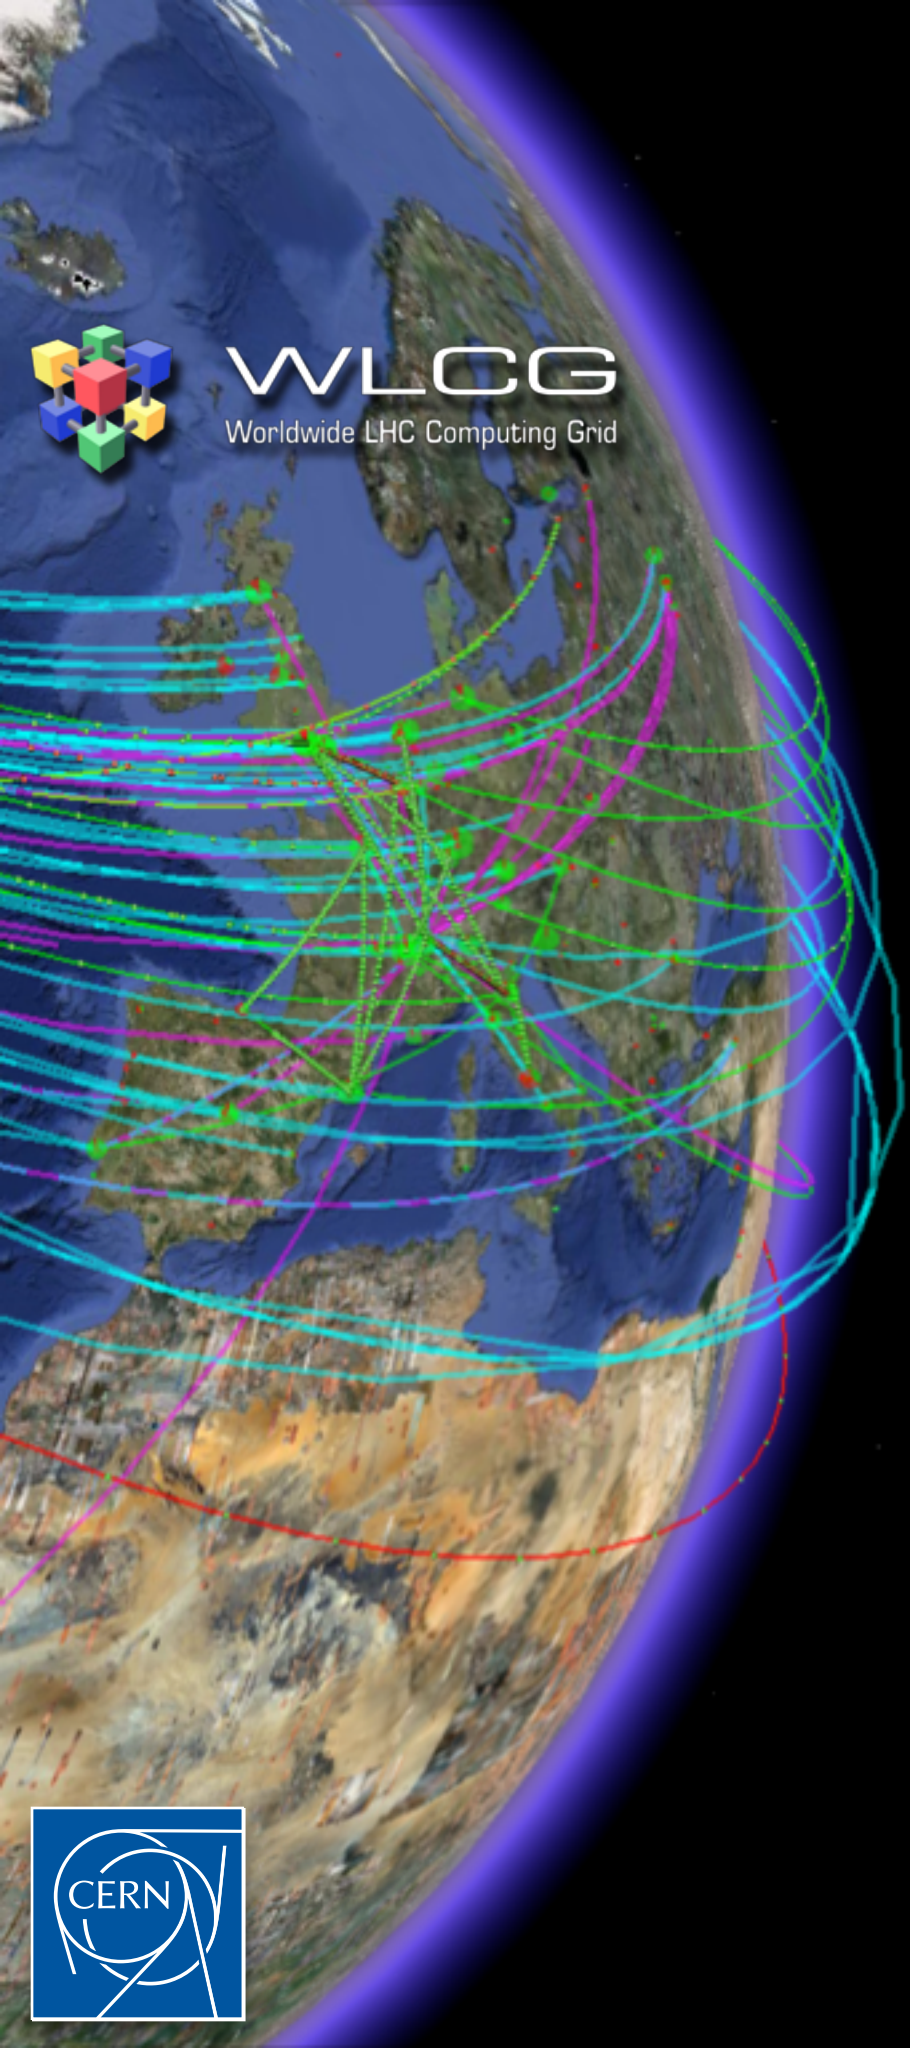

CERN

Markus Elsing

LHC Software and Computing through the Ages

Looking at more than 15 years
of Software and Computing
for the LHC Experiments,
at current developments
as well as at challenges ahead

8th Annual Workshop
1-3 December 2014
DESY, Hamburg

LHC Computing is Big Data

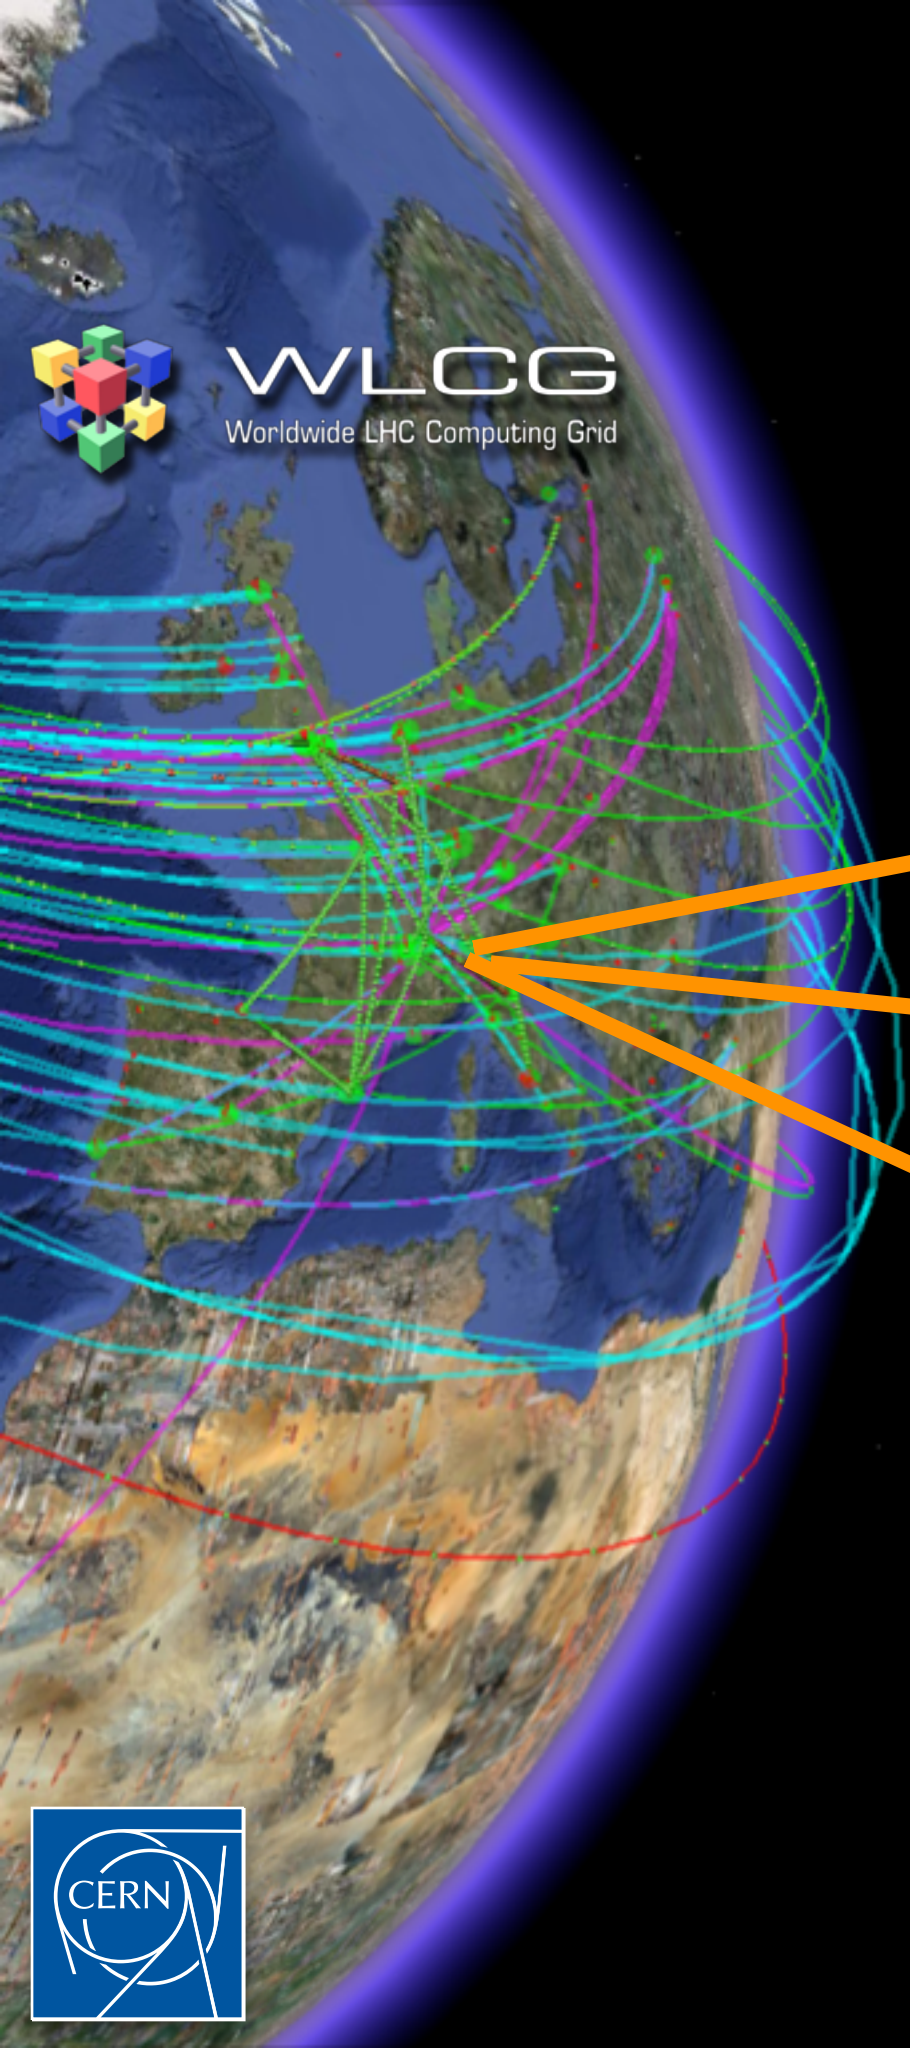

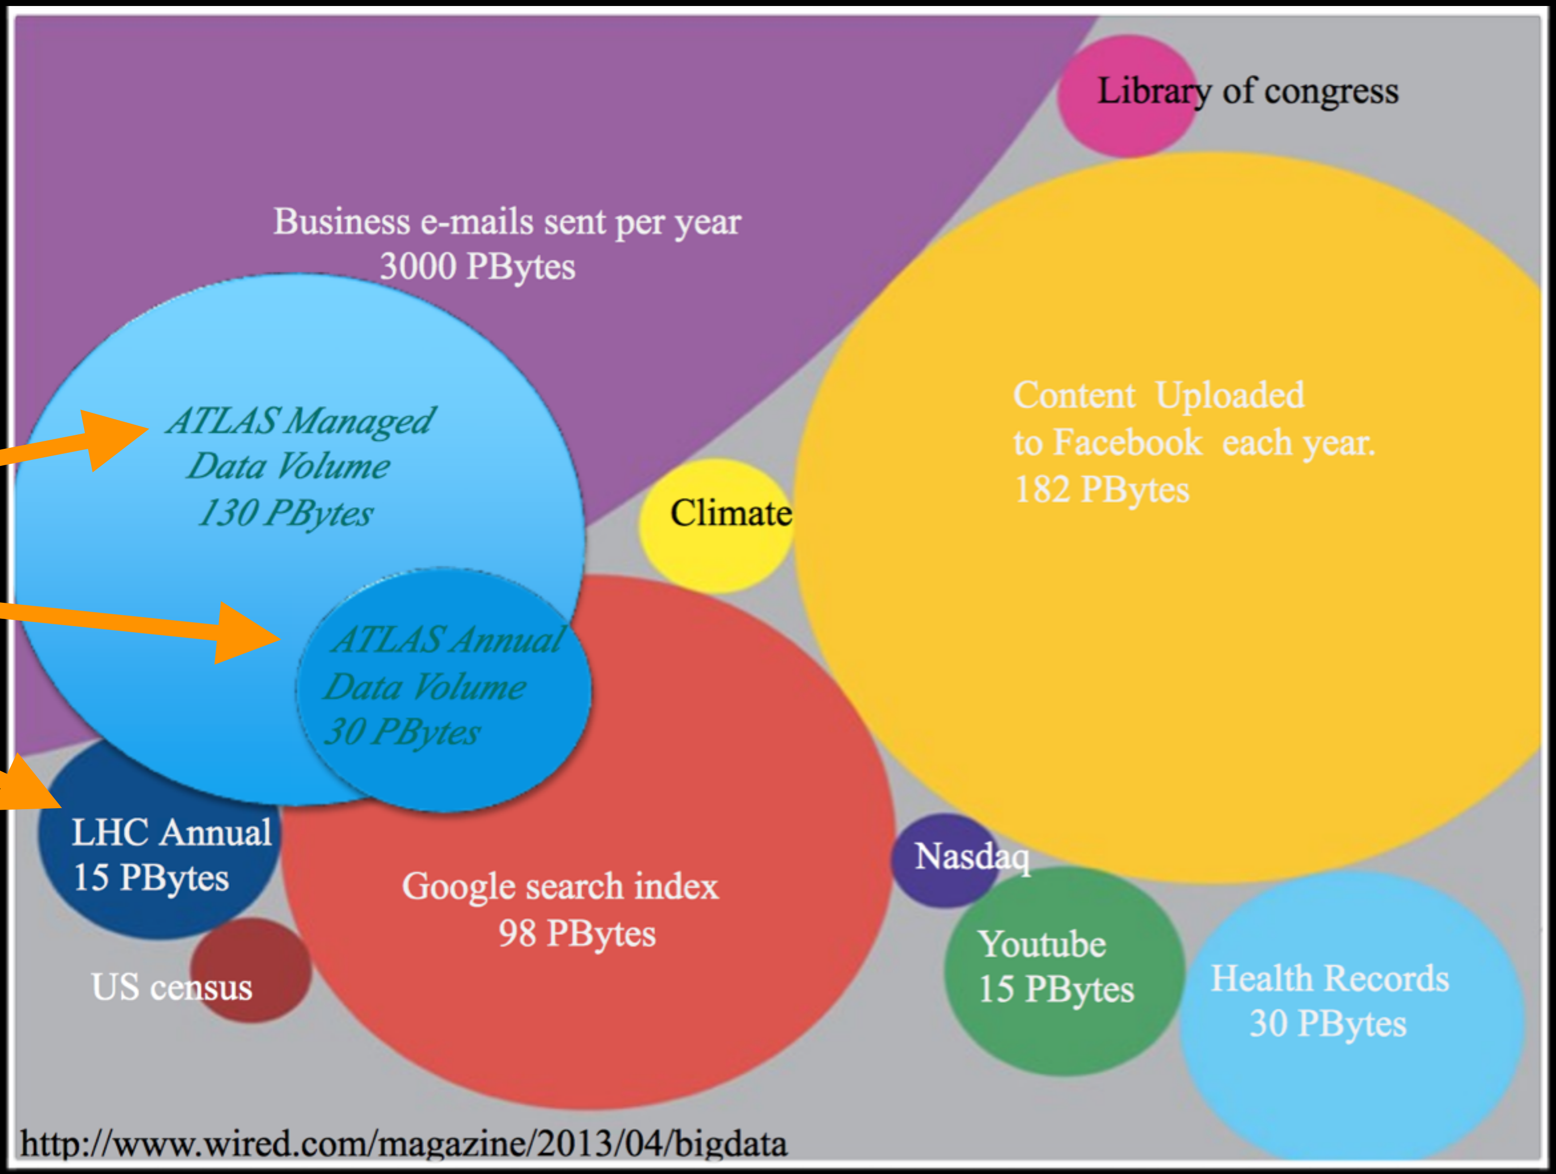

- we started more than a decade before everybody was talking about it !
  - ➡ with a science budget, unlike Google or Facebook

The early Times of the LHC Experiments

- project started during LEP aera in '90s
  - ➡ LoI and TDRs done with infrastructure of the time
    - software in FORTRAN 77, CERNLIB incl. PAW, Geant3
    - general LINUX services at CERN started in 1997
- huge challenges ahead
  - ➡ LHC is a high energy and high luminosity machine
    - unprecedented trigger rates, event sizes, pileup
  - ➡ lots of questions to answer...
    - design the High Level Trigger systems ?
      (can it be done in software, even re-using offline code)
    - how to build up the software infrastructure ?
      (move to C++/OO, learn from BaBar and CDF/D0 Run-2 preparation)
    - a computing infrastructure matching the needs ?
      (building "the" LHC computing centre at CERN wasn't an option)
    - ...
  - ➡ not to forget, LHC startup was supposed to be 2005
    (well, it came different after all)

CERN computing at LEP times

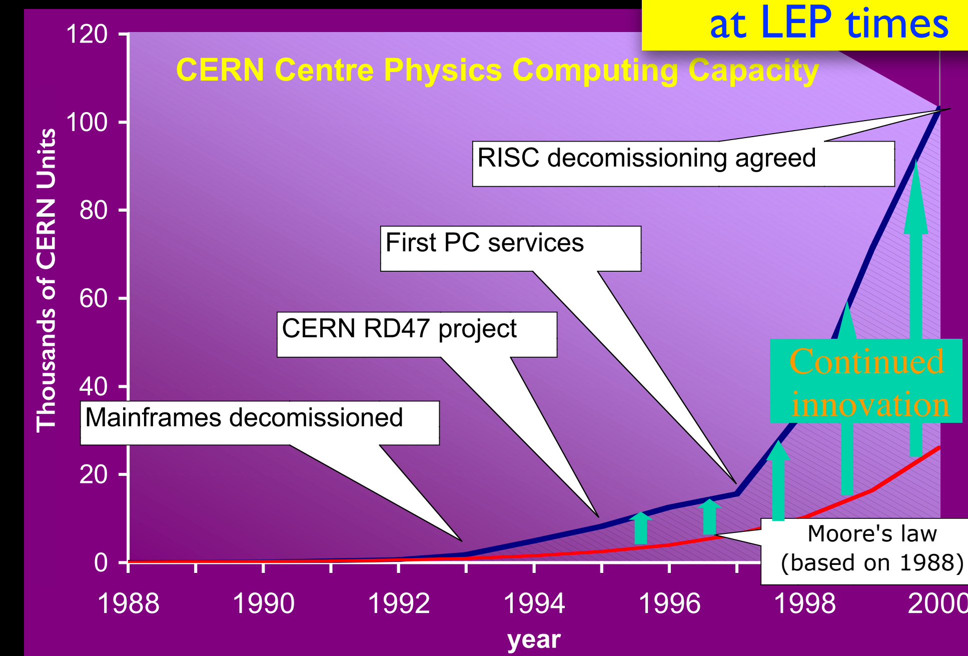

S.Bethke, LHC Computing Review, 2001

log scale

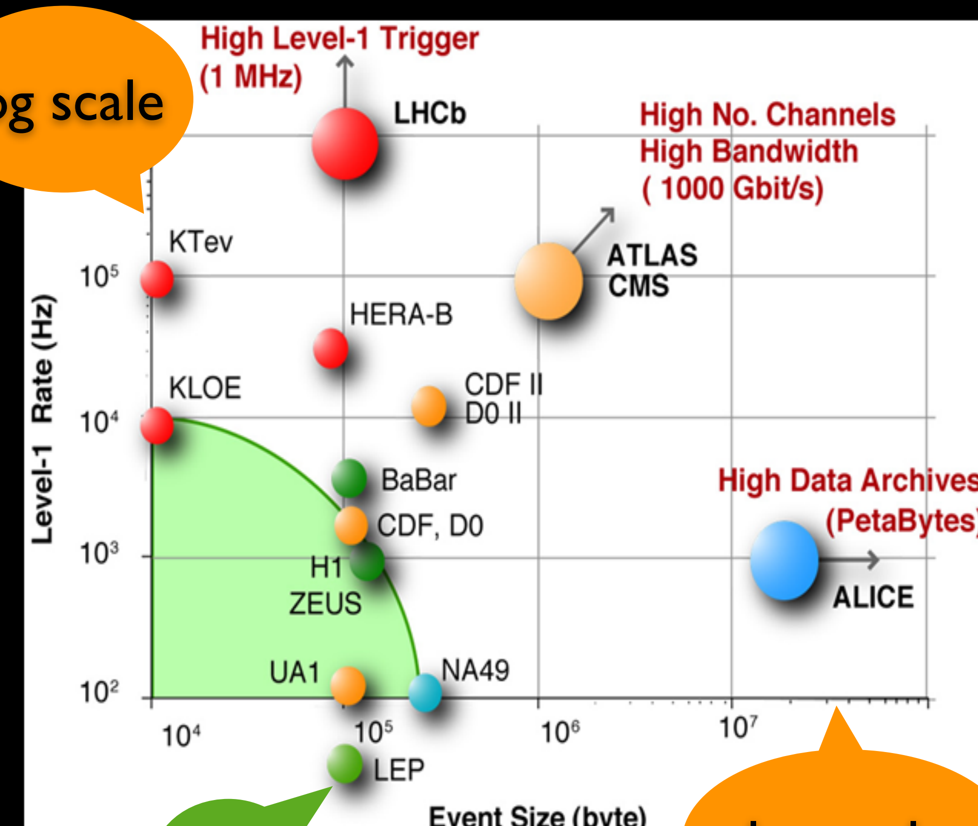

LEP

log scale

Markus Elsing

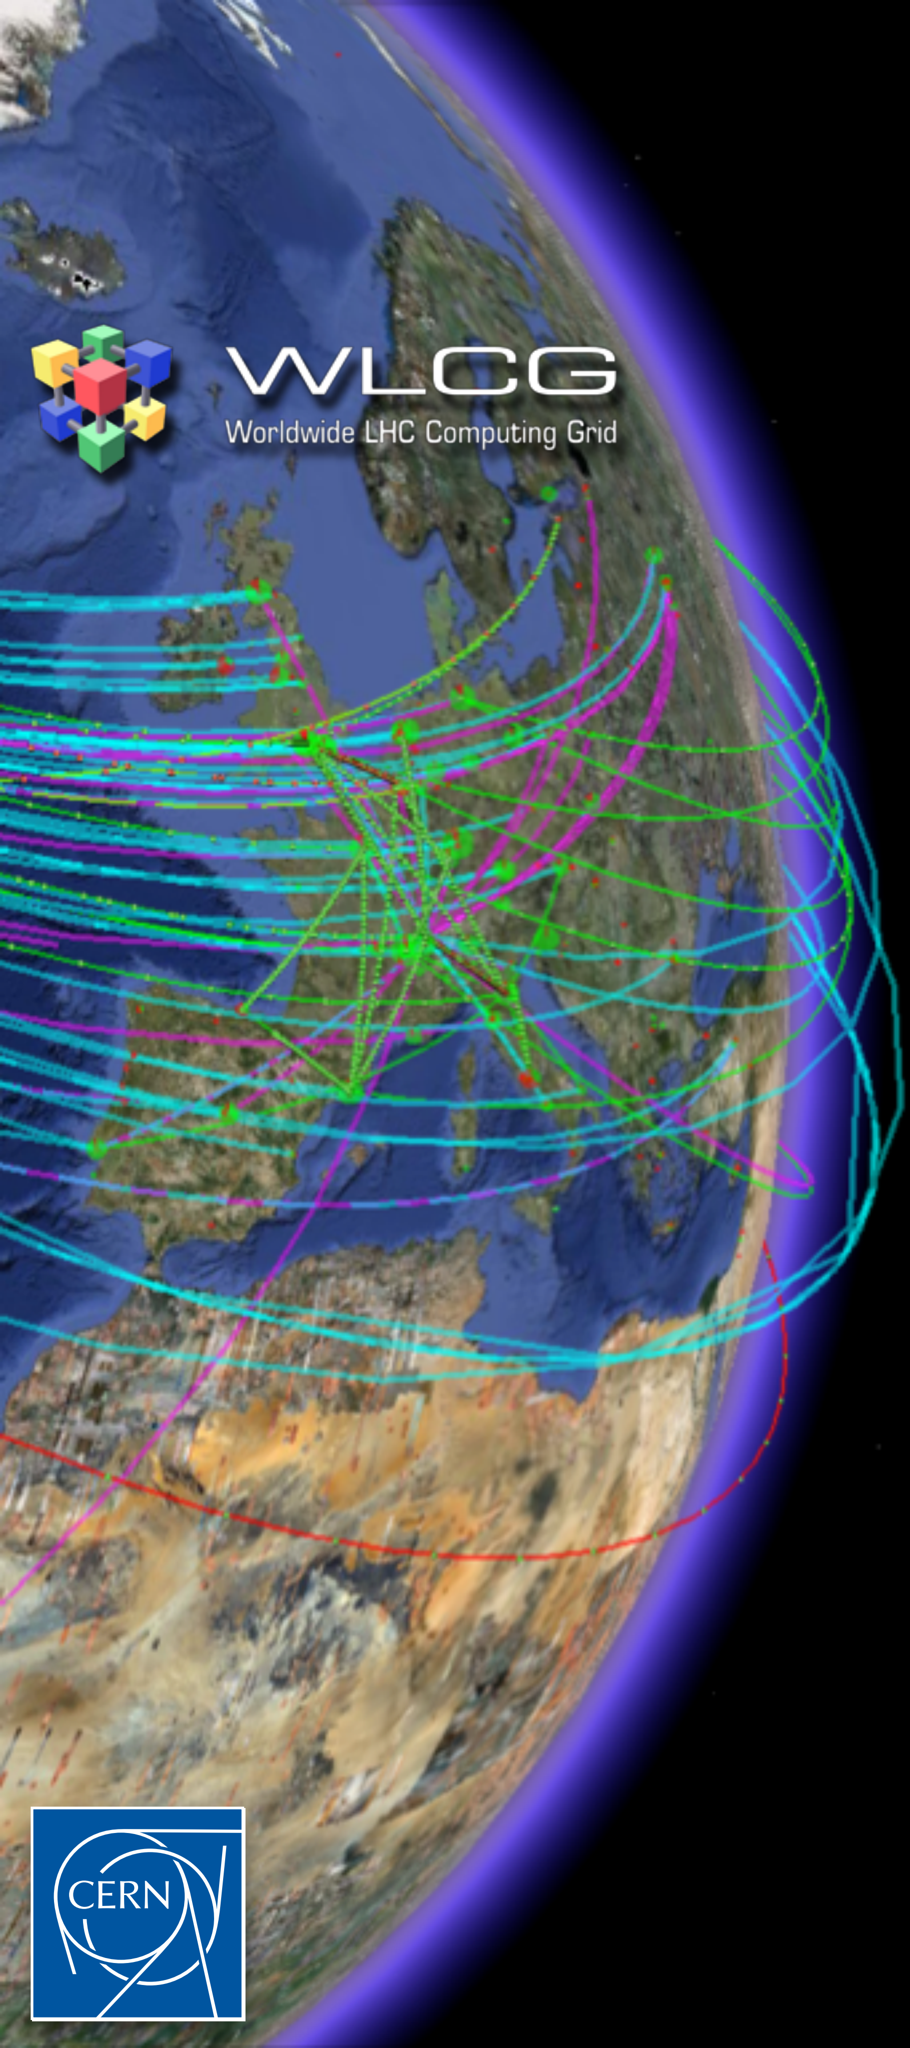

Outline of this Talk

- the LHC Computing GRID
  - ➡ facing the challenge
- Data and Service Challenges
  - ➡ commissioning GRID based computing
- building up the software of the experiments
- early physics and experience from Run-1
- the Higgs discovery
  - ➡ the role of software and computing
- preparing for Run-2
  - ➡ first upgrades of software and computing
- future software and computing challenges
- summary and outlook

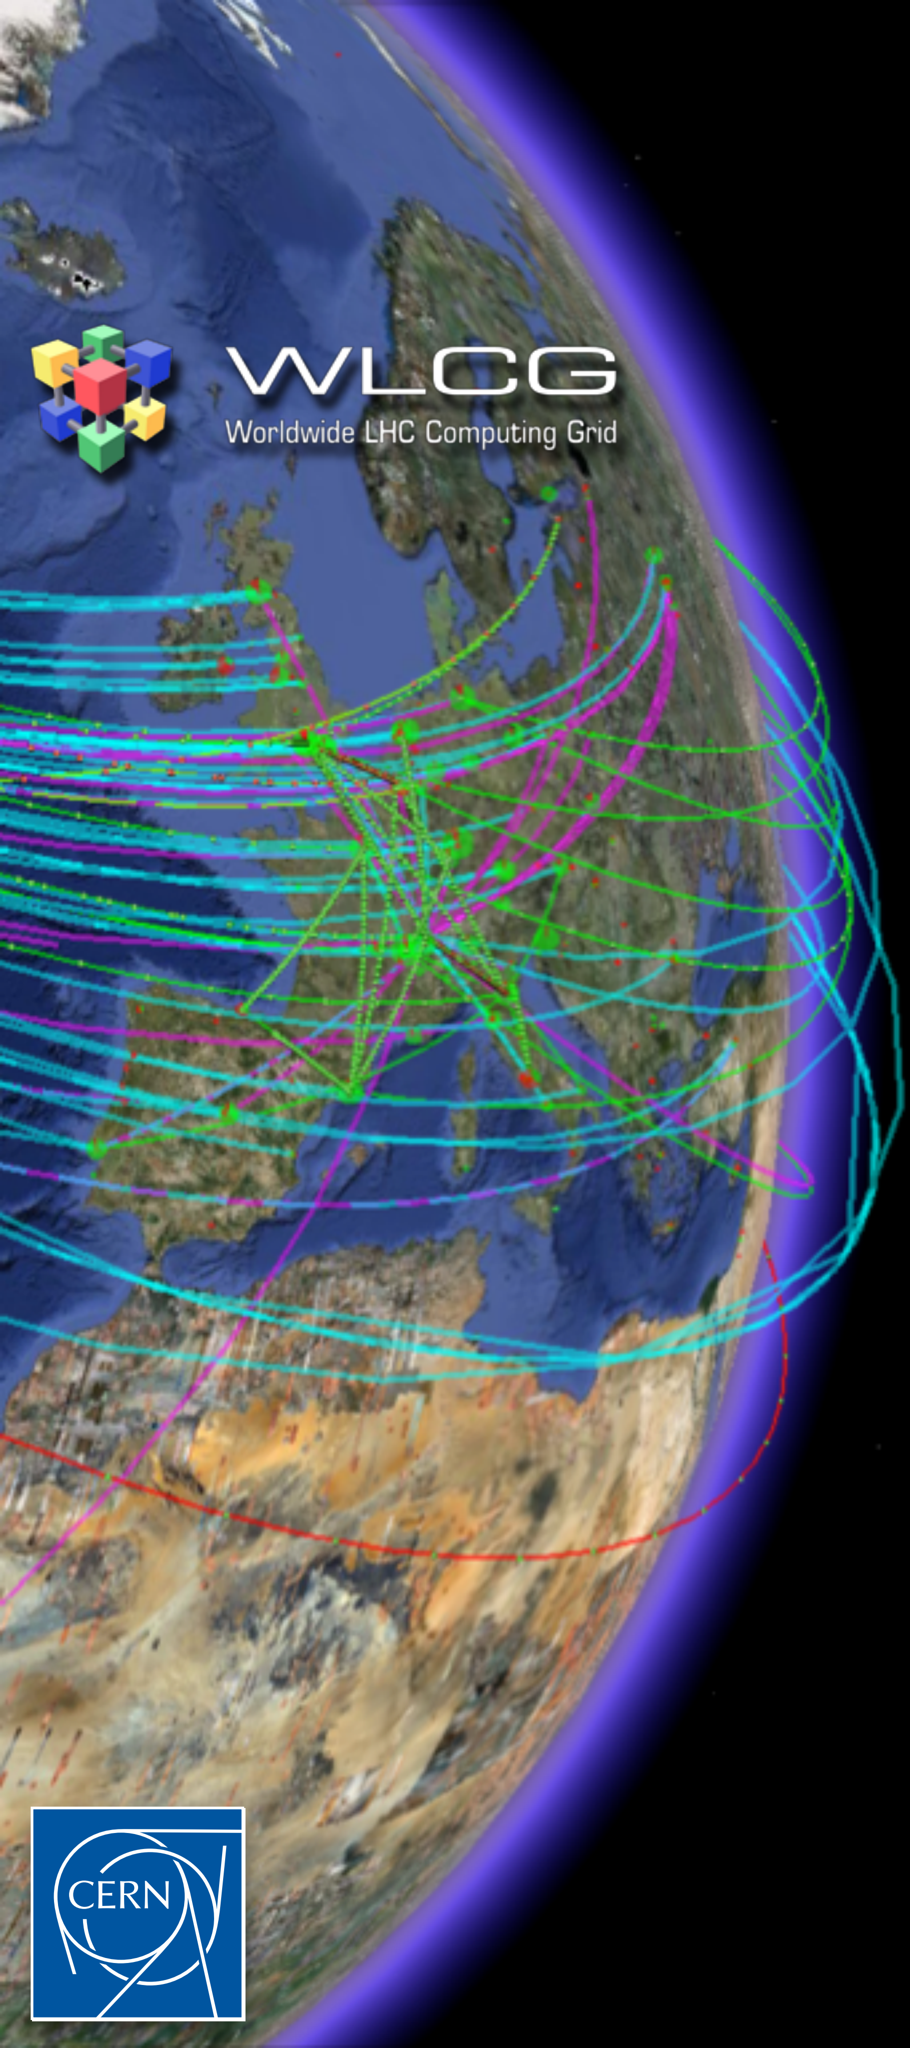

The LHC Computing GRID: Facing the Challenge

The Grid: Blueprint for a New Computing Infrastructure

I.Foster, C.Kesselmann (1998)

"The grid promises to fundamentally change the way we think about and use computing. This infrastructure will connect multiple regional and national computational grids, creating a universal source of pervasive and dependable computing power that supports dramatically new classes of applications."

