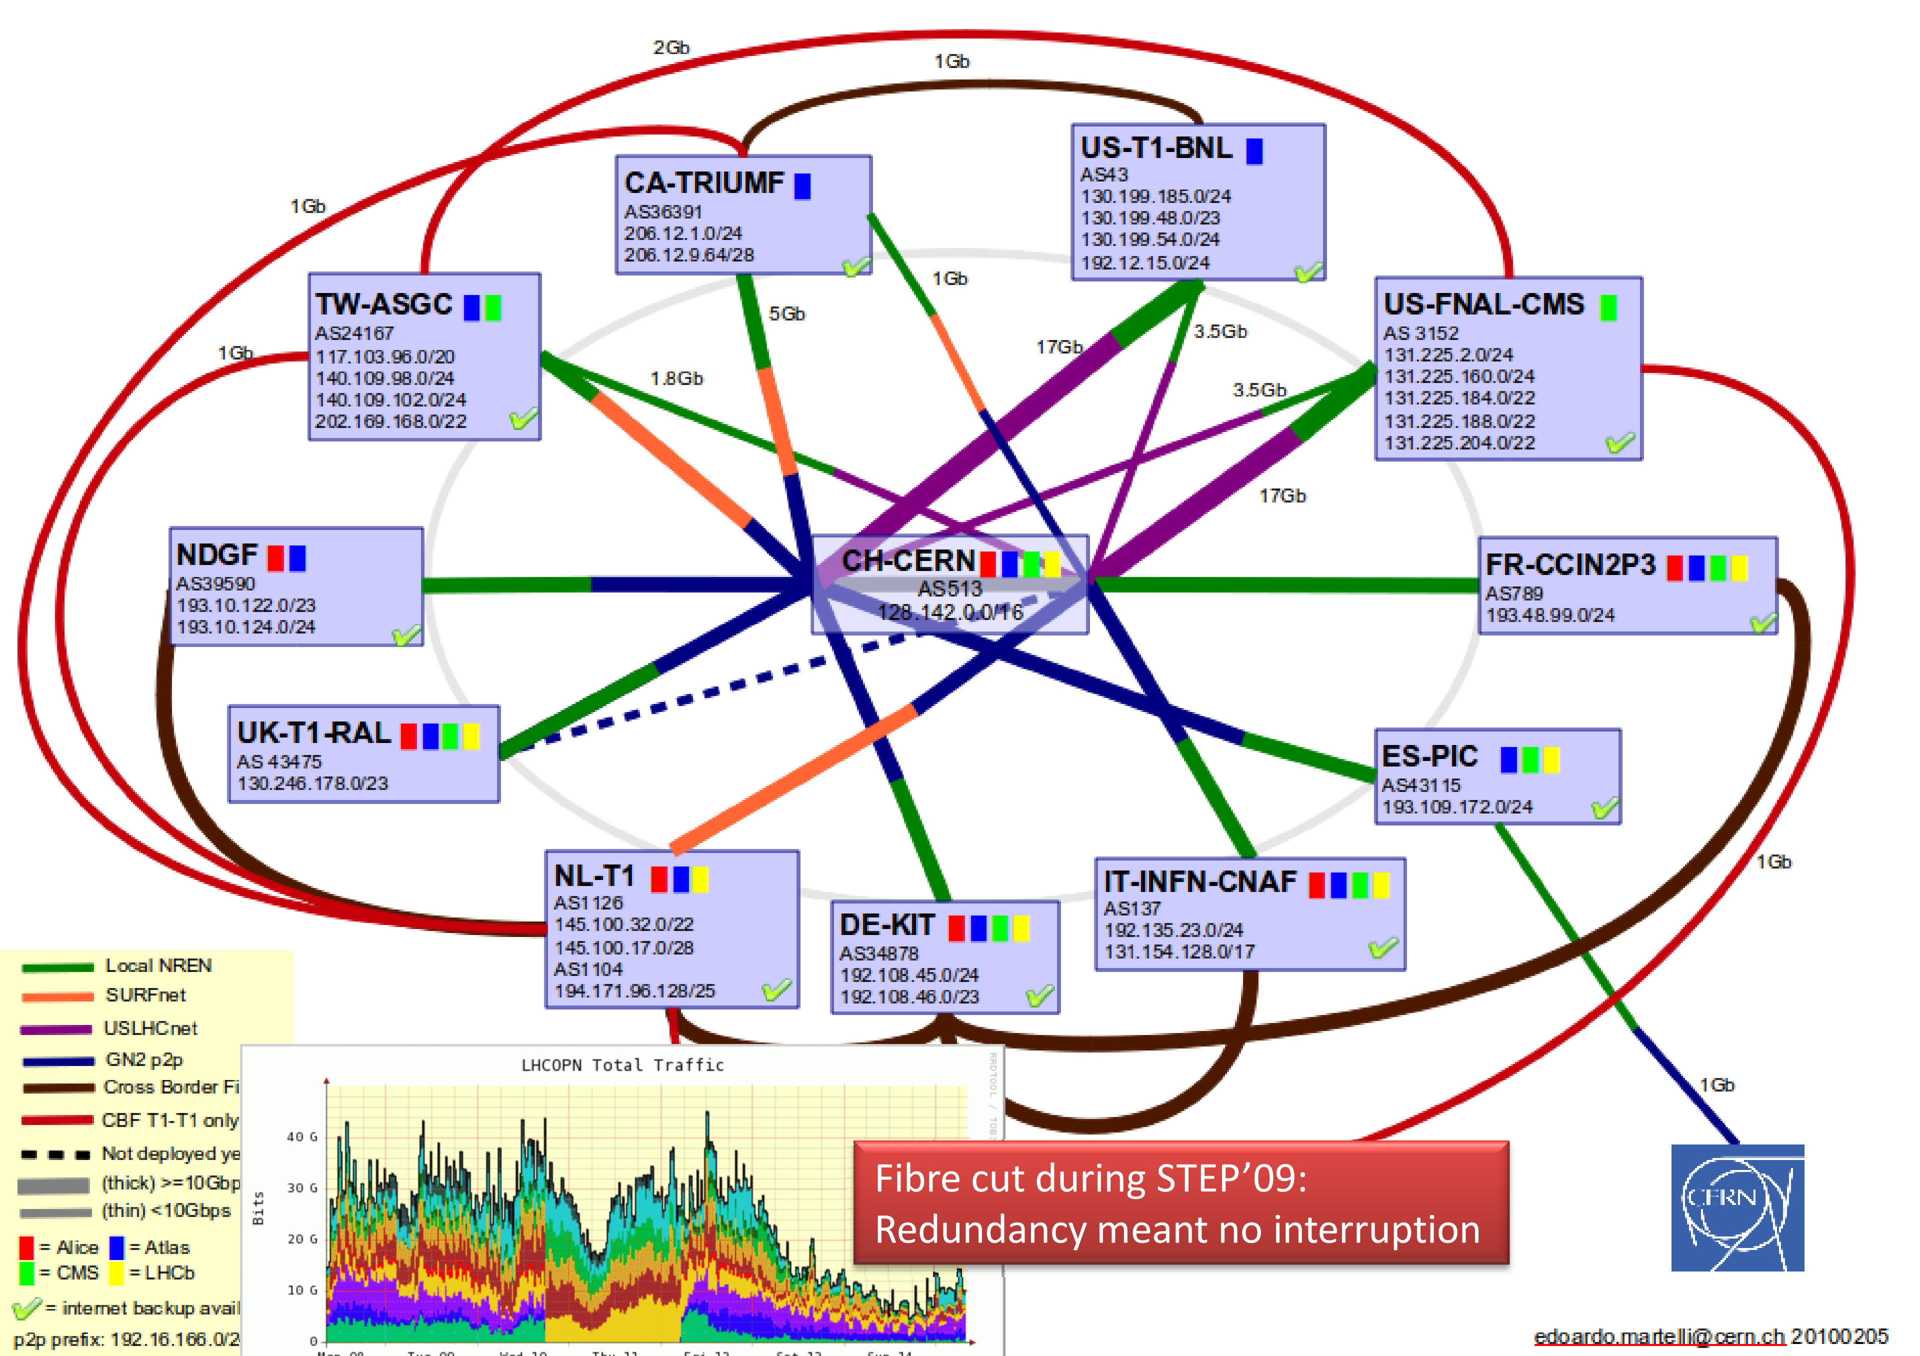

Fibre cut during STEP'09:
Redundancy meant no interruption

edoardo.martelli@cern.ch 20100205

Service quality: defined in MoU

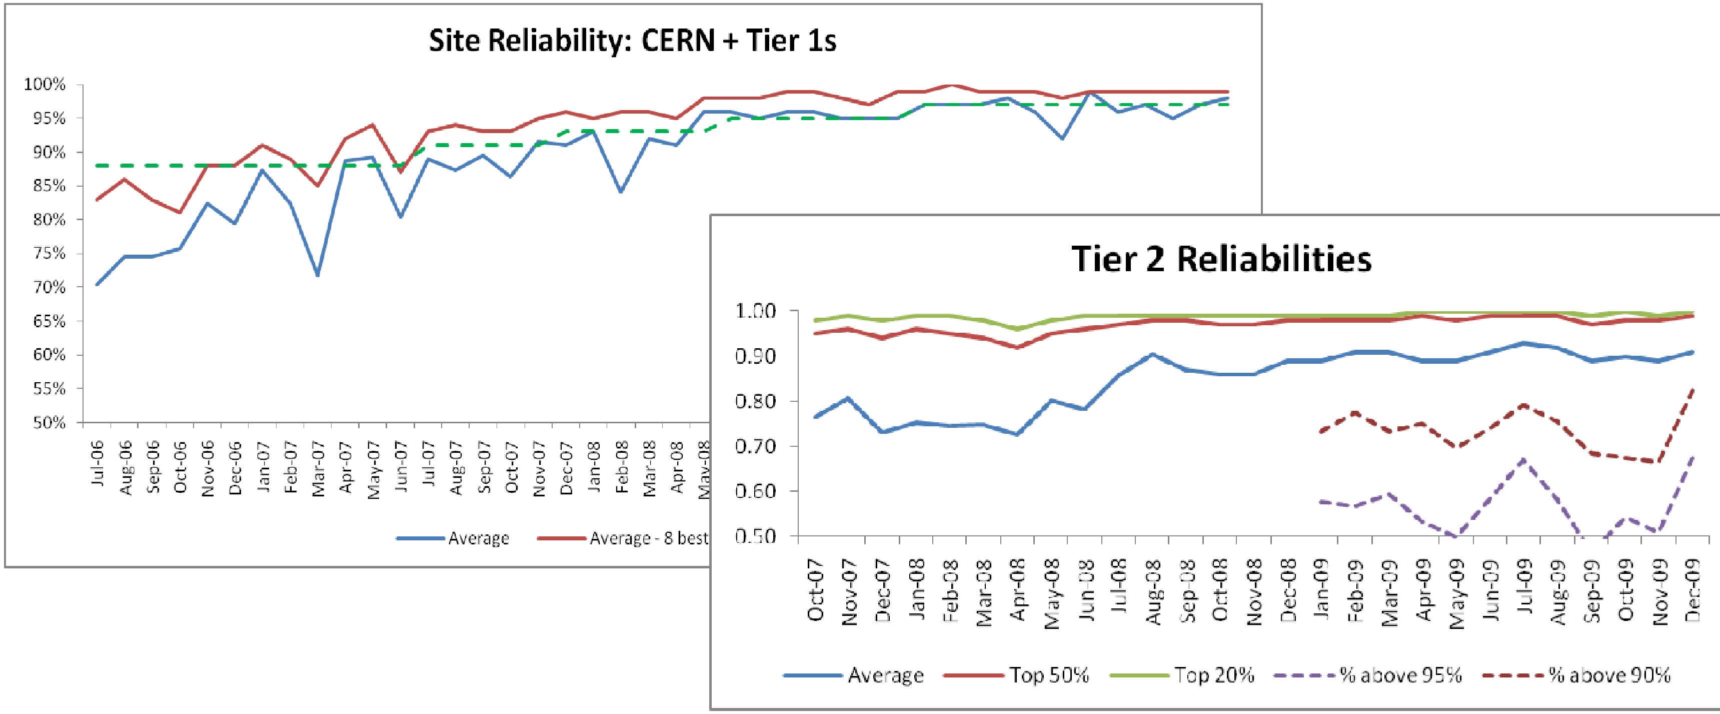

- MoU defines key performance and support metrics for Tier 1 and Tier 2 sites
  - Reliabilities are an approximation for some of these
  - Also metrics on response times, resources, etc.
- The MoU has been an important tool in bringing services to an acceptable level

Success with real data because:

- Focus on real and continuous production use of the service over several years (simulations since 2003, cosmics)
- Data and Service challenges to exercise all aspects of the service – not just for data transfers, but workloads, support structures etc.

- Challenges
  - SC1 → December 2004
  - SC2 → March 2005
  - SC3 → July2005
    - Testing with special emphasis on Data Management
    - Goals largely exceeded for the T2 sites, service reliability and sustained transfer rates
  - SC4 → June 2006
    - Offline data processing requirements can be handled by the Grid to the nominal LHC data rate
    - Large participation of T2 sites, all T1 sites were in
    - Required transfer rates (disk-tape) achieved and in some cases exceeded
  - CCRC'08 → March + June 2008
    - Measurement of the readiness of the Grid services and operations before real data takin
    - All experiments simultaneously stressing the WLCG infrastructure in close to real conditions
    - Experiments running their Full Dress Rehearsals and scheduling key periods together with the CCRC'08 challenge
  - STEP'09 → May 2009
    - Stress and scale testing of all experiment workloads including tape recall and massive end user analysis

Readiness of the computing

- Has meant very rapid data distribution and analysis
  - Data is processed and available at Tier 2s within hours!

