Detector Simulation & GRID Technology

F.Carminati

Erice, September 28, 2003

Alice collaboration

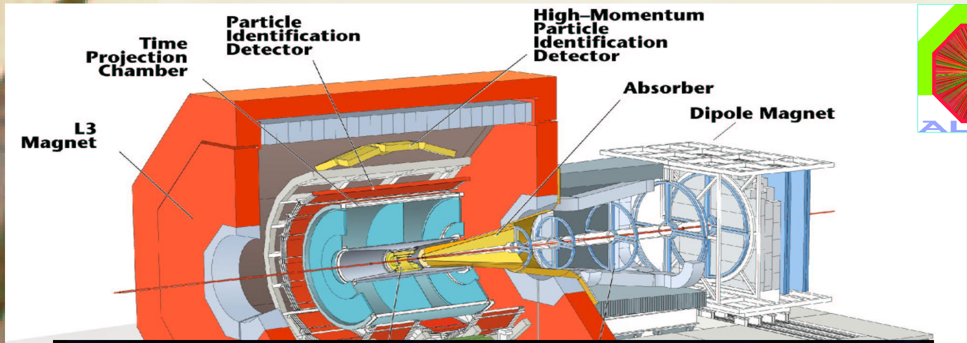

Total weight	2,000t
Overall length	17.3m

Total weight	53,000t
Overall length	270.4m

nd
00t
0m
esla
ure
/sec)
proc
GB/
2 - I
30
(1.25 G
data rec
offline analysis

ALICE Event/100

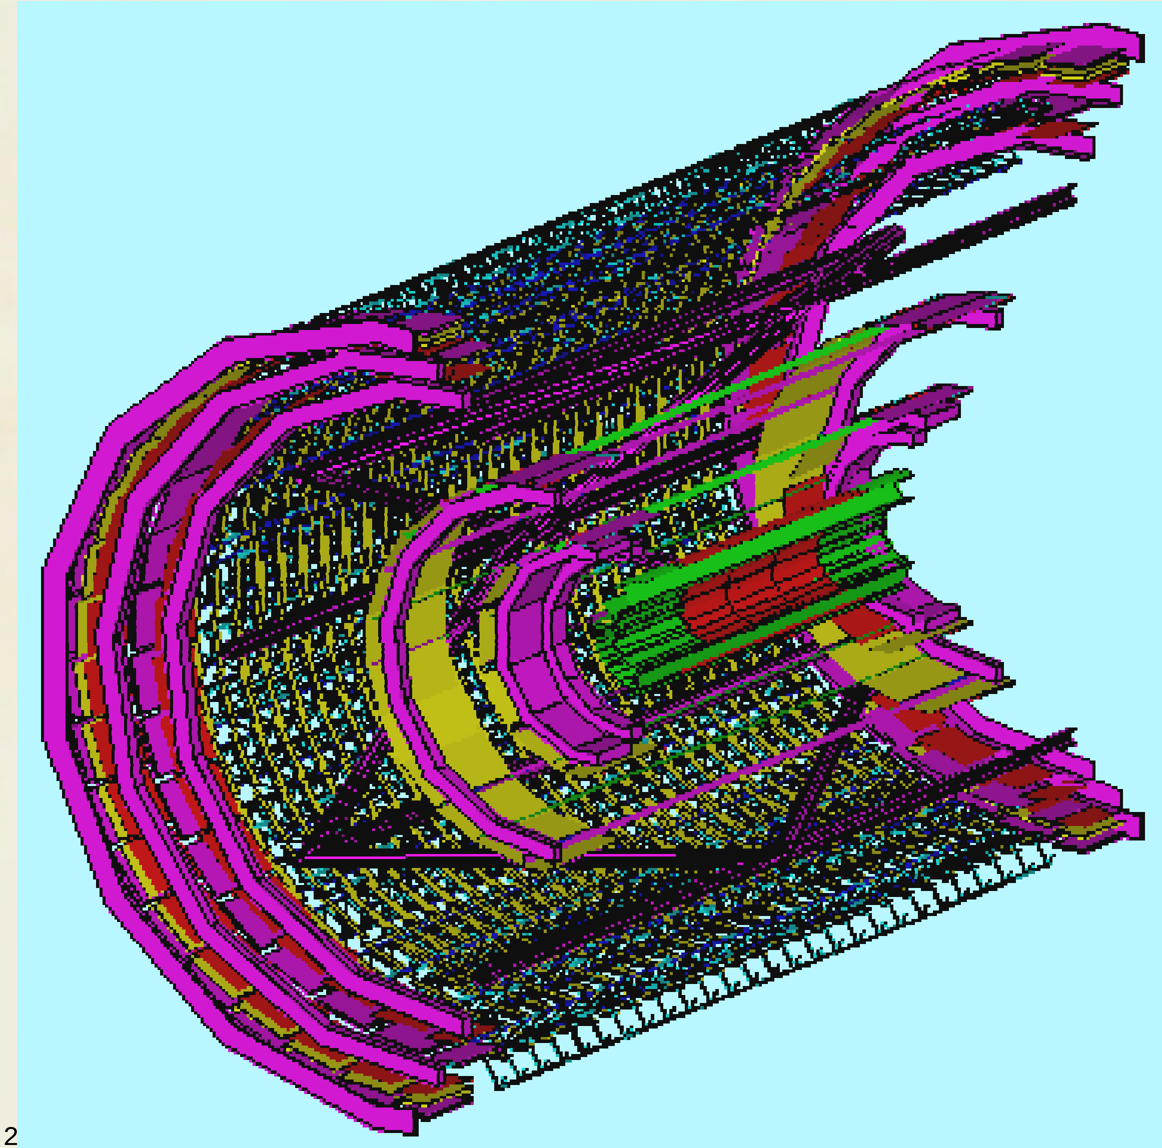

The question of simulation

- Simulation is vital for HEP to evaluate the performance of the detector and estimate background

BUT

- Using GEANT 3.21
  - Stay with FORTRAN, old physics and geometry
- Using GEANT 4
  - Not yet completely validated
- Using FLUKA
  - Good physics but limited user interfa

We are in trouble!

