CERN and Scientific Computing

Massimo Lamanna

CERN
Information Technology Department
Experiment Support Group

1960: 26 GeV proton in the 32 cm CERN hydrogen bubble chamber

1960: IBM 709 at the Geneva airport

1972: Experiment at the CERN PS (East hall)

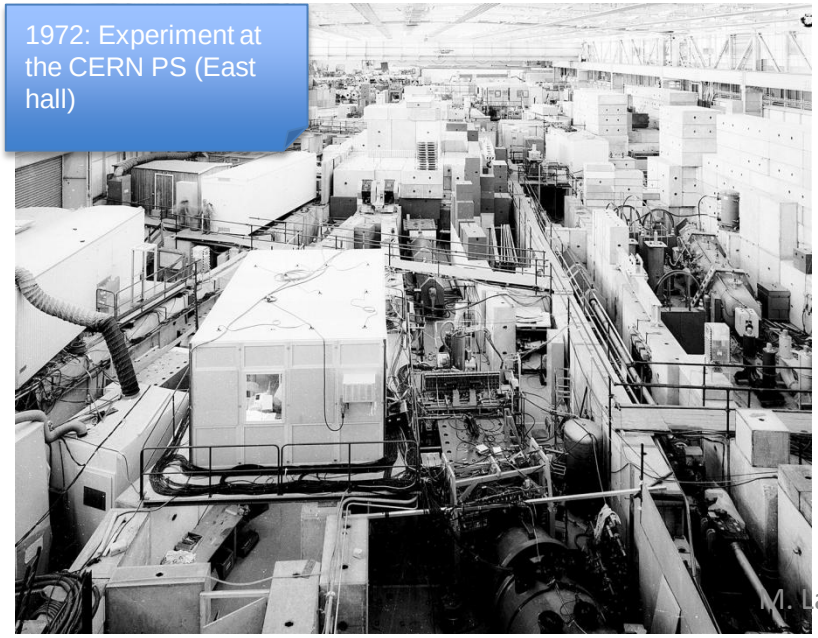

1972: Installation of the CDC 7600

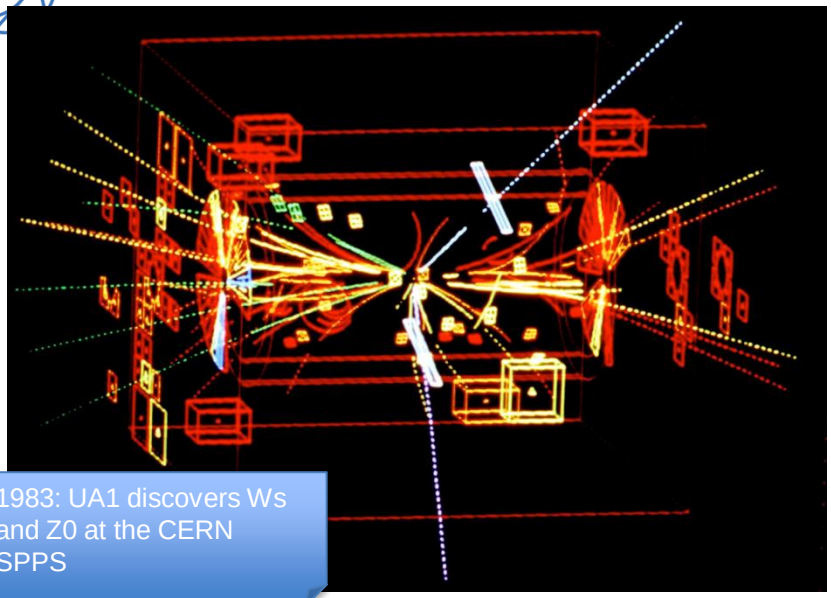

1983: UA1 discovers Ws and Z0 at the CERN SPPS

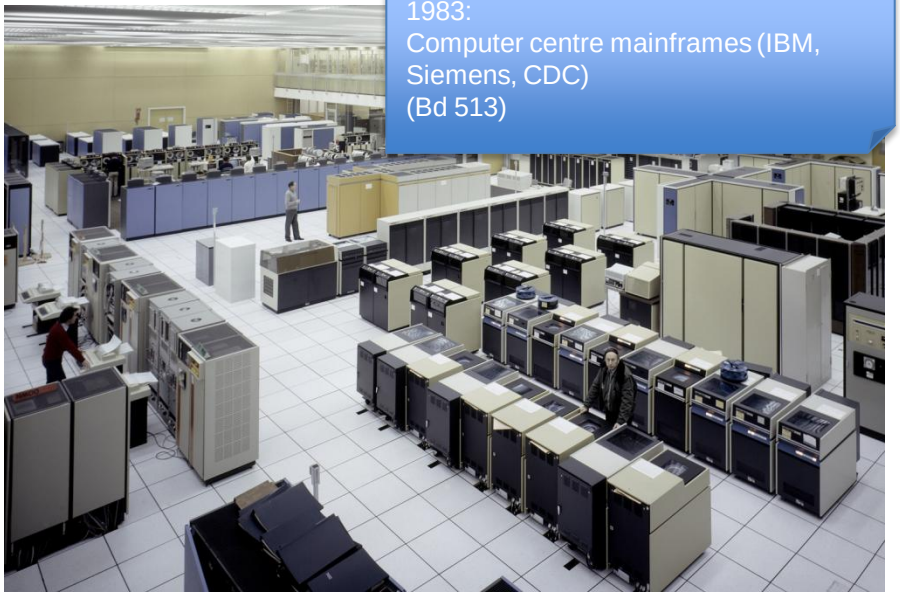

1983:
Computer centre mainframes (IBM, Siemens, CDC)
(Bd 513)

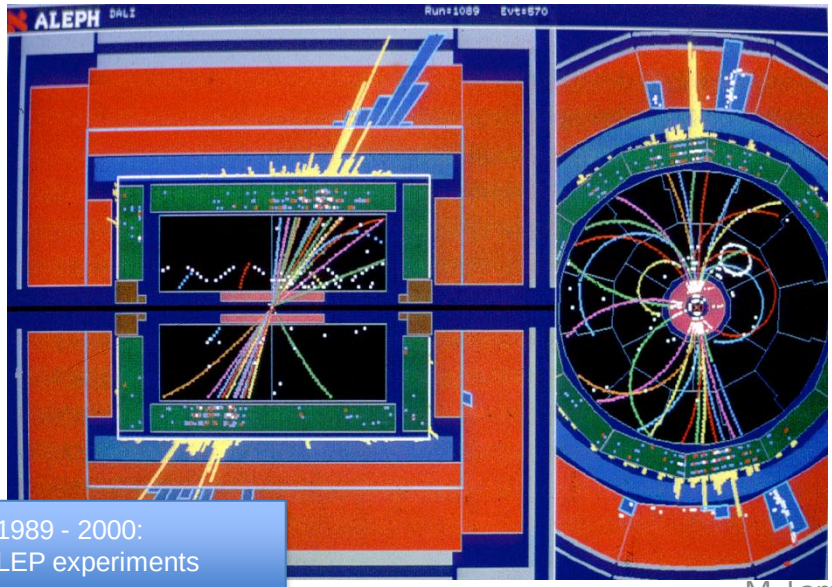

1989 - 2000:
LEP experiments

1990: The first web server: this machine was used by Tim Berners-Lee in 1990 to develop and run the first WWW server, multi-media browser and web editor

Early days of COMPASS (1998)