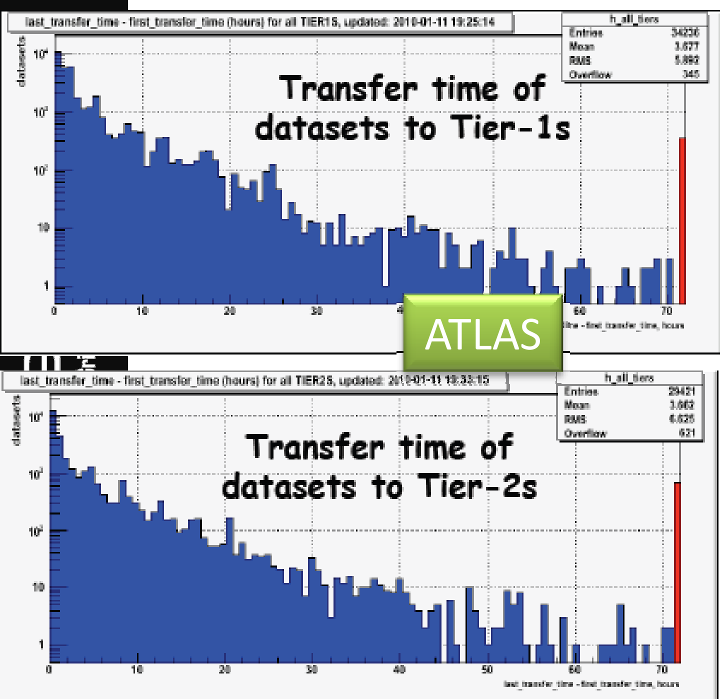

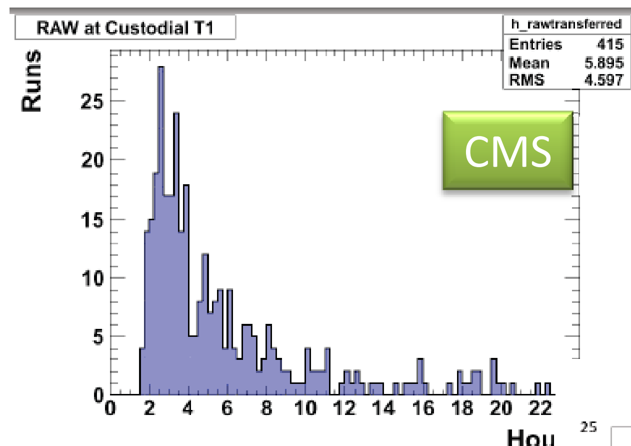

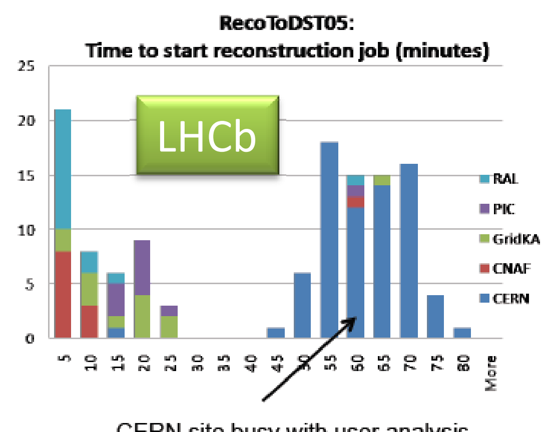

And physics output ...

First paper (submitted 28/11)

- $dN_{ch}/d\eta$ for $|\eta| < 0.5$
- $dN_{ch}/d\eta$ vs η

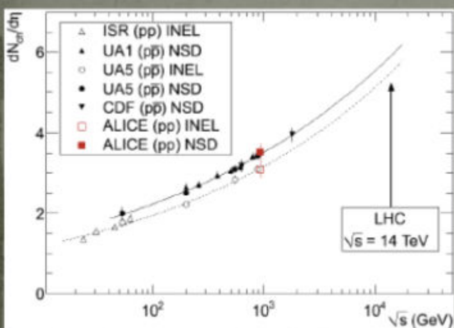

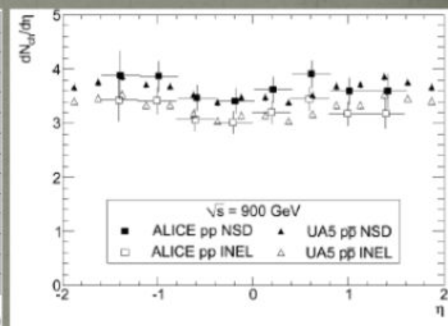

K. Aamodt et al. (ALICE), Eur. Phys. J C 65 (2010) 111

Reconstructed K_s and Λ masses

Tracking without VELO

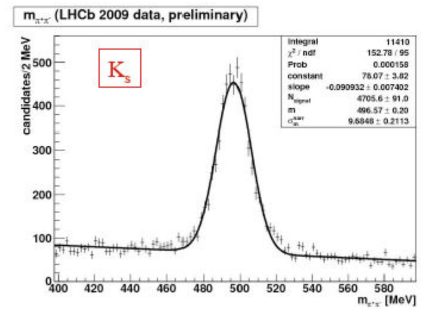

$m = (496.6 \pm 0.2_{stat.})$ MeV/c^2
$\sigma = (9.7 \pm 0.2_{stat})$ MeV/c^2
PDG: 497.61(2) MeV/c^2