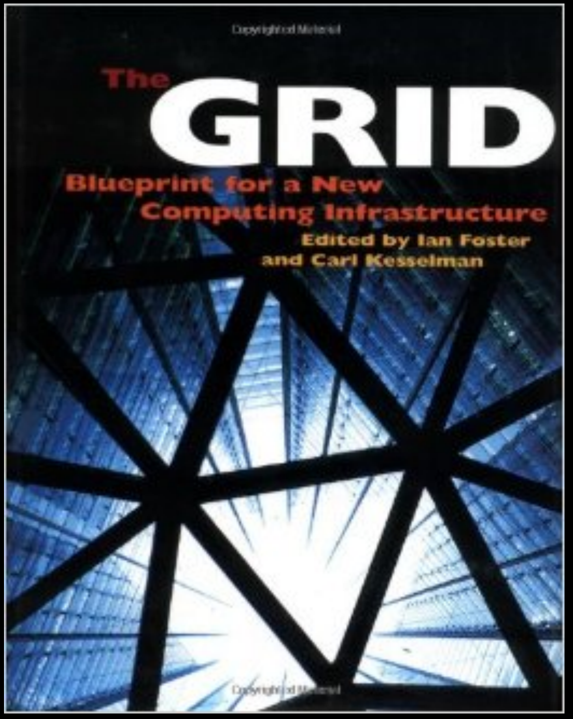

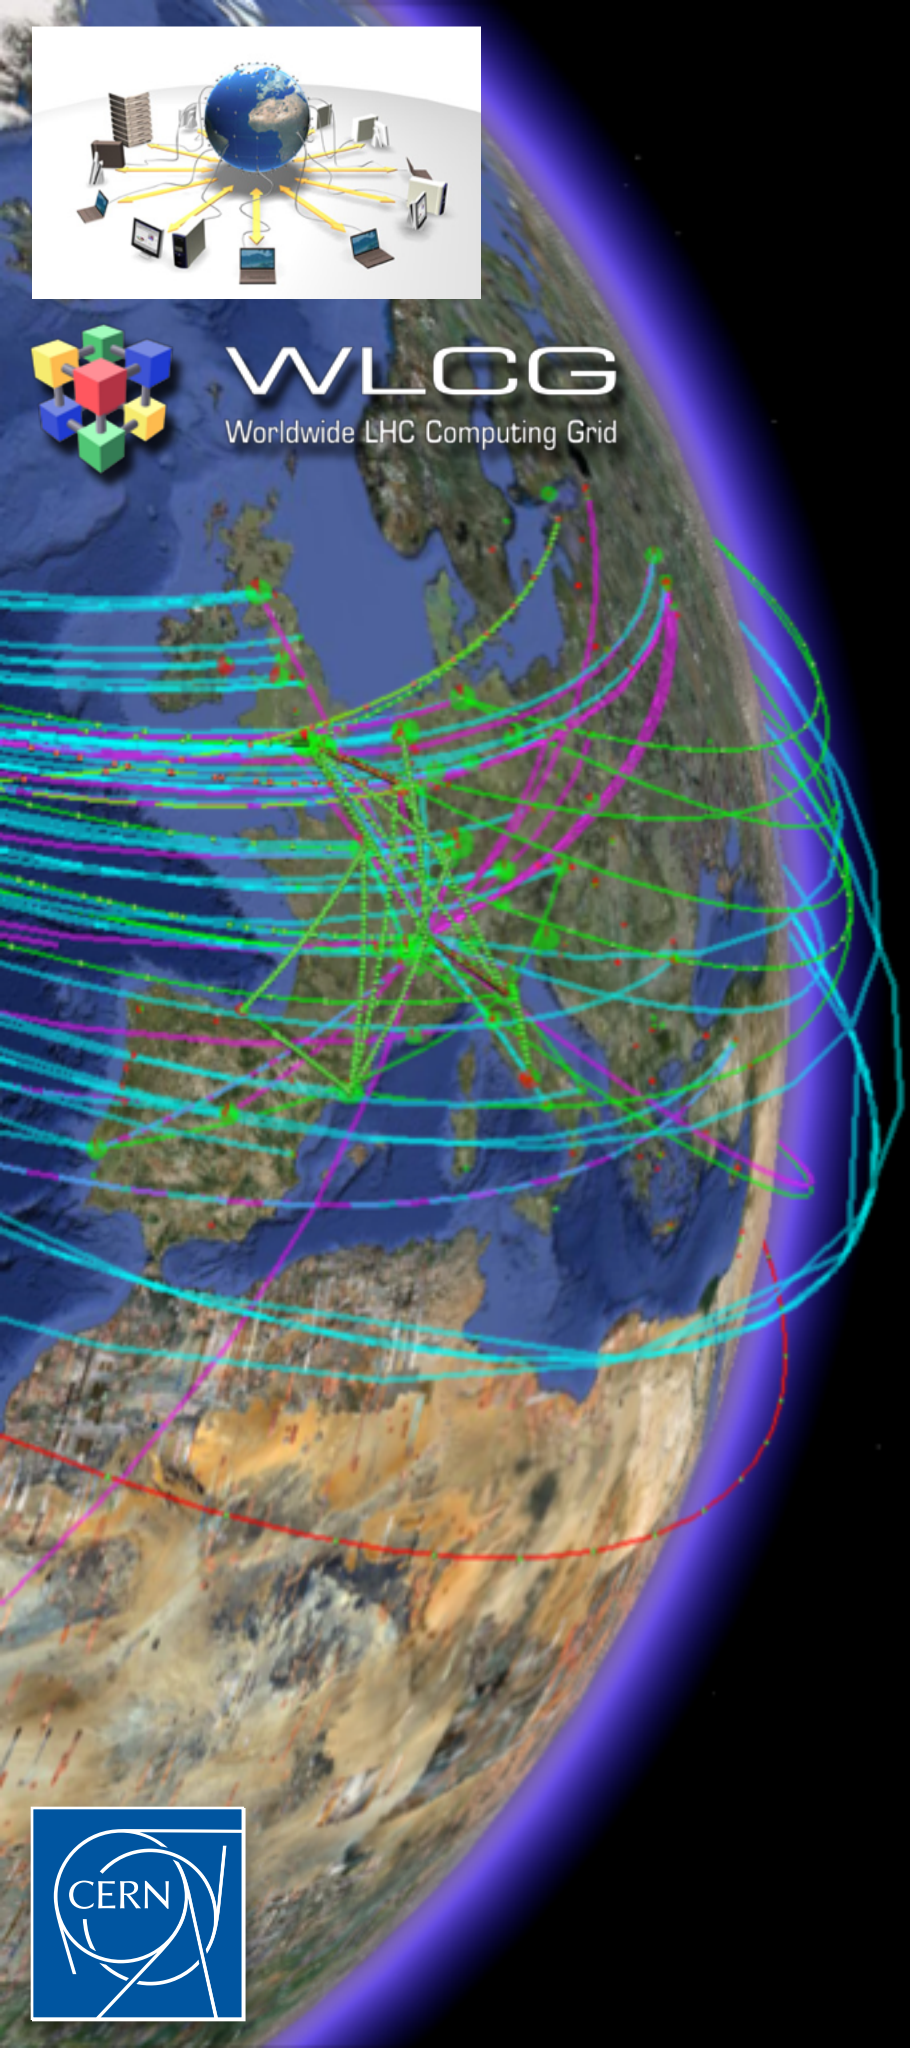

The Middleware

- layer of services to implement a distributed computing GRID
  - ➡ derived from GLOBUS (1998)
    - first middleware widely available (proof of concept)

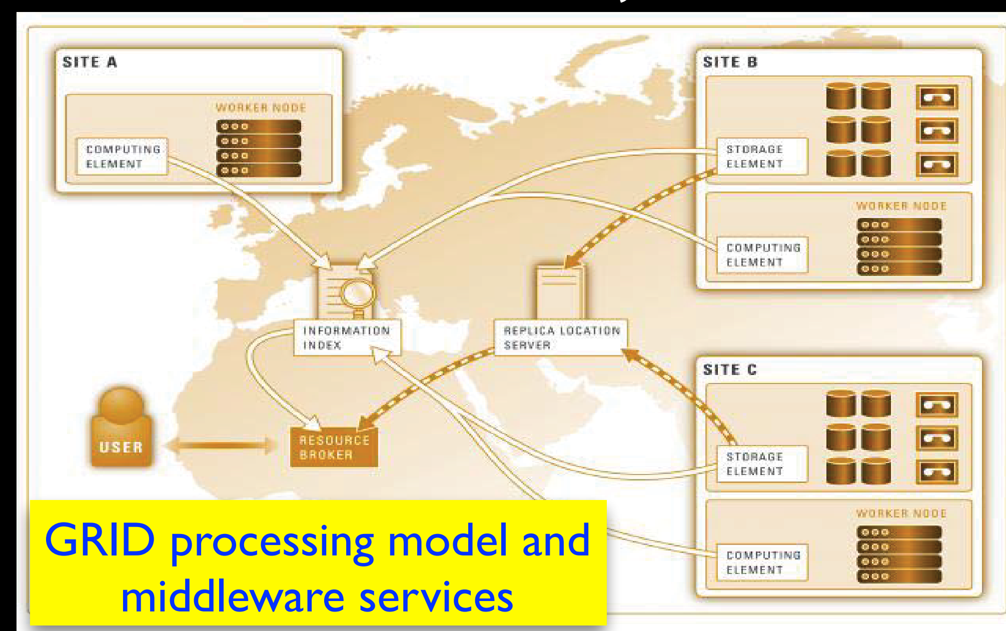

GRID processing model and middleware services

  - ➡ complex software developed in EU and US
    - information system
    - authentication and authorisation system
    - file catalogs and file transfer systems
    - job brokering
    - interfaces to storage and batch systems
    - etc...

The MONARC Model (1999)

- hierarchical model for LHC GRID computing
  - ➡ Models of Network Analysis at Regional Centres (1999)

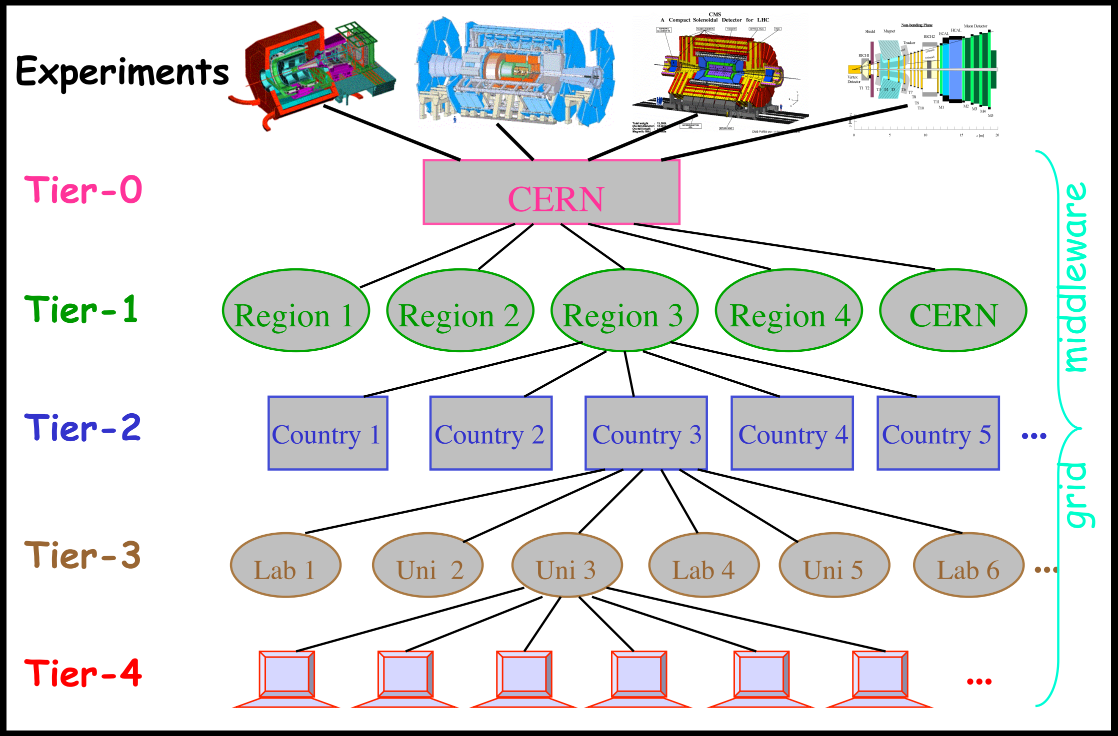

S.Bethke, LHC Computing Review, 2001

  - ➡ hierarchy of functionality and capabilities
    - Tier-0 at CERN, 11 Tier-1s connected via 10 *GB/s* links
    - >100 Tier-2 centres attached by region to Tier-1s
    - data flows along the hierarchy, jobs send to data
    - different tasks assigned to centres according to hierarchy
  - ➡ very structured approach to ease some "fear" of networks and to limit complexity of operation (conservative in a sense)

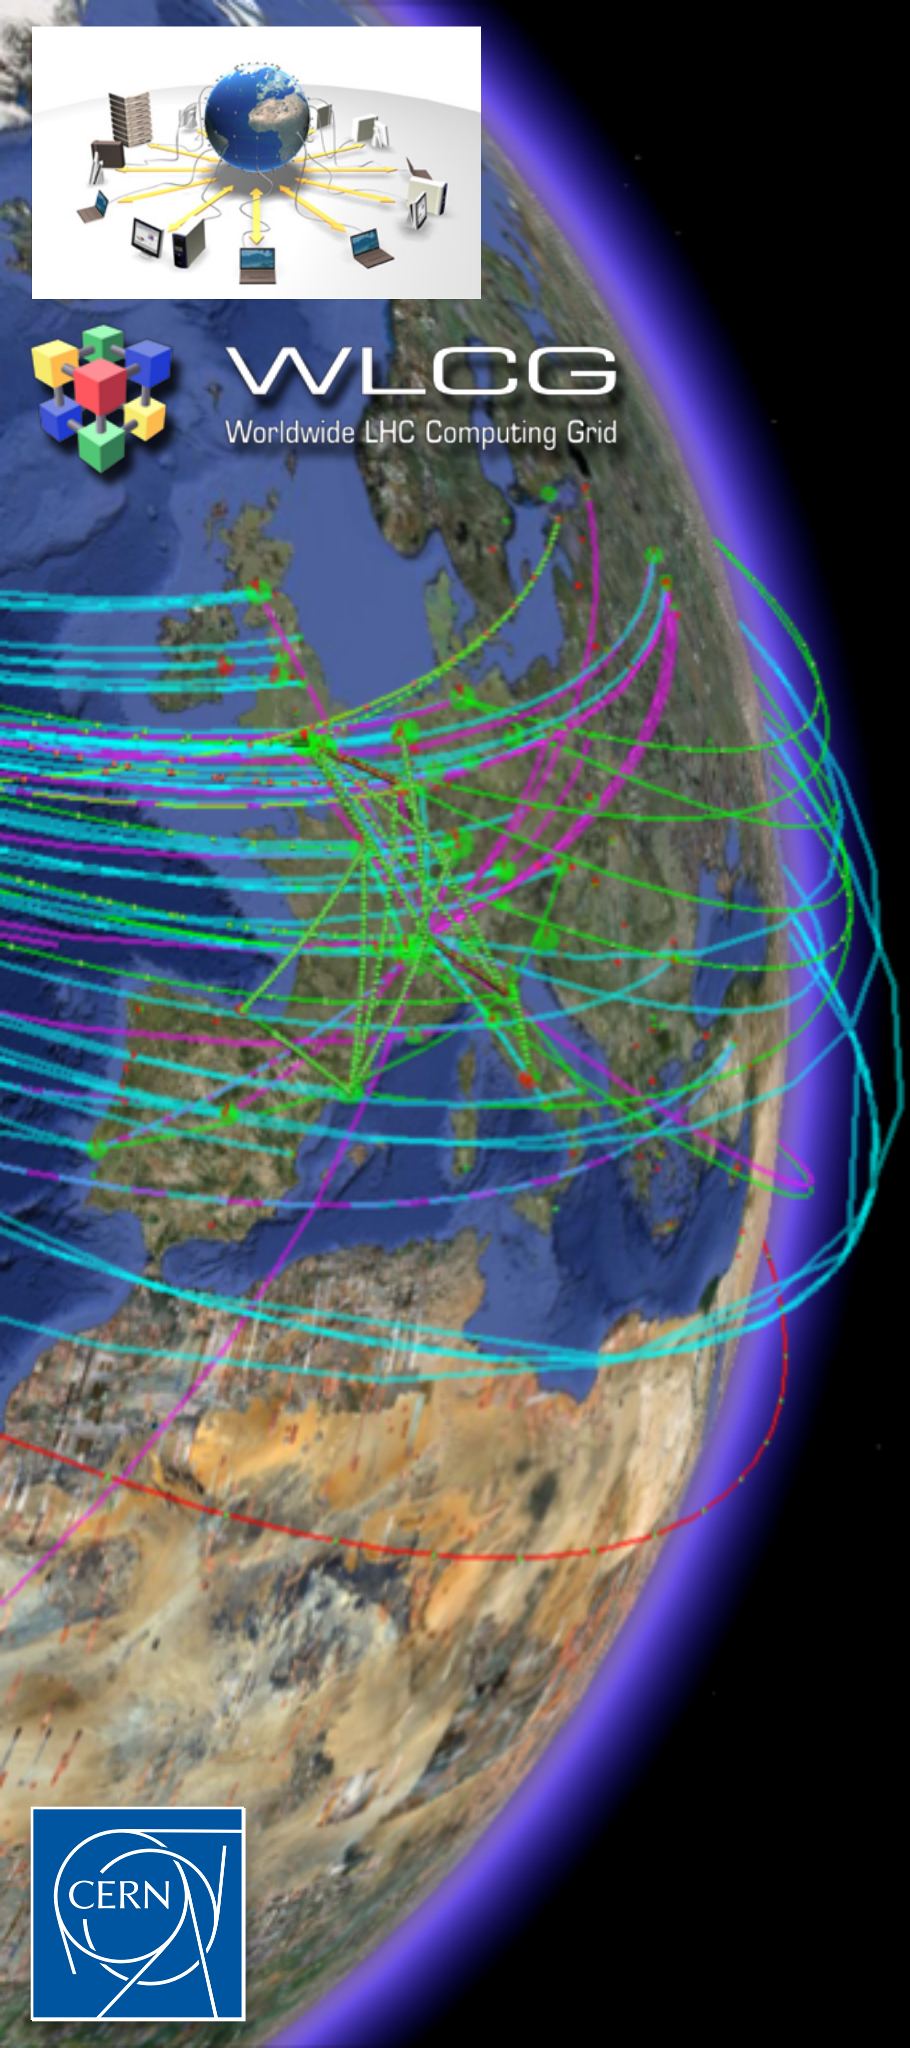

History of WLCG in Europe

(european centric view, ignoring OSG for the moment)

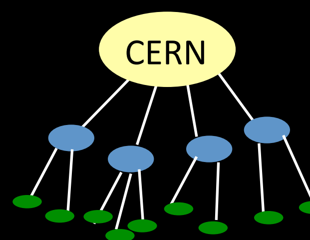

- ➡ 1999 - MONARC project
  - defined the initial hierarchical architecture
- ➡ 2000 - growing interest in Grid technology
  - HEP community main driver in launching the DataGrid project
- ➡ 2001-2004 - EU DataGrid project
  - middleware & testbed for an operational grid
- ➡ 2002-2005 - LHC Computing Grid
  - deploying the results of DataGrid for LHC experiments
- ➡ 2004-2006 - EU EGEE project phase-1
  - a shared production infrastructure building upon the LCG
- ➡ 2006-2008 - EU EGEE project phase-2
  - focus on scaling, stability and interoperability
- ➡ 2008-2010 - EU EGEE project phase-3
  - efficient operations with less central coordination
- ➡ 2010-201x - EGI and EMI
  - sustainability, shared across sciences

see talk of T.Kress

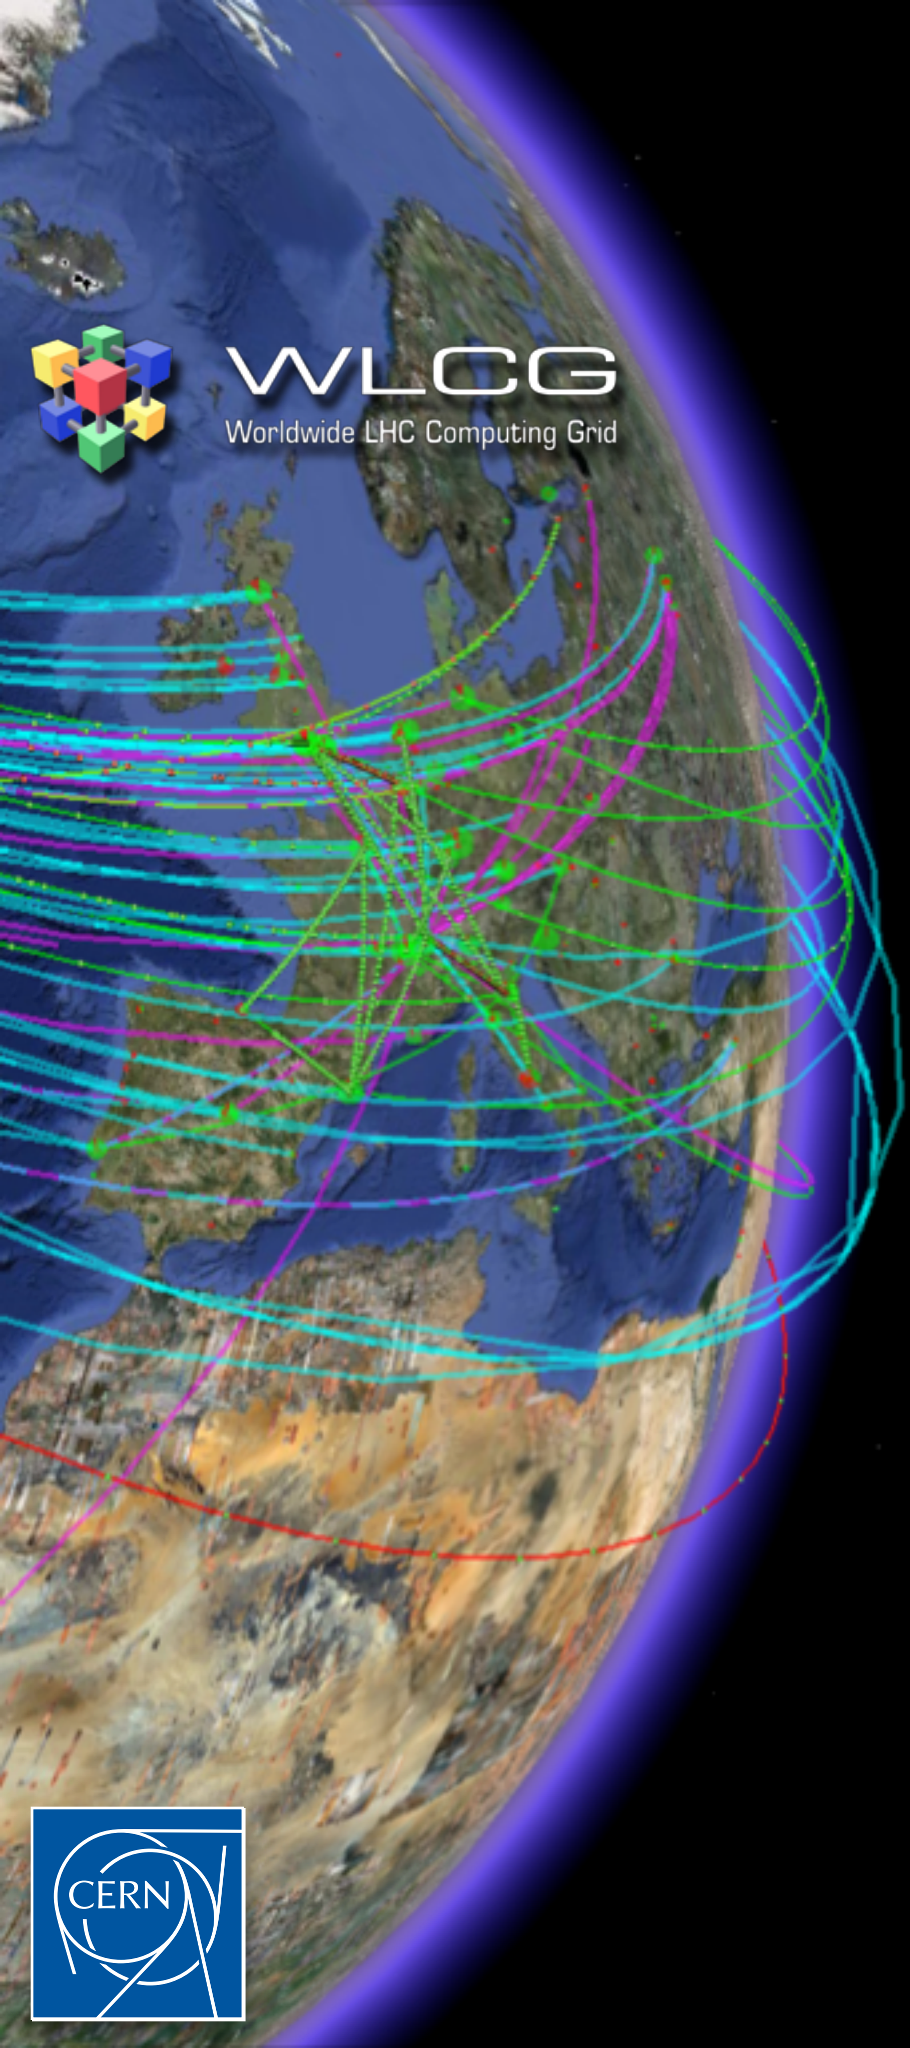

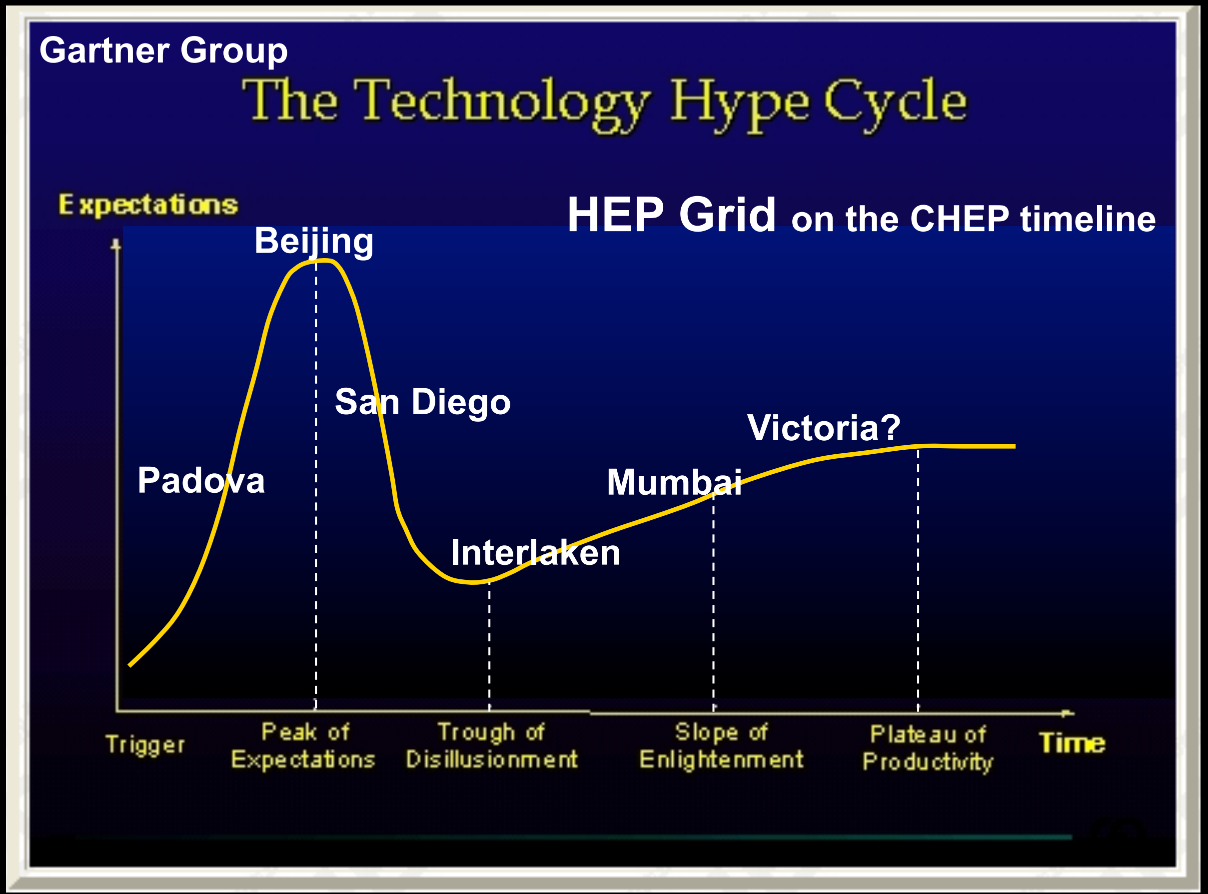

Les Robertson, CHEP Mumbai, 2006

Data and Service Challenges:
Commissioning
GRID based Computing

Role of the GRID Challenges

- experiments and WLCG followed strategy of a series of large scale tests
  - ➡ initially to transition to GRID based computing
  - ➡ later increasing scale and level of complexity
- learning process on all sides
  - ➡ from job success rates to operating site services
  - ➡ with time and operational experience the experiment specific GRID software layers grew:

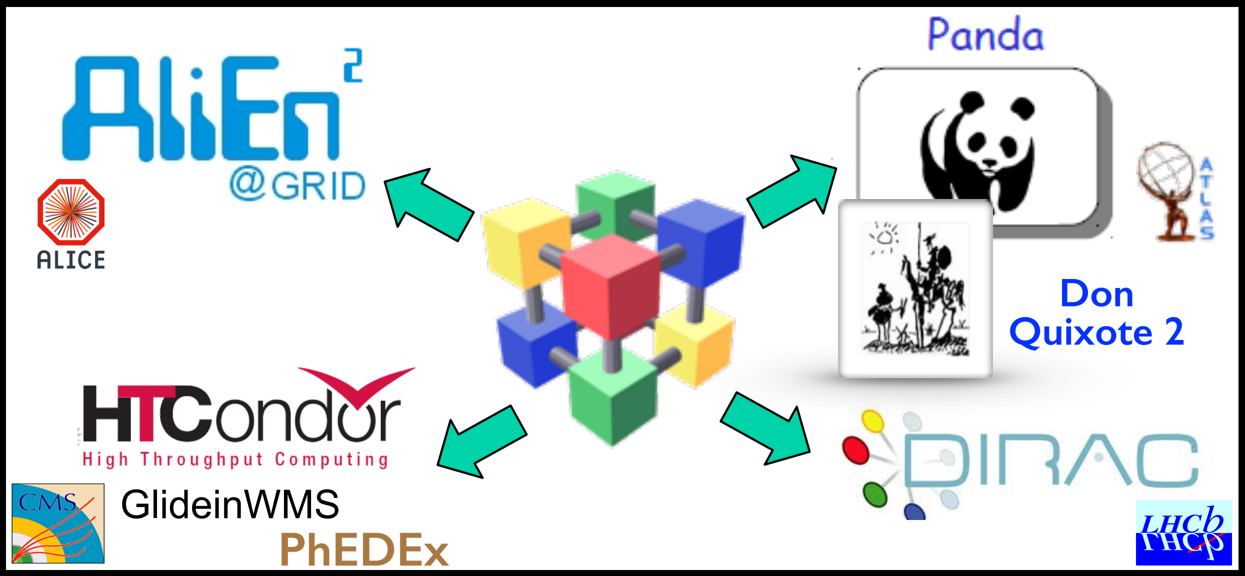

- pilot based production systems (DIRAC...)
- data transfer and data management systems
- etc.

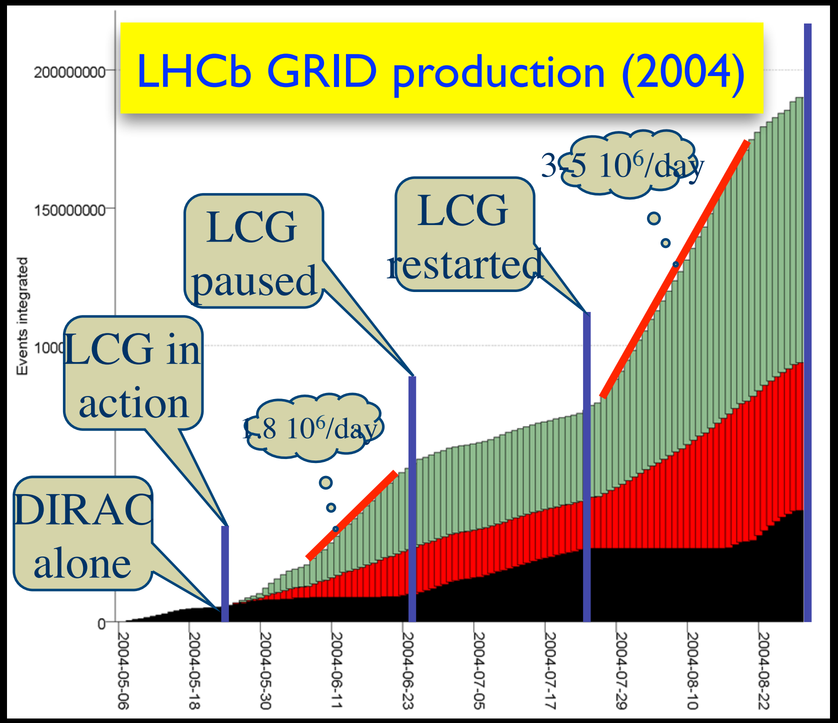

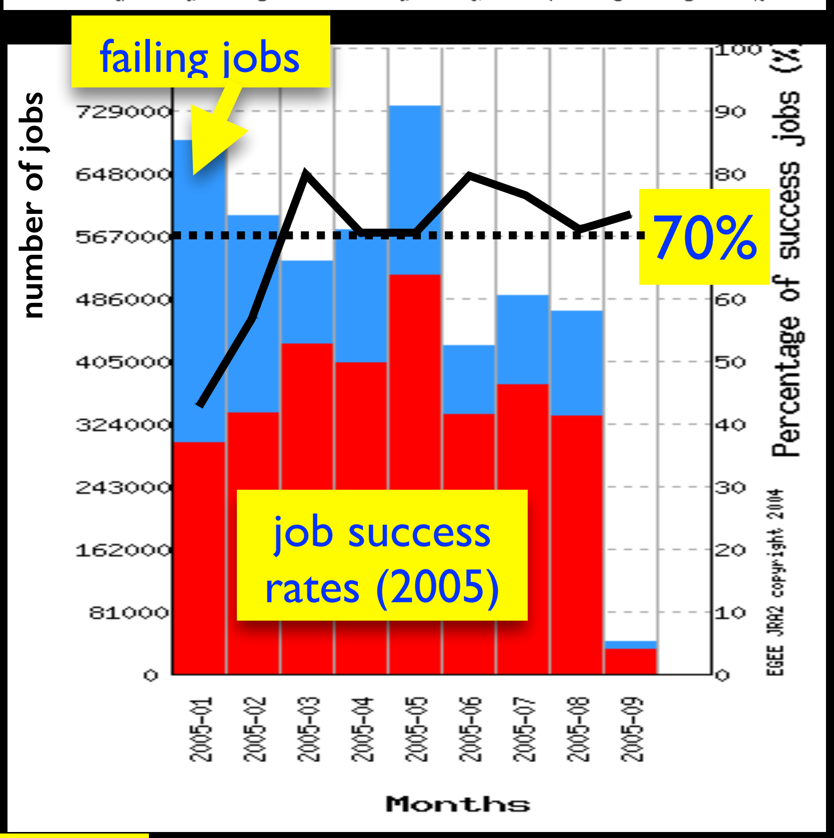

T.Doyle, AHM05 meeting

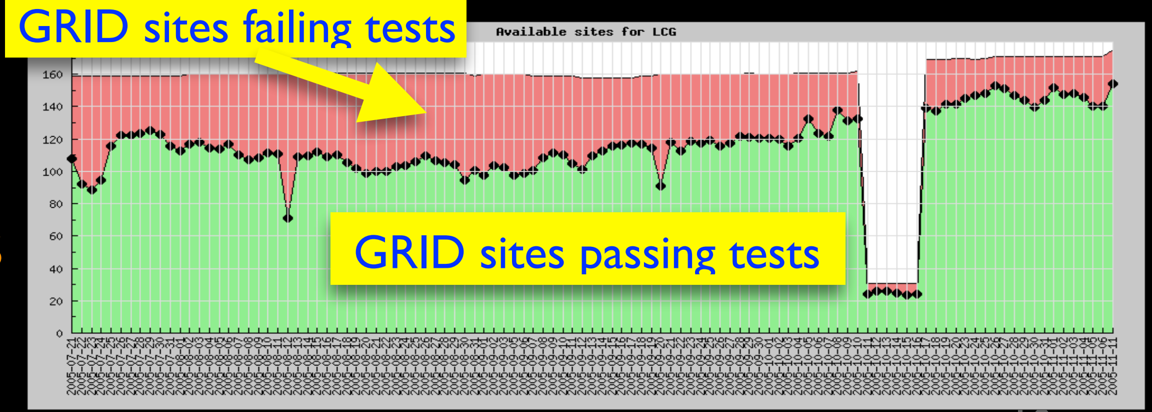

Les Robertson, CHEP Mumbai, 2006

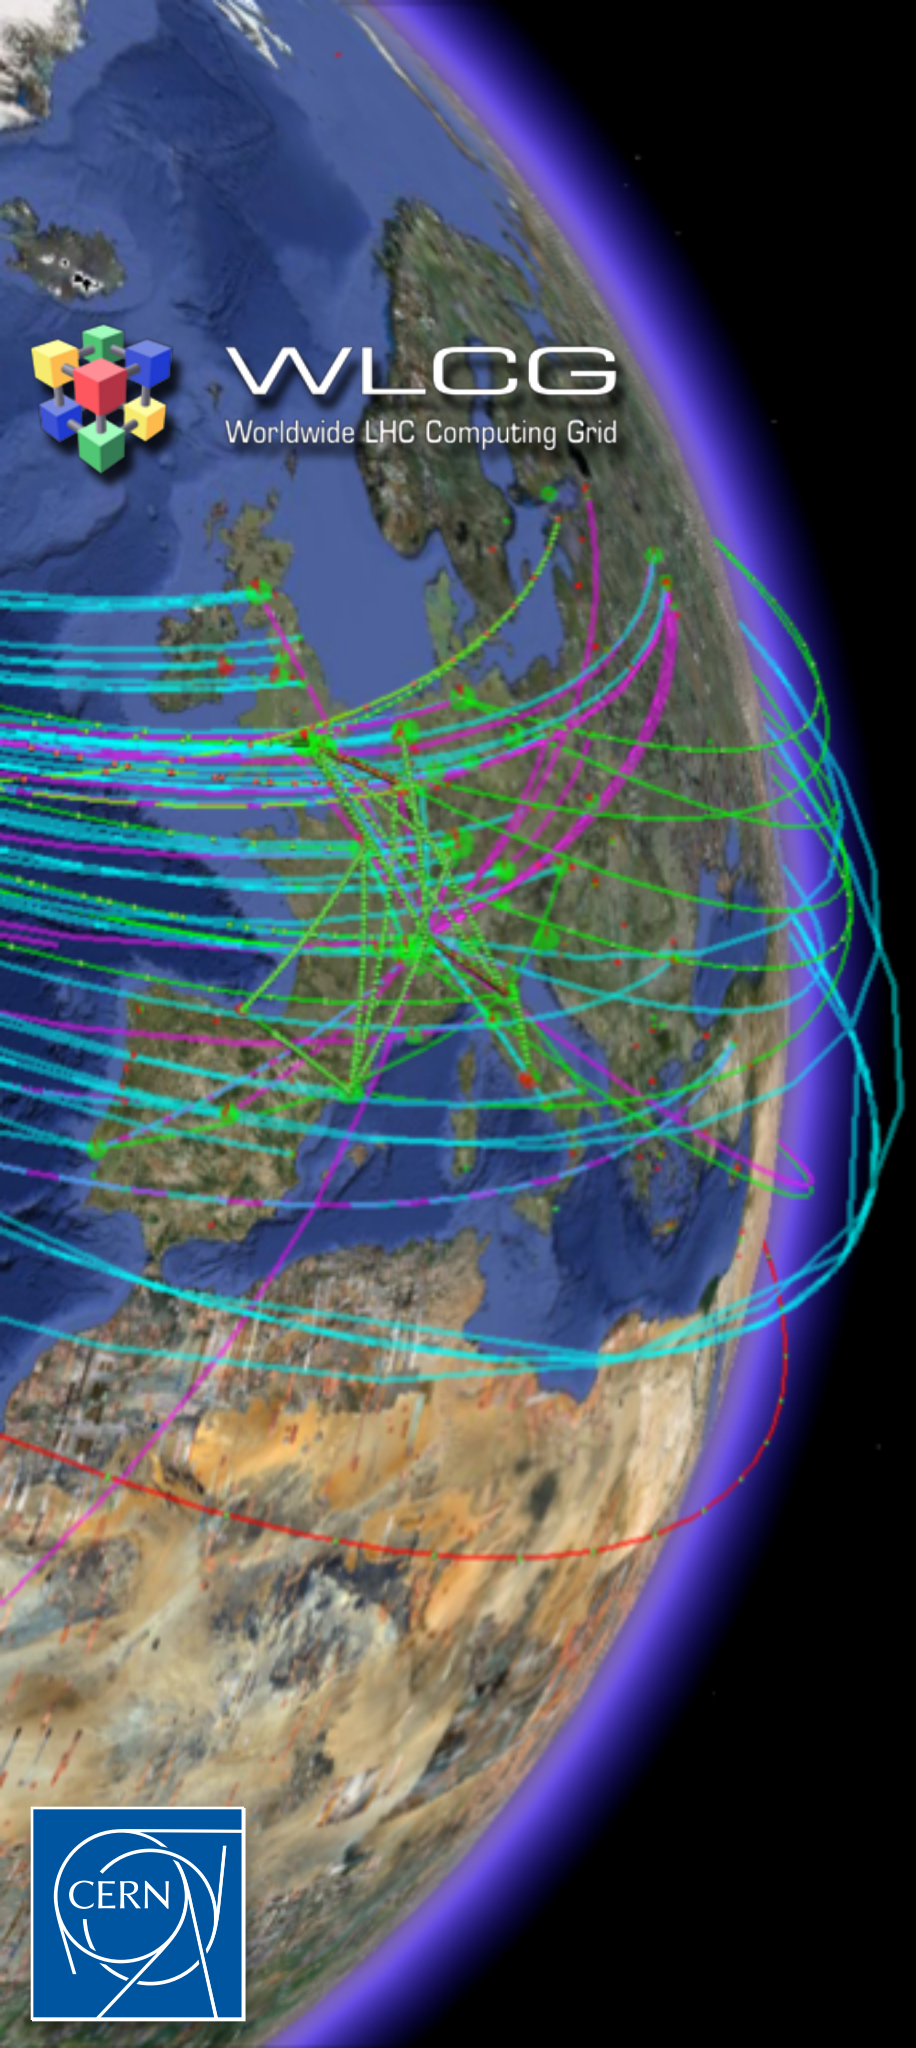

Building up the Software of the Experiments

ROOT (Rene and Rdm OO Technology*)