A new experiment being built…

The SM2 magnet…

… with new computing challenges

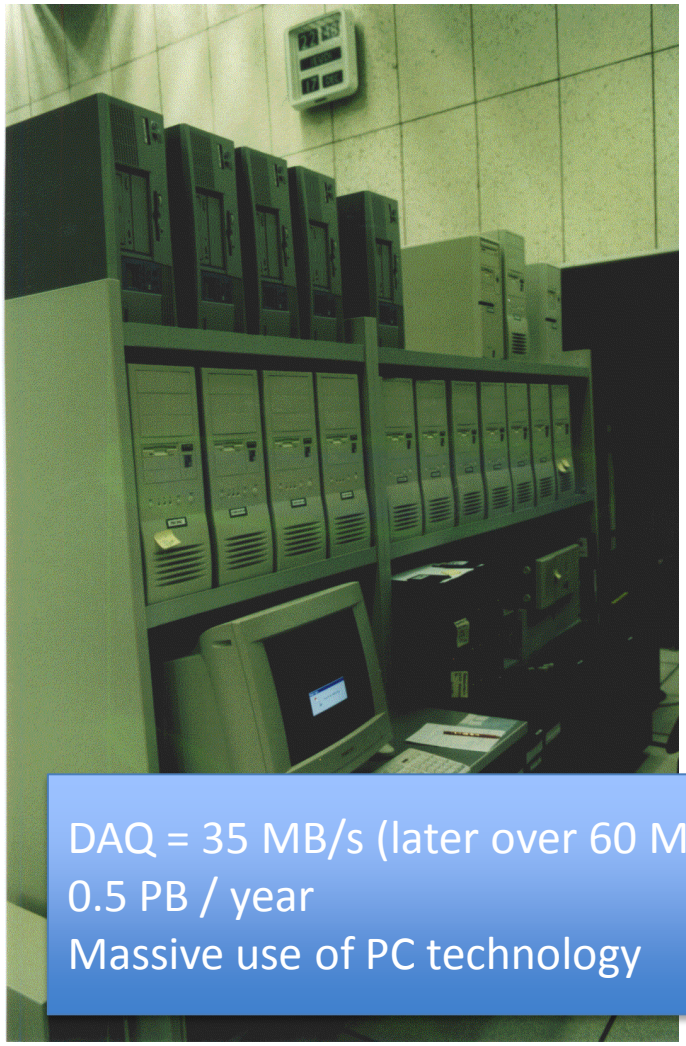

DAQ = 35 MB/s (later over 60 MB/s)
0.5 PB / year
Massive use of PC technology

The COMPASS Computing Farm…

COMPASS Event Display

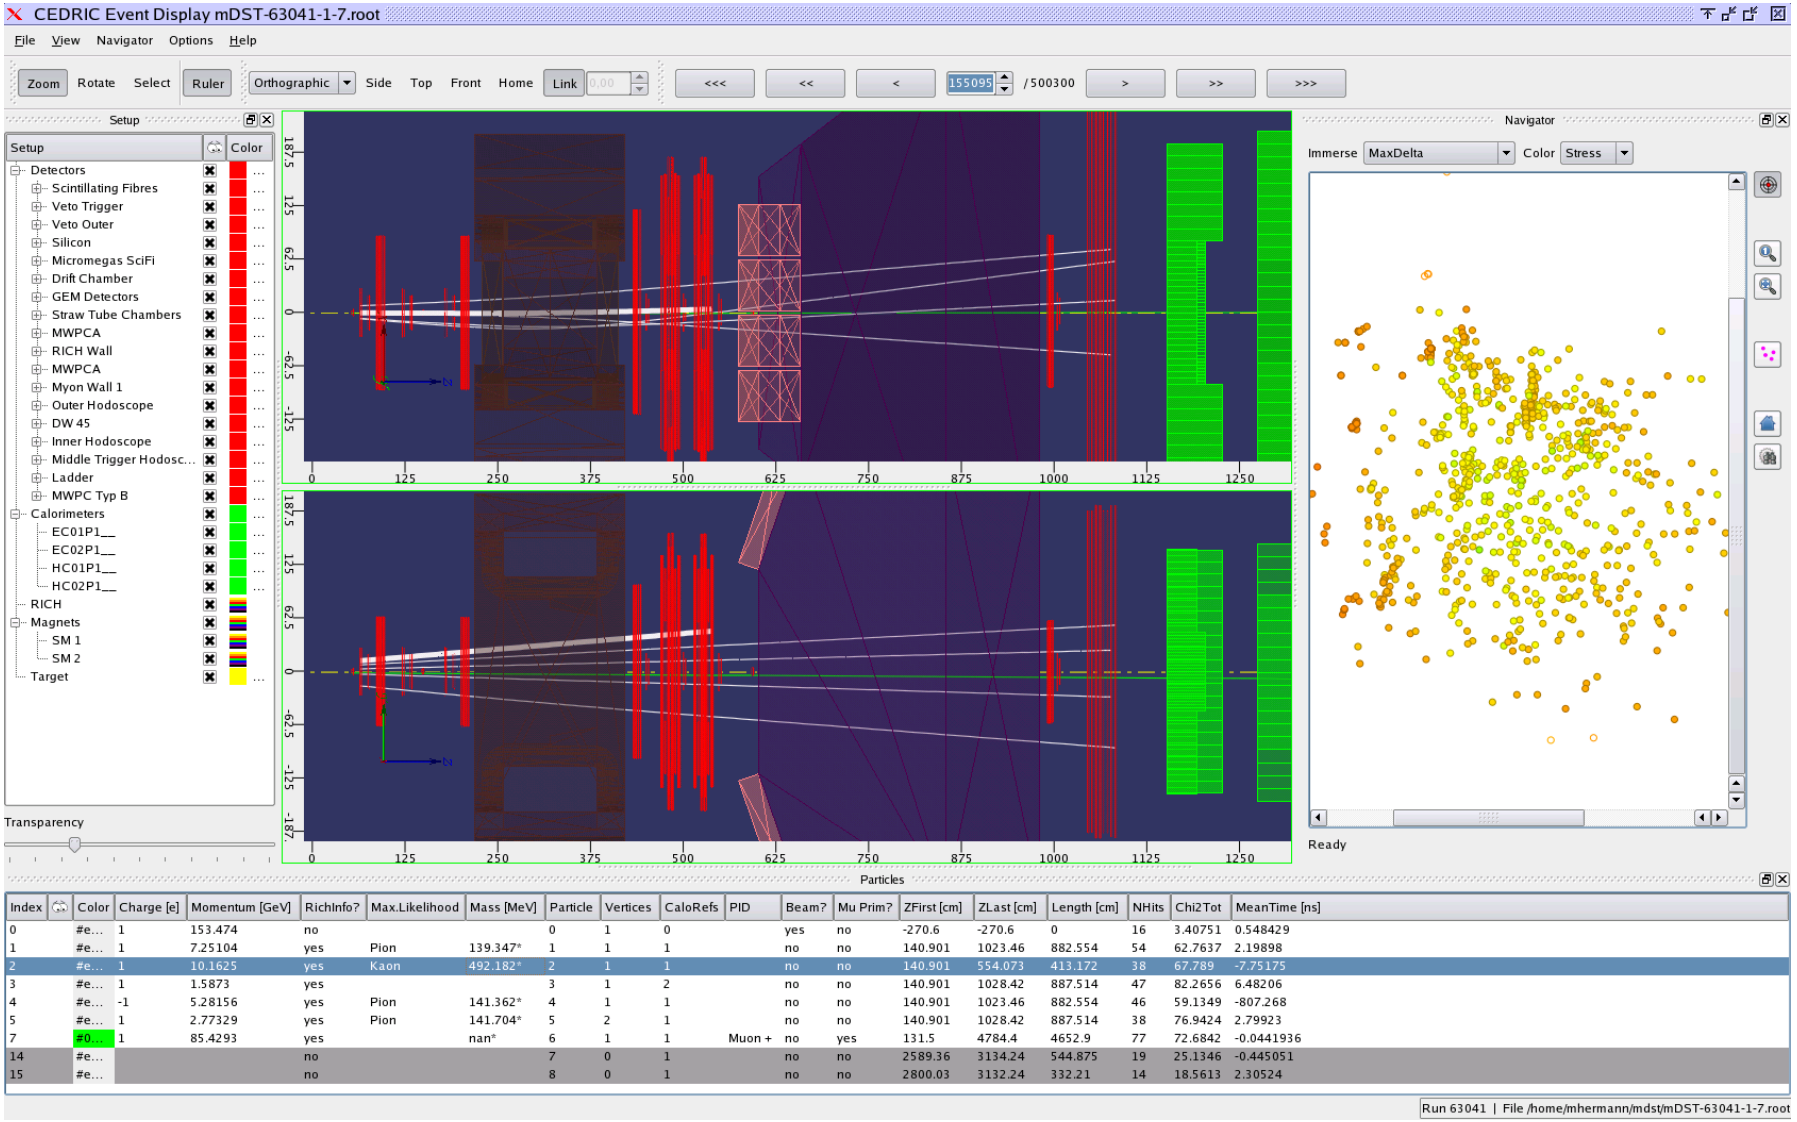

20th of November 2009

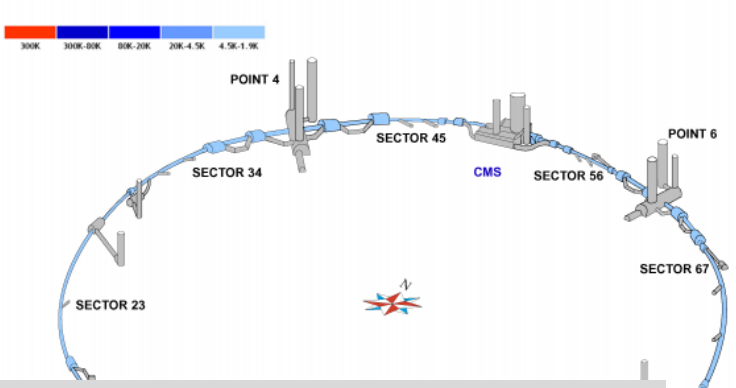

27 km tunnel (2000+ superconductive dipoles)
450 GeV injection energy
7 TeV max beam energy
3.5+3.5 TeV (March the 30th world record)