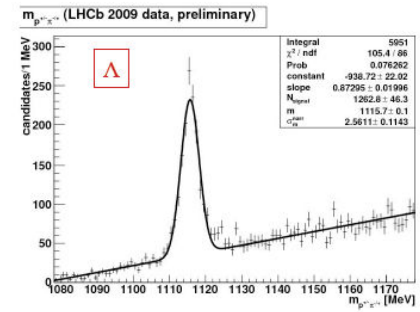

$m = (1115.7 \pm 0.1_{stat.})$ MeV/c^2
$\sigma = (2.6 \pm 0.1_{stat})$ MeV/c^2
PDG: 1115.683(6) MeV/c^2

Analysis overview - Results

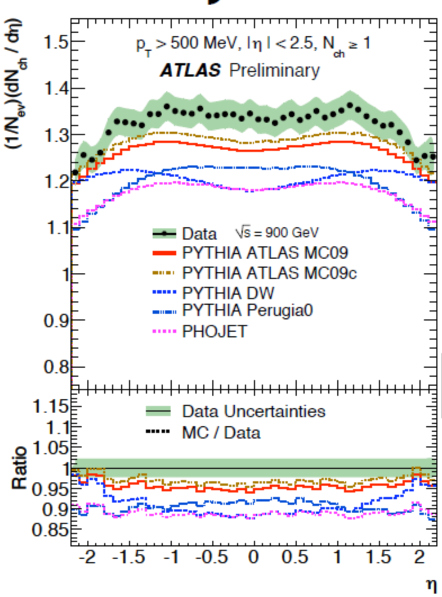

- dN/dη distribution
- Shape agrees well with some PYTHIA tunes
- ATLAS data shows higher value then all MCs
  - MCs tuned in different region of phase space

Note suppressed 0 on y-axis

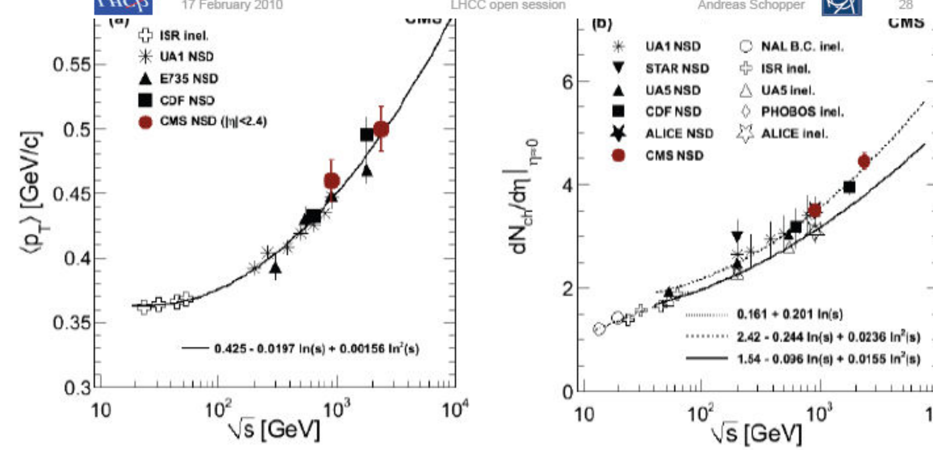

First CMS Paper on pp data
http://arxiv.org/abs/1002.0621
Accepted Feb 7 for publication in the Journal of High Energy Physics (JHEP)

Grids & HEP: Common history

- CERN and the HEP community have been involved with grids from the beginning
- Recognised as a key technology for implementing the LHC computing model
- HEP work with EC-funded EDG/EGEE in Europe, iVDGL/Grid3/OSG etc. in US has been of clear *mutual* benefit
  - Infrastructure development driven by HEP needs
  - Robustness needed by WLCG is benefitting other communities
  - Transfer of technology from HEP
    - Ganga, AMGA, etc used by many communities now