- project started 1995
  - ➡ by R.Brun and F.Rademacher (hence the name)
    - OO framework, having in mind the future LHC needs
    - as well, provided alternative to Objectivity/DB at the time
    - 1998 selected by Fermilab for Run-2 experiments
  - ➡ became "the standard" for HEP and LHC data analysis
    - used by Astrophysics, other sciences and fields
  - ➡ core team at CERN, effort at FNAL and large community input

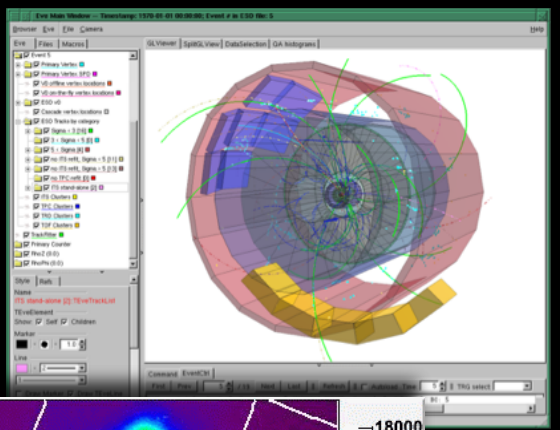

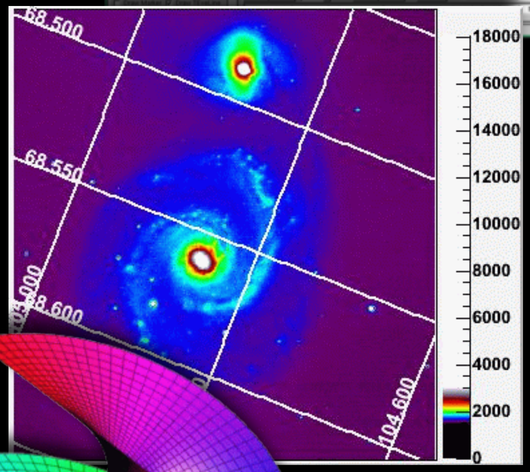

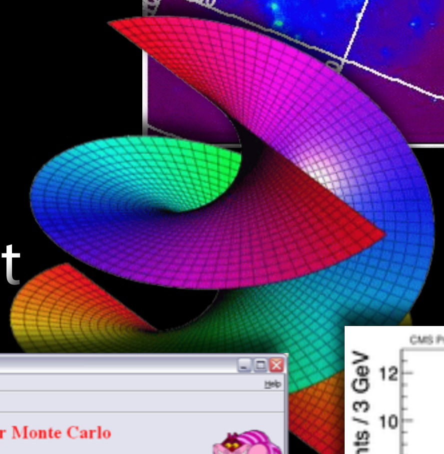

- framework for interactive analysis
  - ➡ visualisation, math libraries, I/O
    - LHC data is based on ROOT persistency
  - ➡ distribution includes suite of other tools
    - xrootd, TMVA, RooFit/RooStats, ...
  - ➡ total about 1.7 *million* lines of code
    - OpenHUB "estimated cost" is 27 M$

https://www.openhub.net/p/ROOT/estimated_cost

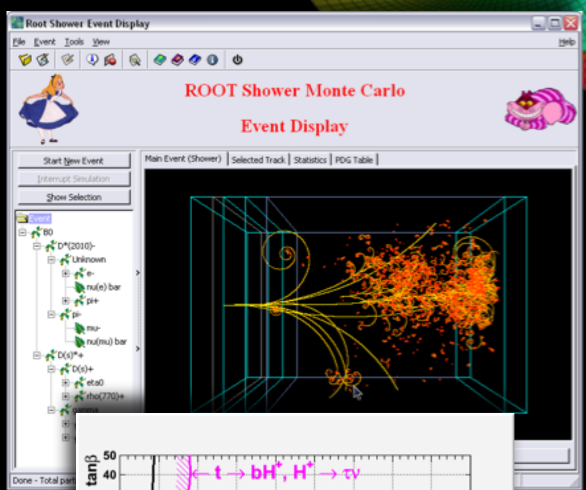

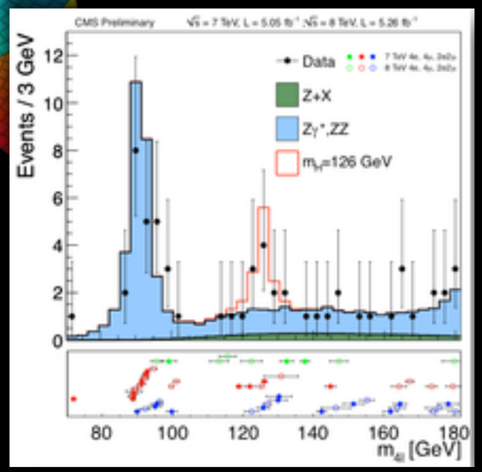

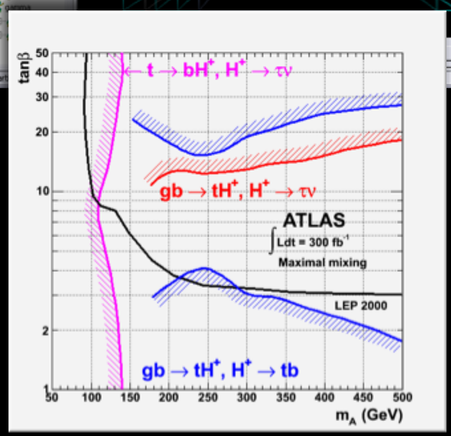

*http://ph-news.web.cern.ch/content/interview-rené-brun

Geant 4

- Geant4 Collaboration started in 1999
  - ➡ successor of Geant series toolkits developed at CERN
    - early studies at CERN and KEK resulted in RD44
    - OO simulation of passage of particles through matter
  - ➡ today effort at many large labs:
    CERN, FNAL, SLAC, KEK, ESA/ESTEC, ...
  - ➡ detector simulation for CMS, LHCb, ATLAS, (ALICE), ...
  - ➡ used by nuclear, accelerator and medical physics,
    as well as space science
  - ➡ about 2.1 *million* lines of code
    - OpenHUB "estimated cost" is 33 M$
      https://www.openhub.net/p/geant4/estimated_cost

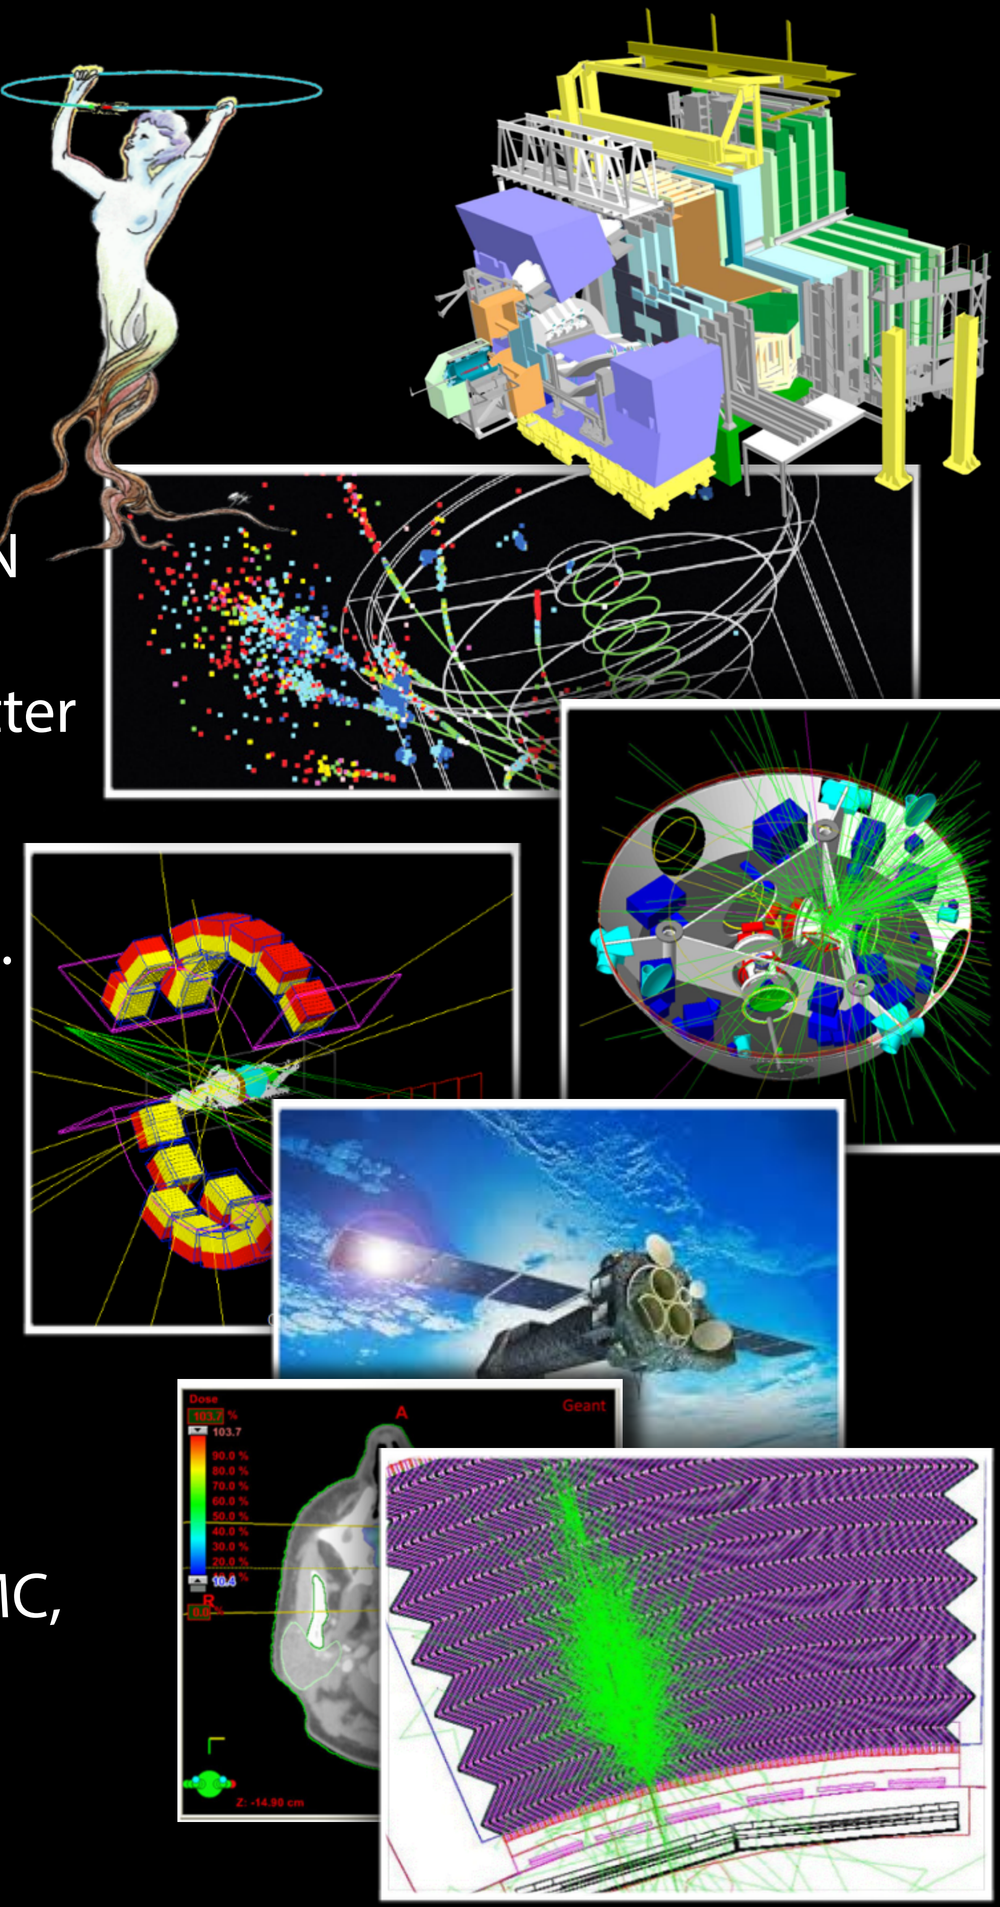

- equally important: event generators
  - ➡ Alpgen, Jimmy, Pythia6/8, Tauola(++), Sherpa, HepMC, Herwig(++), Photos, etc.
  - ➡ C++ and Fortran, about 1.4 *million* lines of code

Software of Experiments

- all developed their own OO frameworks
  - ➡ ORCA (CMS), AliRoot based on ROOT (Alice), GAUDI (LHCb)
  - ➡ ATLAS added its layer to GAUDI and called it ATHENA

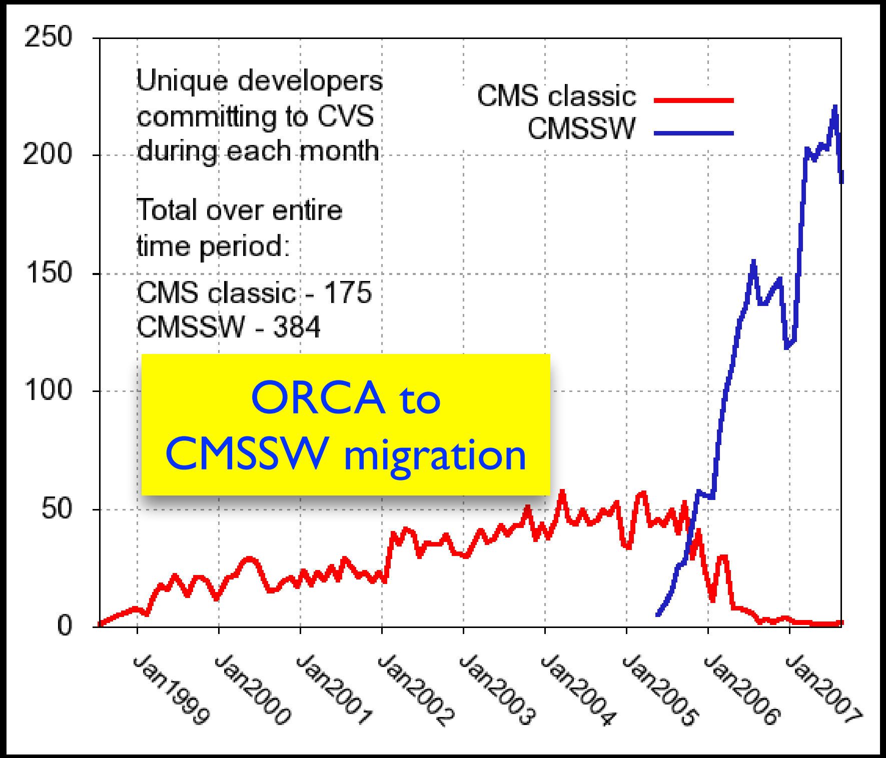

P.Elmar et al.

- CMS started 2005 CMSSW to replace ORCA
  - ➡ based on experience from FERMILAB experiments
    - huge effort, took >3 years
  - ➡ today a full CMSSW release has 7.5 *million* lines of code
    - OpenHUB "estimated cost" is 125 M$

      https://www.openhub.net/p/cms-sw-cmssw/estimated_cost
    - framework itself is only a fraction of this

- software stacks of the experiments
  - ➡ applications implemented in framework
    - detector simulation, trigger, reconstruction, ...
  - ➡ based on common software toolkits
    - development organised within LCG Application Area (Pool, Cool, Coral, Geant4, Root, ...)

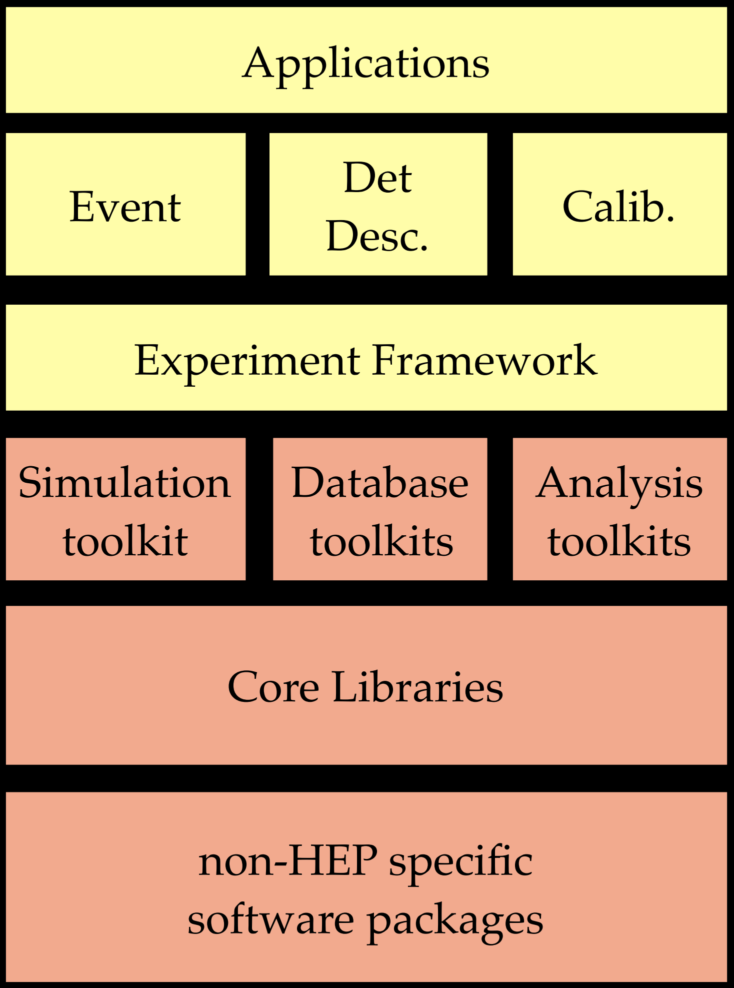

Building the Offline Reconstruction

- migration to C++ based reconstruction
  - ➡ existing FORTRAN algorithmic code often state of the art
    - new ideas from LEP experience, later BaBar and CDF/D0
  - ➡ lot of work (too much) went into OO design
    - "hip" at the time, today we have to back off again (see later)

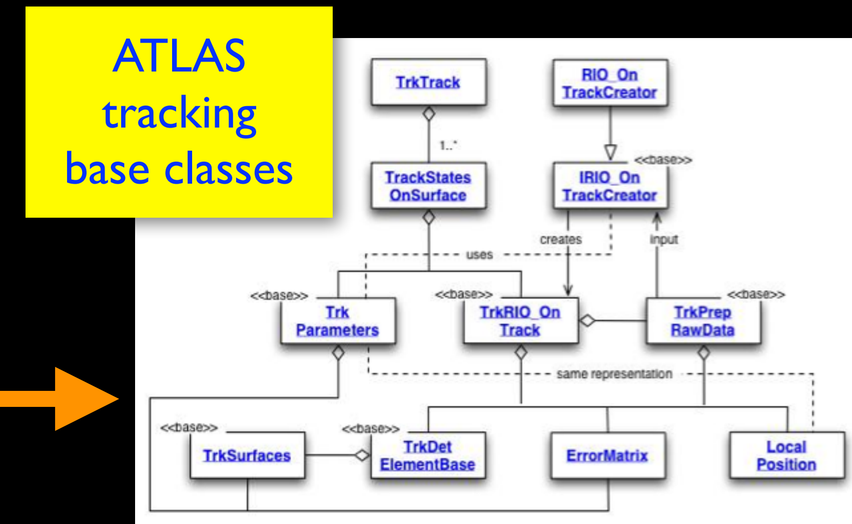

M.E. et al.

- new ideas to meet the LHC challenges
  - ➡ driver for innovation, lots of examples:
    - Deterministic Annealing Filters (Com.Phys.Com. 120 (1999) p.197) ~ tracking in ATLAS TRT at high pileup
    - STEP (J. Instr. 4 (2009) p.04001) ~ Runge-Kutta field integration for ATLAS+CMS muon tracking
    - JetFitter (J.Phys.Conf.Ser. 119 (2008) 032032) ~ novel secondary vertexing in jets for b-tagging
    - FastJet (hep-ph/0512210) ~ fast jet finding
    - Particle Flow (hep-ex/0810.3686) ~ reconstruction in CMS
  - ➡ later significant influx from CDF/D0, example:
    - Jet-Vertex-Fraction (hep-ex/0612040) ~ pileup suppression

A.Strandlie et al.

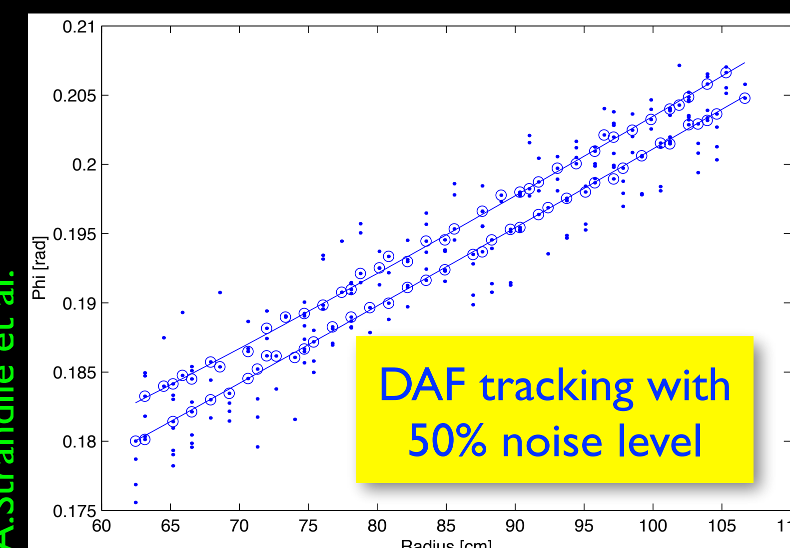

G.Pacquadio

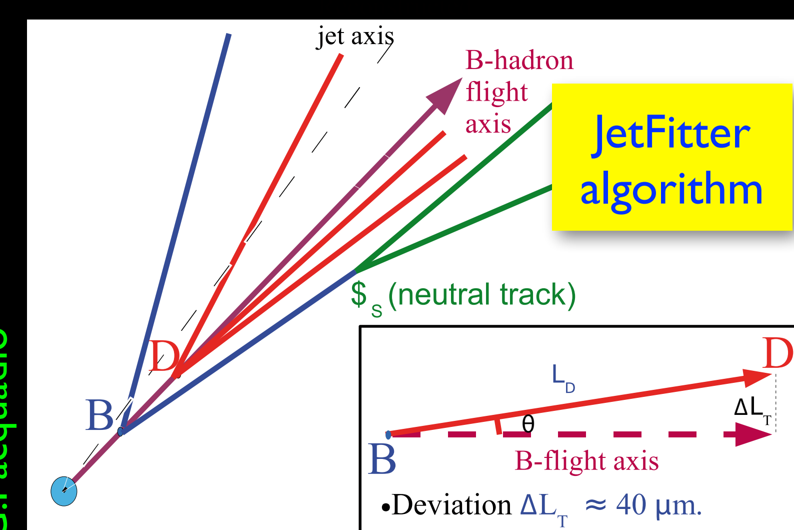

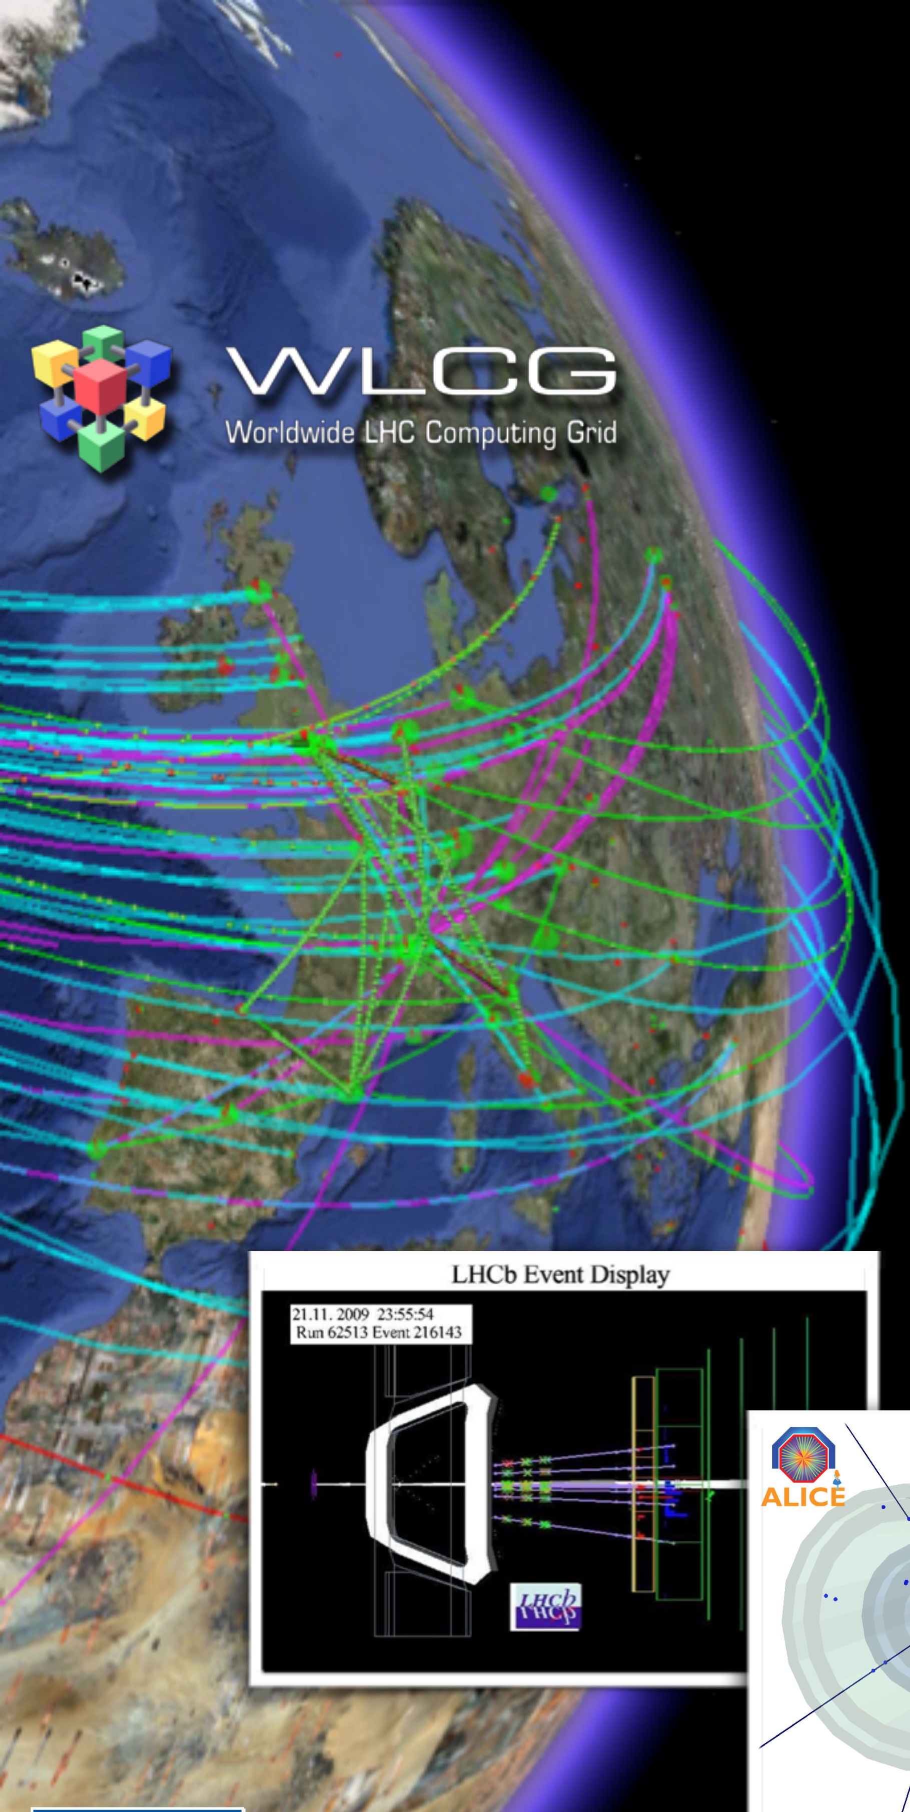

Early Physics and the Experience from Run-1

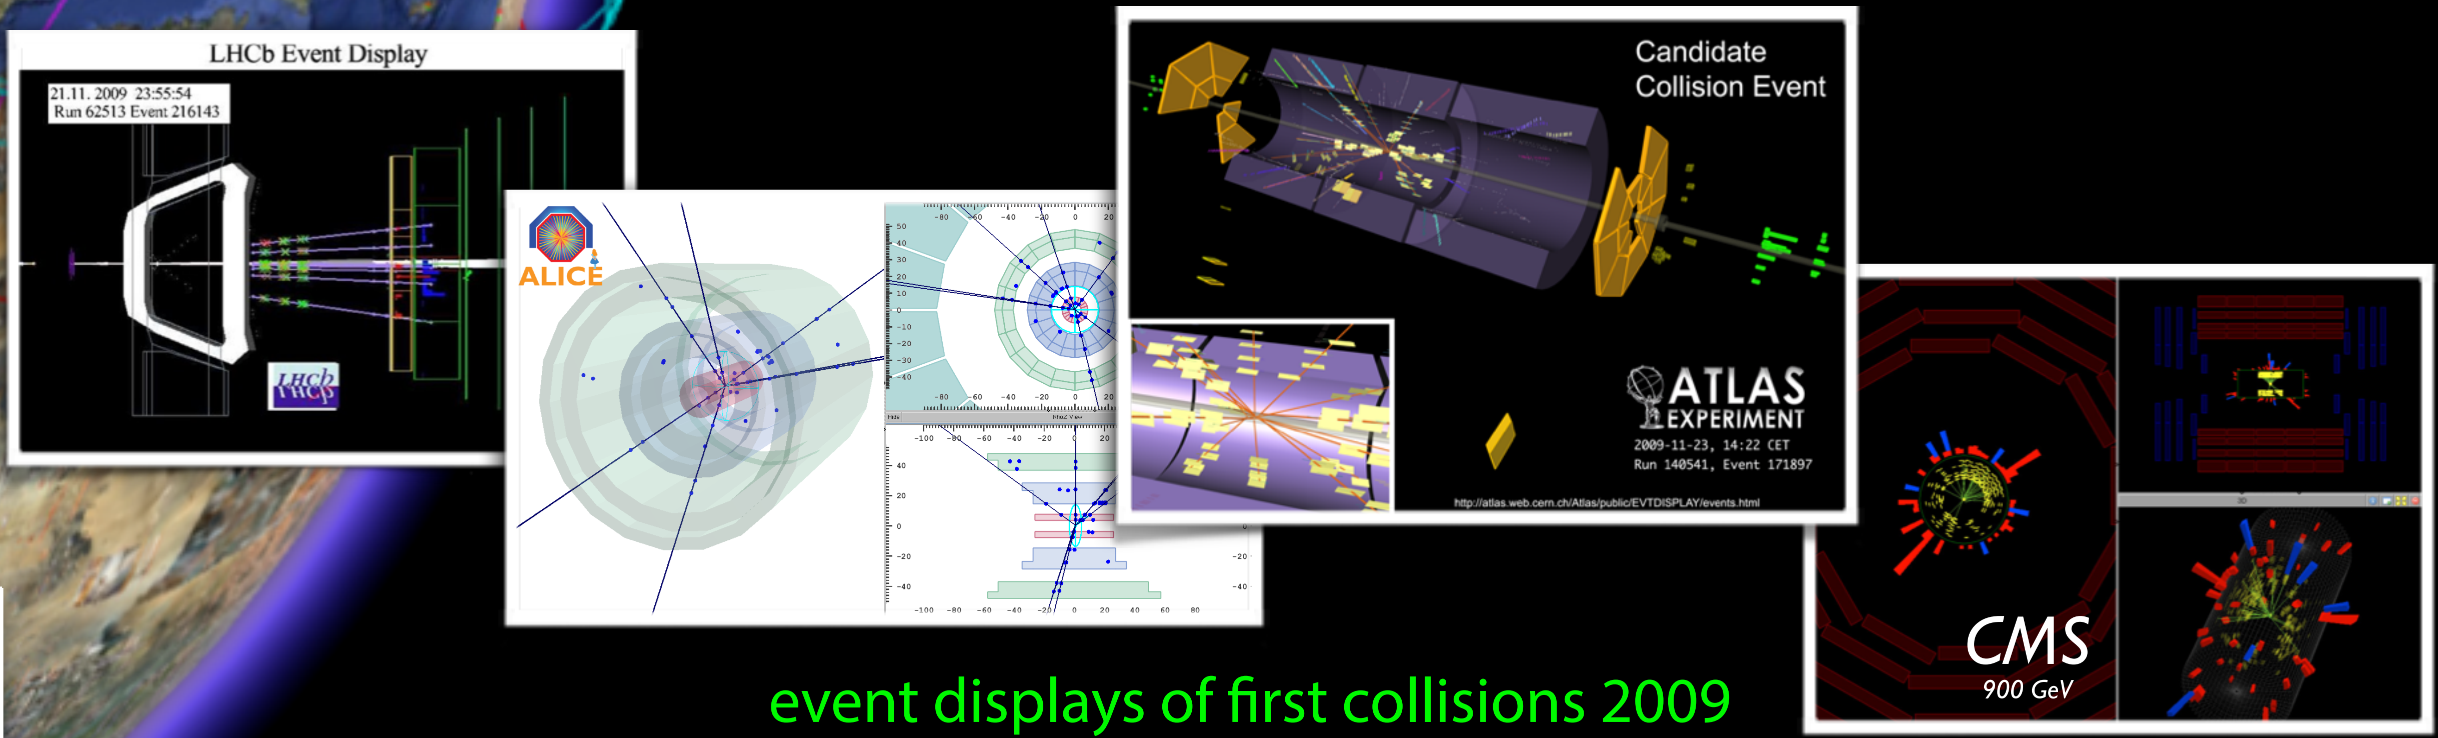

event displays of first collisions 2009

First Data to Physics Results

- a success story all along...
  - ➡ detector, DAQ and trigger worked !
  - ➡ excellent quality of first data
    - fast convergence of calibration and alignment procedures
    - much smoother than many expected
  - ➡ striking level of modelling by simulation
    - thanks to careful preparation work, e.g. excellent model of tracker material
    - helped a lot the fast production of physics results
- with luminosity increasing over the year 2010
  - ➡ quality of data approaching design levels with series of reprocessings
  - ➡ "re-discovered" the standard model particles one-by-one