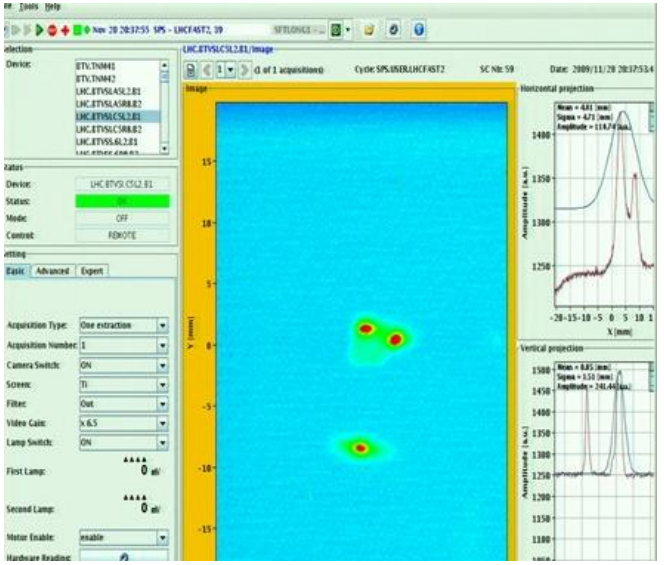

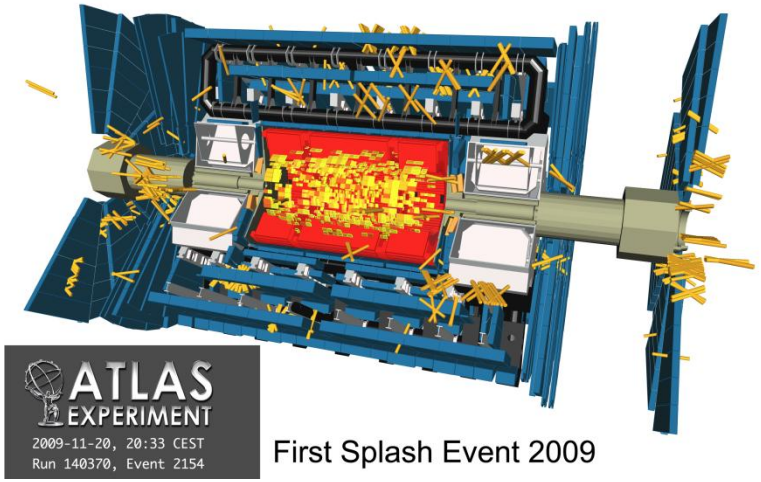

First Splash Event 2009

Very recent LHC pp collisions!