The Virtual MC

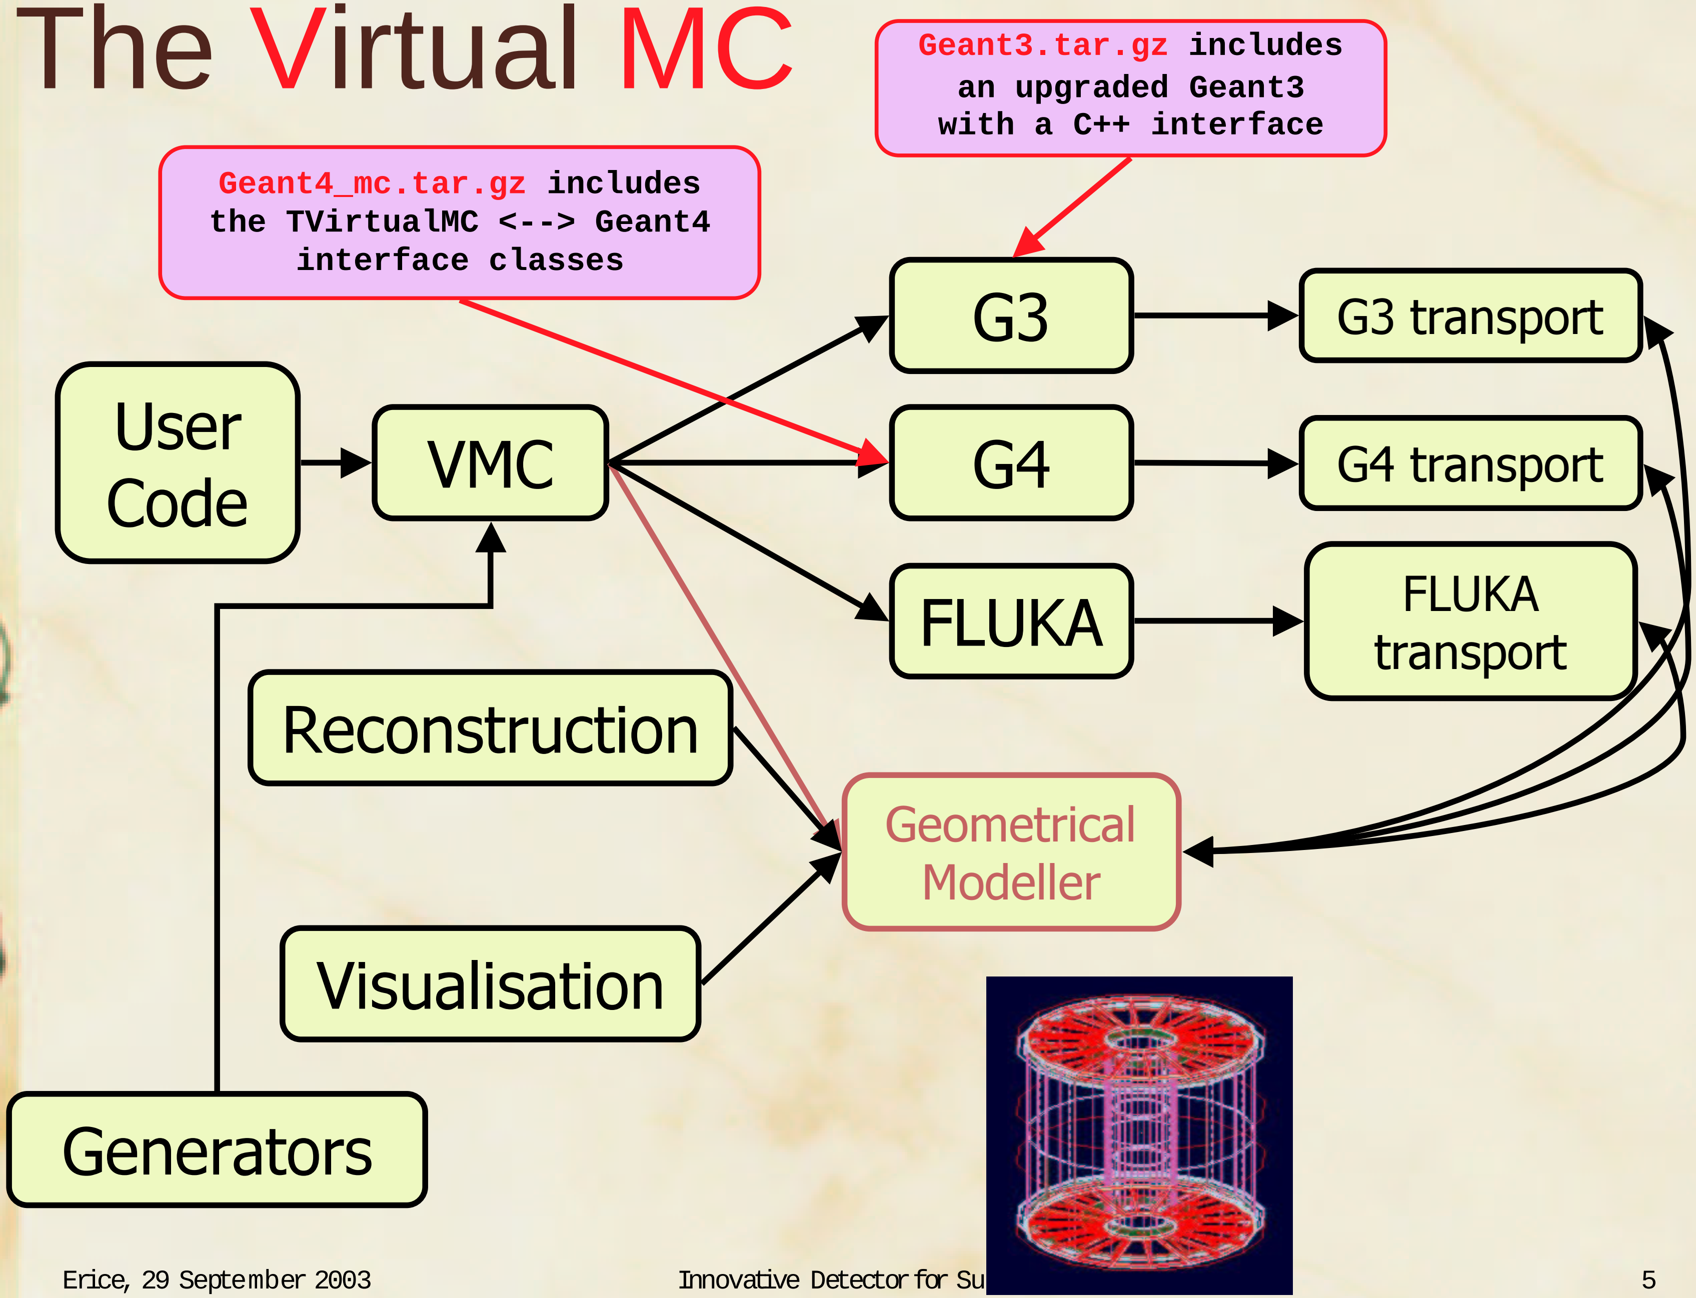

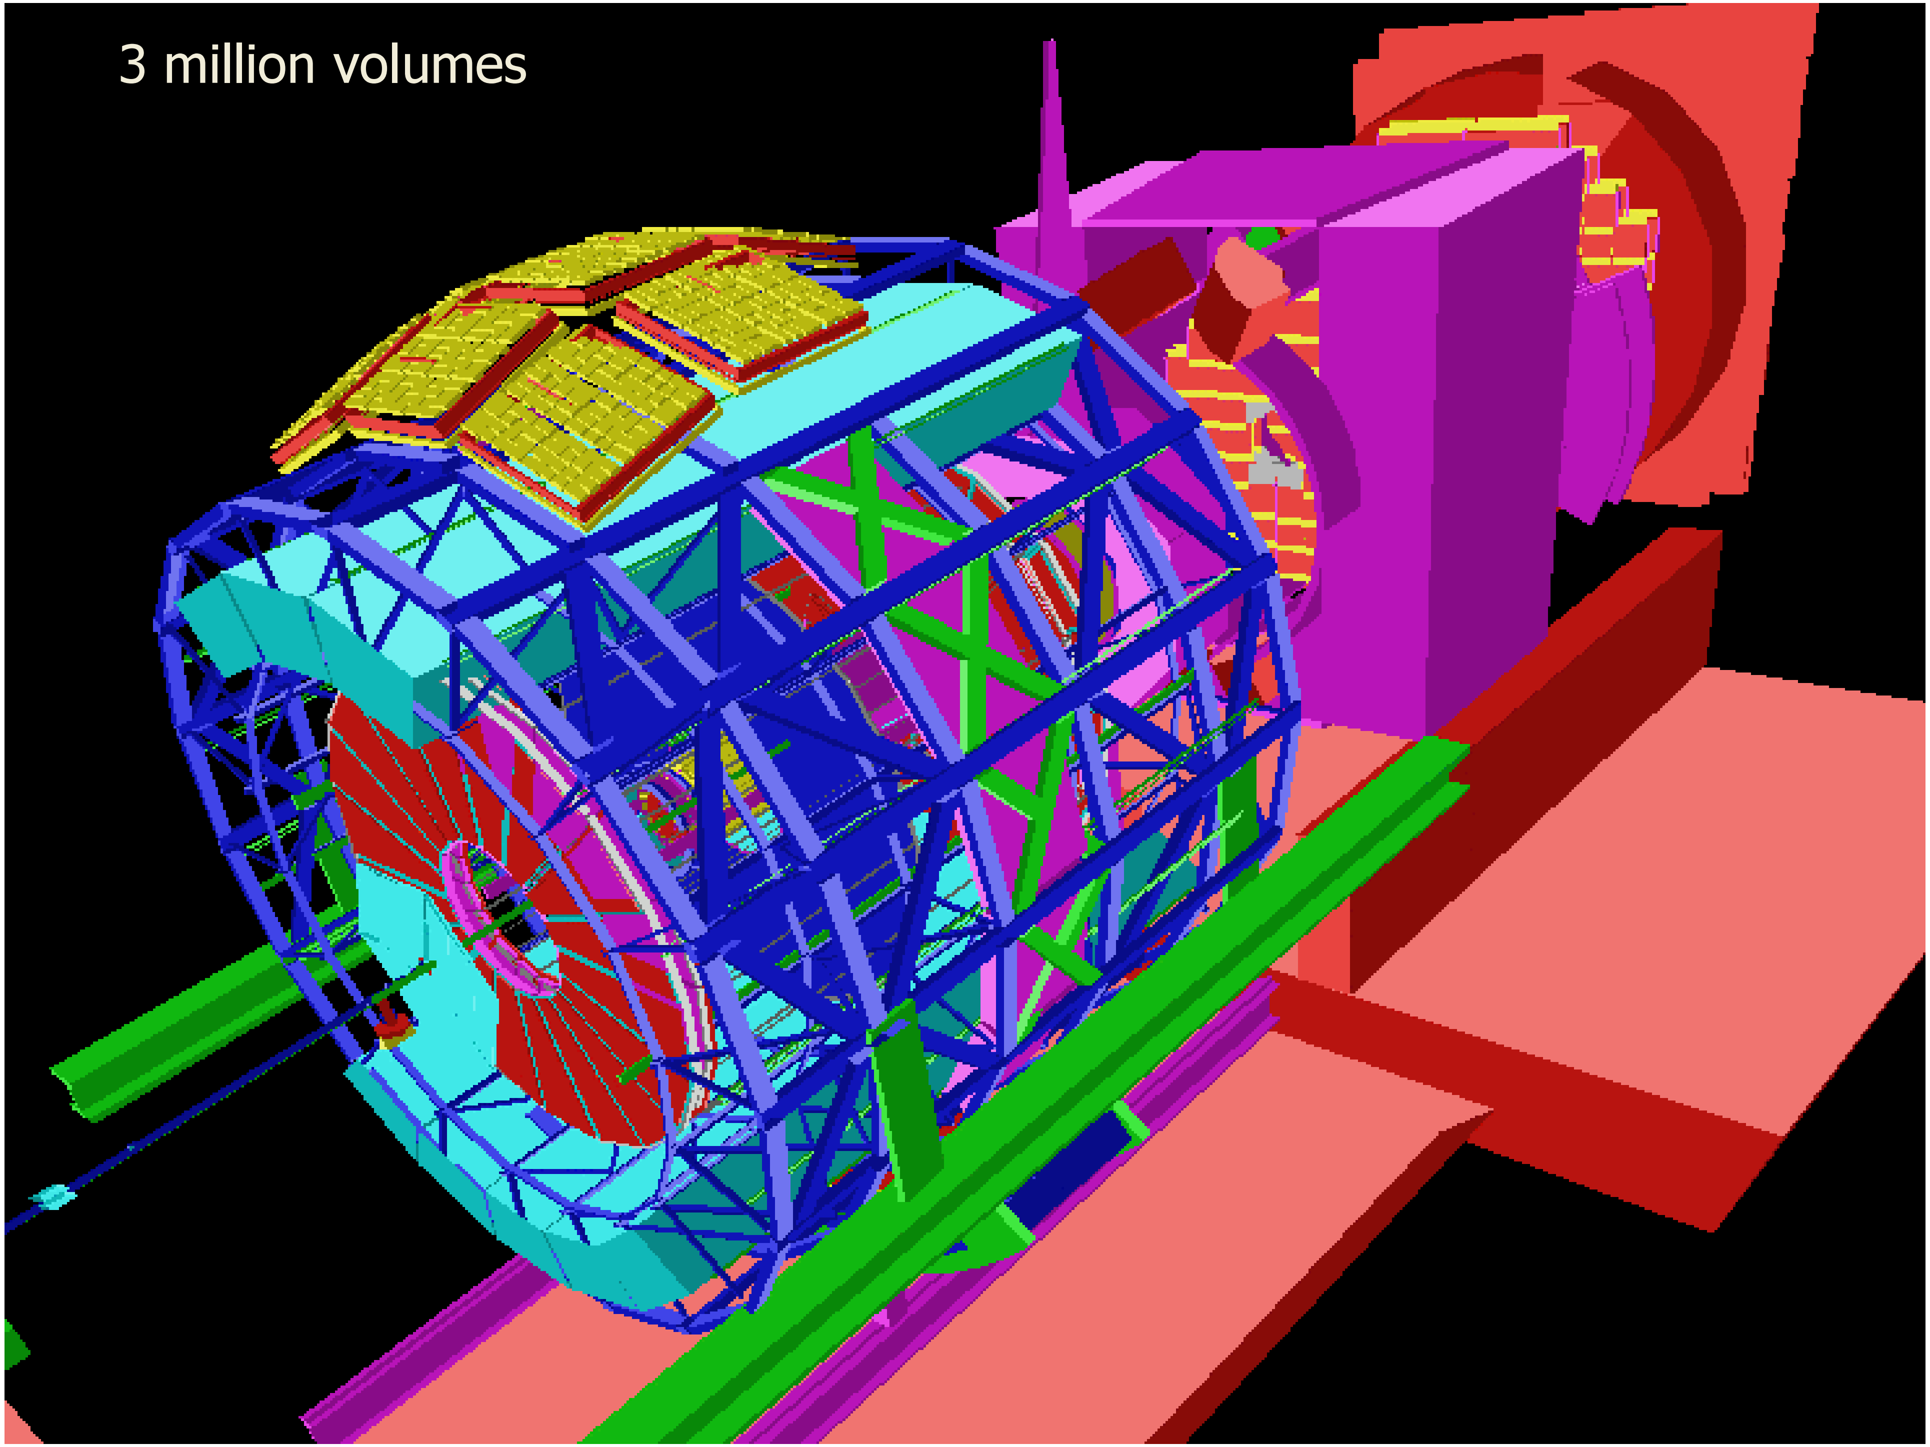
3 million volumes

Issues in simulation

- Reliability of the MC
  - The simulation result are (at best!) as reliable as the MC used!
  - The problem is how well a MC describe a non existing detector in an inexperienced domain of energy
  - Multi-million €choices depend on this
- Availability of adequate computing power
  - Cannot beat the 1/√n law
  - Variance reduction techniques are not commonplace in all HEP MC’s, and they are not of general use

Reliability

- The question usually asked is what is the precision needed for the different processes
- Unfortunately this question is ill posed
  - The interaction between processes is highly non linear
  - It is almost impossible to determine “a priori” which process is important for a given case
- The best strategy is
  - Describe processes “as well as possible” (sic!)
  - Constantly control the quality of the MC with experimental data
- BTW, the same applies with cuts

Do we know everything about multiple scattering?