SCT hits in data and MC in first runs

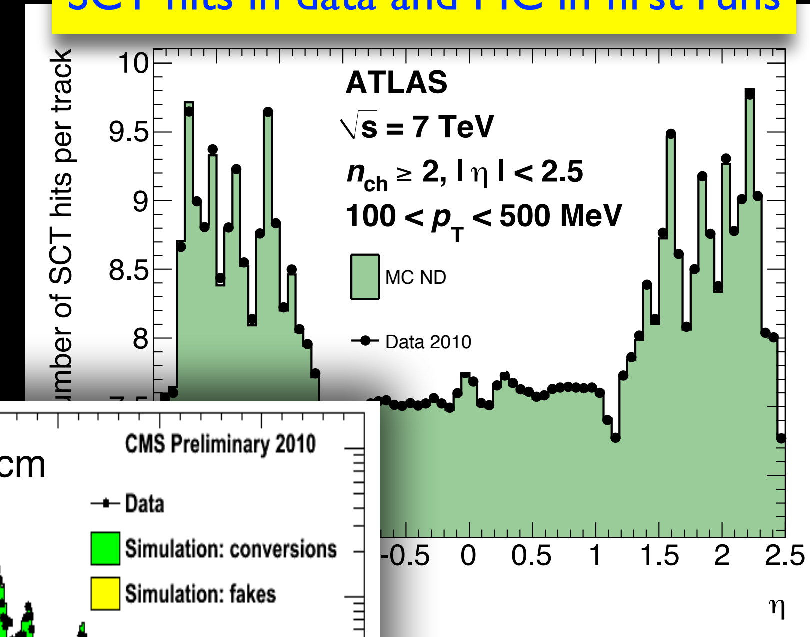

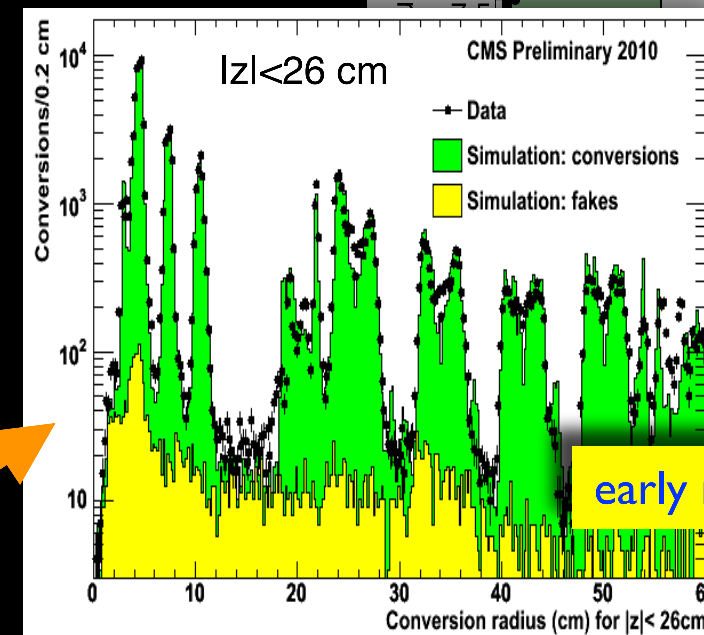

early material studies

F.Gianotti, ICHEP 2010

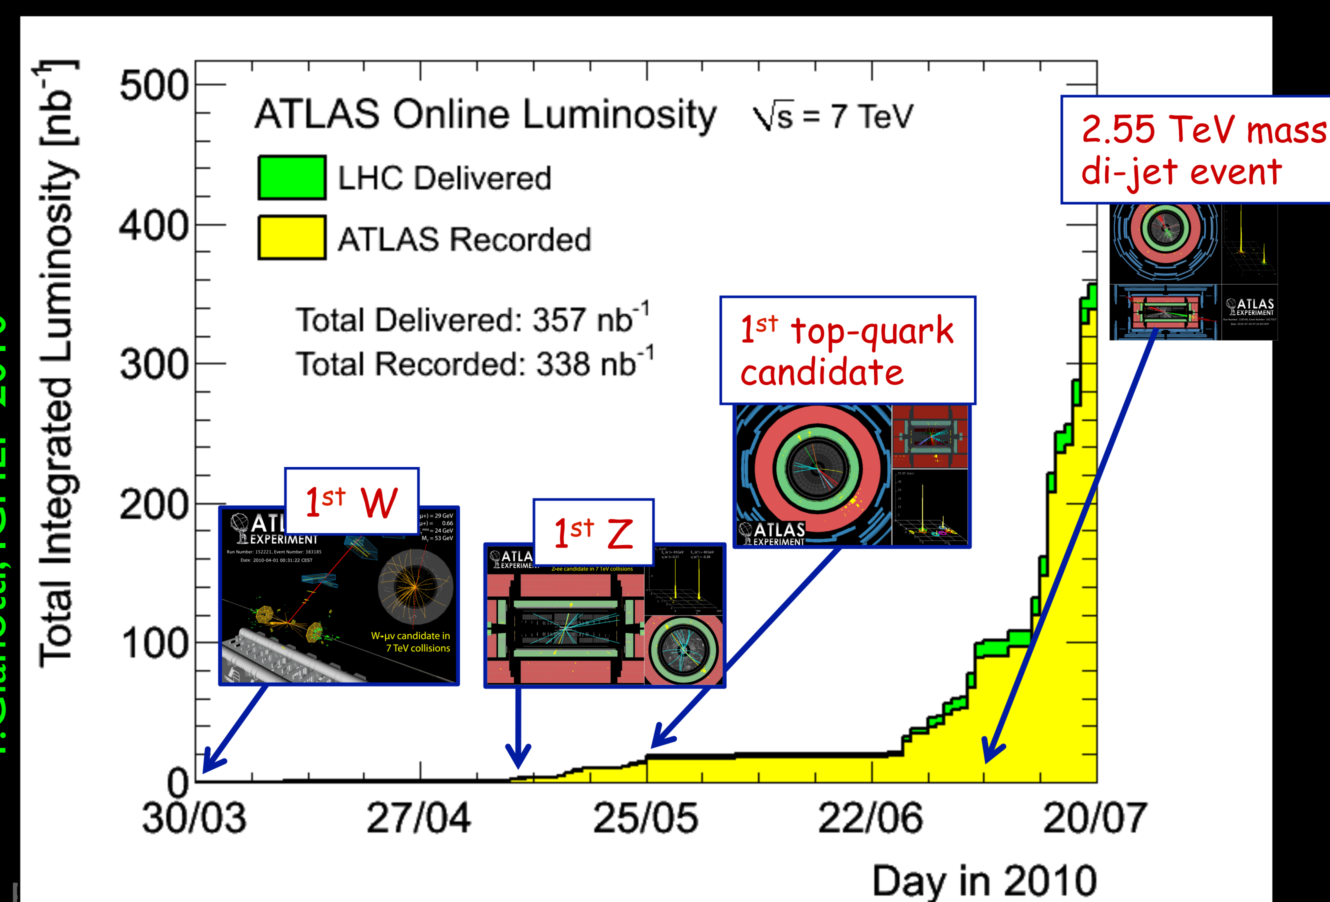

2.55 TeV mass di-jet event

1st top-quark candidate

1st W

1st Z

What about GRID Computing

- it worked !
  - ➡ even beyond expectations
    - Tier-0 processing and GRID distribution
    - MC production and reprocessing
    - distributed analysis
  - ➡ good data available for analysis in timely fashion
    (we talked much less about computing than many expected)

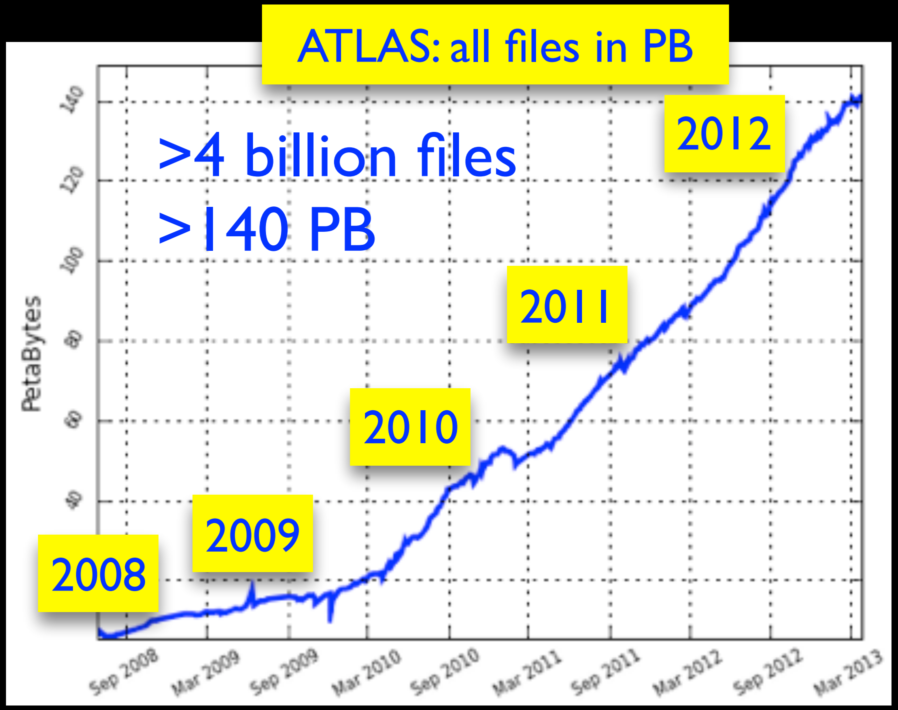

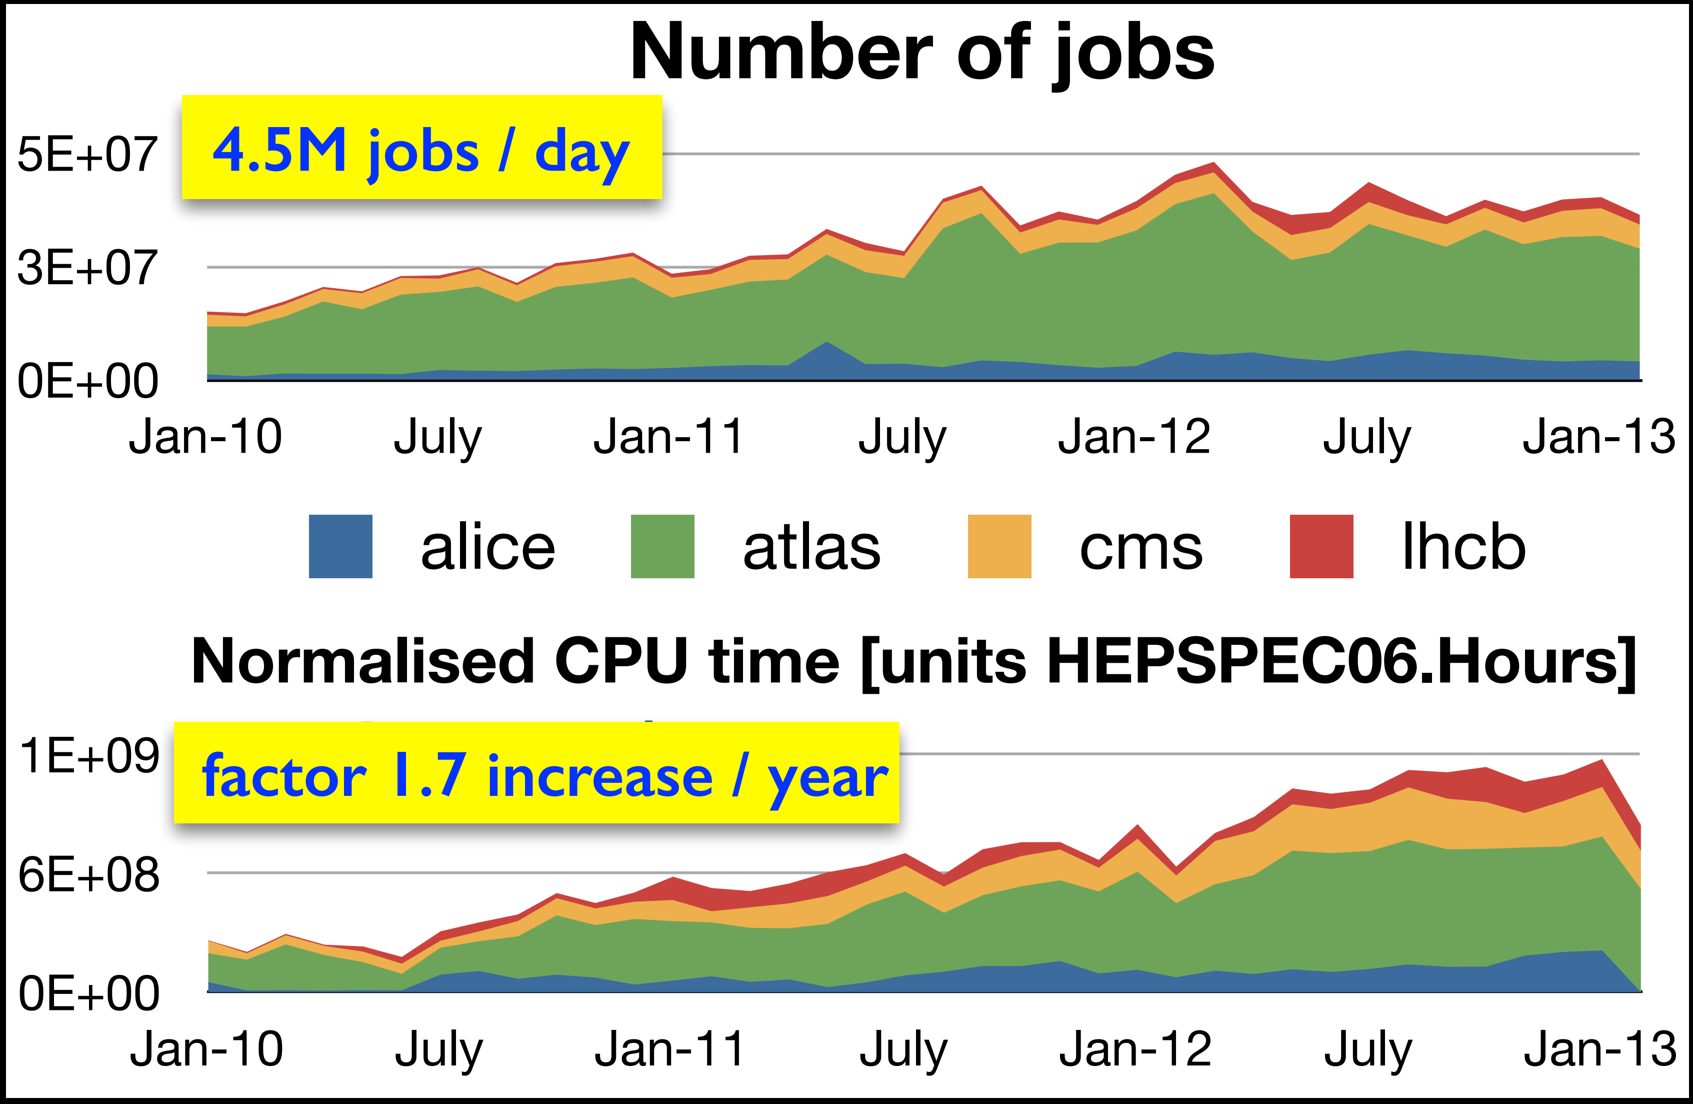

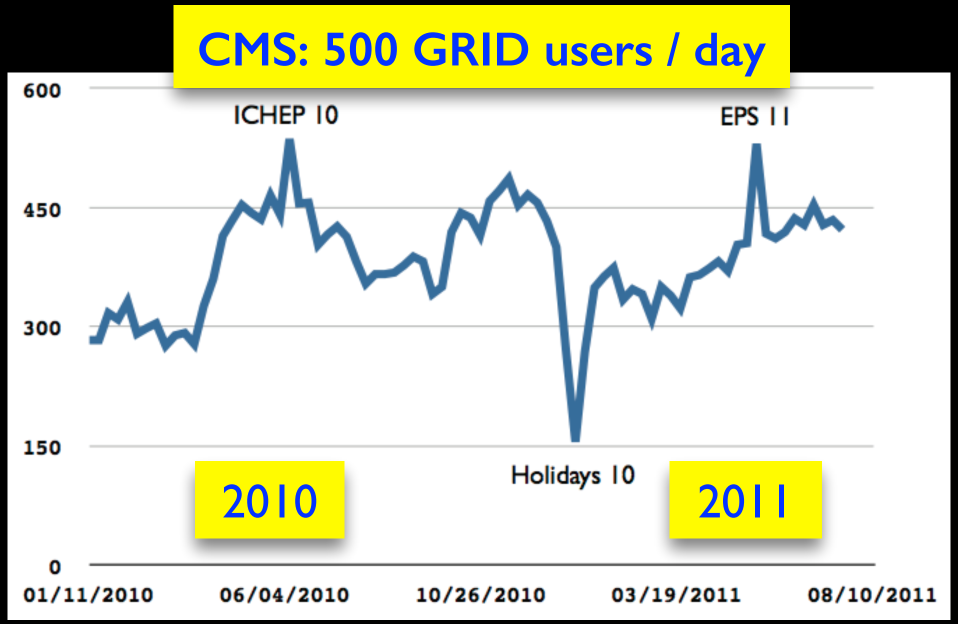

E.Lancon, source EGI portal, 2013

Changes in Computing during Run-1

- with time we made our models more and more flexible
  - ➡ driven by operational experience gained and technology advancements
  - ➡ loosen operational constraints
    - direct transfers between T2s (LHCONE - Tier-2s connect with 10GB)
    - data transfers to jobs (optional)
  - ➡ caching instead of centralisation
    - conditions access from any site (Squid/FronTier, CVMFS)
    - automatic release distribution
  - ➡ popularity based data placement and deletion (e.g. DP2P)
    - less replicas, better disk usage

E.Lancon, 2014

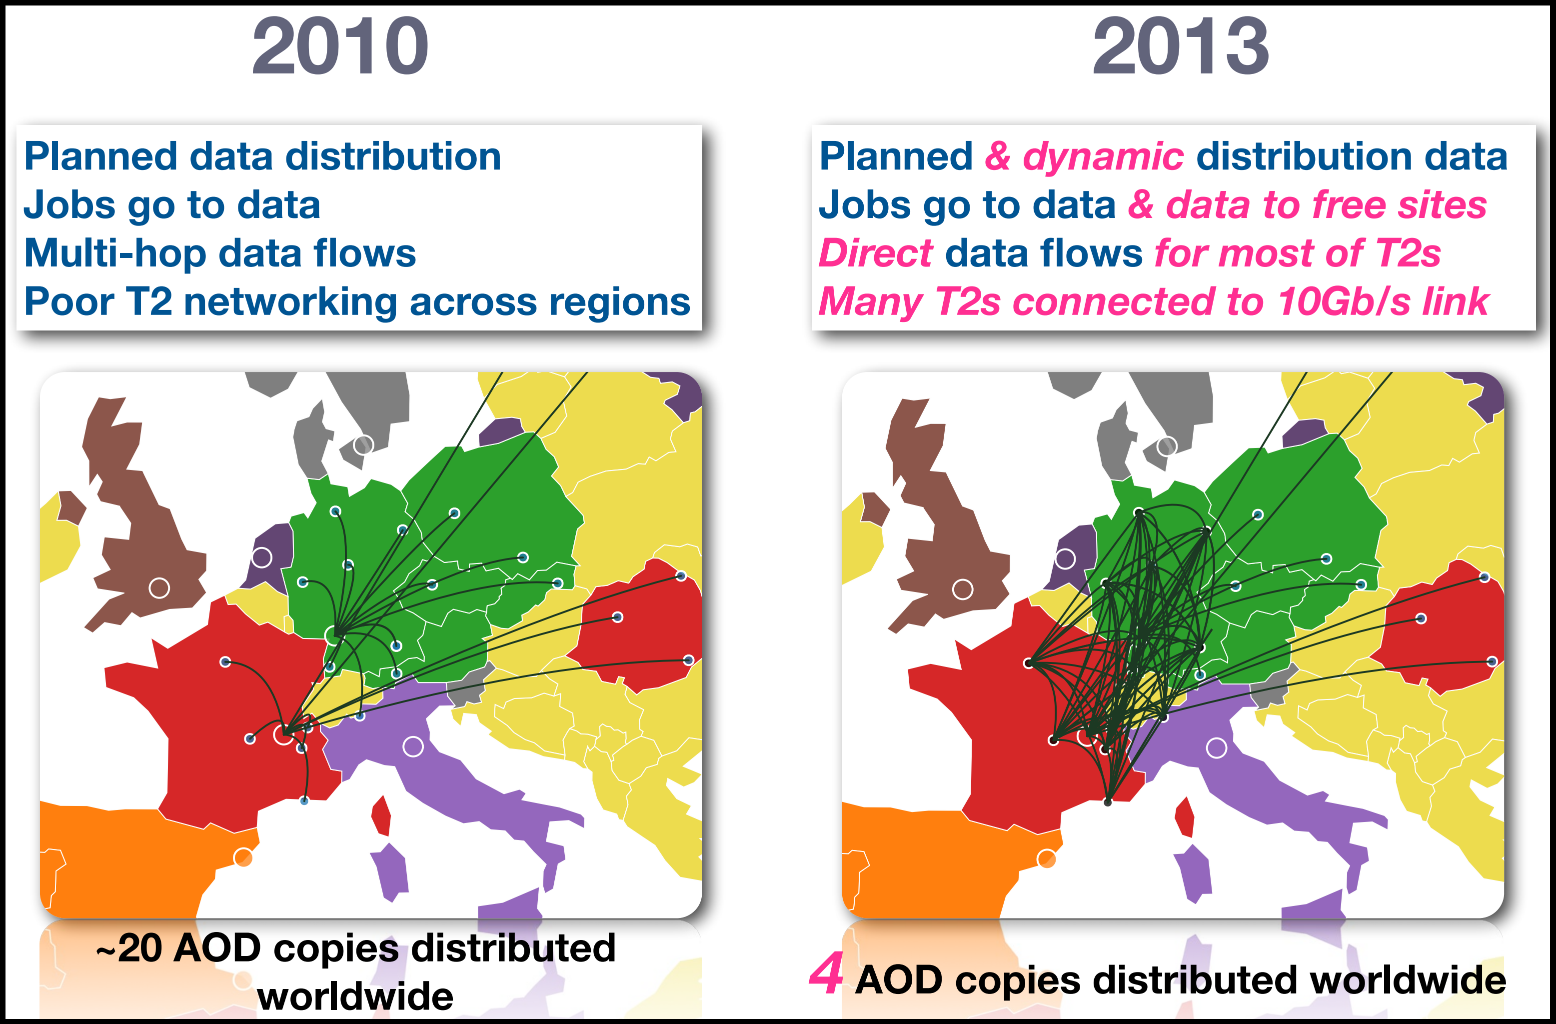

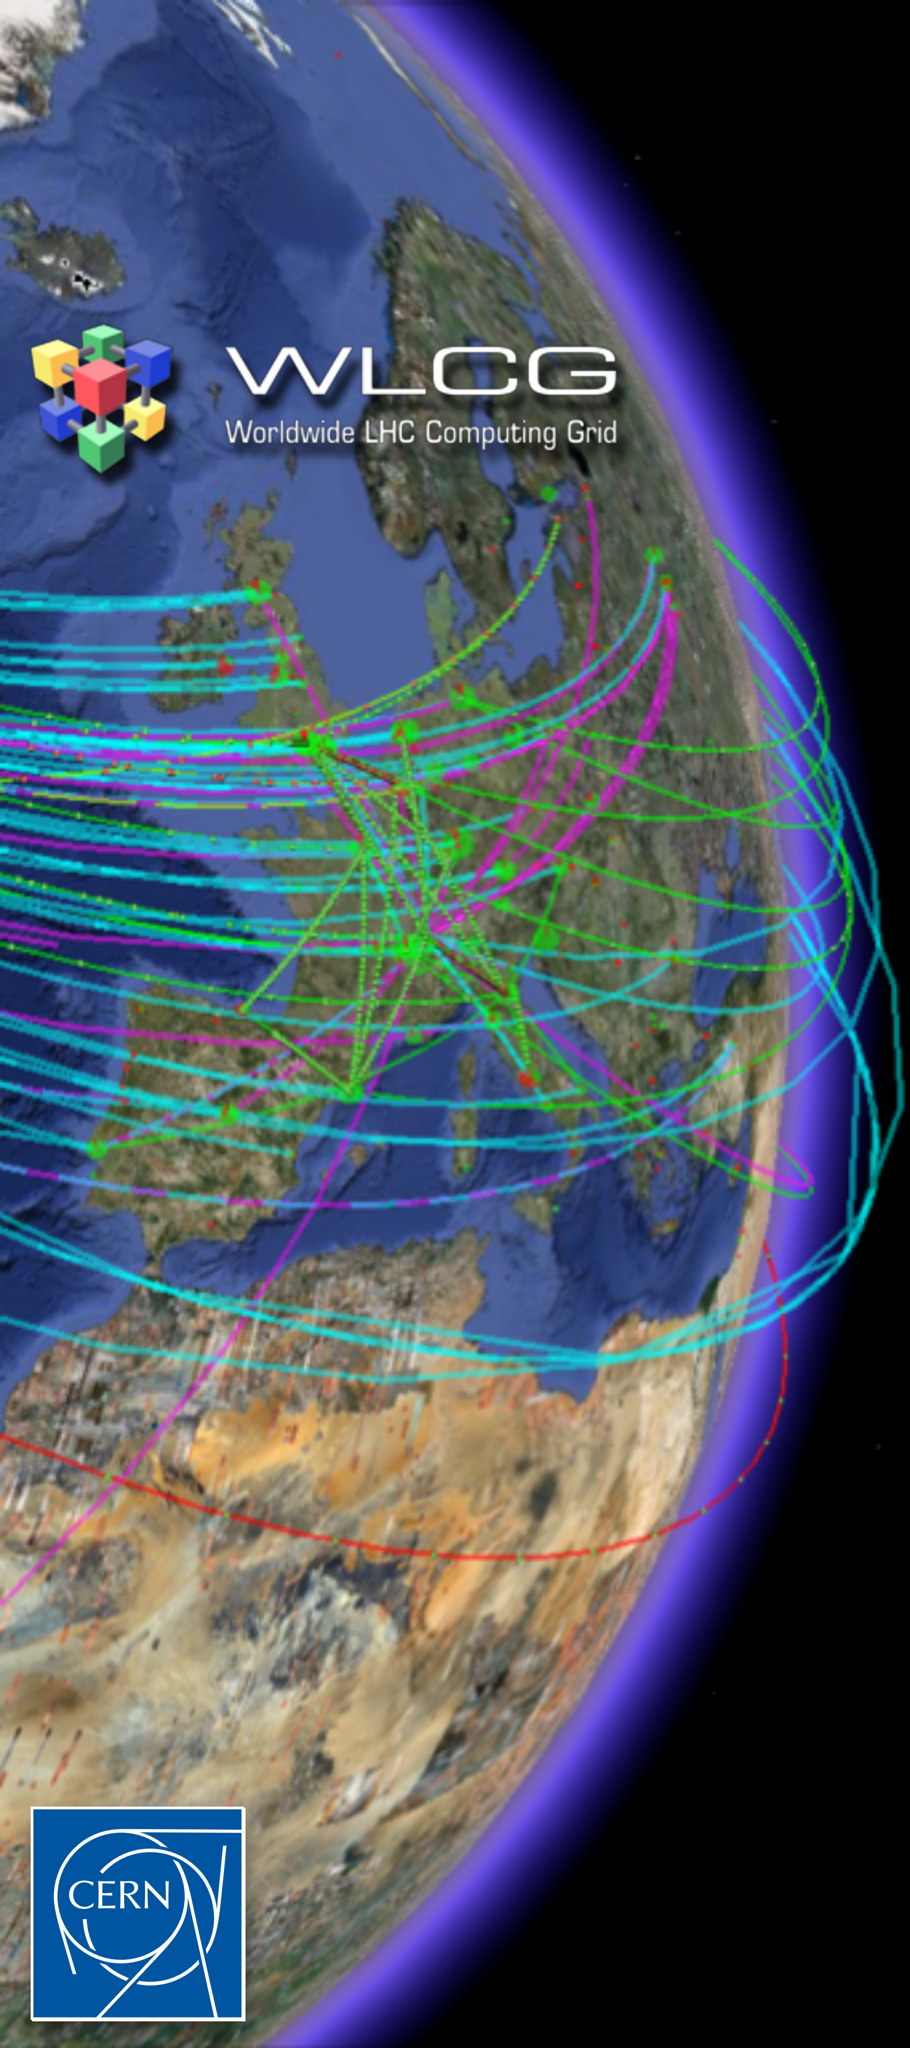

The Higgs Discovery: the Role of Software and Computing

Situation in 2011

- Higgs searches in 2011 data
  - ➡ both experiments saw "hints" for a light Higgs
    - about ~3σ each, ignoring "look elsewhere effect"
    - indications as well in TEVATRON data
  - ➡ low mass region at LHC
    - many decay modes accessible (γγ,ZZ,WW,ττ,bb)
    - γγ and ZZ yield excellent mass resolution (~1%)
  - ➡ detector performance crucial to all analyses (!)
- rapid increase in luminosity
  - ➡ pileup approaching design levels in 2011
    - mainly because of 50 *nsec* operation
    - expectation was to exceed design level in 2012
  - ➡ concerns about pileup robustness and performance of object reconstruction
    - experiments did intensive software development in preparation for 2012 data taking

M.E., for illustration

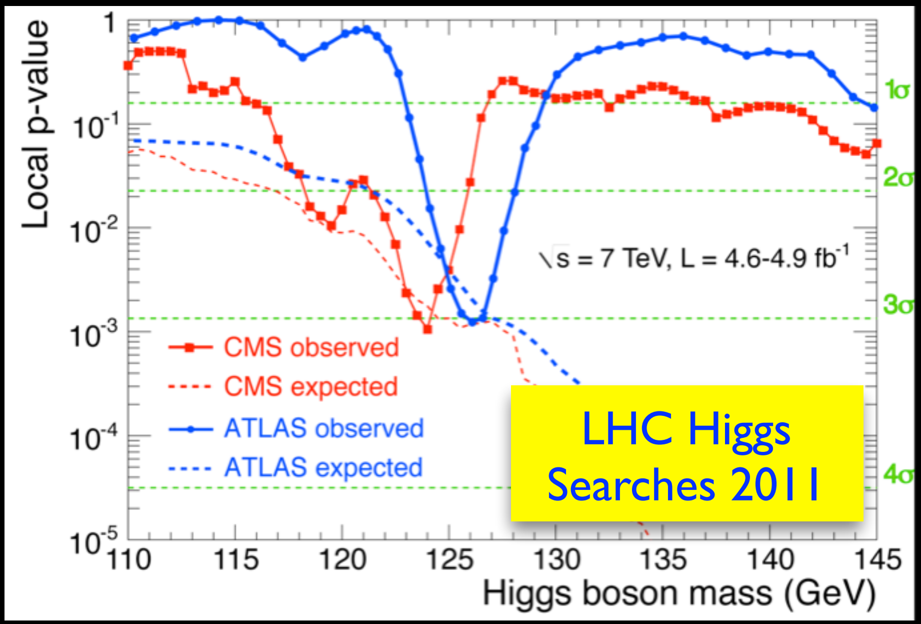

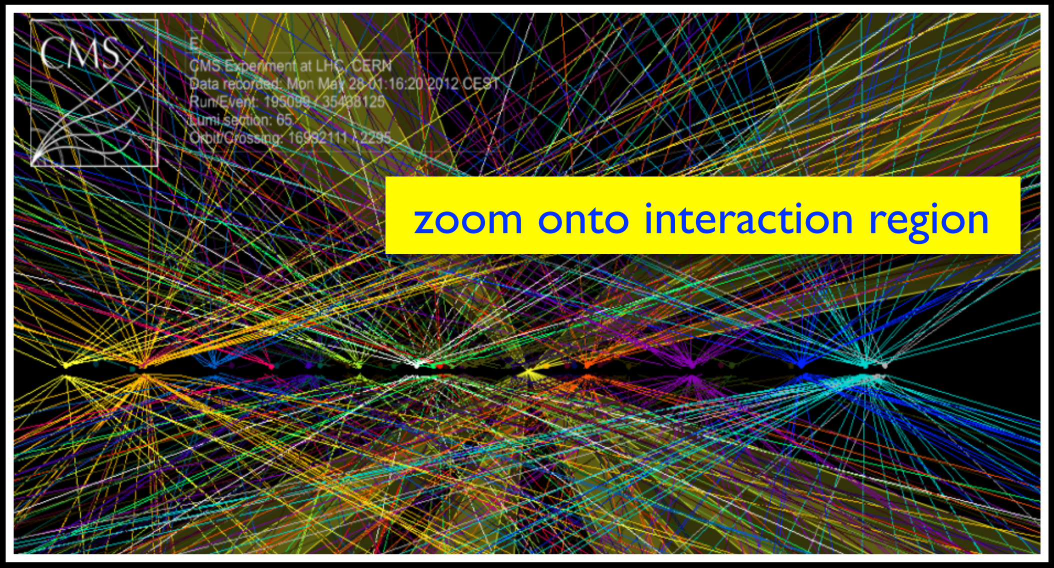

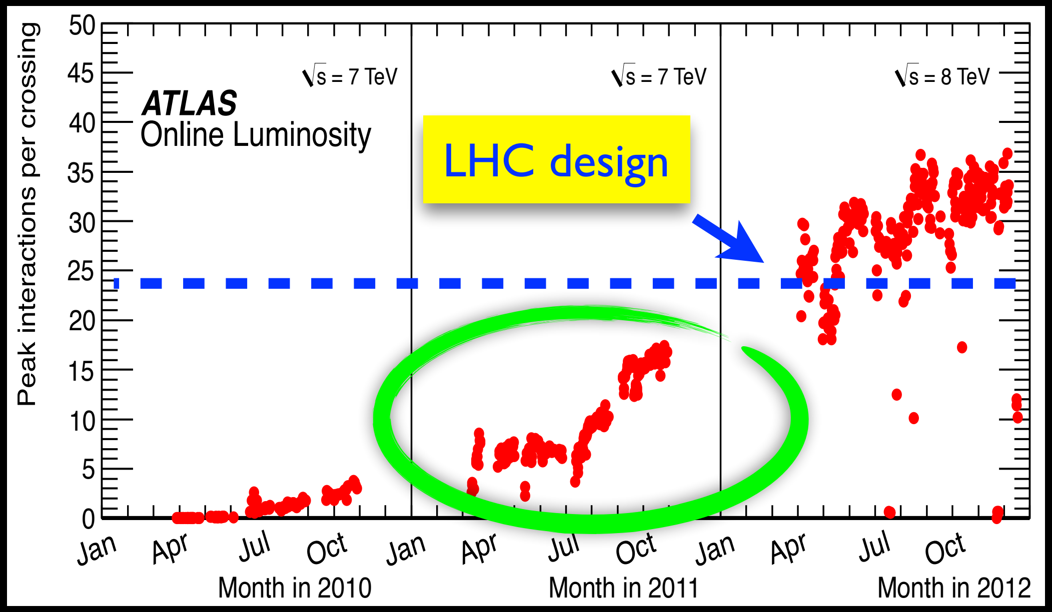

Updates to Tracking

- CPU scales non-linear with pileup
  - ➡ combinatorial explosion
    - CMS ~50% in tracking
      (e/γ dominated by special tracking too)
    - ATLAS ~70% in tracking
  - ➡ e.g. CMS gained factor 2-3 in CPU
    - optimisation of pattern for 30 pileup
    - as well technical optimisation (memory)
- pileup robustness and performance
  - ➡ improve track selections to control fakes and better vertexing cuts
  - ➡ robust tracking cuts for object reconstruction
    - e.g., tracking for conversions in ATLAS optimised to improve pileup stability ($H \rightarrow \gamma\gamma$)

CMS CR-2011-002

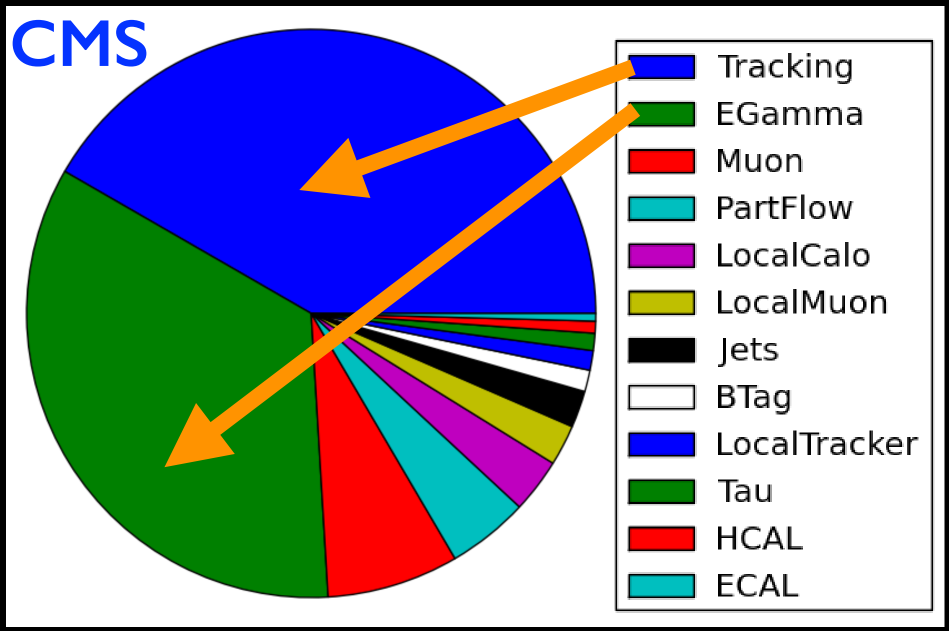

J. Phys.: Conf. Ser. 396 022044

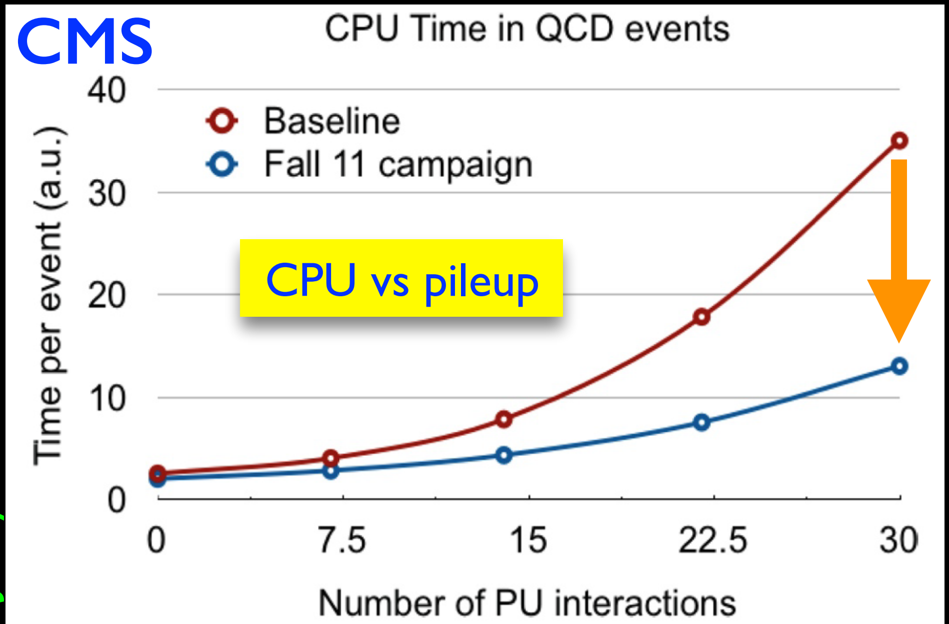

ATLAS-CONF-2012-823

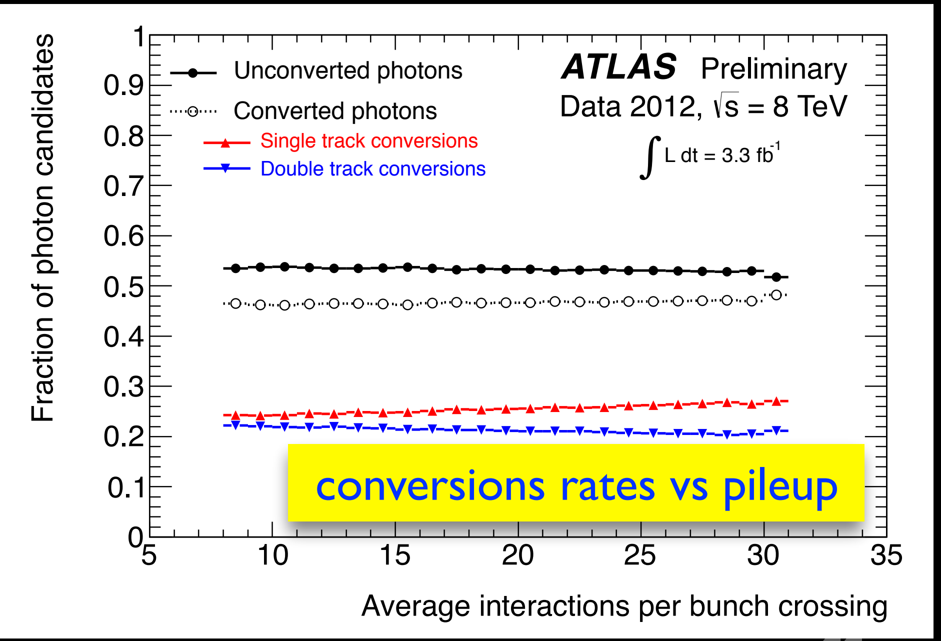

Object Reconstruction Updates

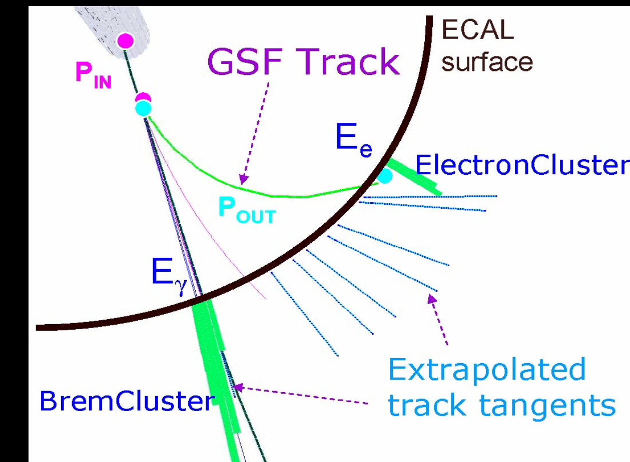

- sophisticated electron brem. recovery
  - ➡ using so called Gaussian Sum Filters
  - ➡ CMS ran dedicated tracking for e/γ
  - ➡ ATLAS introduced Region-of-Interest based tracking
    - brem. recovery for tracks pointing to EM clusters