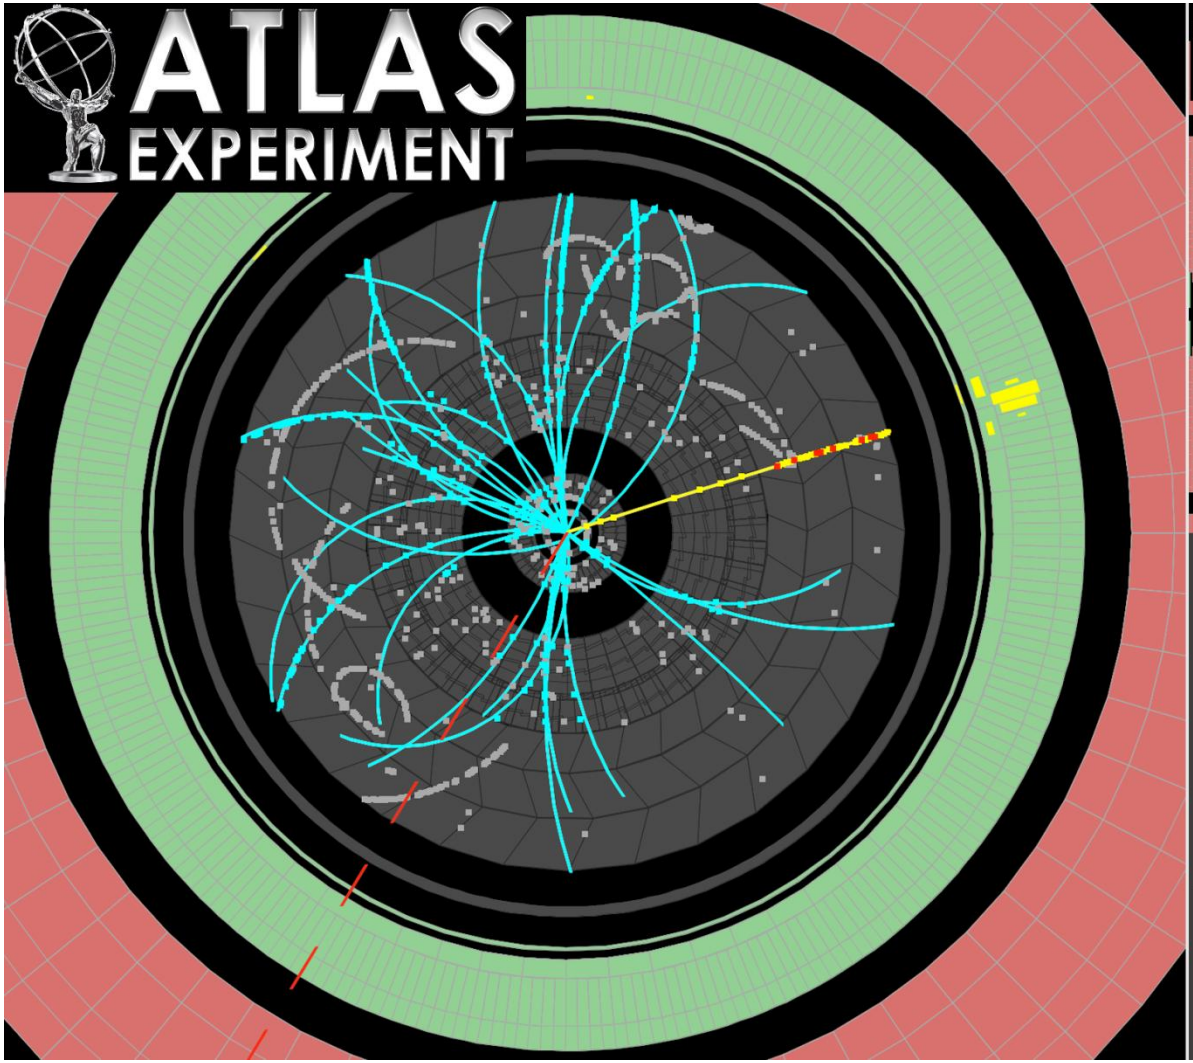

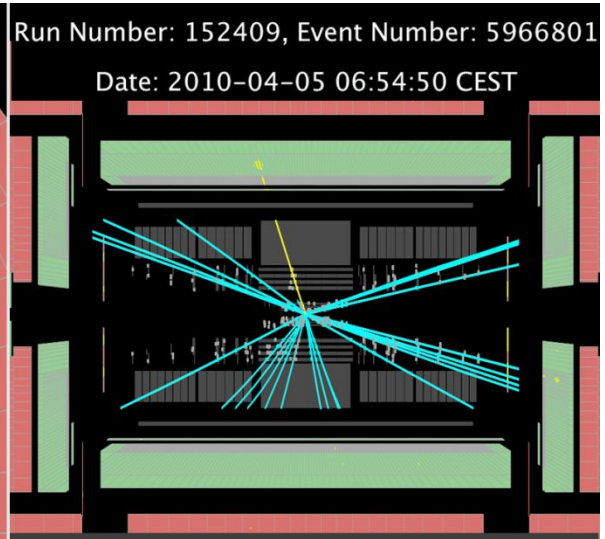

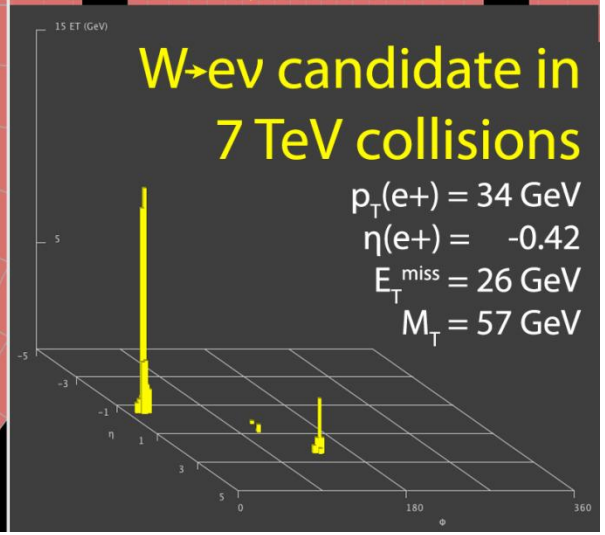

Interesting facts (CERN Computer centre)

- Number of machines
  - About 4,500 batch (18,000 CPUs)
  - About 3,000 disk servers (50,000 hard drives)
  - Several hundred tape servers, console head-nodes, database and Grid servers etc.
- Storage Capacity
  - 5+ PB disk
  - 25+ PB tape (IBM and Sun/StorageTek)
  - There will be an additional 15 PB each year needed for the LHC data (3*10^6 DVDs!)
- Network Capacity
  - Connection at 10 gigabits/second to each Tier 1, plus backup, plus regular (firewalled etc) internet
  - Speed record: 1.1 TB in less than half hour (CERN-CalTech)
- Number of staff
  - CERN: ~2700 ; IT department (computing) ~250 and ~200 on shorter-term Grid projects

Early 2009 data
Fast ramp up going on for the LHC start

Service/Activity	Description	Data
Disk and Tape	Provide storage capabilities	1500 disk servers 5 PB disk space, 16 PB tape storage
Network Campus	Provide local CERN area network service	The Network core has a capacity of 4.8 Tbps The total number of network ports is: >1000 10 Gigabit ports >70000 Gigabit ports feeding 34000 hosts
External Networking	Provide connectivity and infrastructure with other 11 T1s	Out-coming Internet Back bone up to 60GB/s to 11 centers Speed record:1.1 TB transferred in less than half hour (Caltech-CERN)
NICE PC	PC cluster	5,500 active NICE PCs 1,500 Macs 1,100 CMF packages
Computing Facility	Provide local & Grid Computing Power	4500 nodes installed 16K CPUs available Up to 10K concurrent jobs
Computing Center cooling system	The system to cool the centre	more than 500,000 m3/h cold air to cool PCs three 1.5 MW chillers
CVS	Concurrent Version System	over 300 Software Projects and over 3000 users
Messaging services	Email, Ldap, Listbox, News, Fax, Antispam	17,015 mailboxes 7000 lists 2.8TB data 2.5 Million messages/day (~98% SPAM)
Web Services	Web Services , Search Engine, Verisign	- 8852 sites - including 930 sharepoint sites - 6 million hits per day - Bandwith 2 Tbyte per day
WLCG	Coordinating and operating WLCG grid activities	30,000 simultaneously 1 PB /month each ATLAS, CMS 1 GB/s (1.6 GB/s peak) to 12 sites 160 VOs
AFS	Andrew File System (distributed file system for CERN users)	15,000 users 27,000 volumes 25TB allocated

What makes High-Energy Physics so special (also in computing)?

- High rates
  - 100-1,000 MB/s over long periods of time (10+ years, ~10^7 s/year i.e several months a year)
- High data volumes
  - Rates + derived types + replications → o(10) PB/year/experiment
    - Lot of files, relatively small DB (compared to the file storage)
- Several 10^5 PC cores to reconstruct, simulate and analyse the data
  - Reconstruction 10-100 s range on one PC
- Large distributed collaborations
  - O(1000) physicists
  - Everyone ("everywhere") gets in touch with data, in general with custom applications (based on common framework). High concurrency
  - Human activity, "interactive": need to minimise the latency

LHC Computing Grid

Worldwide infrastructure
(EGEE + OSG + NDGF)