Boundary

$$\frac{\Delta\vartheta}{\Delta x} \xrightarrow[x \to 0]{} \infty$$

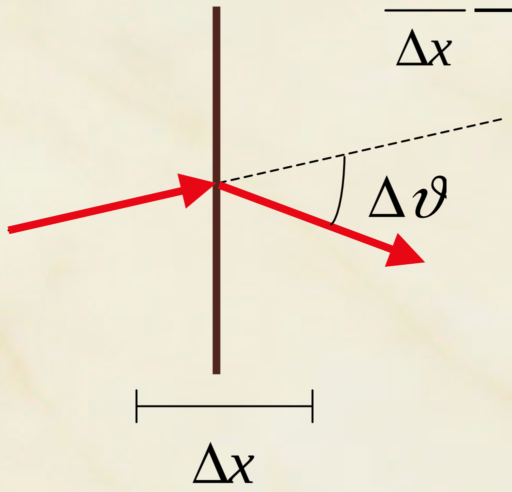

No initial deviation?

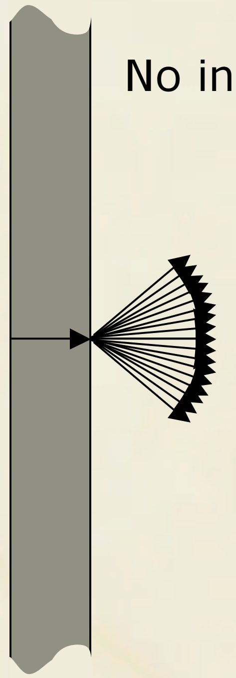

Single scattering through the boundary

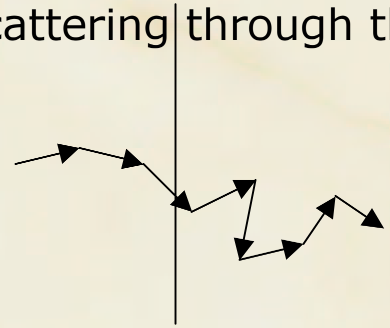

Affects e/π in finely segmented media

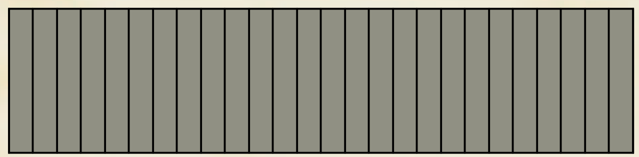

Parametrisation and physics

- All MCs contain parametrisations to be tuned… however
  - If these are the parameters of a spline, no application to a set of data different from the original is allowed
    - This is the case of GHEISHA (page 3 of the GHEISHA manual), still largely used for LHC simulations!
  - If these are parameters of a physic model, a little more optimism is allowed
    - This is the case for FLUKA and the latest developments in GEANT4

MC validation

- Future supercolliders will need predictive simulations
  - Good parametrisations can only be derived from good simulations!
- To have some hope that a MC is predictive two kind of tests have to be performed exhaustively
  - Single process tests
  - Test beam validations
- Experimental errors have to be considered attentively
  - Do not confuse systematic and statistical errors!
- Very hard work…

An example with GEANT4

- Low momentum hadrons are important for ALICE
  - Open geometry (no calorimeter to absorb particles)
  - Small magnetic field (0.4 T)
  - Account for most of the energy deposit
- Particles "leaking" through the front absorbers and beam-shield generate background which limits the performance in central Pb-Pb collisions
- In the forward direction also the high-energy hadronic collisions are of importance