ATLAS-CONF-2013-1287

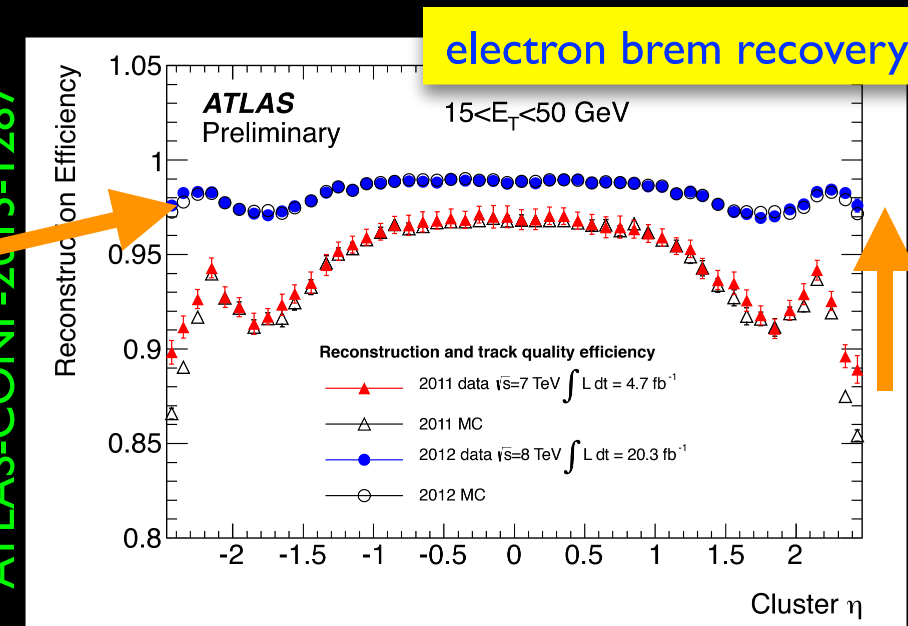

- pileup suppression for jets, τ, E_T-mis ...
  - ➡ combining calorimeter and tracking information
  - ➡ ATLAS pileup jet tagging (JVF and variants of it)
  - ➡ full fledged particle flow in CMS

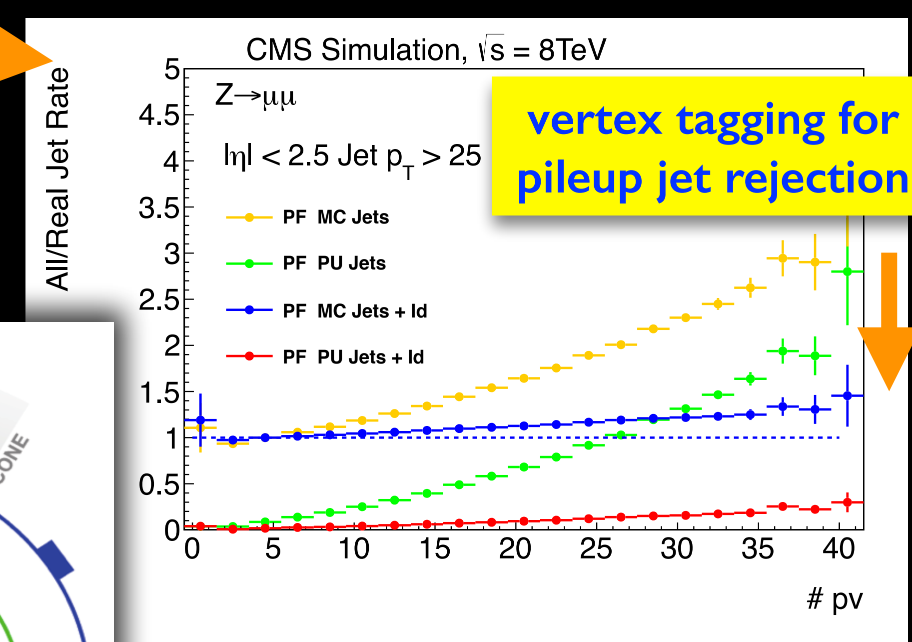

- more MVA based object identification
  - ➡ optimally combining all available information

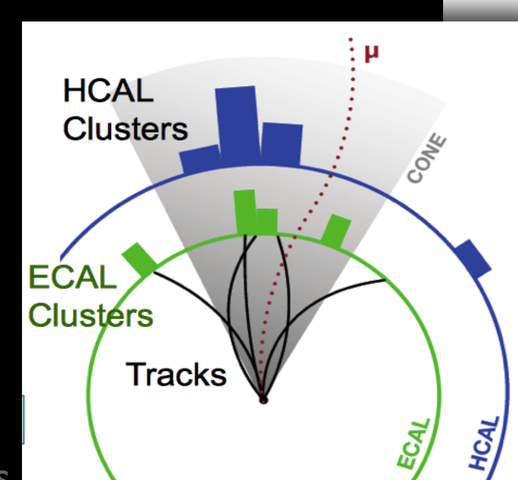

Distributed Computing

- analysis preparation for 2012
  - ➡ flexible and effective GRID operations
    - massive production of 8 TeV Monte Carlo
    - distribution of data samples across Tier-1 and Tier-2 centres
  - ➡ e.g. ATLAS used GRID resources continuously beyond pledges
  - ➡ >1500 active GRID analysers in ATLAS

F.Gianotti, ICHEP 2012

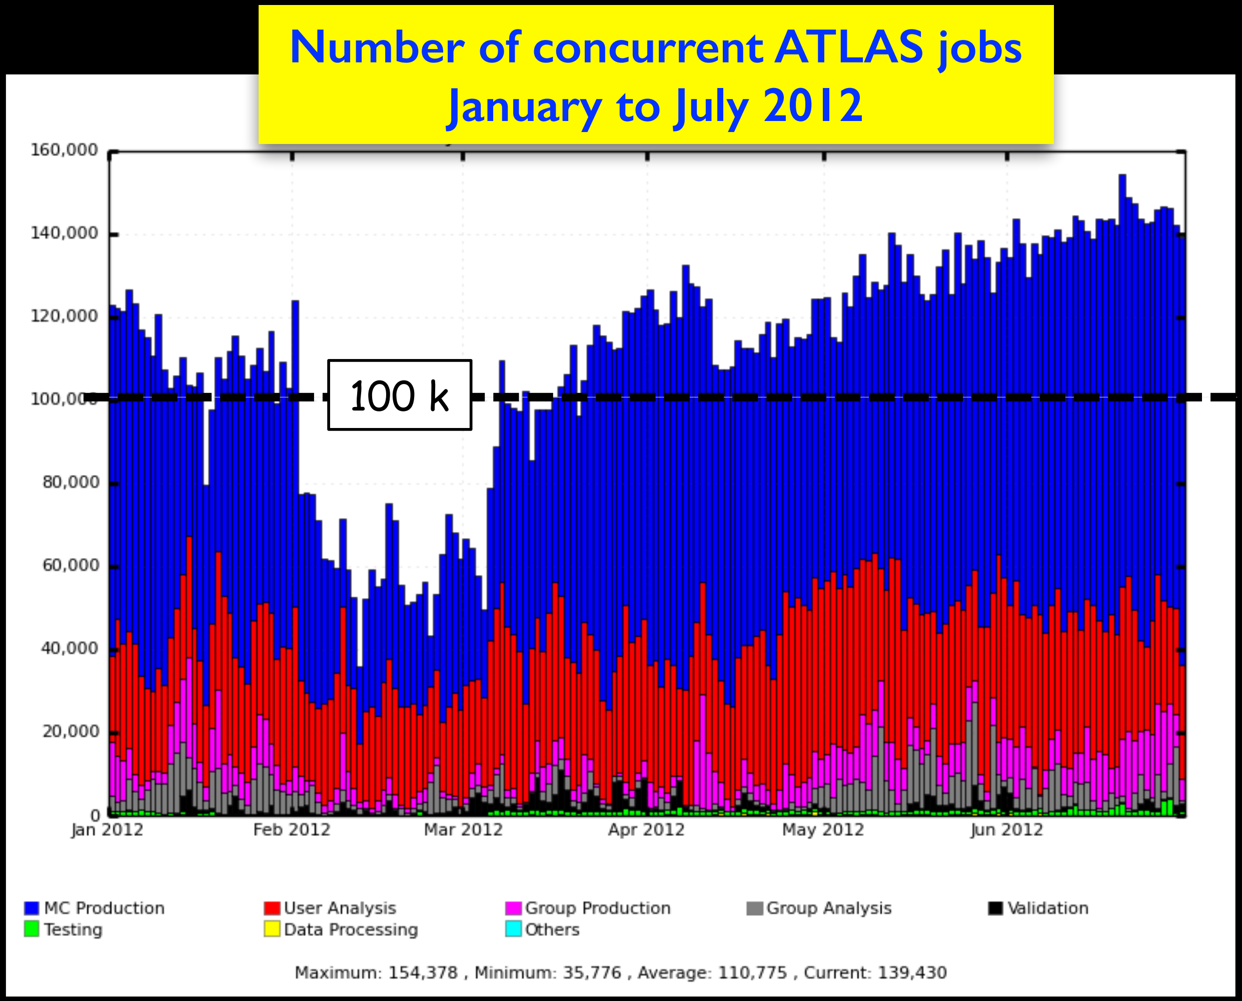

- fast updates of preliminary results using latest data for ICHEP 2012
  - ➡ relied on Tier-0 prompt data processing
    - required excellent quality of fast calibration
  - ➡ only final Higgs results used reprocessed data
    - reprocessing campaign takes few months

CMS CR-2012-145

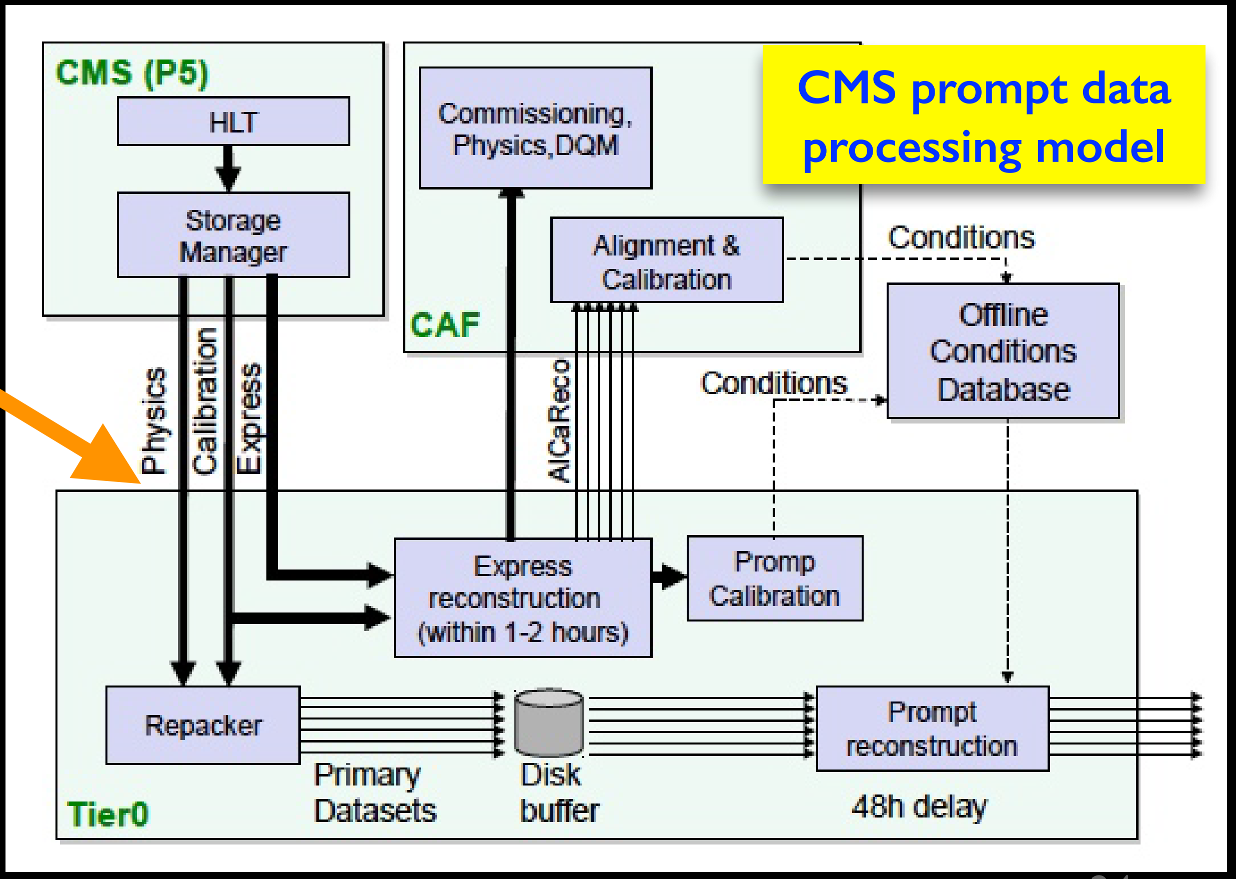

CERN Seminar July 4th, 2012: the Higgs

F.Gianotti, ICHEP 2012

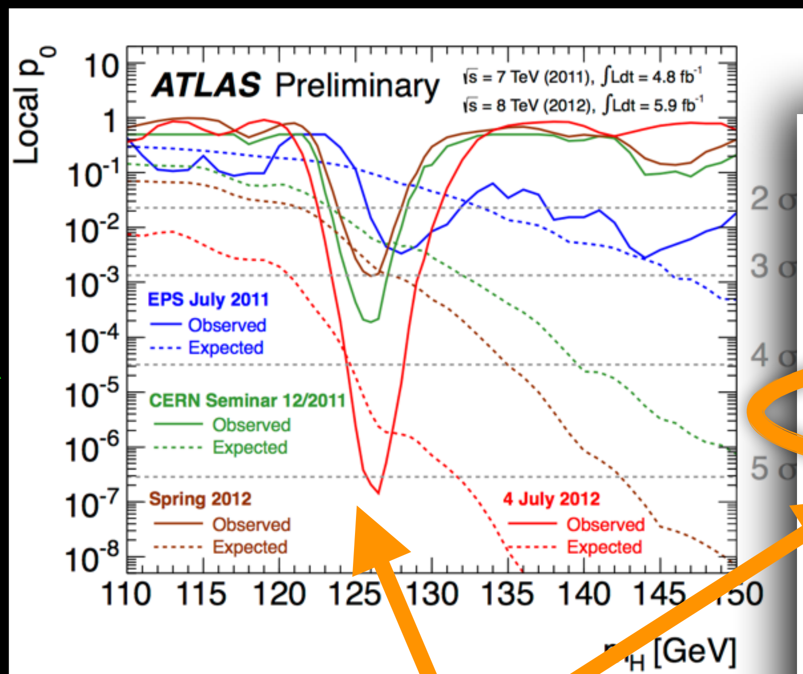

We present updated results on SM Higgs searches based on the data recorded in 2011 at $\sqrt{s}$=7 TeV (~4.9 fb^{-1}) and 2012 at $\sqrt{s}$=8 TeV (~5.9 fb^{-1})

Results are preliminary:
- ❑ 2012 data recorded until 2 weeks ago
- ❑ x2 larger event pile-up
- ❑ new, improved analyses deployed for the first time

H → γγ and H→ 4l: high-sensitivity at low-m_H; high mass-resolution; pile-up robust
- ❑ analyses improved to increase sensitivity → new results from 2011 data
- ❑ all the data recorded so far in 2012 have been analyzed
- → results are presented here for the first time

Other low-mass channels: H→ WW$^{(*)}$→ lνlν, H→ ττ, W/ZH→ W/Z bb:
- ❑ E_T^{miss} in final state → less robust to pile-up
- ❑ worse mass resolution, no signal "peak" in some cases
- ❑ complex mixture of backgrou
- → understanding of the detecto advanced, but results not yet
- → 2011 results used here for th

ATLAS: Status of SM Higgs searches, 4/7/20 3

- fantastic success (!!!)
  - ➡ software and computing had its share in it ...
  - ➡ full chain worked excellent:
    - from detector + trigger to
    - prompt calibration,
    - Tier-0 reconstruction,
    - GRID distribution and
    - fast distributed analysis !

CERN

We all know what happened next ...

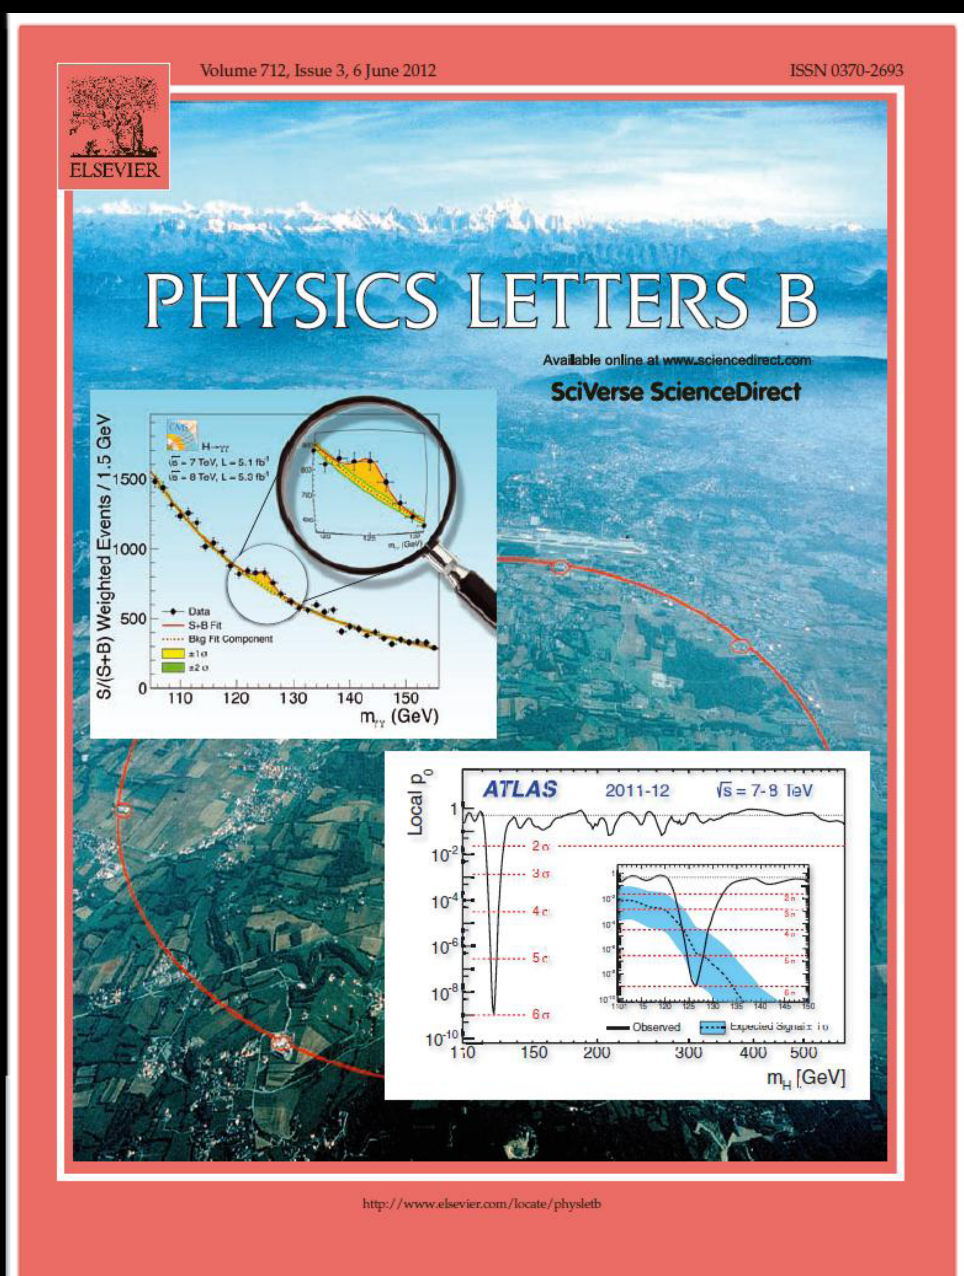

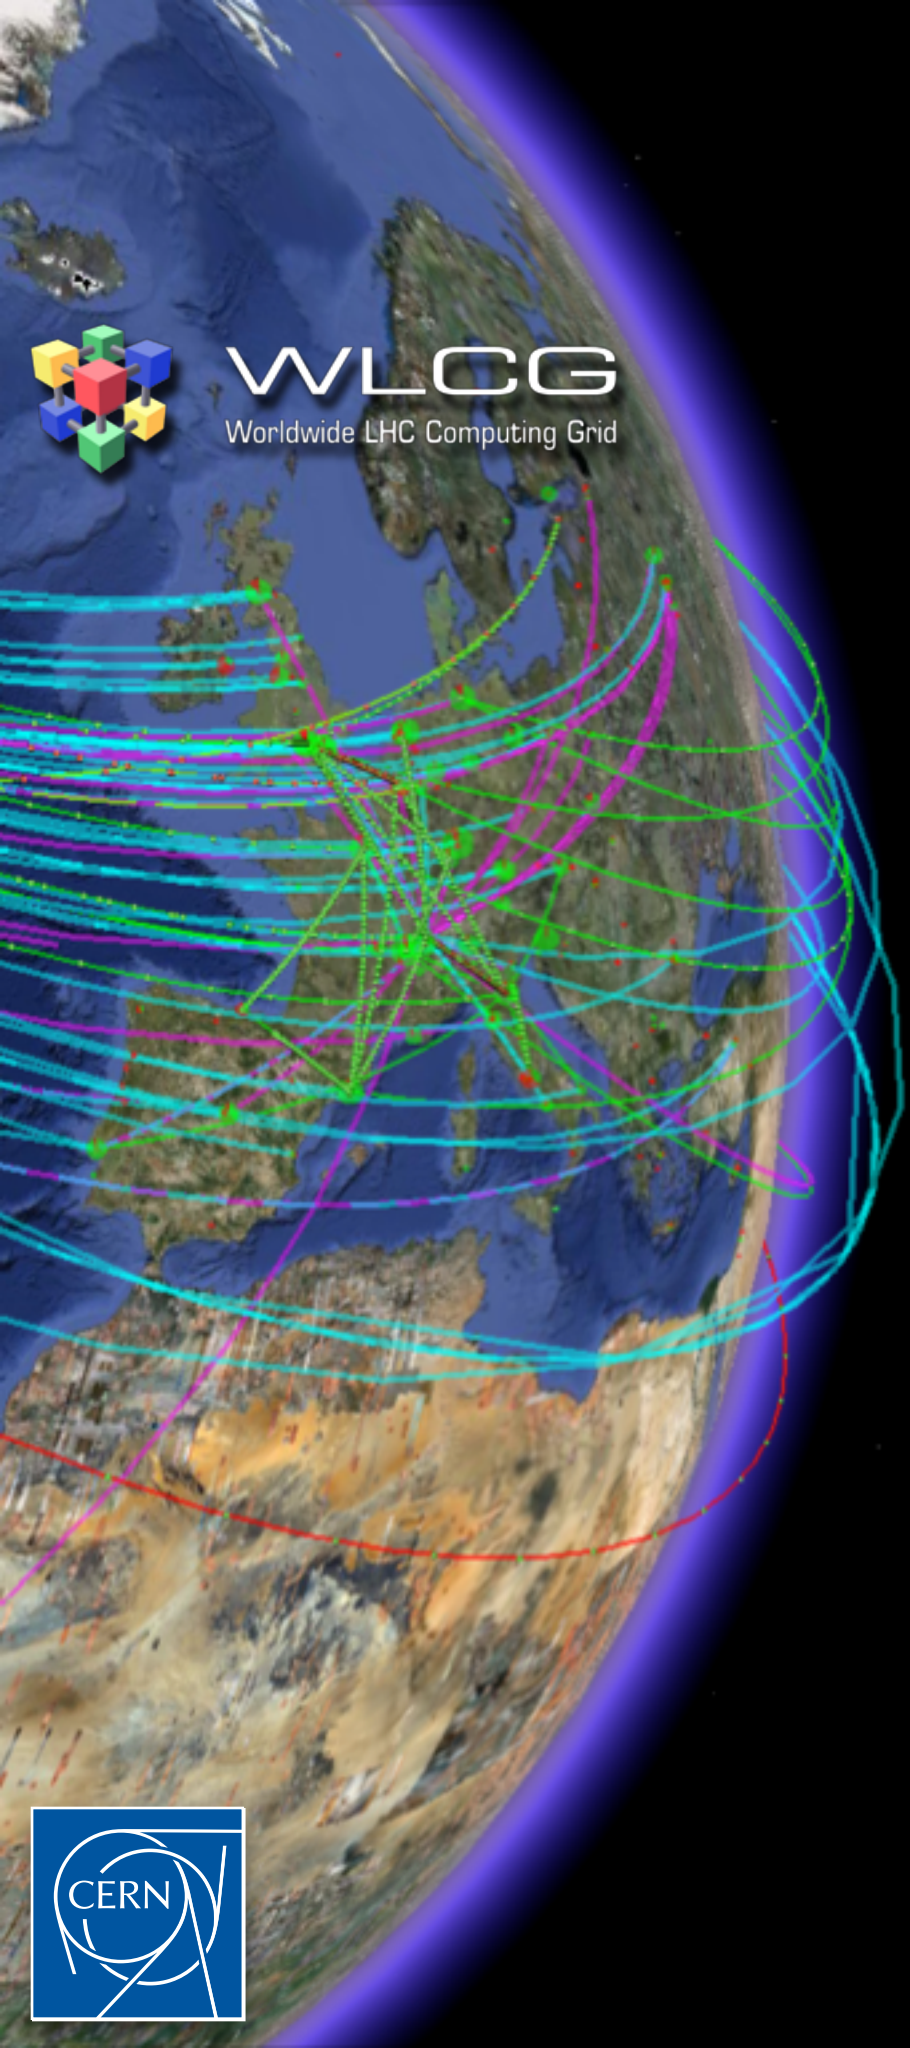

Preparing for Run-2: First Upgrades of Software and Computing

Computing Constraints for Run-2

- unlike Run-1, computing resources will be limited !
  - ➡ assumption is a constant computing budget
  - ➡ interplay of technology advancement, market price and needed replacements

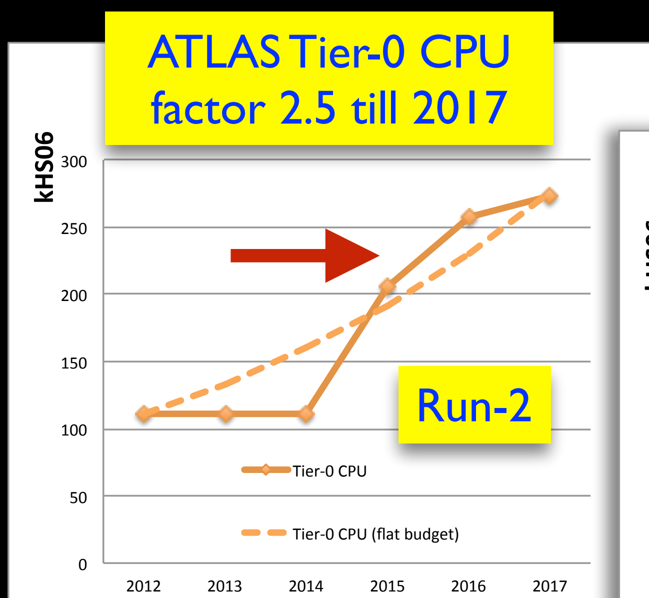

B.Kersevan et al.

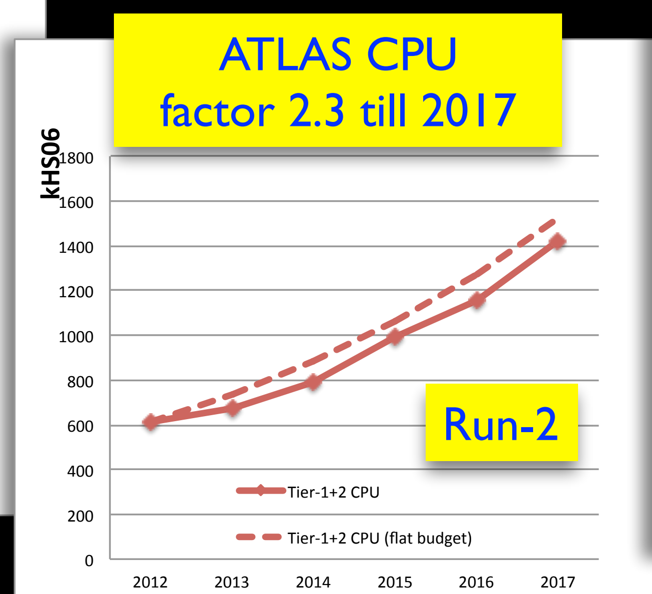

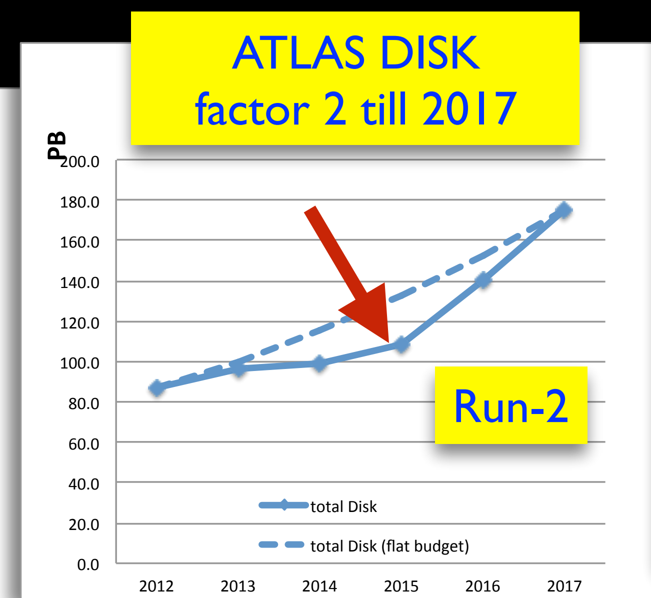

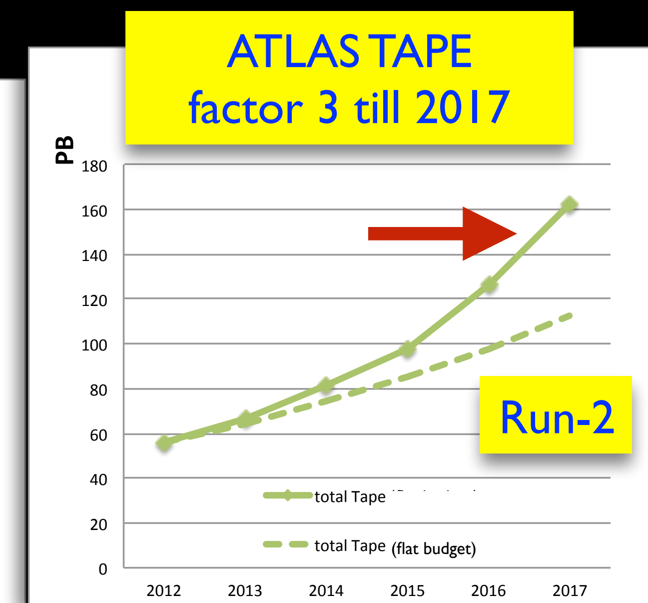

- motivation for LS1 software upgrades
  - ➡ ensure that Tier-0 can process 1kHz trigger rate
  - ➡ optimise disk usage (e.g. ATLAS new Analysis Model)
- biggest problem will be disk !

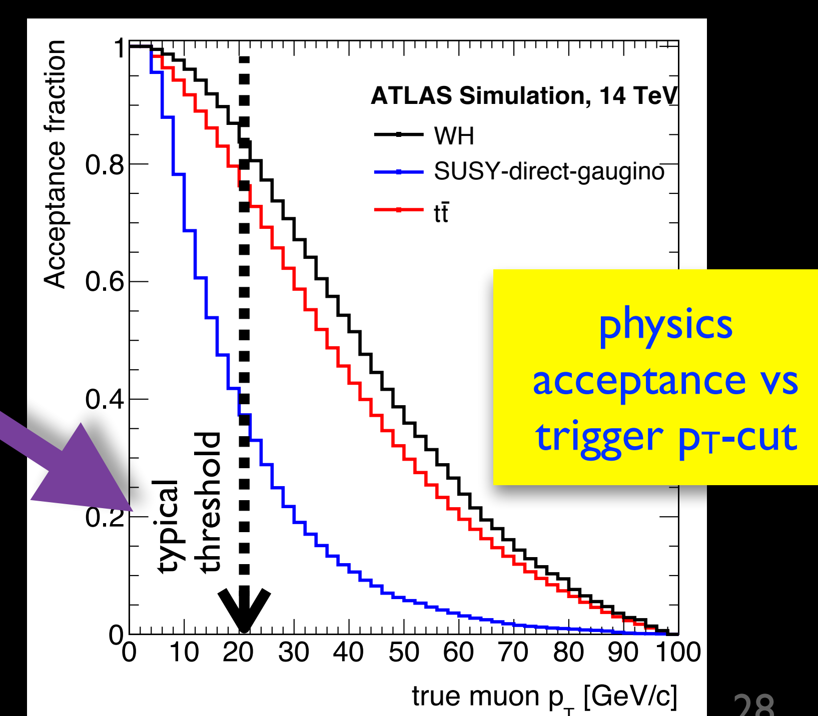

CPU for Reconstruction

- focus on software technology and improve current algorithms
  - ➡ improve software technology, including:
    - simplify EDM design to be less OO ("hip" 10 years ago)
    - faster vector+matrix algebra libs (Eigen)
    - vectorised trigonometric functions (VDT, Intel)
    - work on CPU hot spots
  - ➡ tune reconstruction strategy (very similar in ATLAS and CMS)
    - optimise track finding strategy for 40 pileup
    - modify track seeding to explore 4th Pixel layer

- huge gains achieved !
  - ➡ ATLAS reports overall factor 3 in CPU time
    - touched >1000 packages for factor 4 in tracking
  - ➡ CMS reports overall factor 2 in CPU time
    - as well dominated by tracking improvements
  - ➡ both experiments within 1 *kHz* Tier-0 budget
    - required to keep single lepton triggers

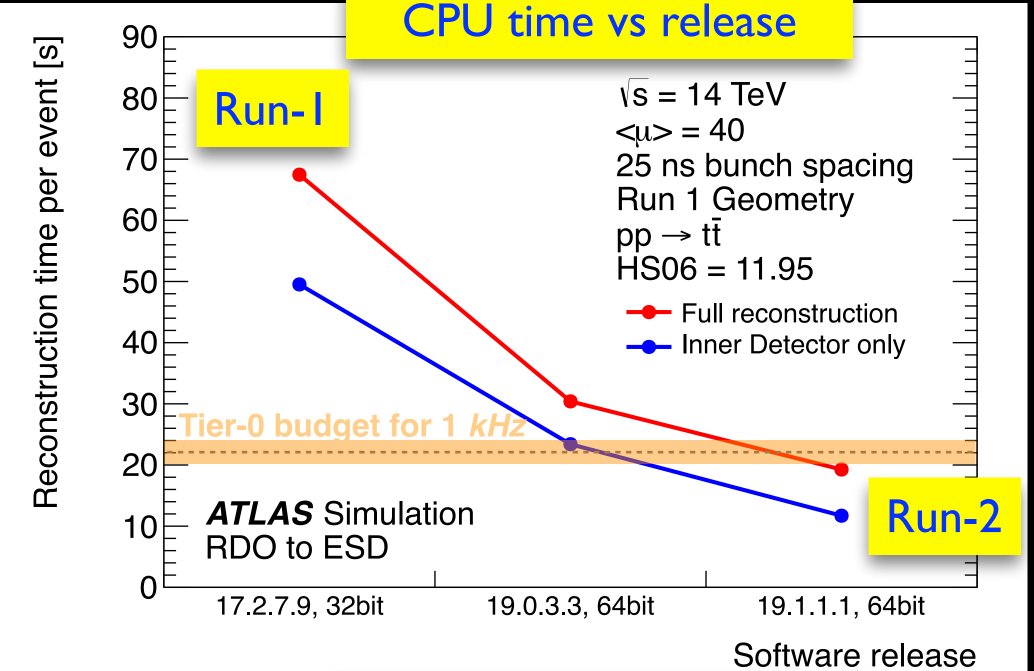

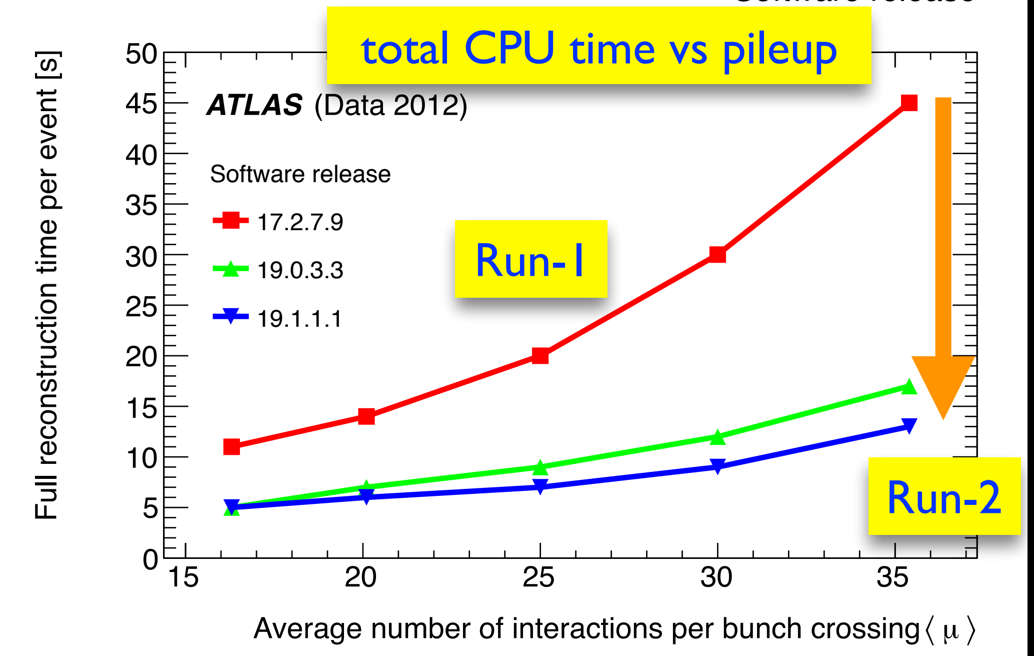

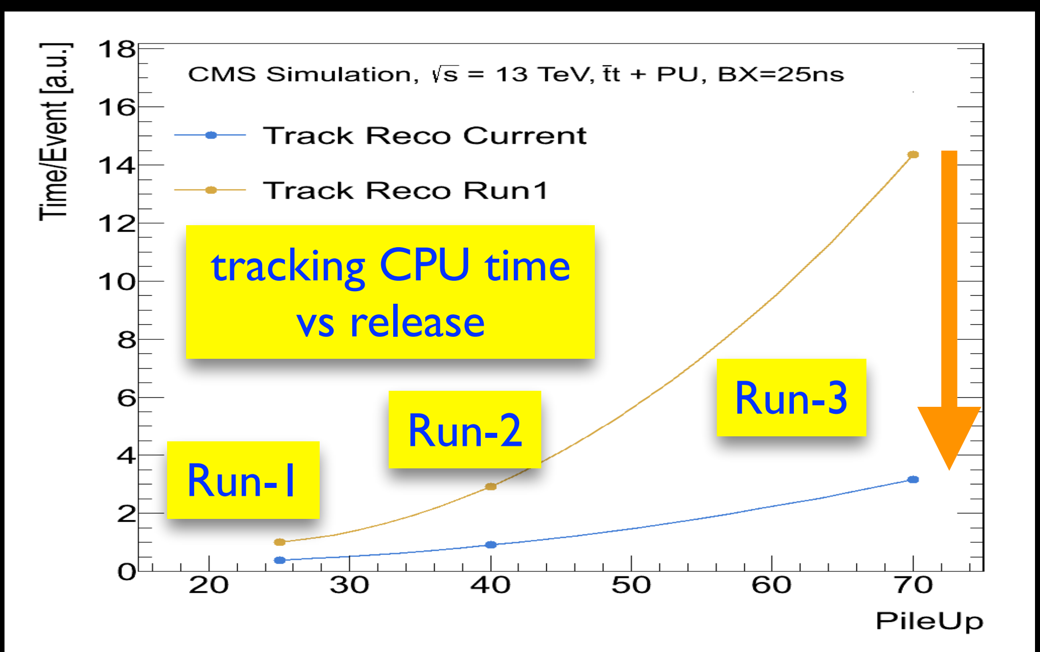

ATLAS New Analysis Model for Run-2

- several issues with Run-1 model
  - ➡ analysis ntuples duplicate AOD (disk !)
  - ➡ production of ntuples costly (time !)
  - ➡ analysers develop in ROOT (compatibility !)