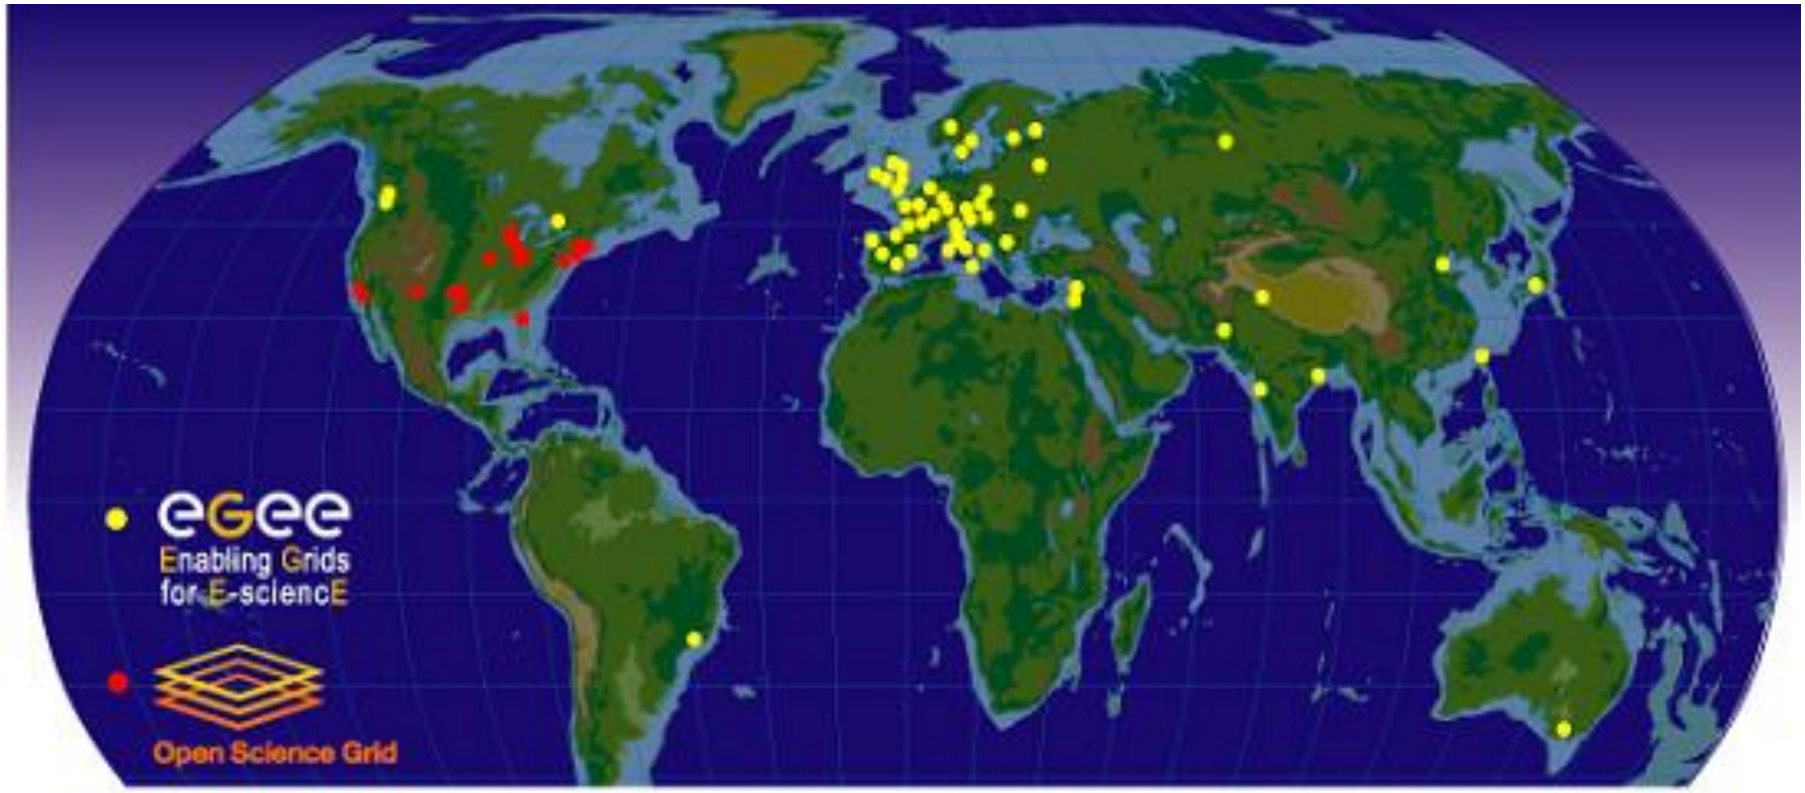

CPU

- "Number crunching" boxes
  - No resident scientific data
  - Shared facility for all CERN users (basically every physicists participating in an experiment at CERN)
- Faster PC means more SPECs per box
  - SPEC_INT_2k still very much used
    - 1,000 SI2K ~ Pentium4 @ 3GHz (~3GFlops)

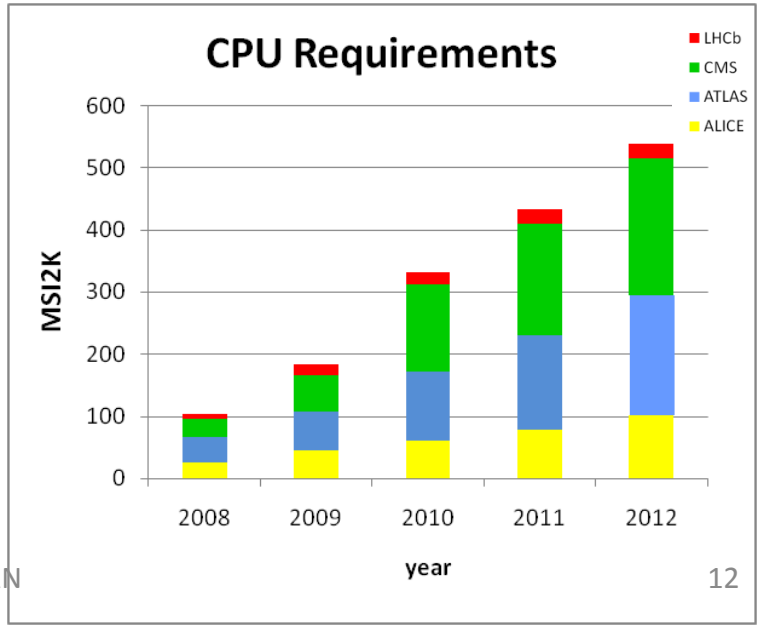

Disks

- "Staging" area: they keep "hot" data
- Access to these disks is managed by dedicated PCs (serving/receiving data to the PC crunching numbers)
- "Moore"'s law at work here!
  - Gb per € goes up
  - Disk capacity goes up
    - GB/cm^3 goes up
  - Mass-market items!

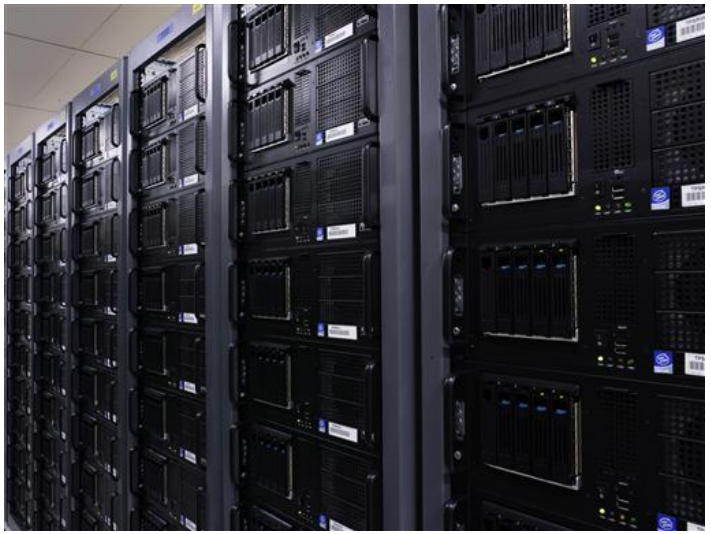

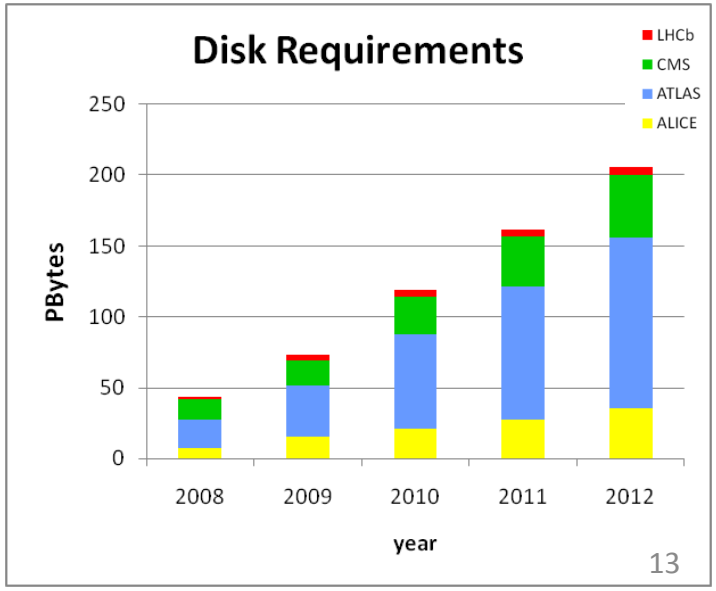

Tapes

- Data custodial
  - We build accelerator and experiments to collect scientific data
- Write-once / Read-many
  - At variance with backups (Write-once / Read-never)
- Evolving (a' la Moore's law) but more "gently" than PCs
  - No surprise, none of us has a tape library at home, I guess ☺

- Expect to store 40 PB (40,000 TB) of data per year
  - Scientific data, corresponding derived data (reconstruction, analysis), simulation data