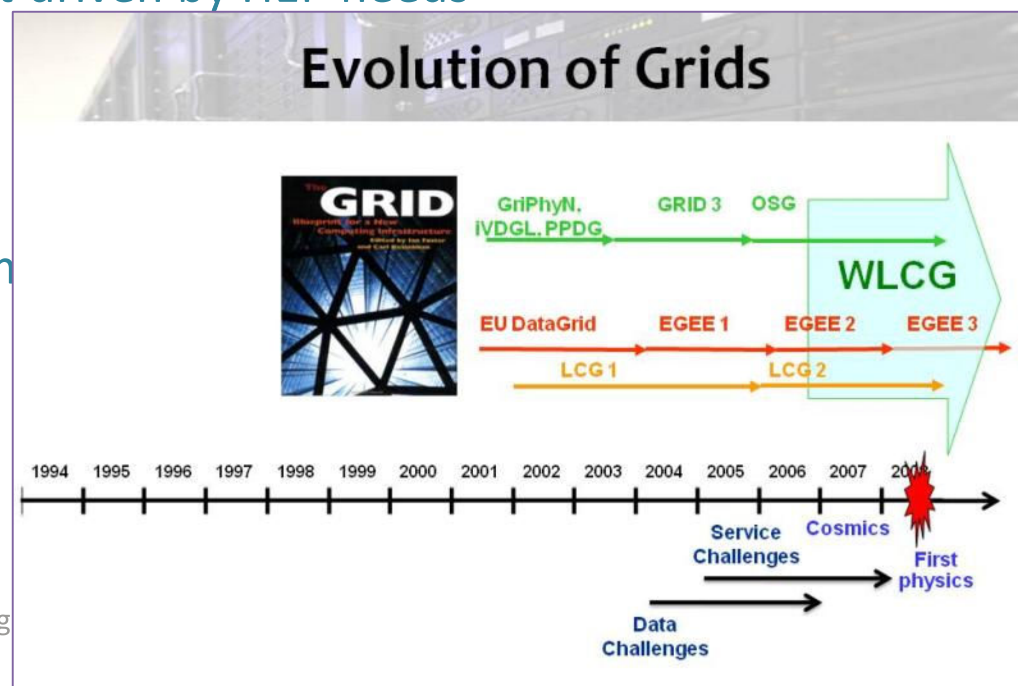

Large scale = long times

- LHC, the experiments, & computing have taken ~20 years to build and commission
- They will run for at least 20 years
- We must be able to rely on long term infrastructures
  - Global networking
  - Strong and stable NGIs (or their evolution)
    - That should be eventually self-sustaining
  - Long term sustainability - must come out of the current short term project funding cycles

WLCG Worldwide LHC Computing Grid

LCG

Longer term future

- Long term sustainability of the infrastructure

We have achieved what we set out to do – provide an environment for LHC computing;
And we have spun-off significant general science grid infrastructures
BUT: is it sustainable in the long term???

- Need to adapt to changing technologies
  - Major re-think of storage and data access
  - Virtualisation as a solution for job management
  - Complexity of the middleware compared to the actual use cases
- Network infrastructure
  - This is the most reliable service we have
  - Invest in networks and make full use of the distributed system (i.e. Leave data where it is)?