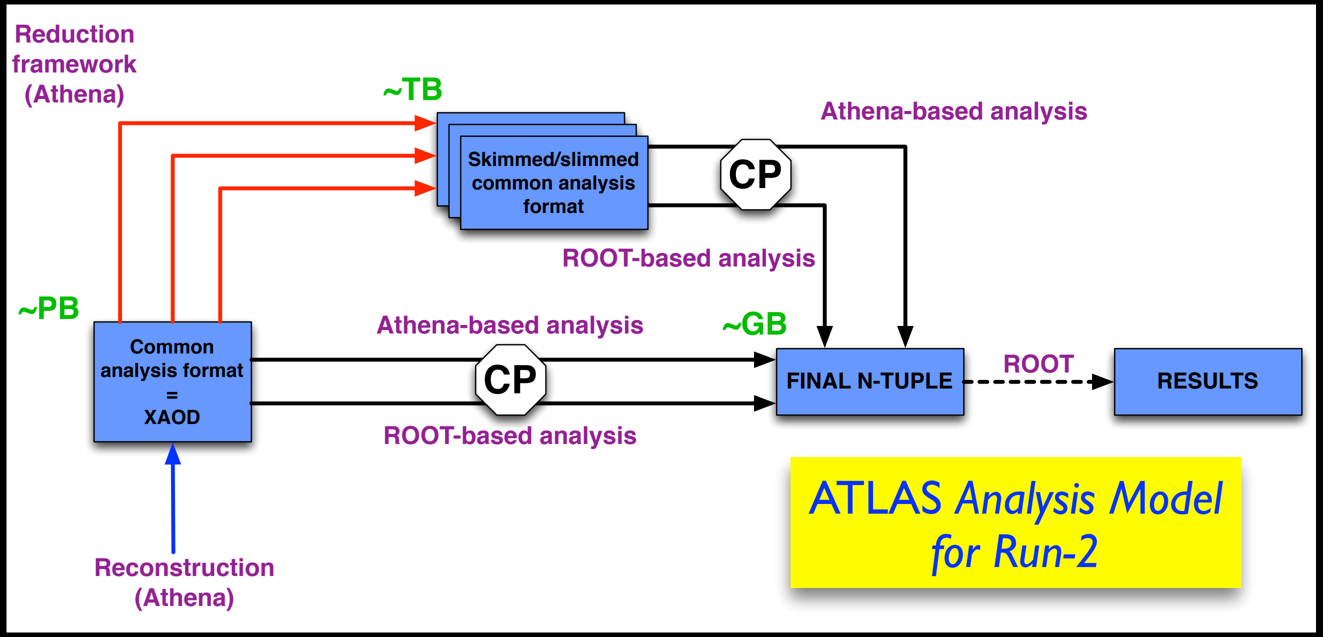

- "small" revolution for ATLAS
  - ➡ new format (xAOD) readable in ROOT
    - branch-wise reading at ROOT speed
    - object decoration with user data
  - ➡ centrally produce skims for analysers
    - train production model
    - smart slimming of xAOD objects
  - ➡ analysis tools transparently usable in ROOT and ATHENA
    - ROOT based and ATHENA based analysis software releases

- changes for other experiments are less extreme
  - ➡ similar pressure to reduce resource needs

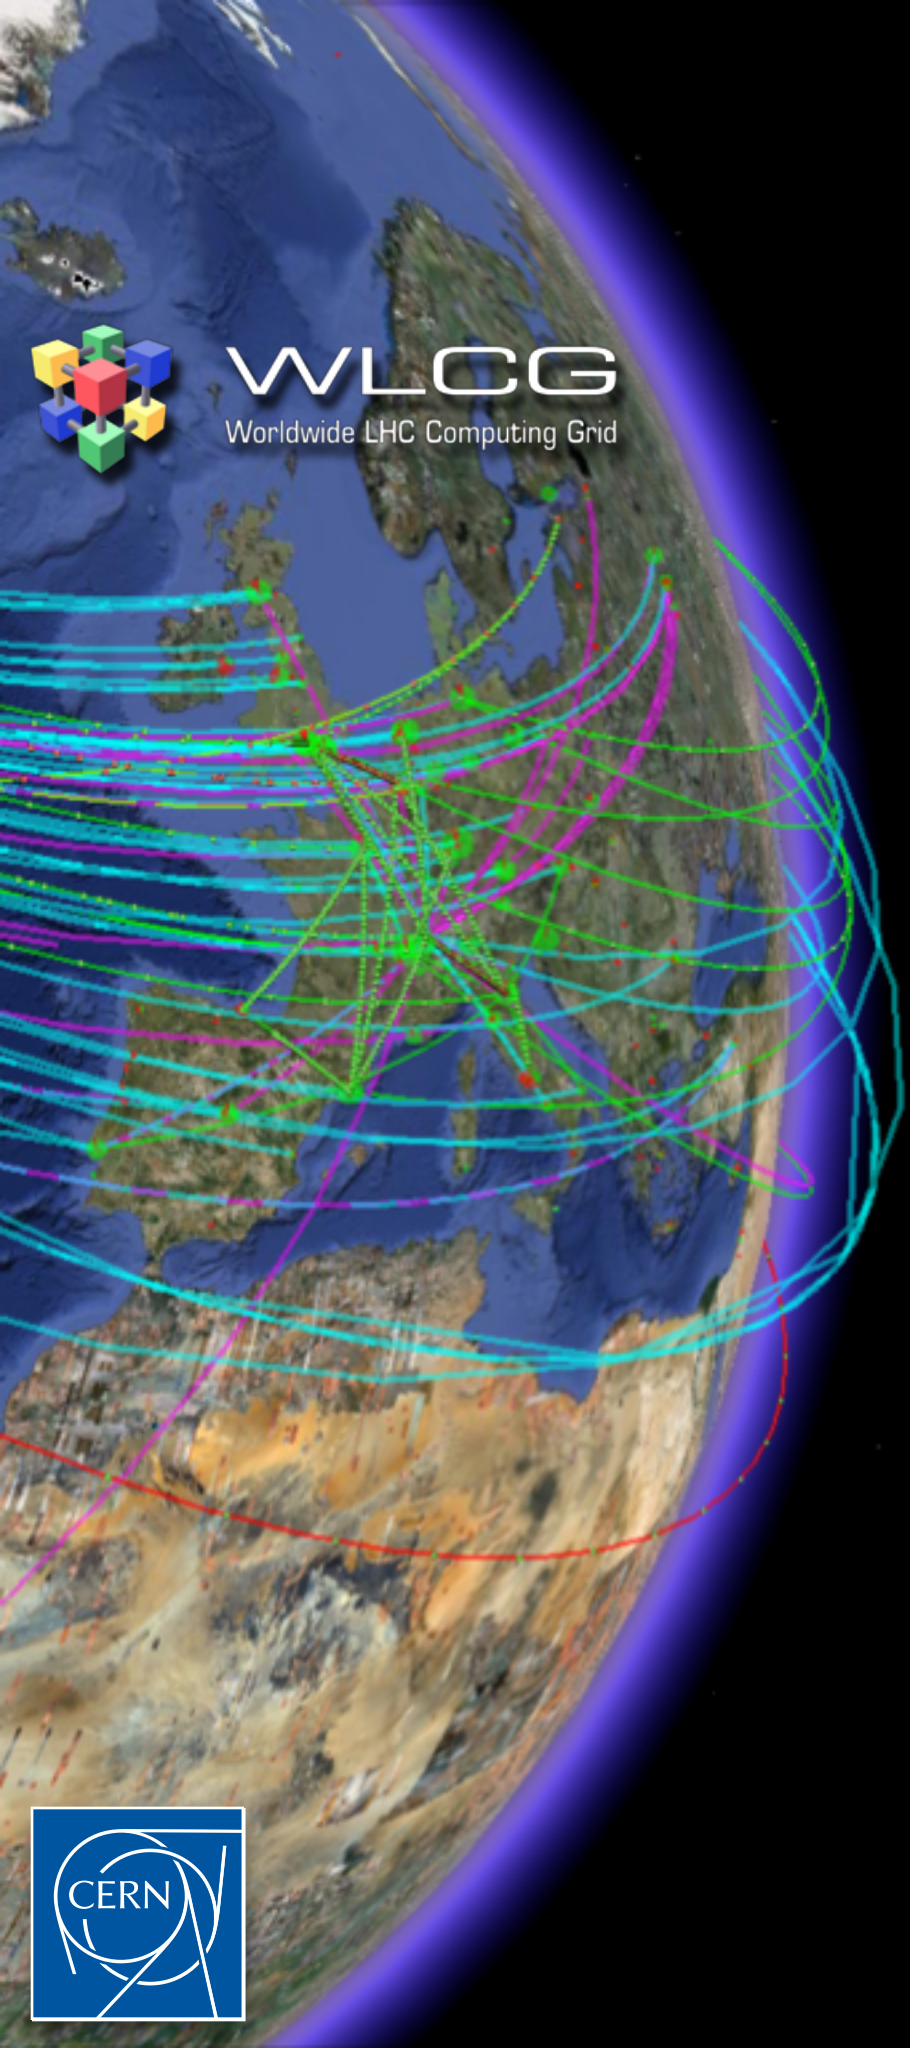

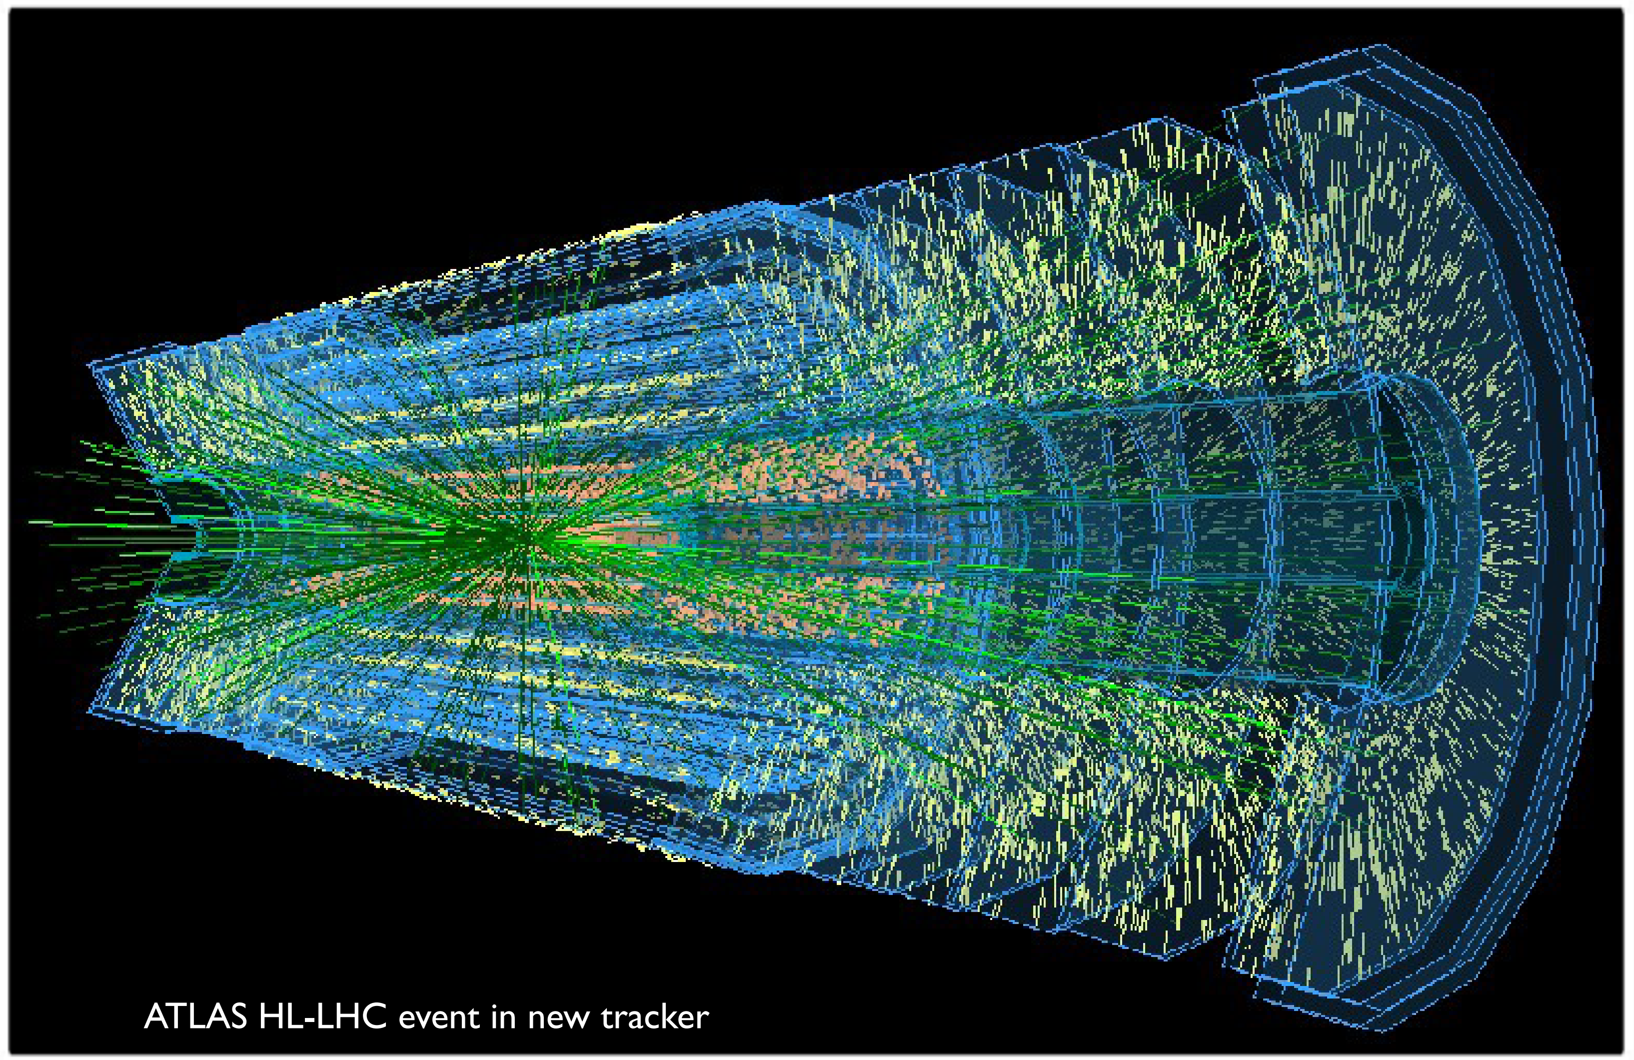

ATLAS HL-LHC event in new tracker

Software for Detector Upgrades

LHC Upgrade Program

- Phase-1 upgrades (2018→)
  - ➡ LHCb and ALICE trigger-less readout
  - ➡ CMS and ATLAS ready for 350 fb^{-1}
- Phase-2 upgrades (2023→)
  - ➡ HL-LHC upgrades for CMS and ATLAS for 3000 fb^{-1}
- software plays key role in this program
  - ➡ physics prospects, detector design, TDRs...
  - ➡ preparing offline and trigger for detector upgrades itself

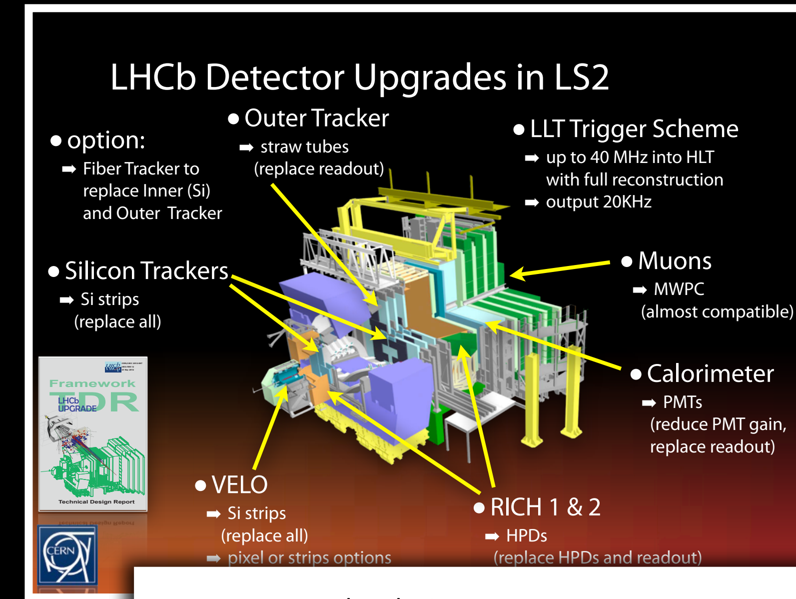

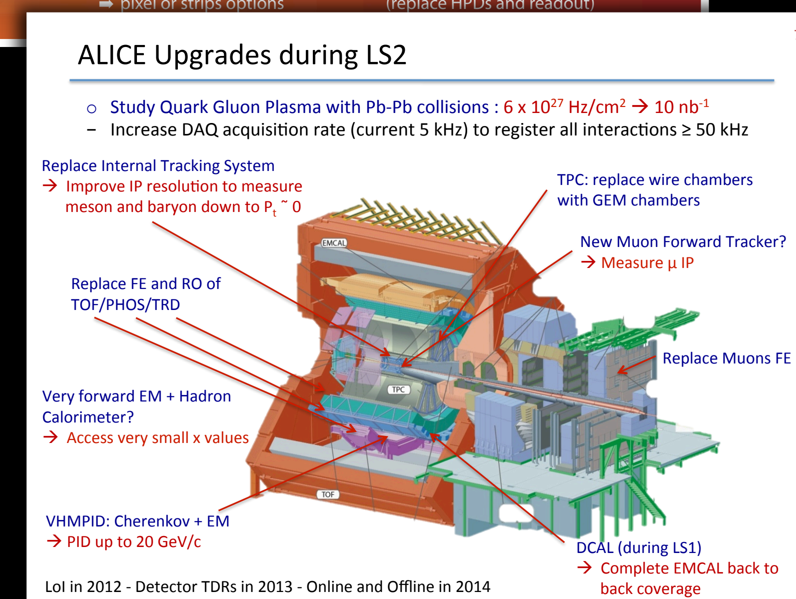

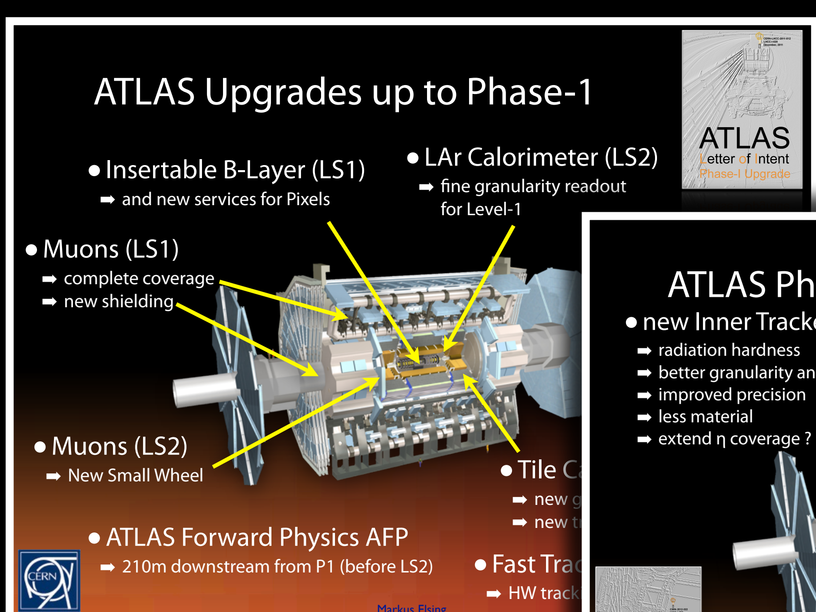

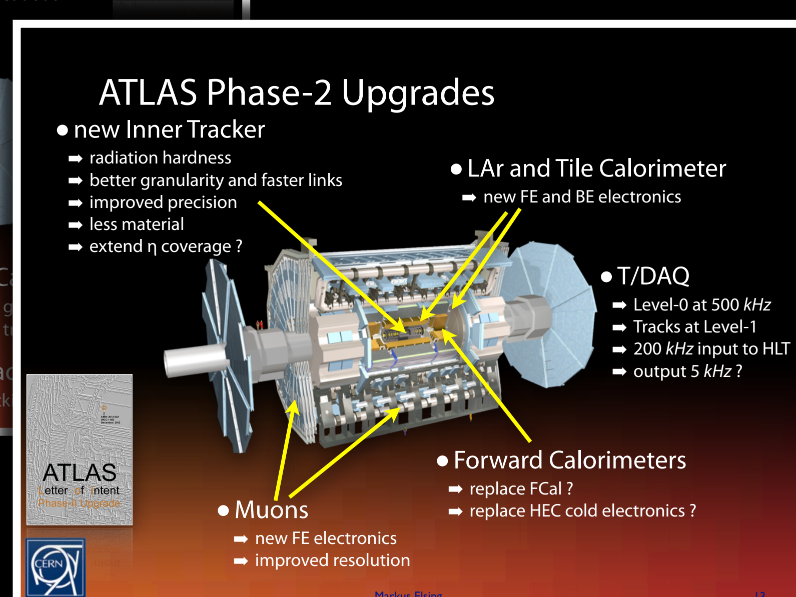

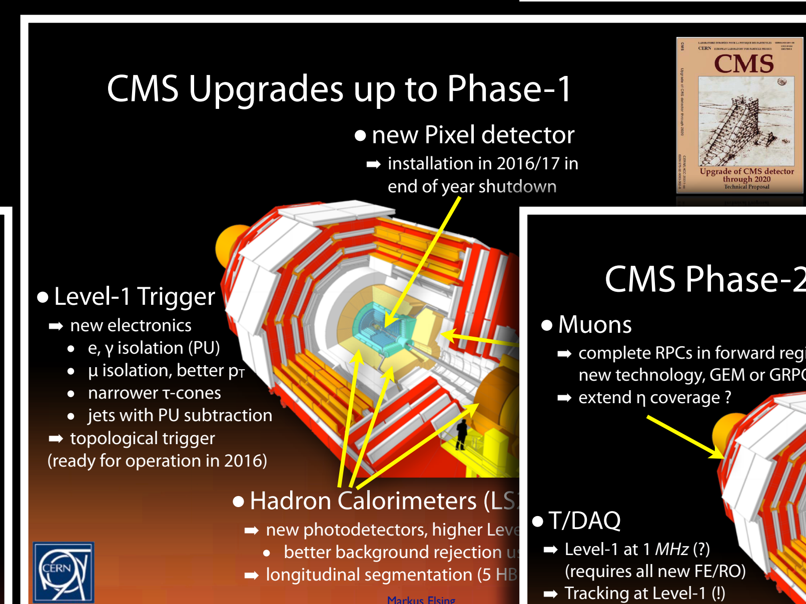

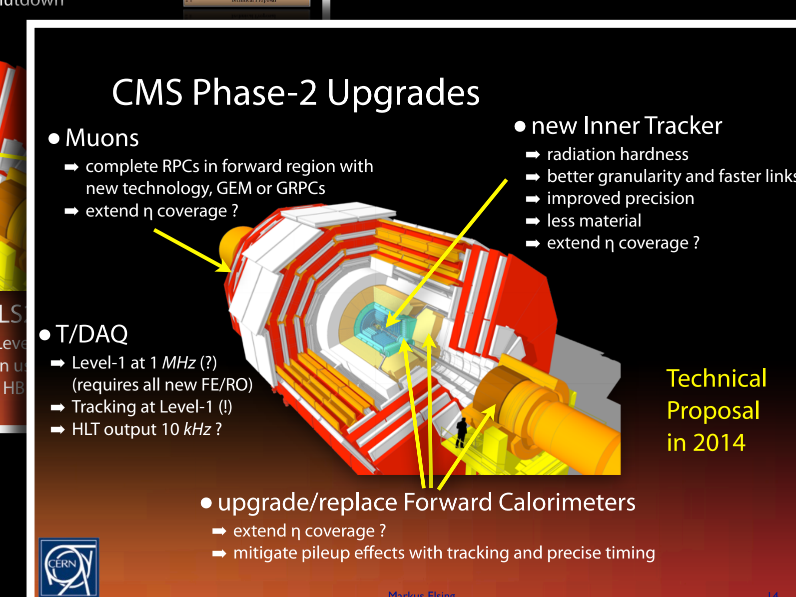

Software and Manpower

- software follows a natural life cycle
  - ➡ building up the software for an experiment
  - ➡ start of operations and data taking
  - ➡ data analysis and detector upgrades
- loss of software manpower in ATLAS/CMS
  - ➡ (mostly) students and postdocs moved on to do physics
    - same trend like in previous experiments
  - ➡ like CDF/D0 Run-2, LHC upgrade program is ambitious
    - need to find sufficient manpower to develop the software for the upgrade

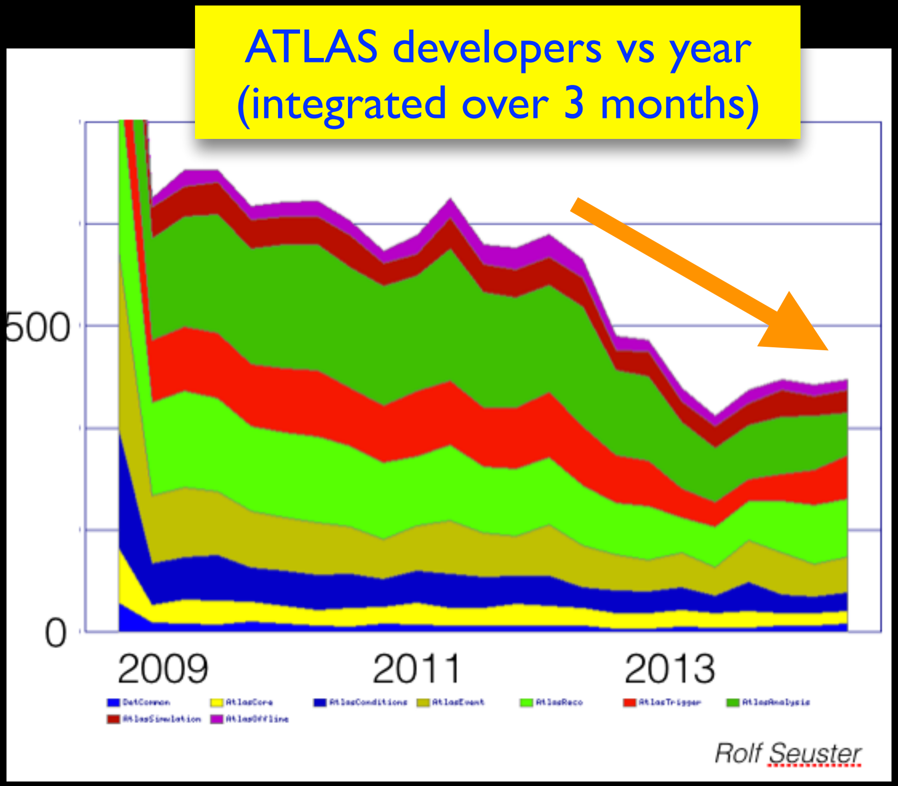

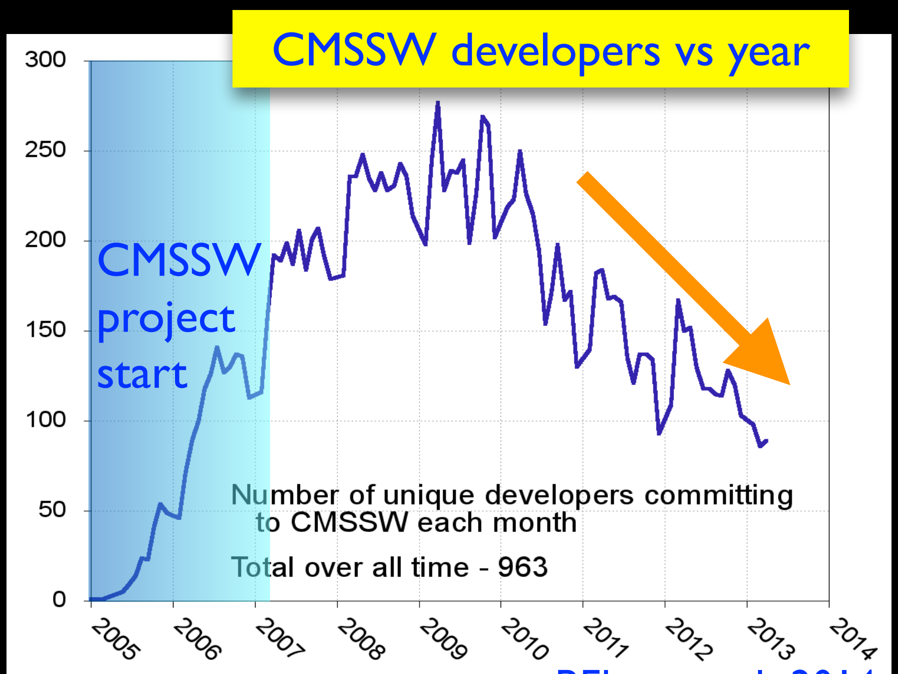

P.Elmar et al., 2014

P.Elmar, L.Sexton-Kennedy, C.Jones, ICHEP 2007

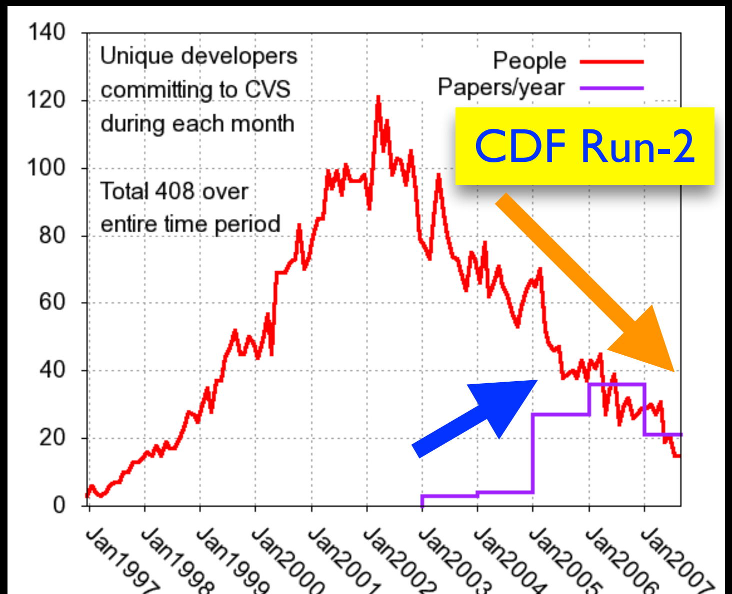

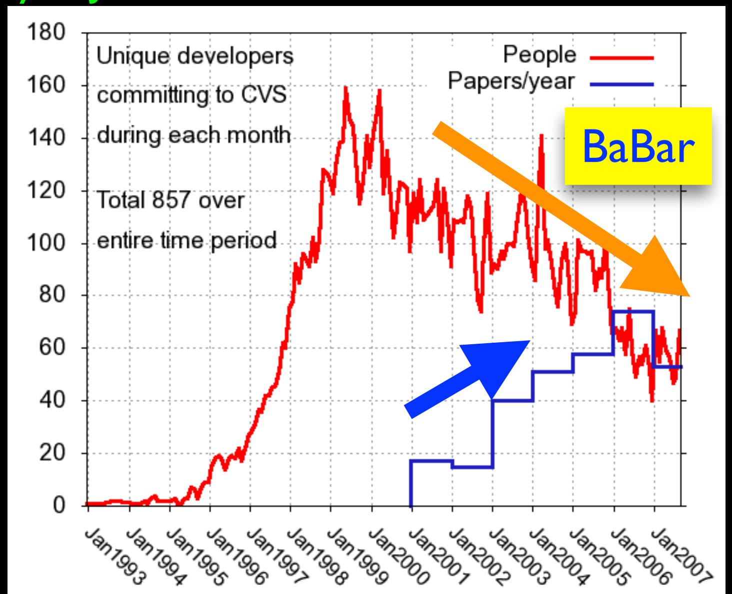

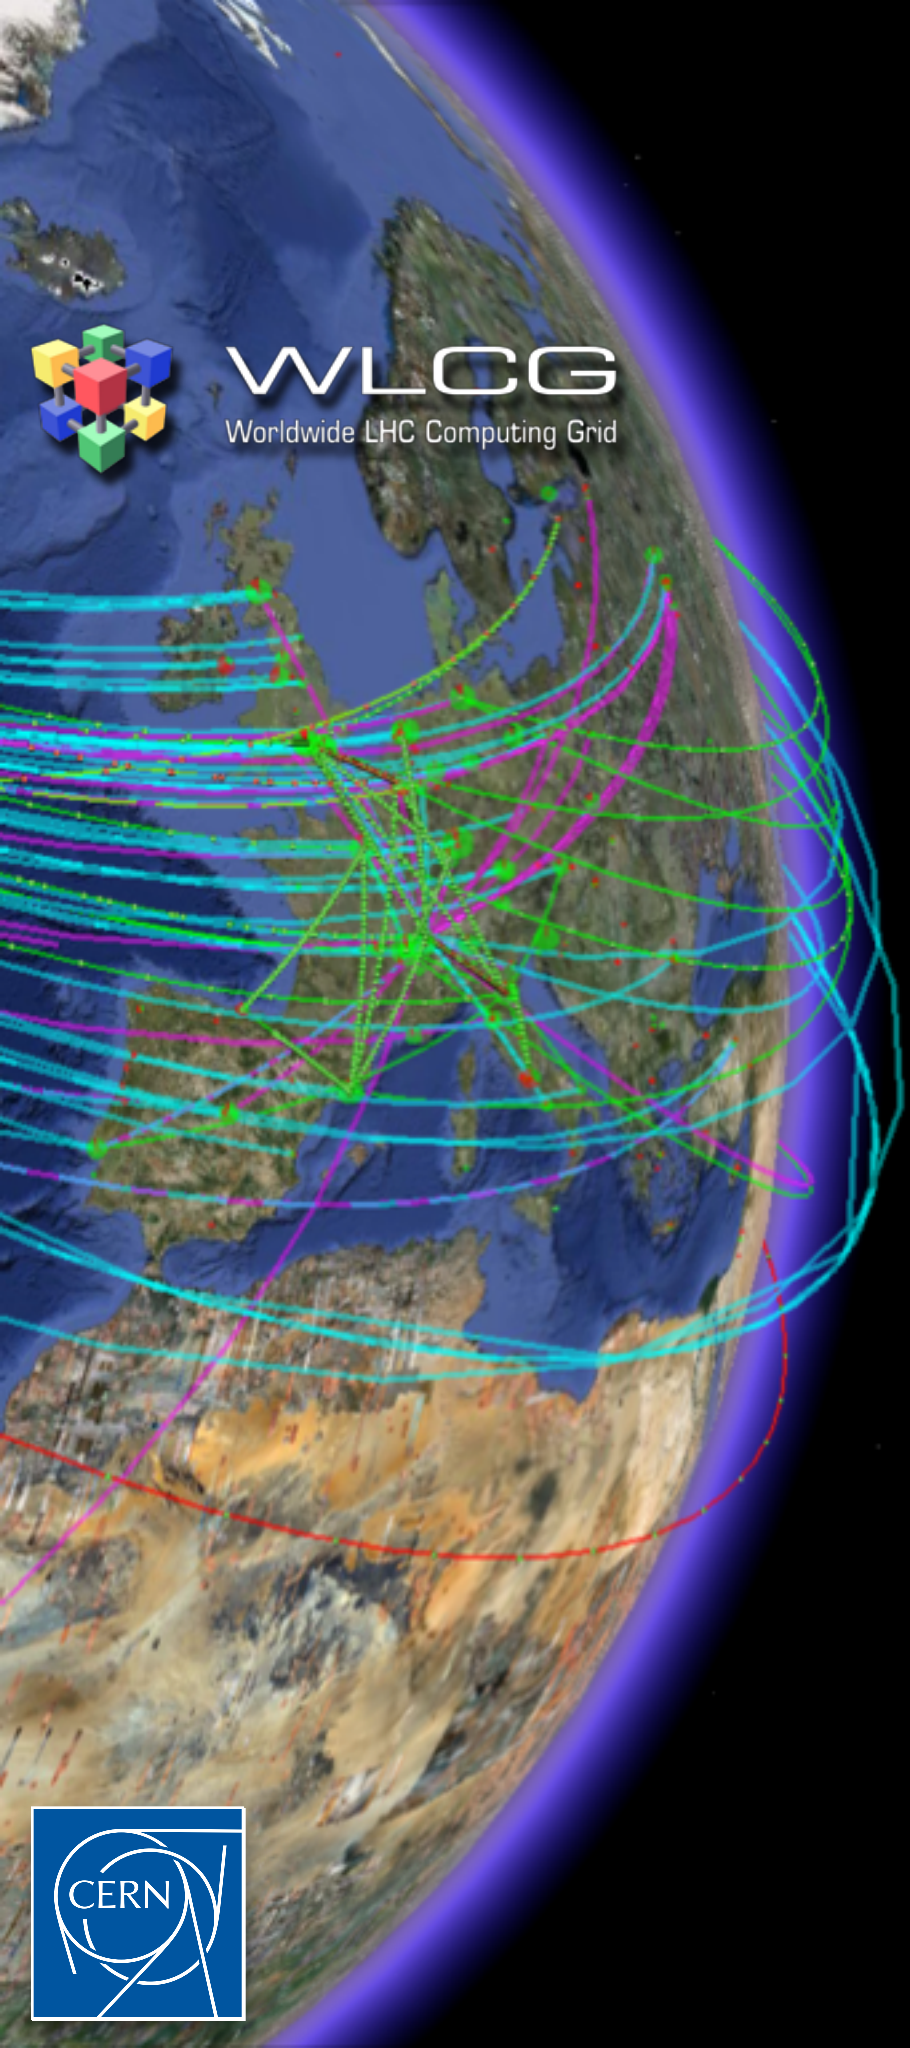

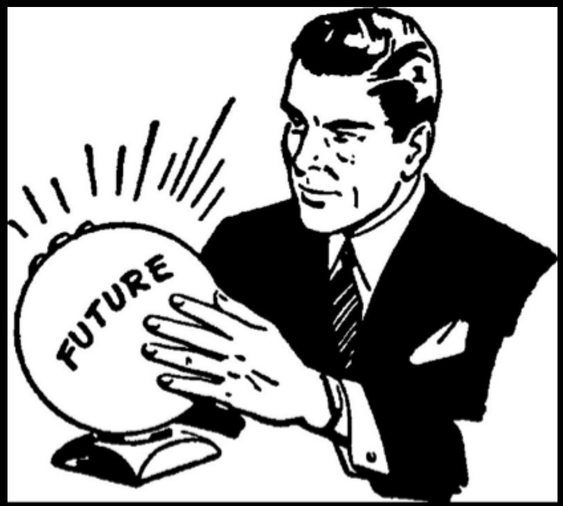

Future Software and Computing Challenges

the million dollar question:
how to process HL-LHC events

Future Computing Needs

- increase in raw data samples
  - ➡ driven by ALICE trigger-less readout
    - mostly for their online disk buffer
  - ➡ ATLAS and CMS increase of trigger rate and event size (pileup)
- total disk needs scales with raw
  - ➡ current models are above constant budget, hence need:
    - smaller data formats
    - new analysis models
    - use more tape (cheaper, continues to scale)
    - less replicas (use growing network bandwidth)
- CPU needs less certain
  - ➡ best estimates are factors above budget
    - based on current applications and models