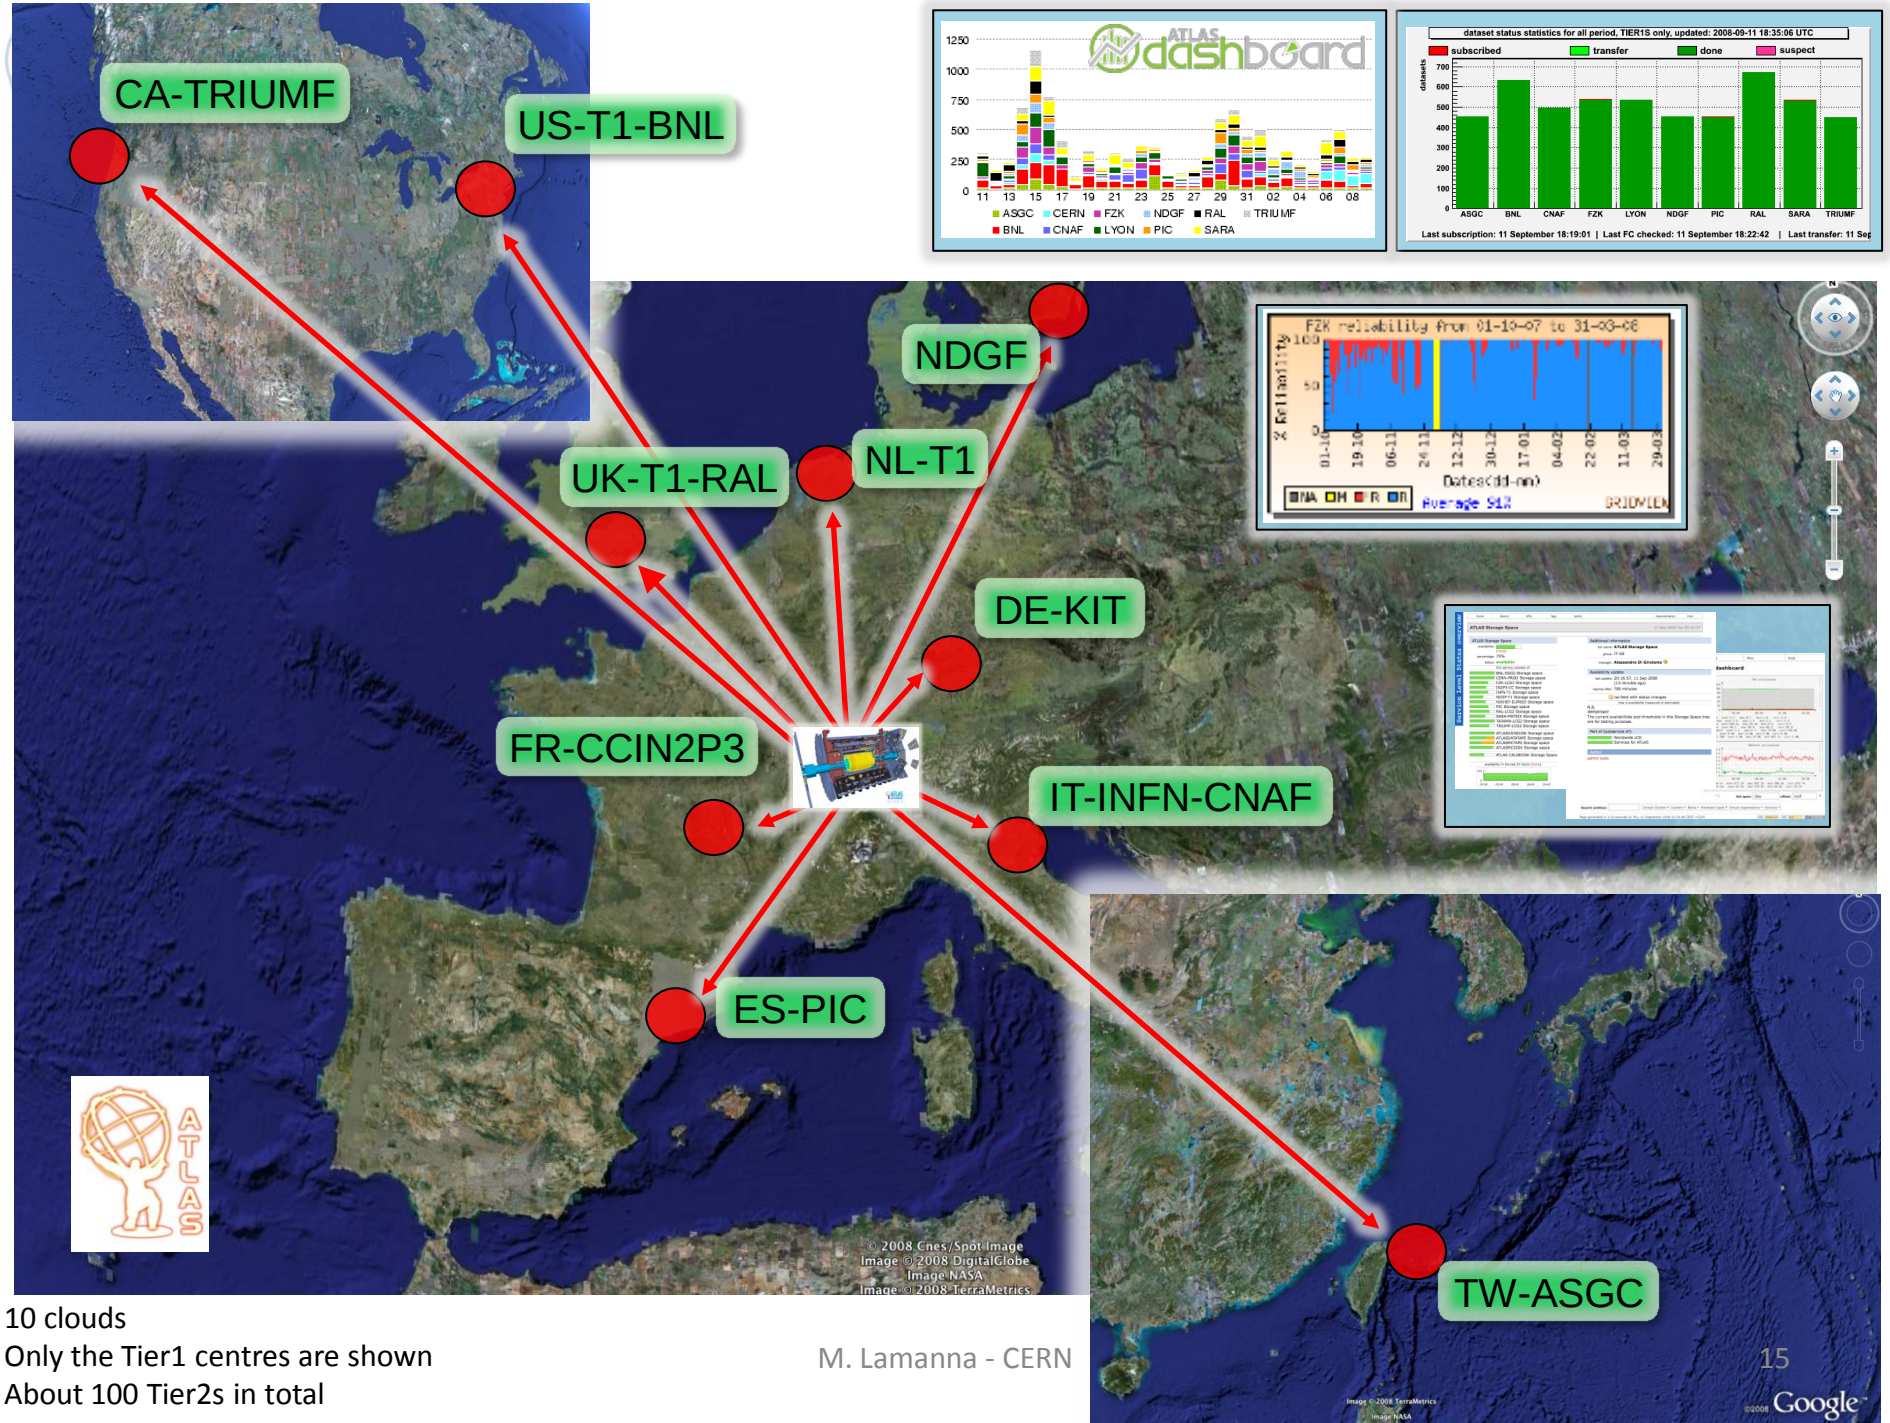

10 clouds
Only the Tier1 centres are shown
About 100 Tier2s in total

EXAMPLES OF OTHER APPLICATION (CERN COLLABORATIONS)

ITU conference (2006)

The problem:
Assign frequencies for digital radio and television (international treaty)

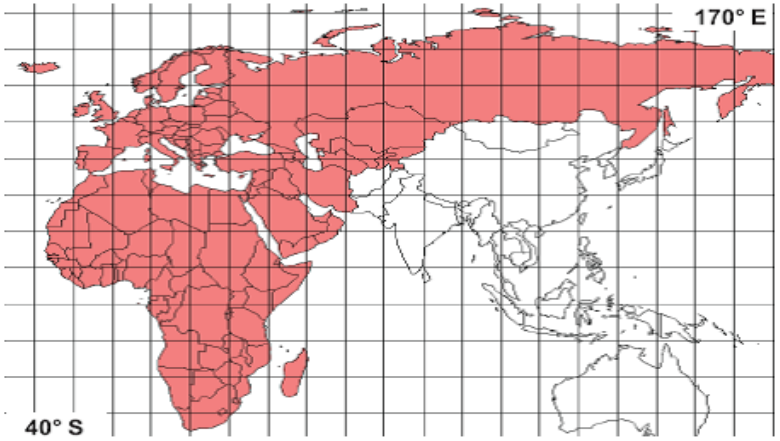

Figure 1
The extent of the planning area for the RRC-06

Critical point:
Need on dependability:
verify (iteratively) the compatibility between radio stations

Solution:
Use the EGEE grid + a system used in ATLAS and LHCb to increase the reliability of the Grid