WLCG Worldwide LHC Computing Grid
LCG

Sustainability

- Grid middleware
  - Is still dependent upon project funding – but this is a very risky strategy now
  - Limited development support in EMI (for example)
- Must (continue) to push for mainstream, industrial solutions:
  - Messaging, Nagios for monitoring are good examples
  - Fabric and job management are good candidates for non-HEP-specific solutions
- Because .... Data Management is not solved
  - And we must invest significant effort here to improve the reliability and overall usability; must reduce complexity (e.g. SRM – functionality and implementations)
  - But – we are not alone – other sciences expect to have significant data volumes soon
  - Must take care not to have special solutions

CERN, WLCG and EGI – the future

- WLCG needs to be able to rely on strong and stable global e-science infrastructures
  - In Europe this means the NGIs and EGI
- WLCG is a very structured large user community
  - It can serve as a model for others – they can also learn from our mistakes
- CERN has connections to the other EIROs which are also large scientific communities, several of which are associated with ESFRI projects
  - Can play a role in bringing these to EGI
- CERN also supports other visible communities:
  - E.g. UNOSat

LHC is not alone

- HEP has been a leader in needing and building global collaborations in order to achieve its goals
- It is no longer unique – many other sciences now have similar needs
  - Life sciences, astrophysics, ESFRI projects
  - Anticipate huge data volumes
  - Need global collaborations
- There are important lessons from our experiences,
  - HEP was able to do this because it has a long history of global collaboration; missing from many other sciences
- We must also collaborate on common solutions where possible

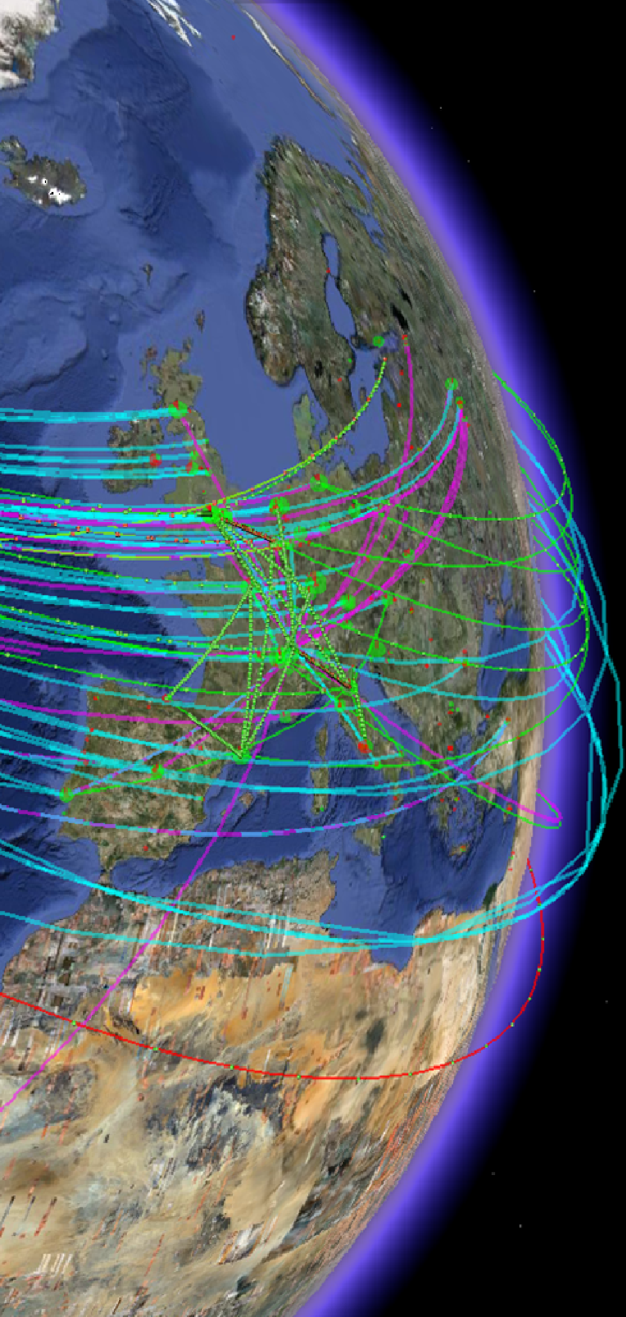

Summary

- LHC is operational and producing physics!
- Collaborative science on a global scale is a reality and LHC can act as a model for others