plots M.Krzewicki, ECFA WS 2014

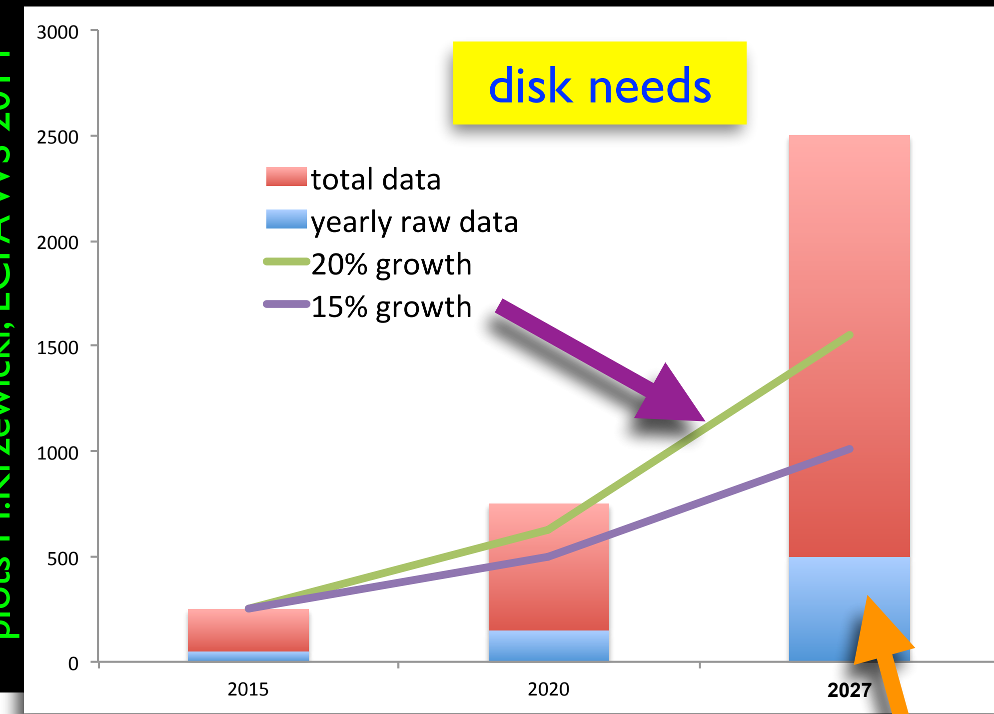

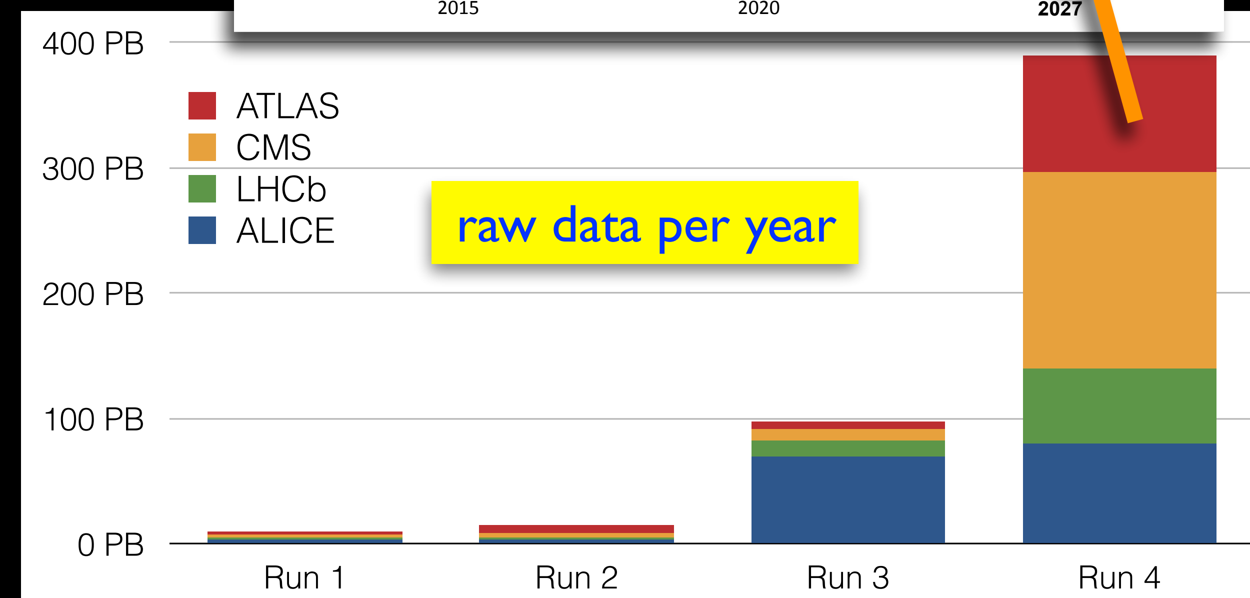

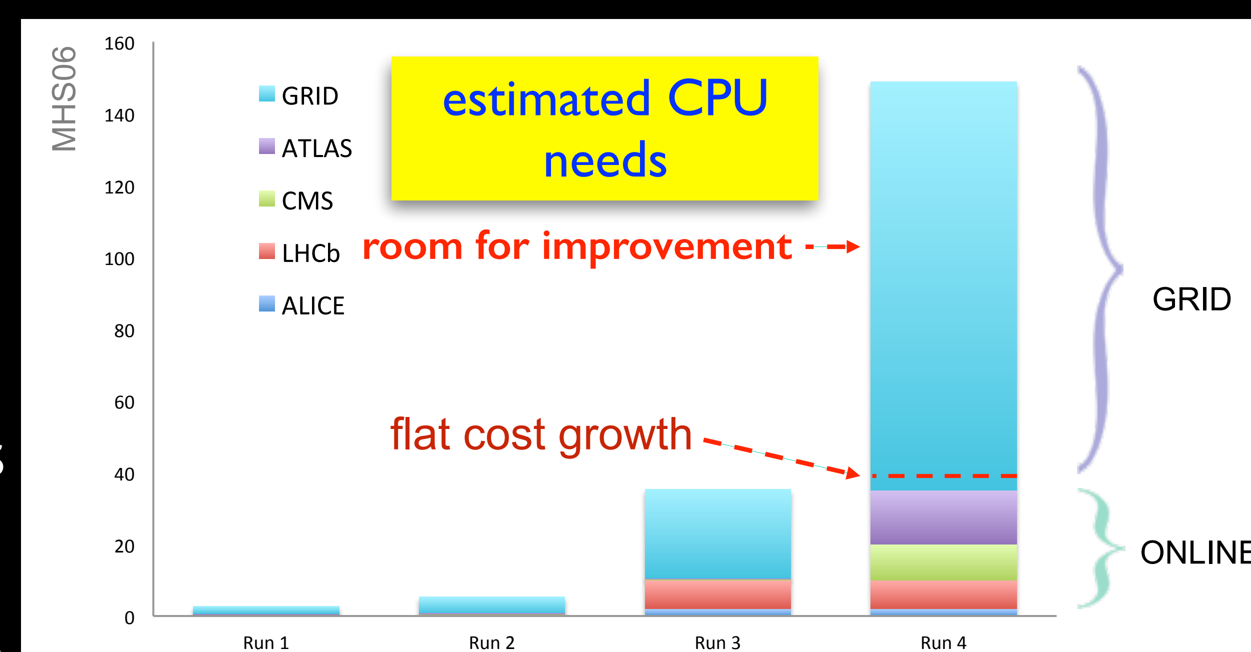

Processor Technology

- Moore's law is still alive
  - ➡ number of transistors still doubles every 2 years
    - no free lunch, clock speed no longer increasing
  - ➡ lots of transistors looking for something to do:
    - vector registers
    - out of order execution
    - hyper threading
    - multiple cores
  - ➡ increase theoretical performance of processors
    - hard to achieve this performance with HEP applications

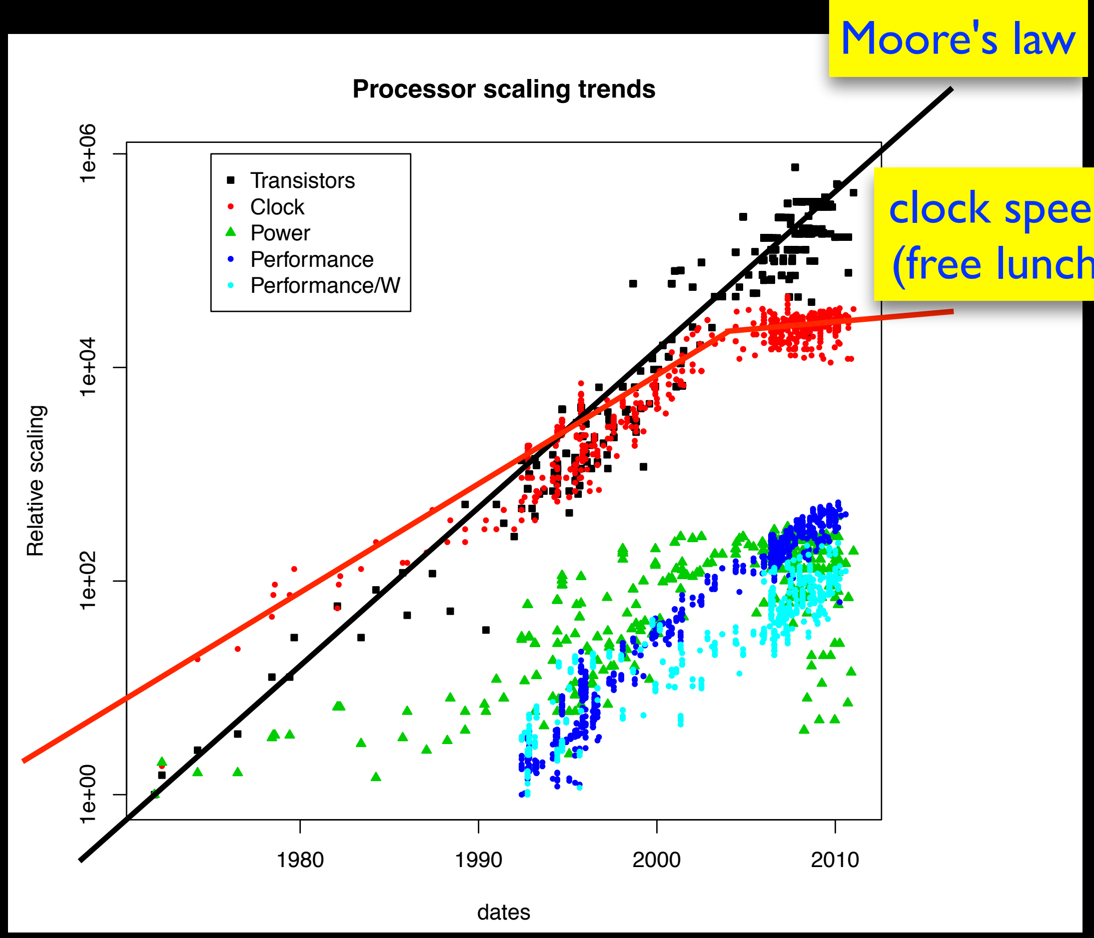

- many-core processors, including GPGPUs
  - ➡ e.g. Intel Xenon Phi, Nvidia Tesla
  - ➡ lots of cores with less memory
    - same for ARM or ATOM based systems
  - ➡ challenge will be to adapt HEP software
    - need to parallelise applications (multi-threading)
      (GAUDI-HIVE and CMSSW multi-threading a step in this direction)
    - change memory model for objects, more vectorisation, ...

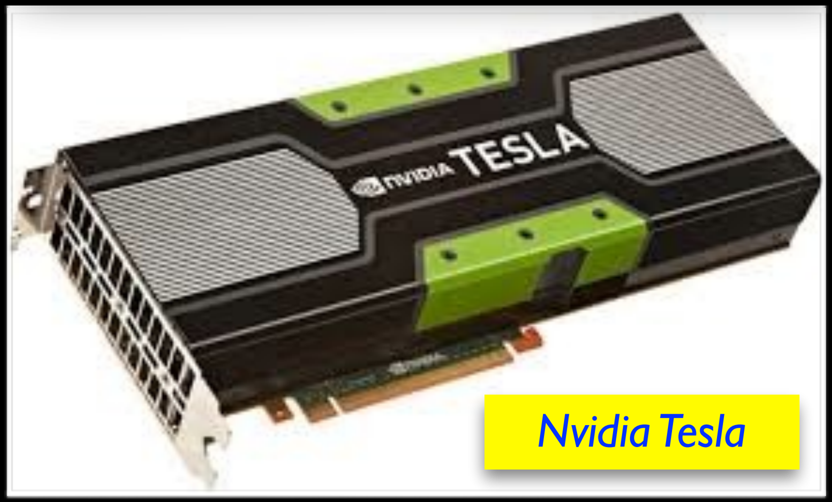

Trends in LHC Computing

- pledged GRID resources indispensable
  - ➡ will continue to be basis for LHC computing
  - ➡ make full use of resources (e.g. HLT farms outside data taking)

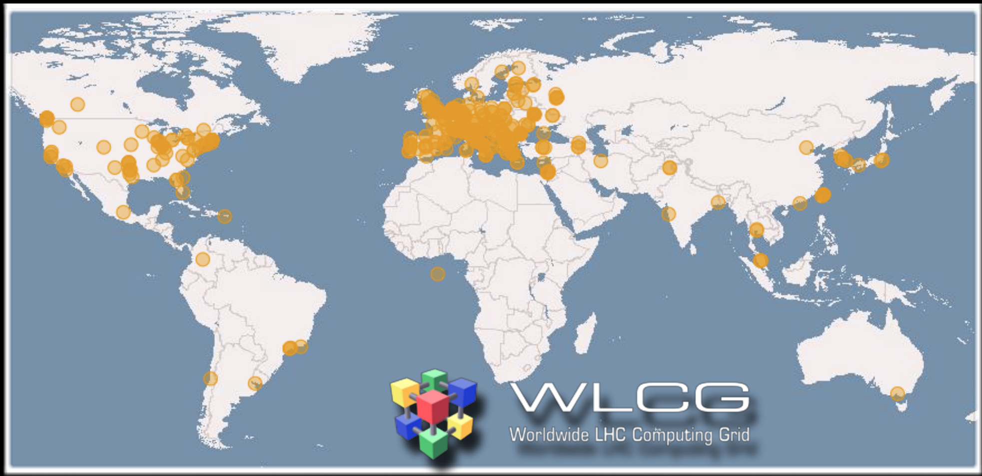

- more heterogeneous infrastructure
  - ➡ opportunistic usage of additional resources
    - commercial Cloud providers (i.e. Google, Amazon)
    - free CPU in High Performance Computing centres
      (big HPC centres outperform WLCG in CPU)
  - ➡ storage will not become opportunistically available

- GRID services become (even) more flexible
  - ➡ global data federations serve data to jobs at remote sites
    (FAX - ATLAS, AAA - CMS, AliEn - ALICE)
  - ➡ ATLAS "event service"
    - short payloads for opportunistic remote computing

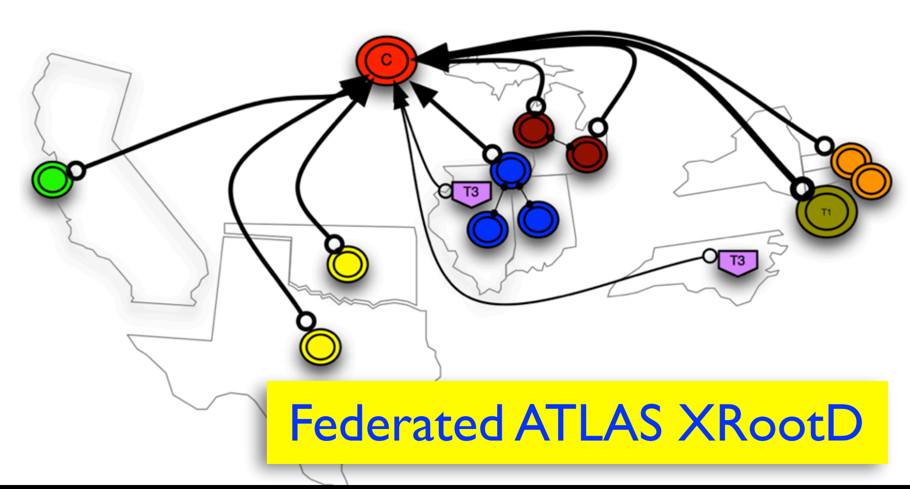

CERN

Detector Simulation

- simulation limited by CPU
  - ➡ avoid MC limiting physics precision
  - ➡ need to increase GRID "MC luminosity"
- major software technology developments in simulation
  - ➡ Geant 4.10 introduces multi-threading support
  - ➡ Geant V redesign to explore vectorisation
- ATLAS Integrated Simulation Framework (ISF)
  - ➡ mixes fast and full sim. in one event
    - spend time on important event aspects
  - ➡ towards complete fast software chain
    - avoid digit. and reco. bottleneck
    - directly produce analysis formats (disk)

A.Gheata, ACAT 2014

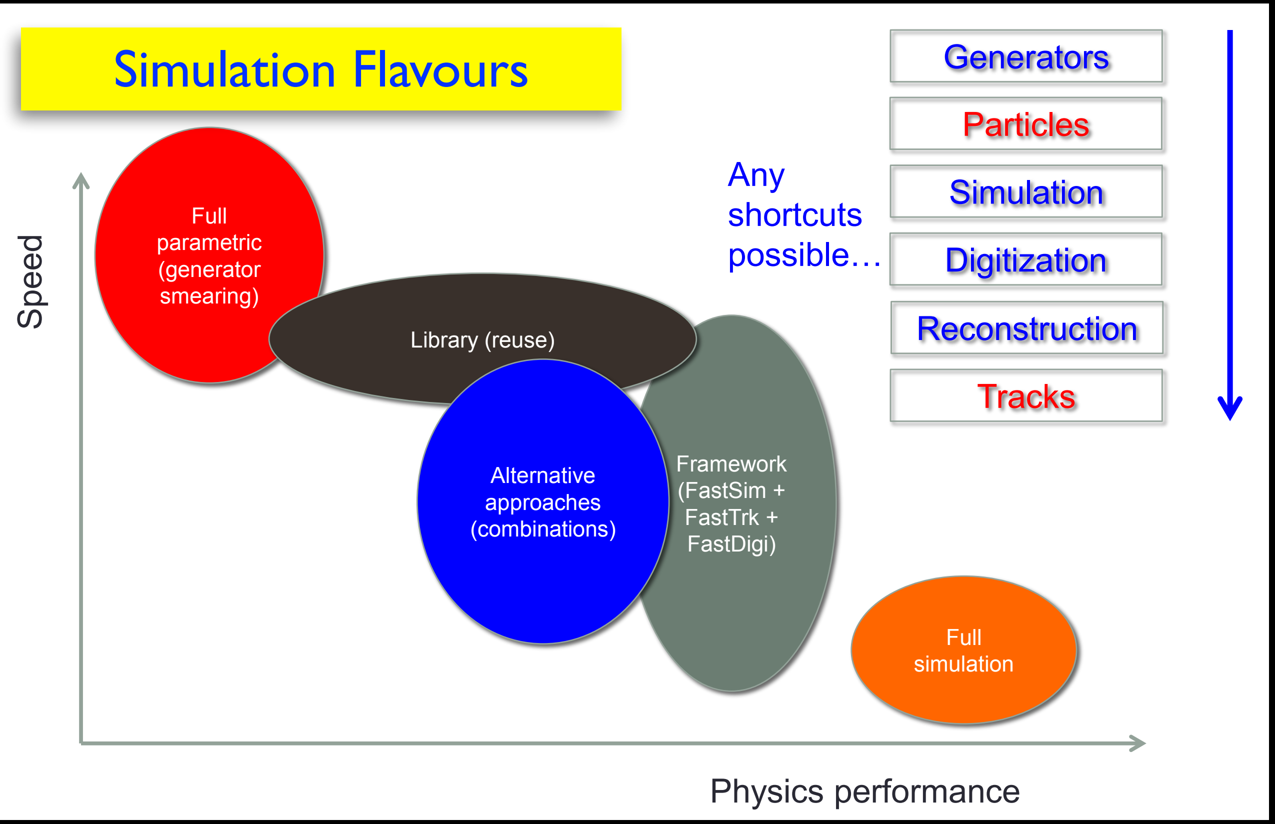

A.Salzburger

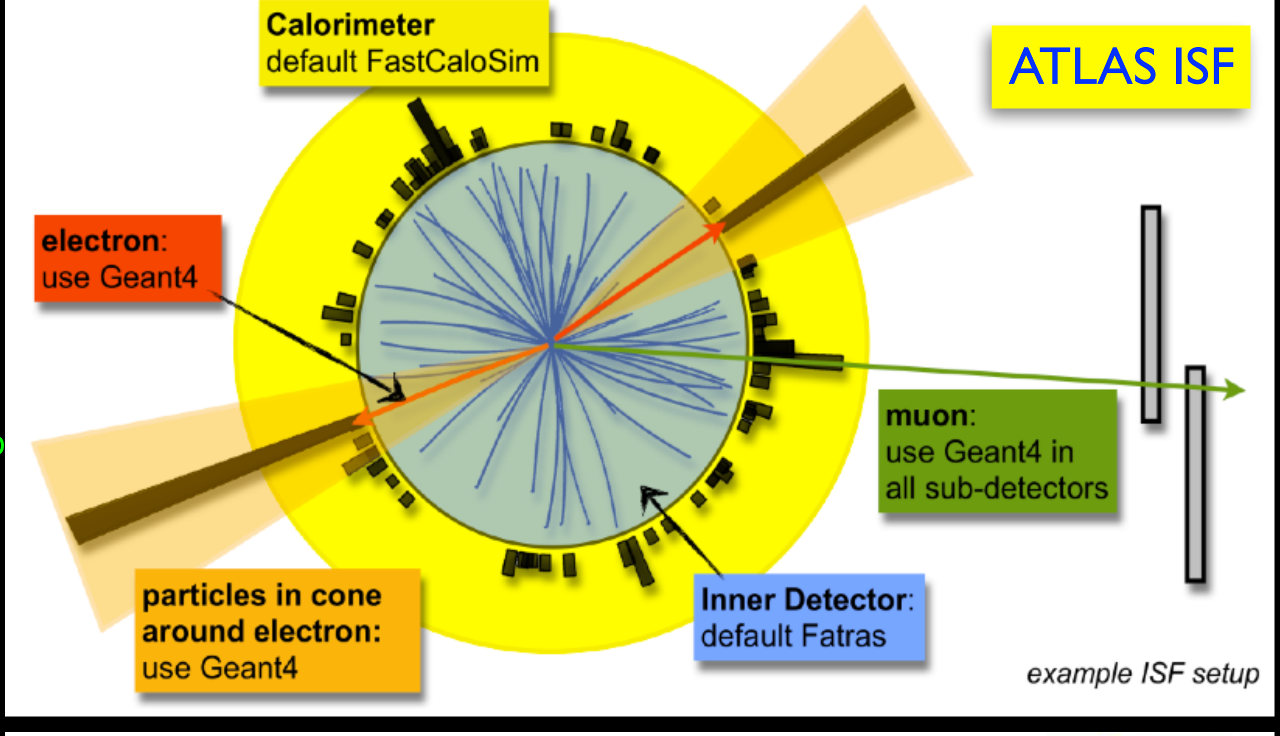

ATLAS fast software chain

HEP Software Foundation

upcoming workshop

http://hepsoftwarefoundation.org

- initiative to raise profile of HEP software projects
  - ➡ building upon existing and previous initiatives
    - hepfroge.org
    - Concurrency Forum
    - (less known) US HEP Forum for Computational Excellence
    - previous LCG Application Area
  - ➡ as well, existing HEP SW projects
    - Geant4, Root, ...
  - ➡ hopefully as well GRID software

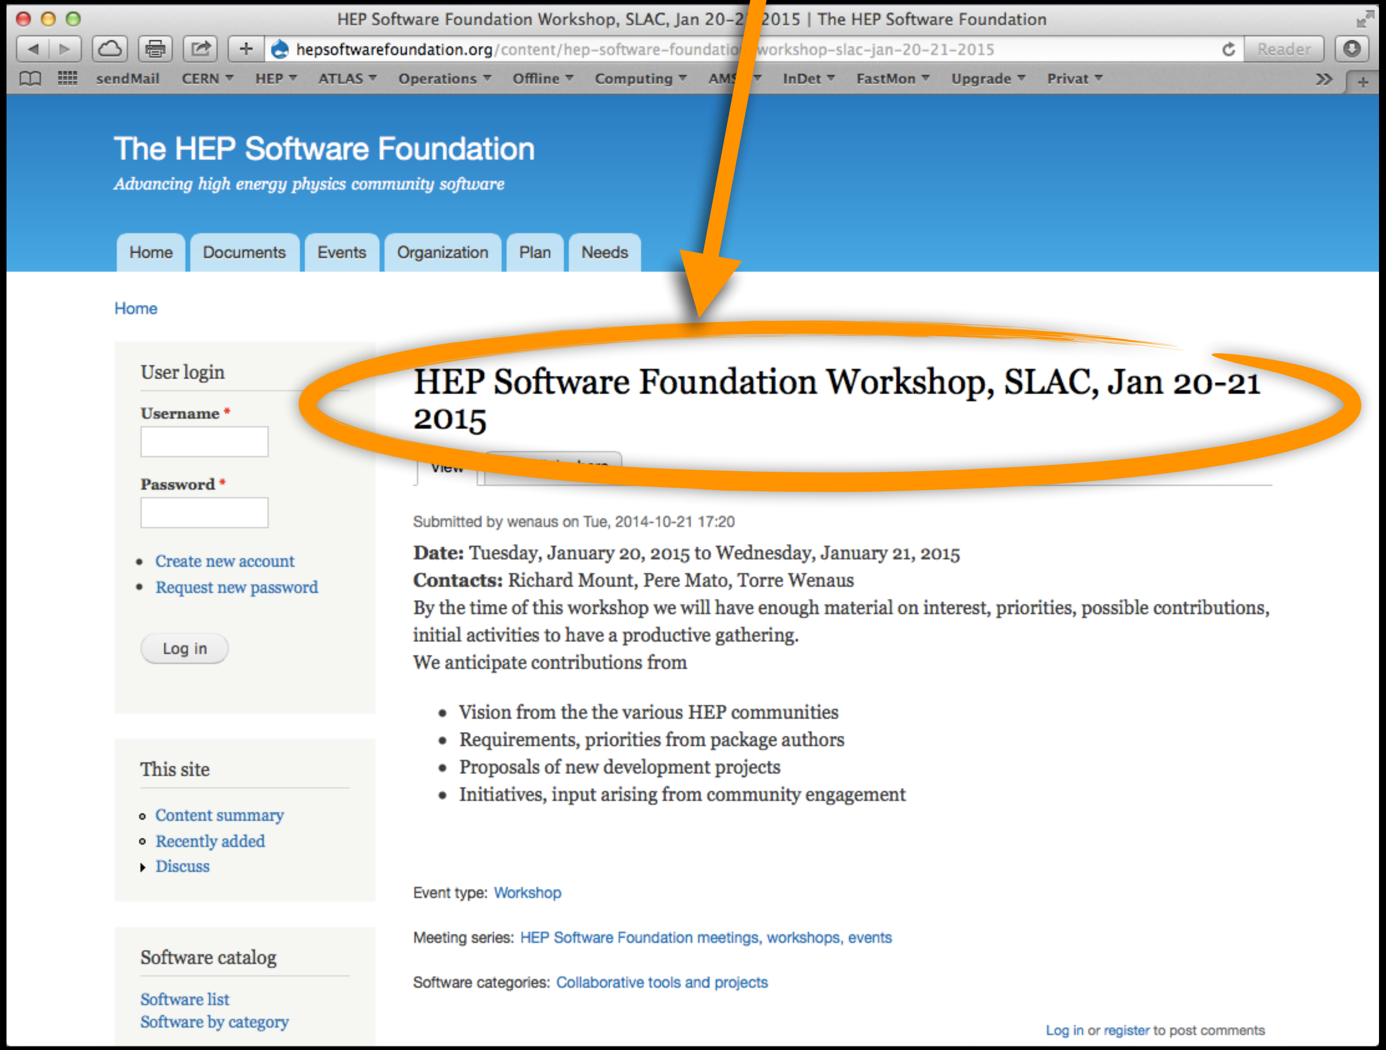

- foundation as a bottom-up approach
  - ➡ invite participation in projects across experiments and collaboration beyond HEP
  - ➡ hope to achieve synergies and bundle expertise on crucial technology developments

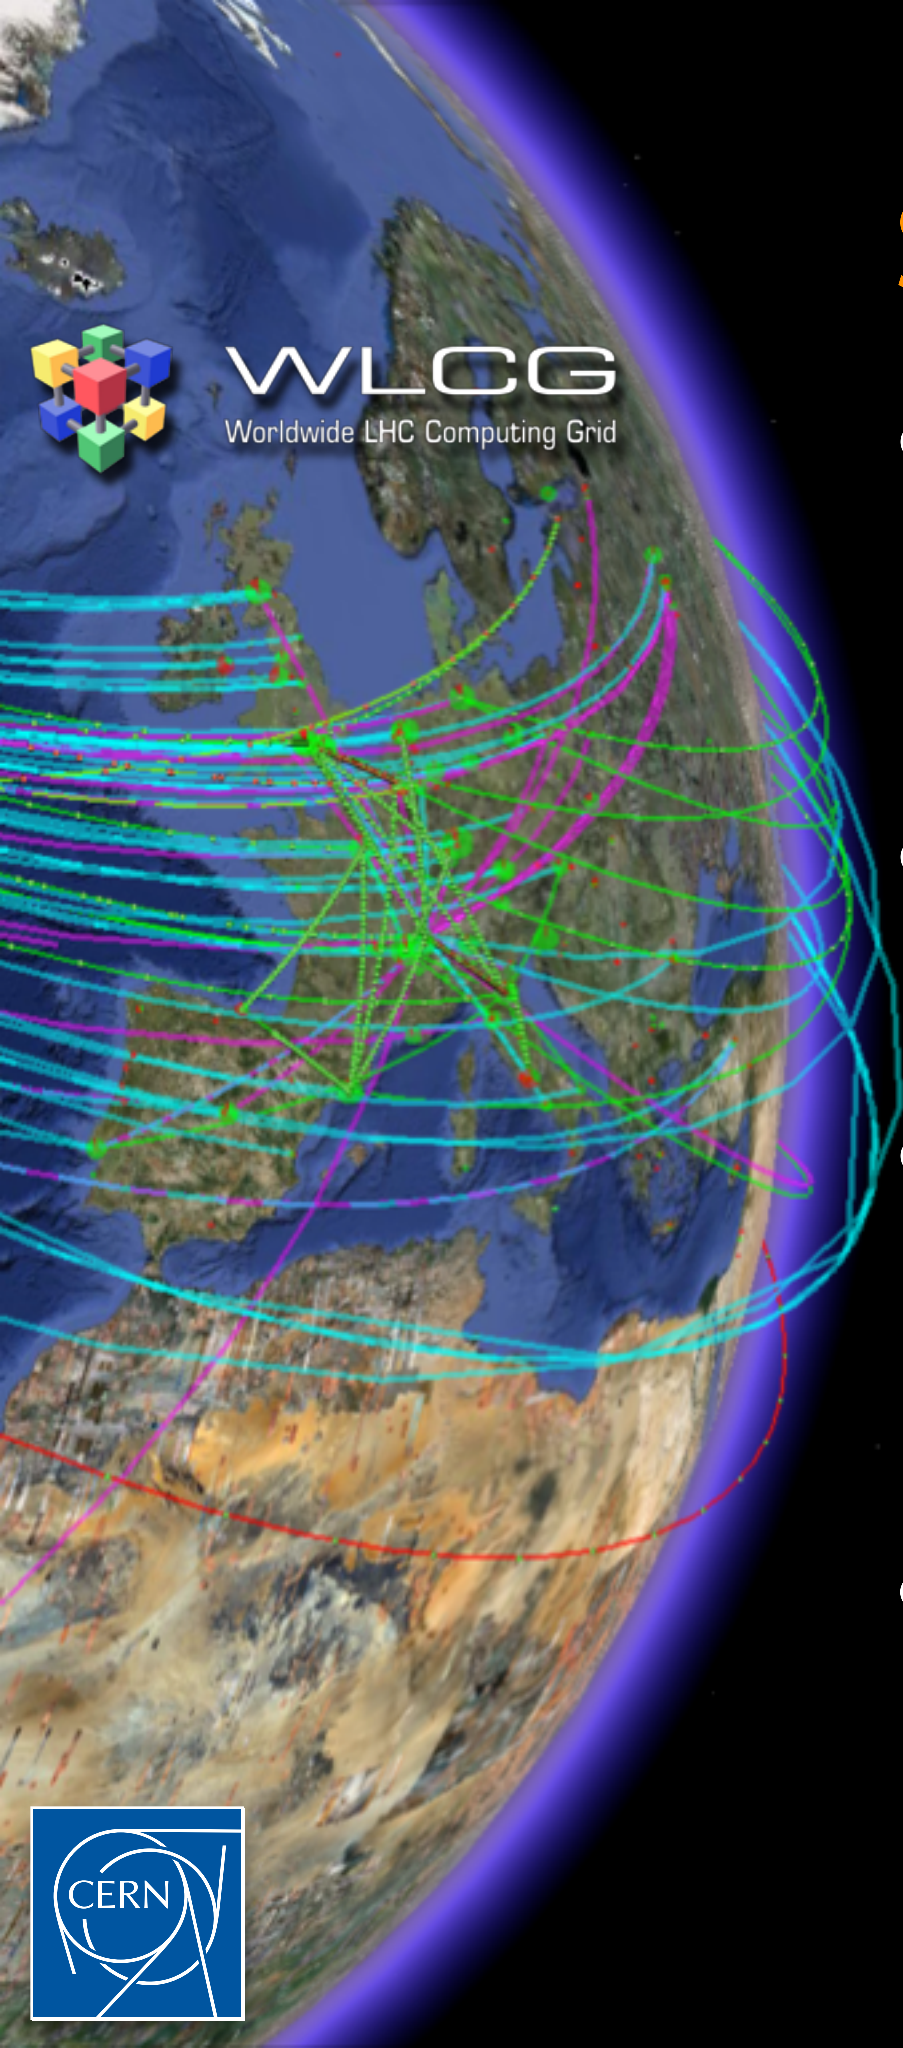

Summary

- facing the LHC computing challenge
  - ➡ the voyage started nearly 2 decades ago
    - from FORTRAN to GRID computing
  - ➡ it was a success story !
    - computing & software worked extremely well, enabling LHC physics program
- shutdown preparations for Run-2
  - ➡ first round of upgrades to software and computing
  - ➡ even higher pileup and limited computing resources
- many more challenges ahead
  - ➡ Phase-1 and Phase-2 detector upgrades
    - pileup will rise further, up to 140-200 for HL-LHC
  - ➡ IT technologies are changing dramatically
    - more heterogeneous, more complicated to program
- and finally:
  - ➡ to fully explore the LHC potential we not only need the computing resources, but our community at large needs as well to find adequate manpower for the necessary software developments !