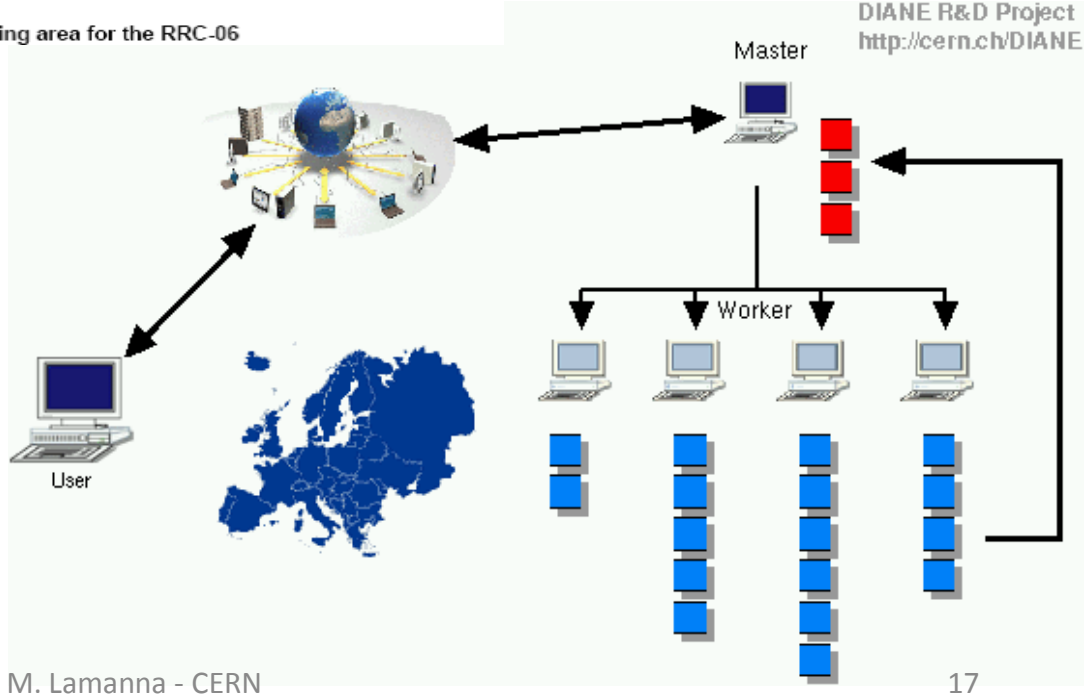

Bird Flu

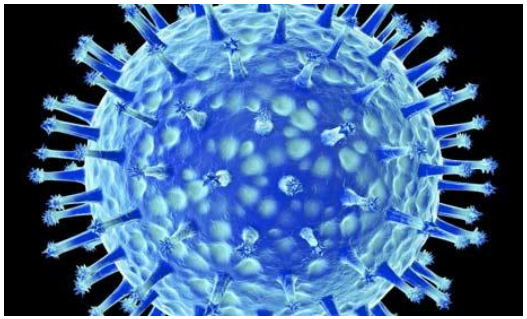

- Basic idea:
  - Compute how a given chemical interacts with a protein (e.g. belonging to a virus)
  - High affinity means the chemical is a potential drug against the virus
- In silico (i.e. use your PC):
  - Scan millions of chemicals (~10^3 s per chemical-protein pair)
    - With 1,000 PCs, 1 docking per second
  - Good candidate given to biologist (verification longer –and more expensive- than in silico docking)
    - In practice, you enrich the initial sample saving time (and money)
    - Essential to fight to pandemia (H5N1) or to fight neglected diseases (like Malaria)
- WISDOM collaboration
  - Malaria
  - H5N1 (Bird Flu)

Molecular docking

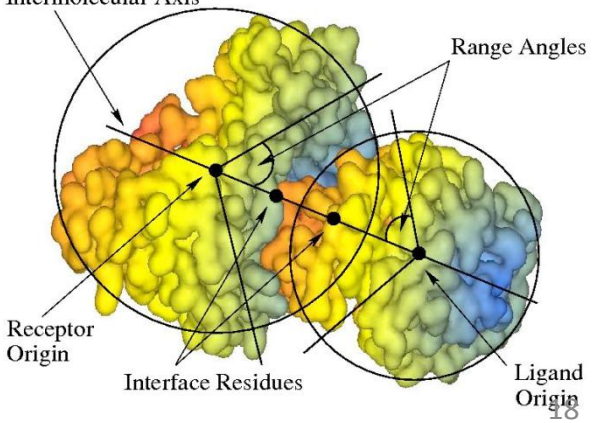

Image search and tagging

Innovation and commercial use.

powered by GANGA

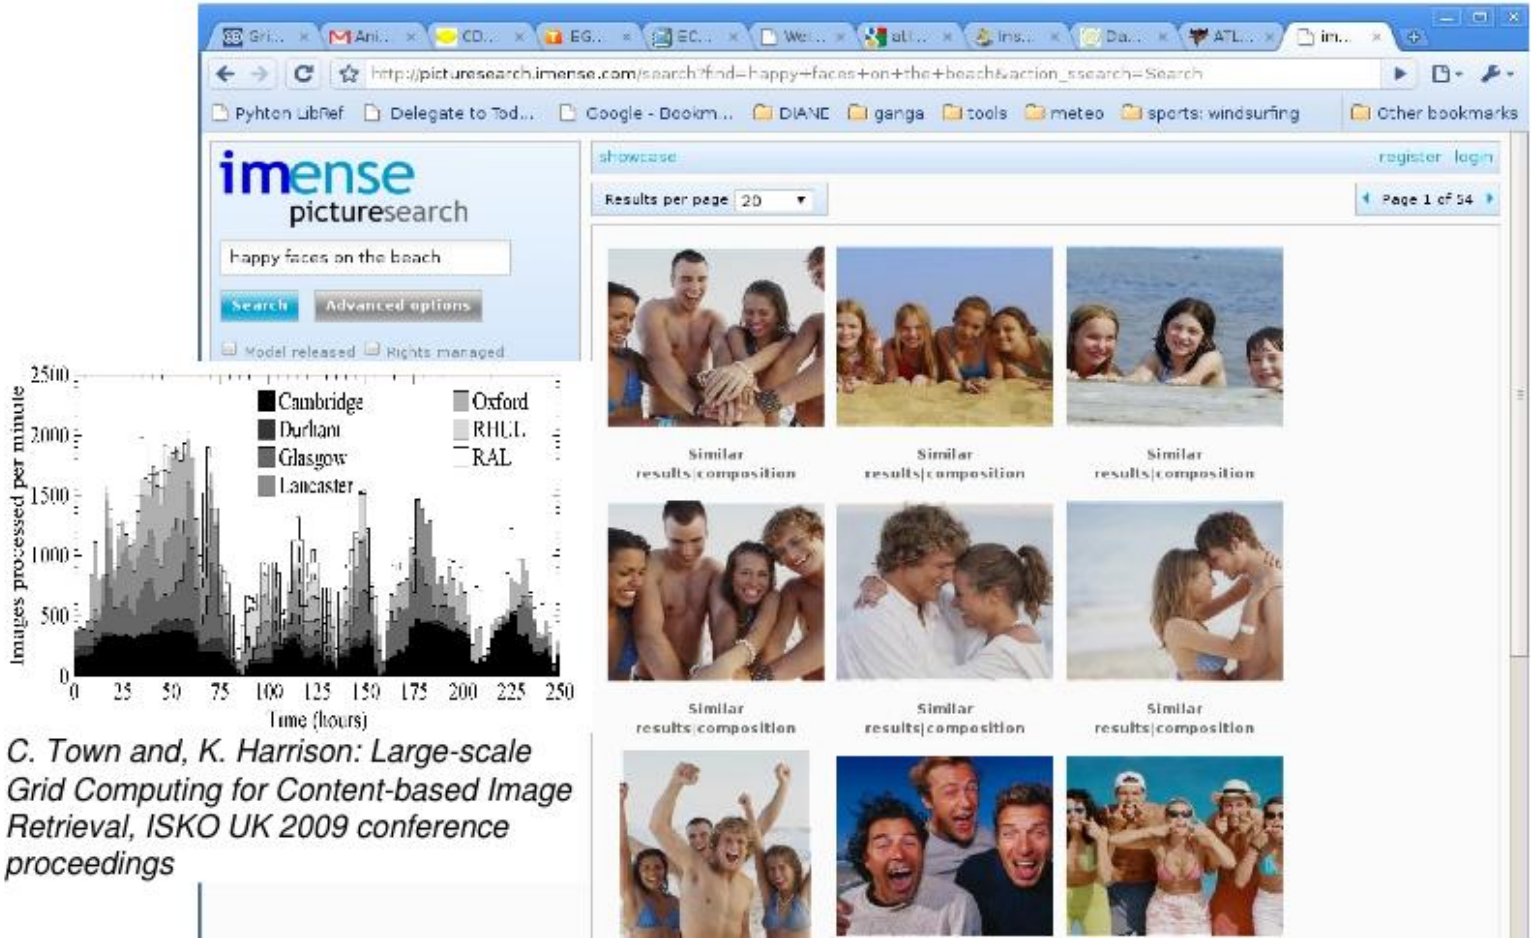

C. Town and, K. Harrison: Large-scale Grid Computing for Content-based Image Retrieval, ISKO UK 2009 conference proceedings