Proton Thin Target Experimental Set-Up

Data from Los Alamos in: Nucl. Sci. Eng., Vol. 102, 110, 112 & 115

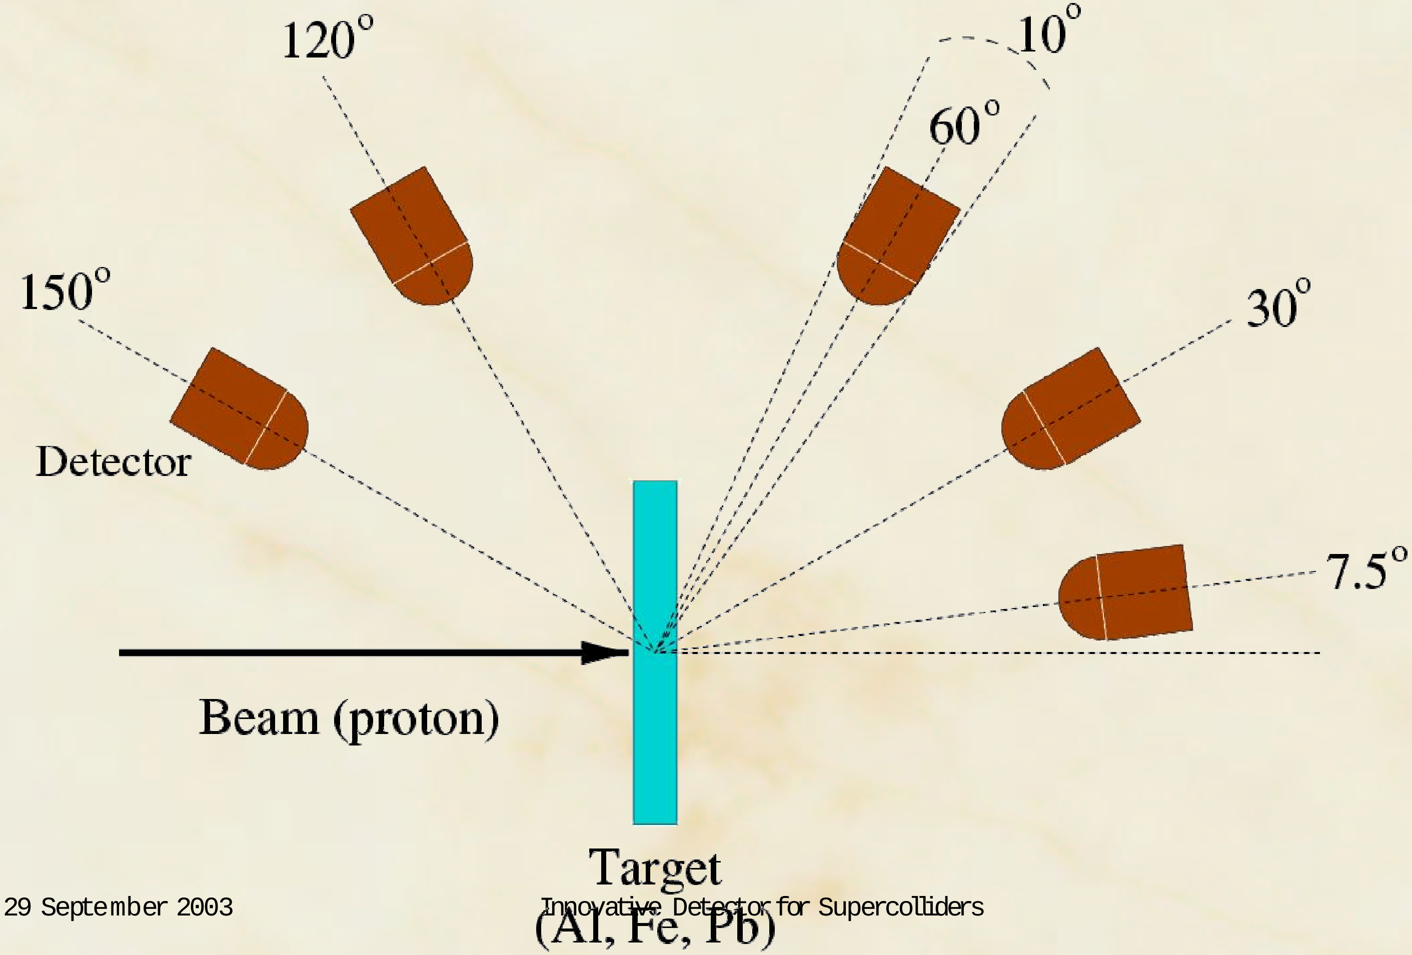