Acknowledgements

- I'd like to thank:
  G.Duckeck, G.Quast, G.Stewart, E.Lançon, D.Froidevaux, P.Clark, K.Jakobs

Backup

LHC history

- 1980: LEP not yet built, but physicists think about the possibility to re-use the tunnel for a hadron collider;
- 1984: Glimmerings of the LHC (2x5...9 TeV, symposium in Lausanne) and SSC (2x20 TeV);
- 1988: SSC approved (Waxahachie, Texas);
- 1989: First collisions in LEP and SLC, R&D for LHC detectors begins;
- 1993: SSC construction cancelled;
- 1994: LHC approved (start in 2005)
- 1995: Discovery of the top quark at Fermilab;
- 1996: ATLAS and CMS approved. 1997: ALICE, 1998 LHCb;
- 2000: end of LEP running, no Higgs yet;
- 2005: first cosmic seen in the ATLAS pit;
- 2006: new CERN accelerator control centre ready;
- 2007, June: the last dipole magnet lowered to the tunnel, first sector @-271 deg;
- 2008: LHC start;
- 2008, 10. September - 10:28: first full turn of a proton bunch
- 2008, 19. September - failure during powering tests
- 2009, 23. November: protons collide again! (30. November: 1.2 TeV collisions)
- 2010, 30. March: first high energy proton collision (3.5 TeV)
- 2012, 4. July: Higgs-like particle seen!
- 2012, 8. November: First observation of $B_s^0 \rightarrow \mu^+\mu^-$; the Standard Model rules....

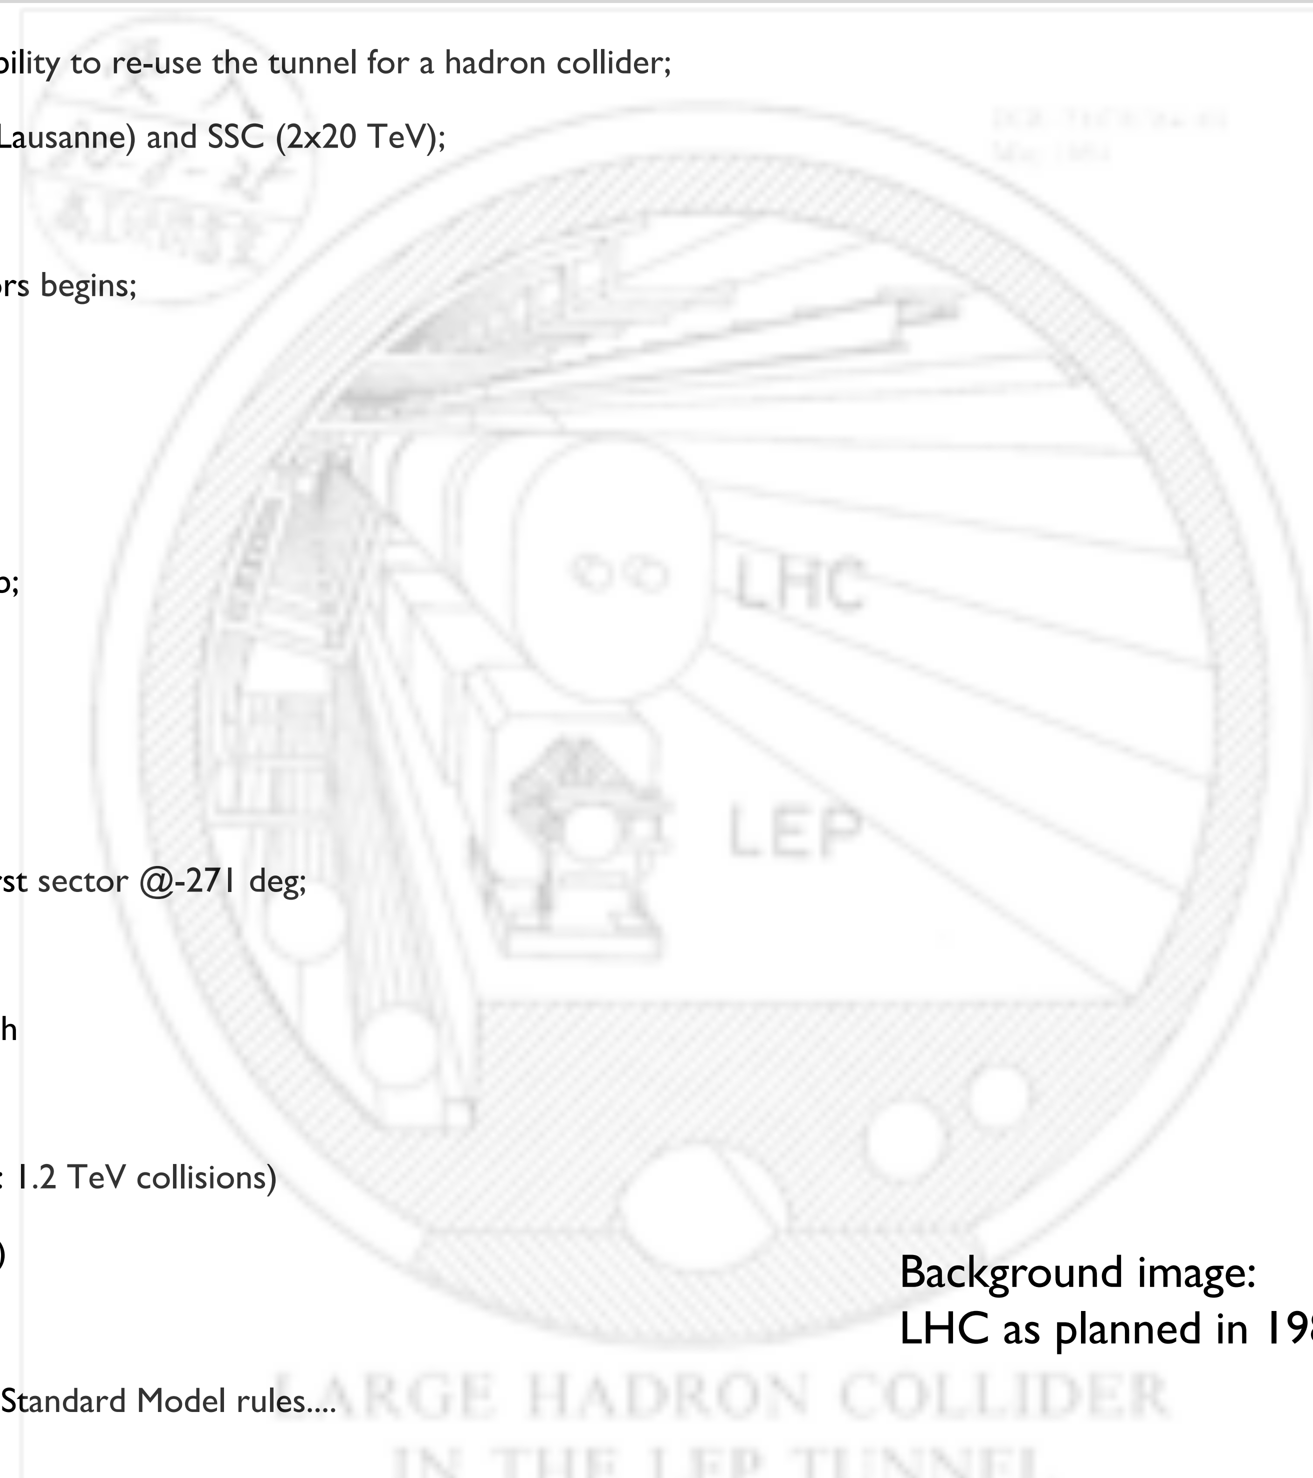

Background image:
LHC as planned in 1984

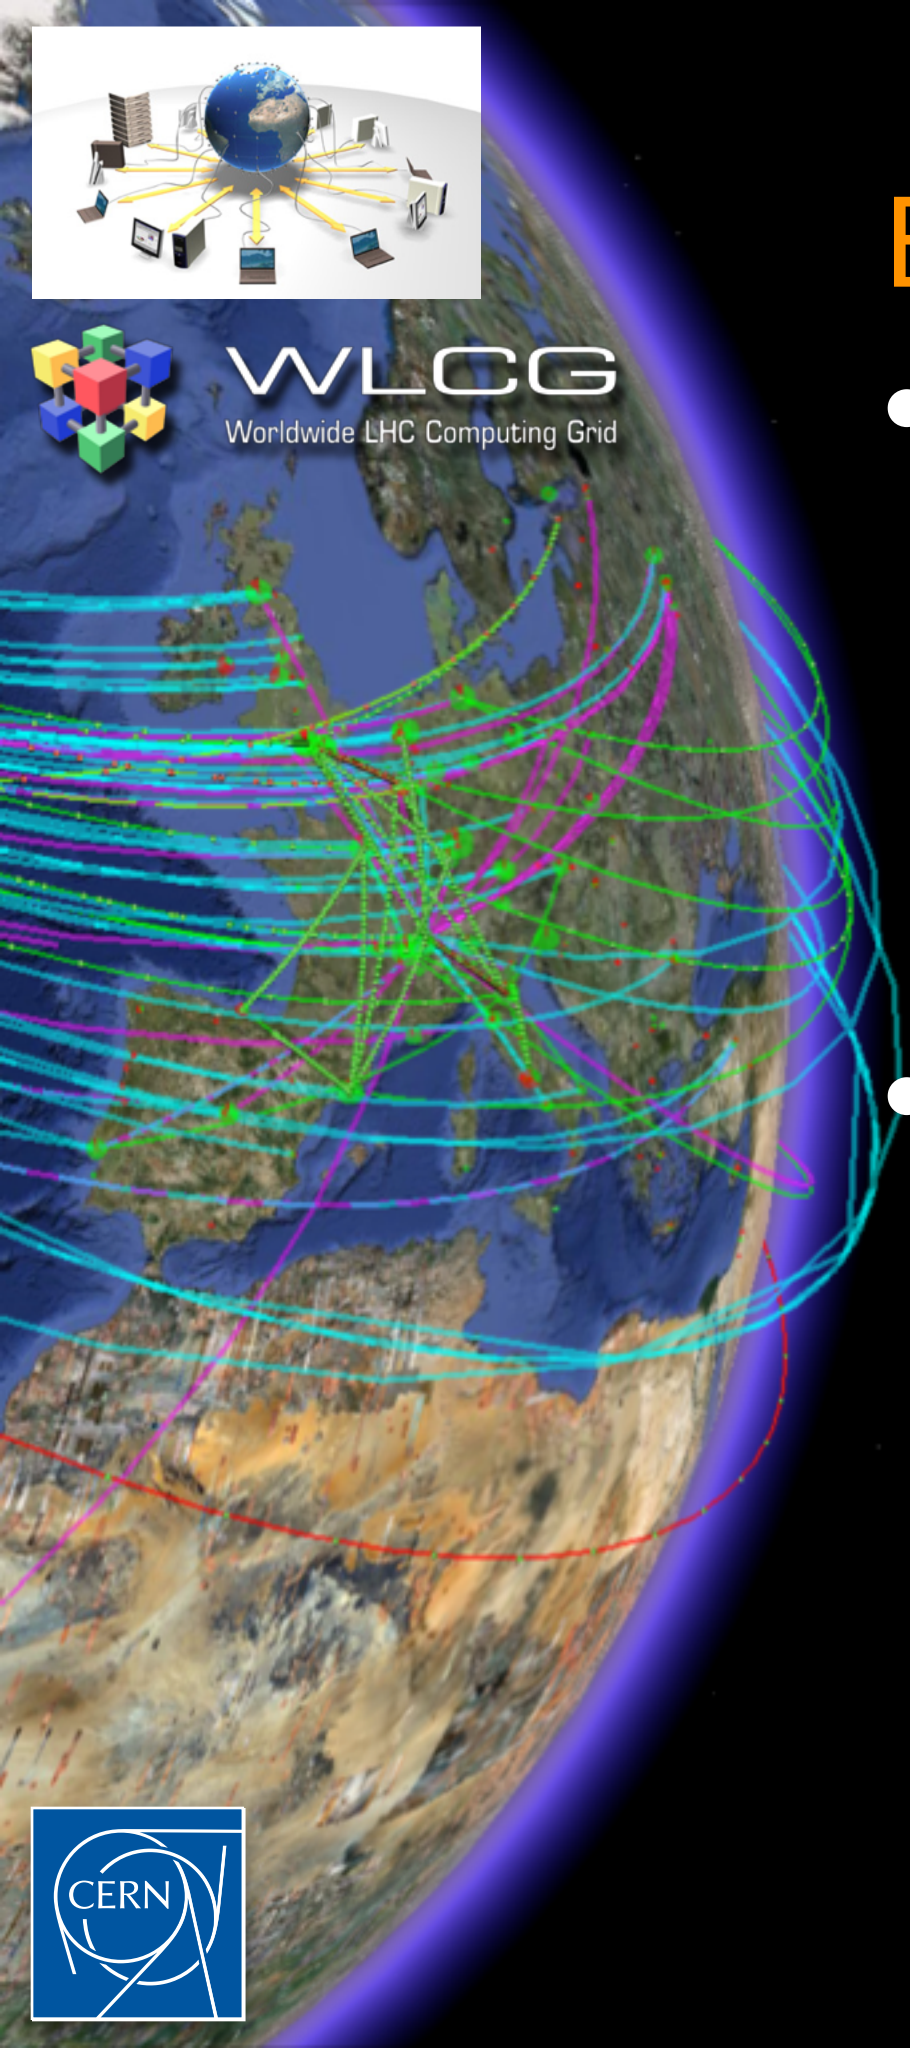

WLCG Worldwide LHC Computing Grid

SKIP SLIDE

European Grid Initiative (EGI)

- EGI federation participants:
  - ➡ National Grid Initiatives (NGIs)
    - funding via NGIs
  - ➡ international research organisations
- GRID in Germany
  - ➡ contributions to worldwide WLCG:
    - 15% to Tier-1s (KIT)
    - 10% to Tier-2s
      (DESY, GSI, MPI München, 5 Universities: Aachen, Freiburg, Göttingen, LMU, Wuppertal)
  - ➡ within Germany
    - 40% at Tier-1
    - 60% at Tier-2

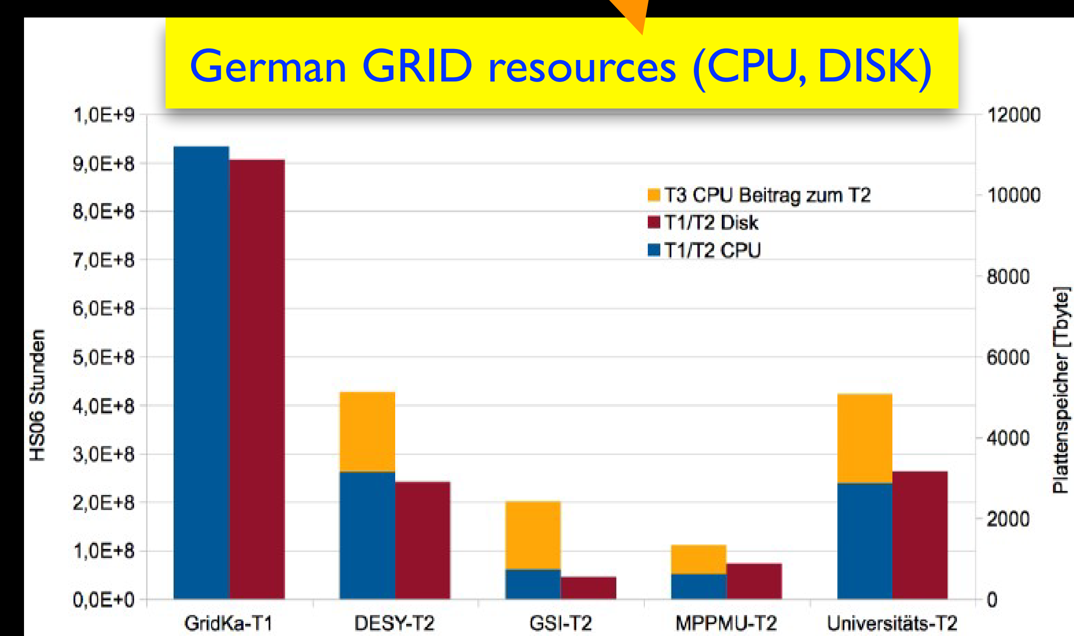

G.Quast. DPG 2014

see talk of T.Kress

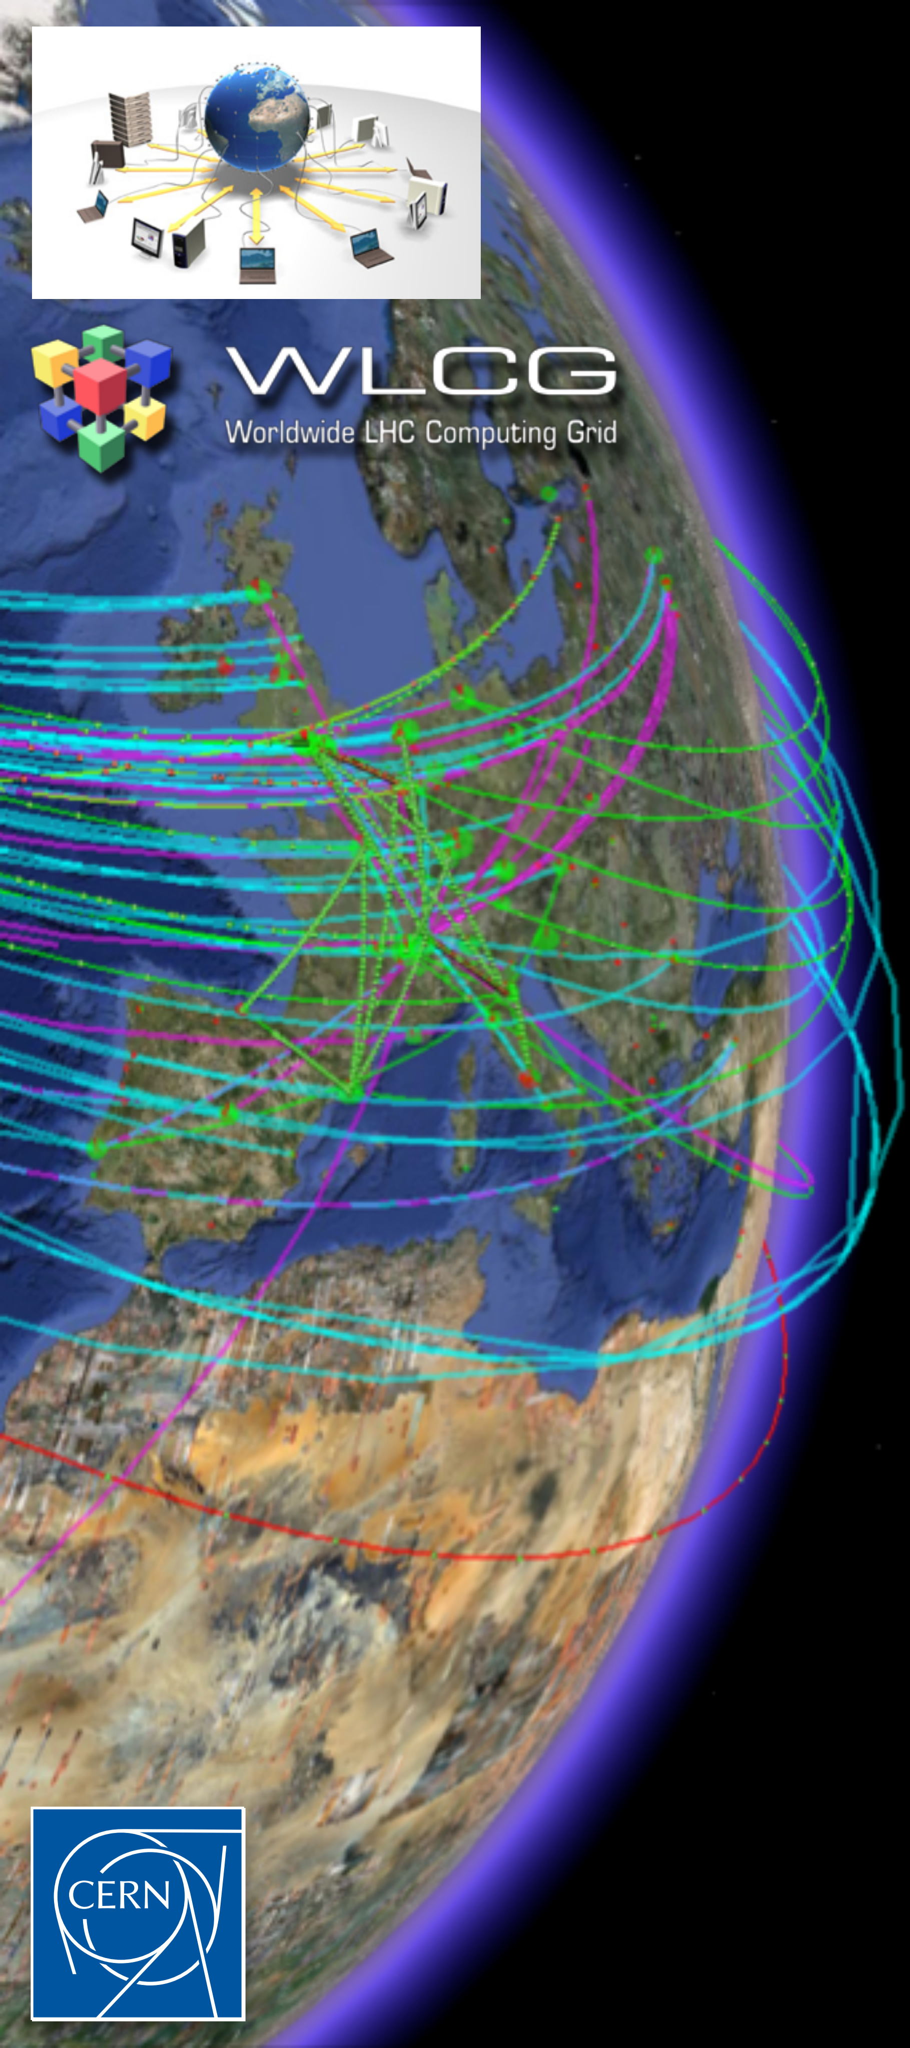

SKIP SLIDE

EGI is a shared Infrastructure

- a few hundred virtual organisations (VOs) from several scientific domains:
  - ➡ astronomy & astrophysics
  - ➡ civil protection
  - ➡ computational chemistry
  - ➡ comp. fluid dynamics
  - ➡ computer science/tools
  - ➡ condensed matter physics
  - ➡ earth sciences
  - ➡ fusion
  - ➡ high energy physics
  - ➡ life sciences
  - ➡ ...

(http://operations-portal.egi.eu/vo/search)

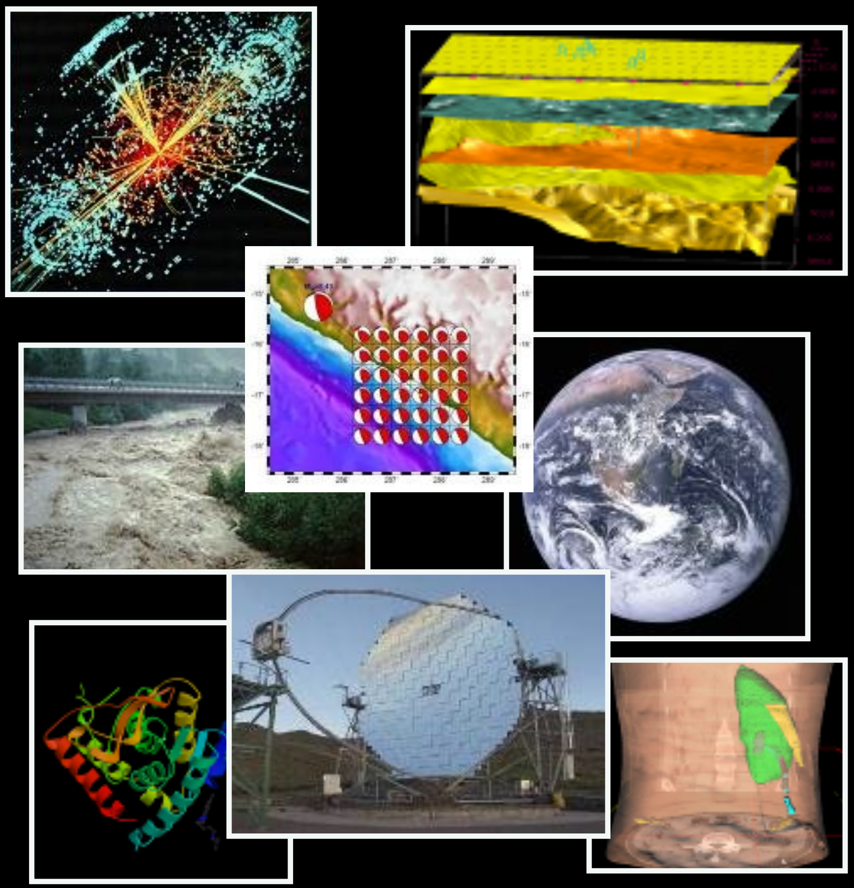

- organisations are joining continuously
  - ➡ e.g. fishery (I-Marine)

SKIP SLIDE

Teddy Todorov (1966-2014), et al.

Tracking Software Concepts

- tracking for LHC luminosities
  - ➡ early years informal collaboration by CMS and ATLAS
    - R&D on fitting techniques, STEP propagation, ...
    - later series of LHC alignment workshops
  - ➡ novel tracking geometries with embedded navigation
    - reduced volume complexity
    - bended material on simple surface shapes
    - much faster than generic voxelisation a la Geant4
  - ➡ speed up reconstruction and fast tracker simulation

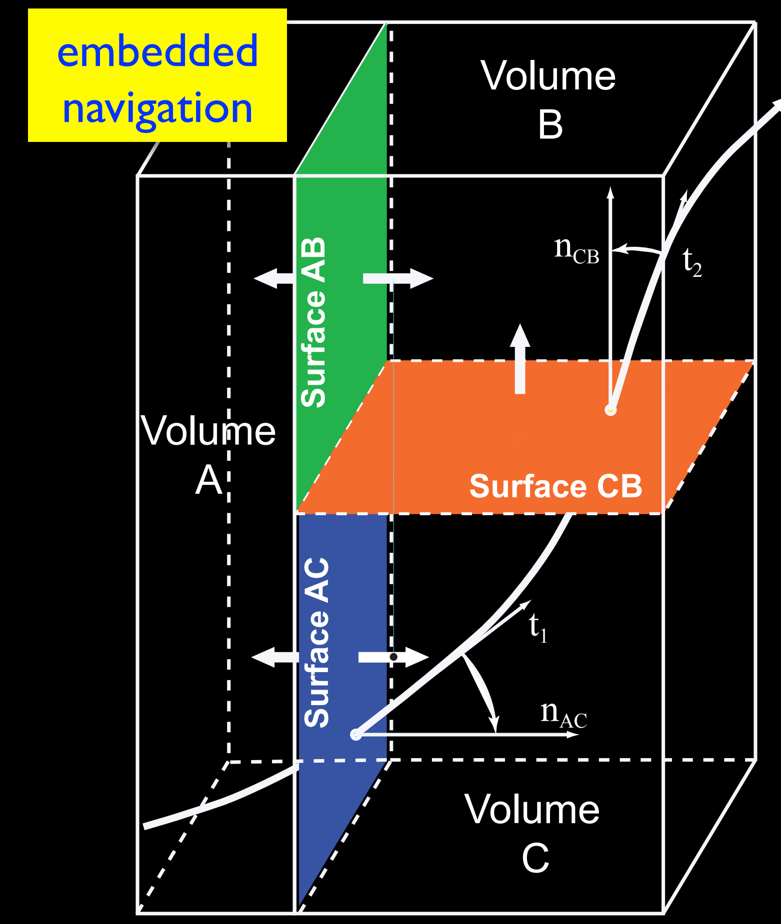

A.Salzburger

- material description of LHC detectors
  - ➡ we knew ATLAS and CMS trackers would be heavy
  - ➡ measure components precisely
    - interplay hardware and software people
    - we will see, it payed off later !

CMS	measured	simulation
active Pixels	2598 g	2455 g
full detector	6350 kg	6173 kg

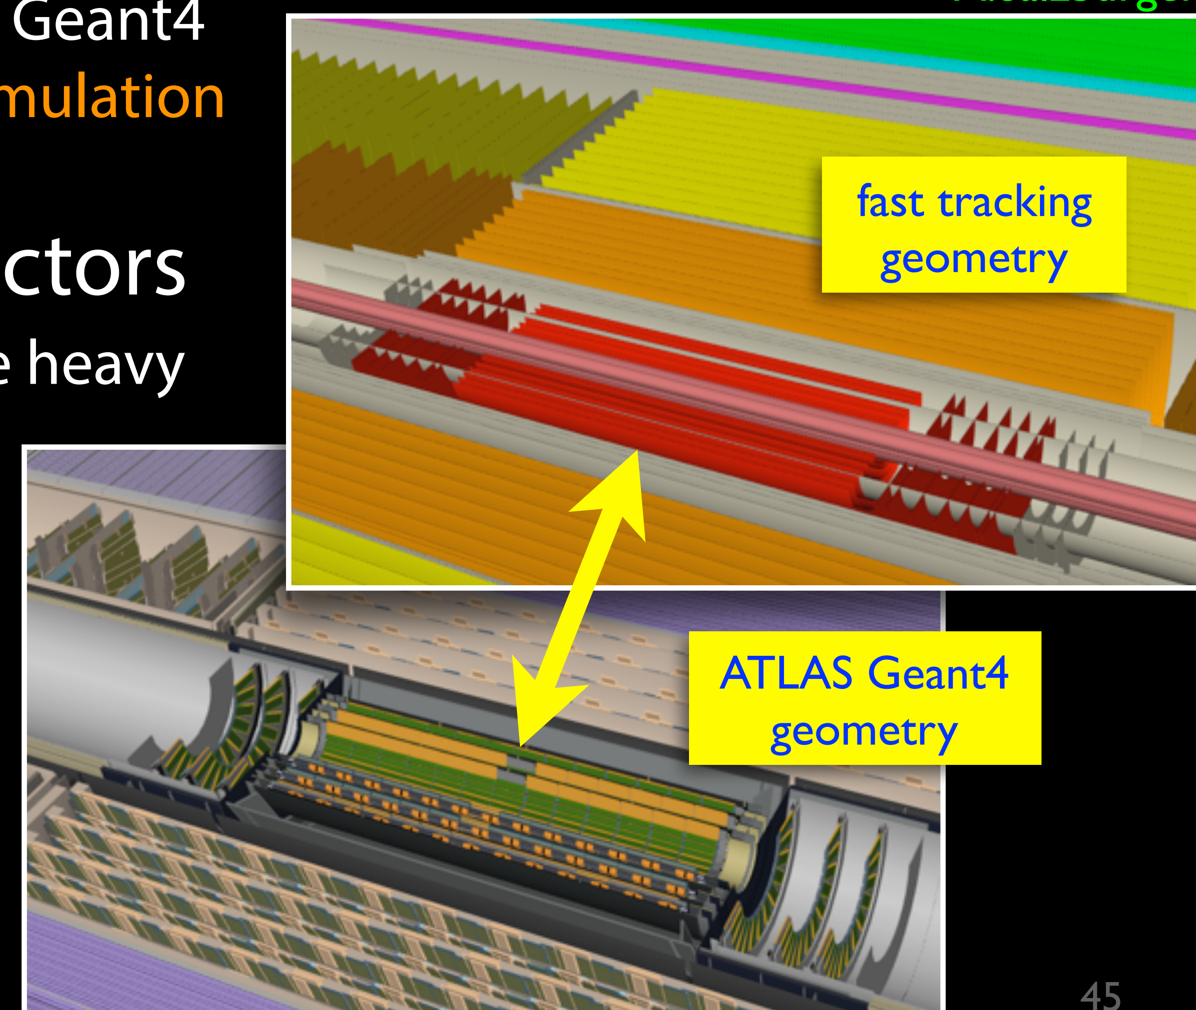

SKIP SLIDE

Software and Upgrade

- ECFA HL-LHC workshop series
  - ➡ software and computing part of the process
  - ➡ across all 4 experiments

- numerous upgrade goals
  - ➡ boost physics reach, including
    - LHCb all software trigger
    - online data compression for ALICE
  - ➡ keeping physics acceptance at higher pileup
    - ATLAS and CMS will increase trigger rates, especially for single lepton triggers
    - even higher pileup will require more resources (CPU, memory, disk)
  - ➡ upgrade software and computing itself
    - follow technology evolution

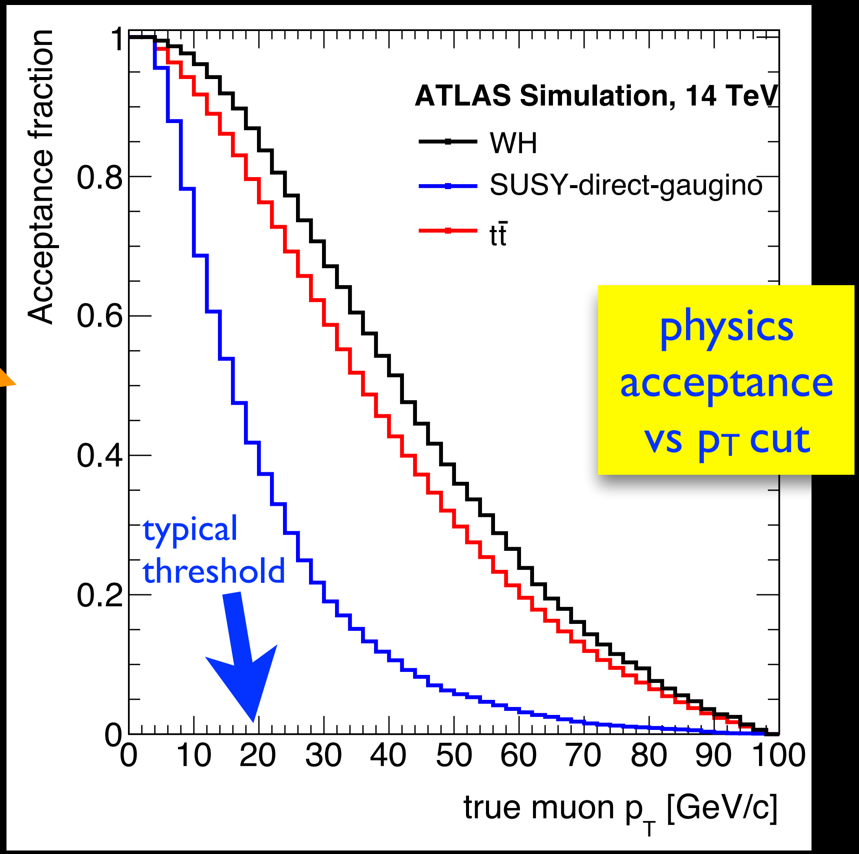

SKIP SLIDE

Framework Support for Concurrency

- Gaudi-Hive
  - ➡ parallelism for the Gaudi framework
    - used by LHCb and ATLAS
  - ➡ Intel TBB toolkit for multi-threading support
    - event and algorithm level parallelism
  - ➡ demonstrators show encouraging results
    - but tracking needs finer-grained parallelism

Gaudi-Hive scheduling model

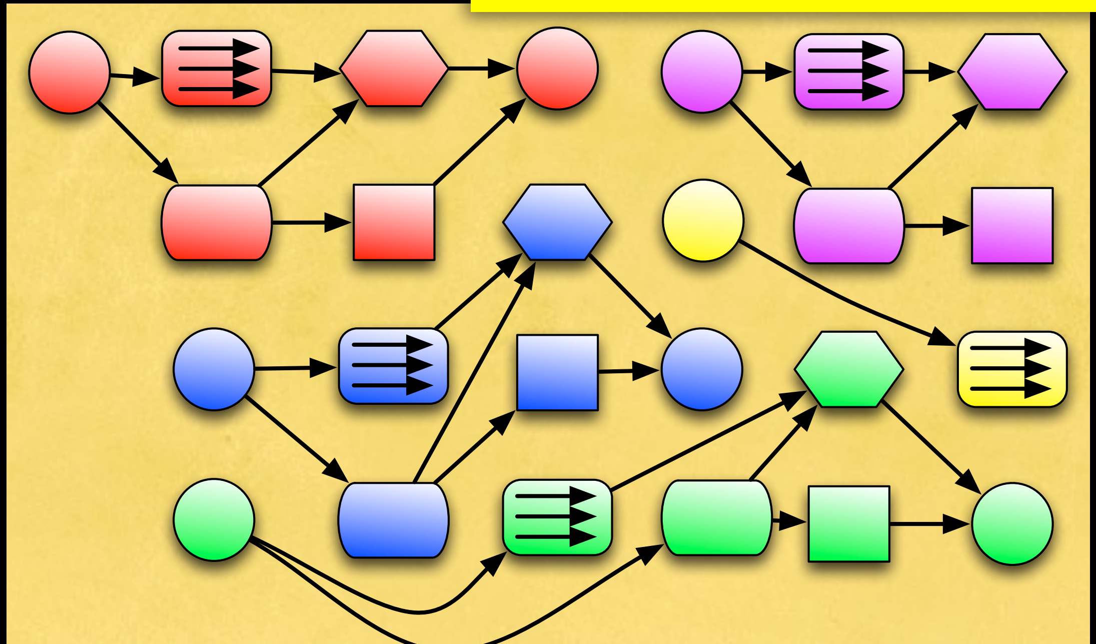

colours represent events, shapes different algorithms

- CMSSW multi-threading
  - ➡ framework splits into global (transitions) and multiple streams (event processing)
    - underlying toolkit is as well Intel TBB
  - ➡ excellent scaling and memory improvements observed on 16 core machines
    - 99% of CMS reconstruction is now thread safe

CMSSW multi-threading