Hydrological modeling

Observations		SWAT	Scenarios1
Observations	Climate	SWAT	Scenarios2
Observations	Land cover	SWAT	Scenarios3

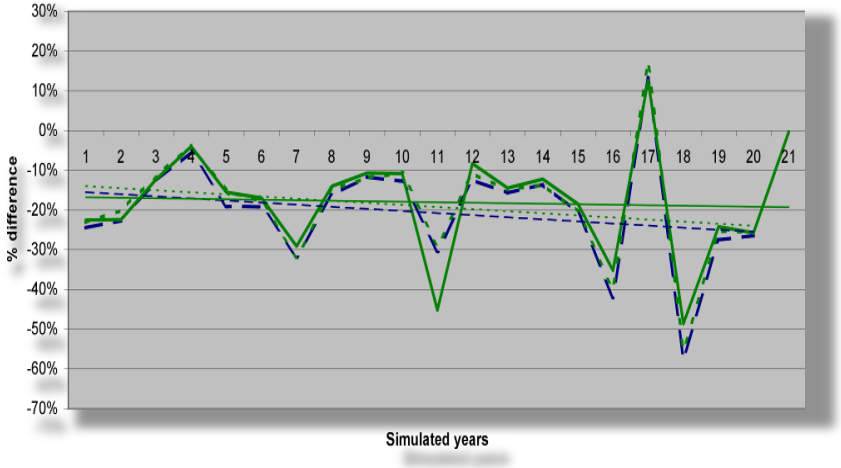

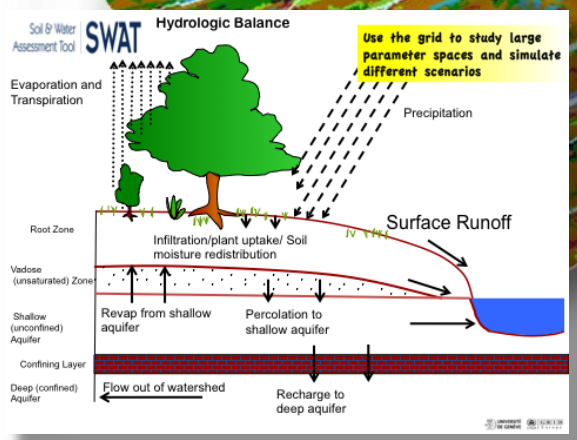

CERN contribution:
"gridification" of SWAT
Using Ganga/Diane

Questions?

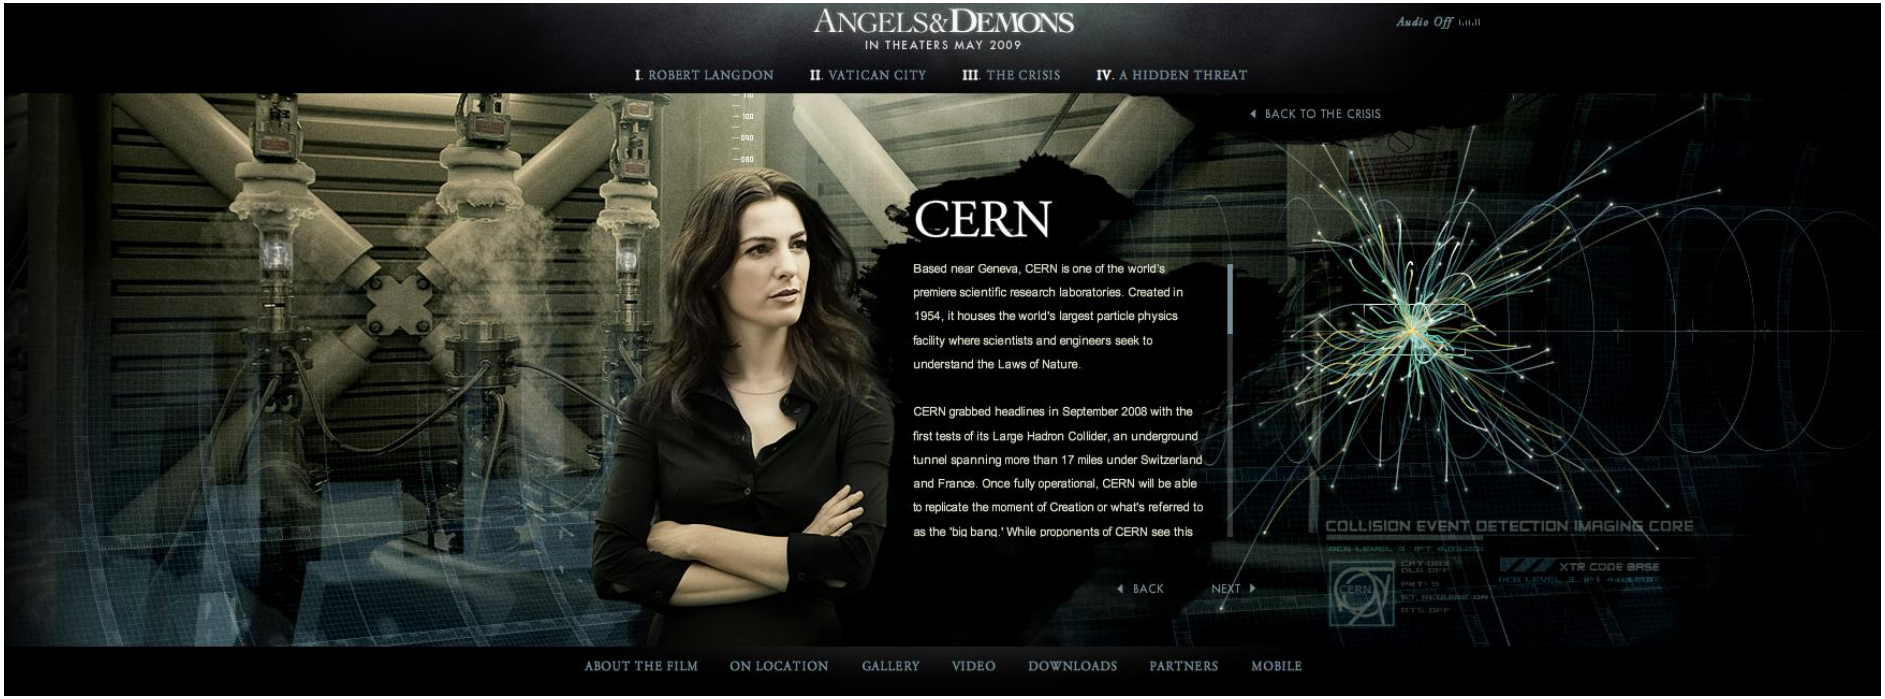
ANGELS&DEMONS
IN THEATERS MAY 2009
Audio Off
I. ROBERT LANGDON
II. VATICAN CITY
III. THE CRISIS
IV. A HIDDEN THREAT
BACK TO THE CRISIS
CERN
Based near Geneva, CERN is one of the world's premiere scientific research laboratories. Created in 1954, it houses the world's largest particle physics facility where scientists and engineers seek to understand the Laws of Nature.
CERN grabbed headlines in September 2008 with the first tests of its Large Hadron Collider, an underground tunnel spanning more than 17 miles under Switzerland and France. Once fully operational, CERN will be able to replicate the moment of Creation or what's referred to as the 'big bang.' While proponents of CERN see this
COLLISION EVENT DETECTION IMAGING CORE
XTR CODE BASE
BACK
NEXT
ABOUT THE FILM
ON LOCATION
GALLERY
VIDEO
DOWNLOADS
PARTNERS
MOBILE