Parameterised model:
φ_{pions}: (p,Al) @ 597 MeV

Before

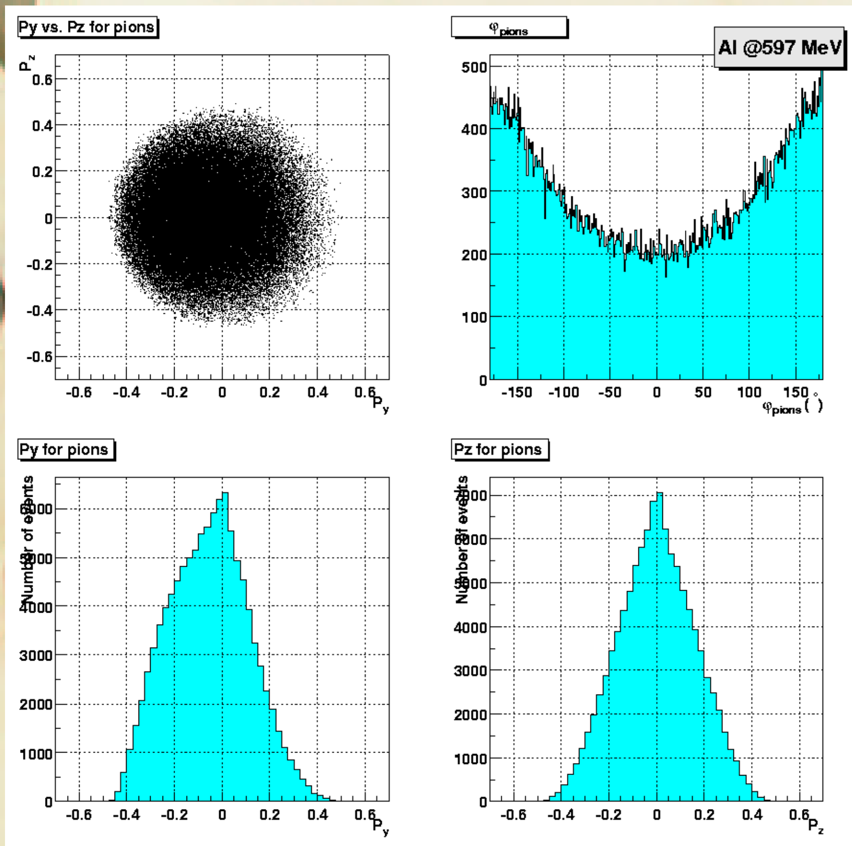

Now

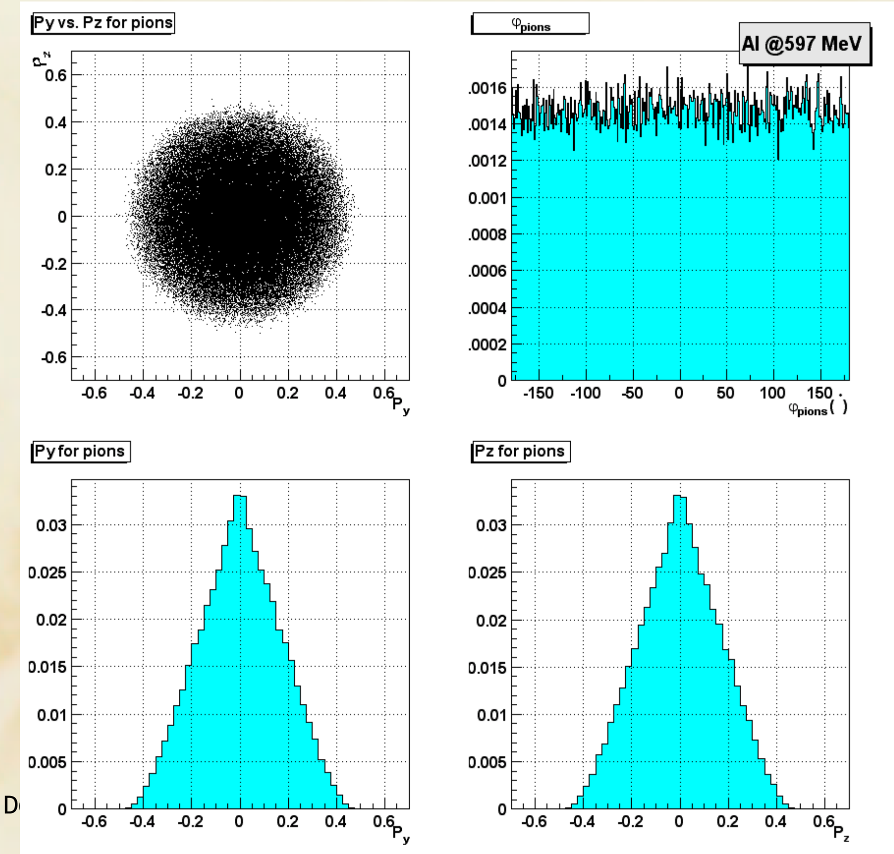

Double differentials

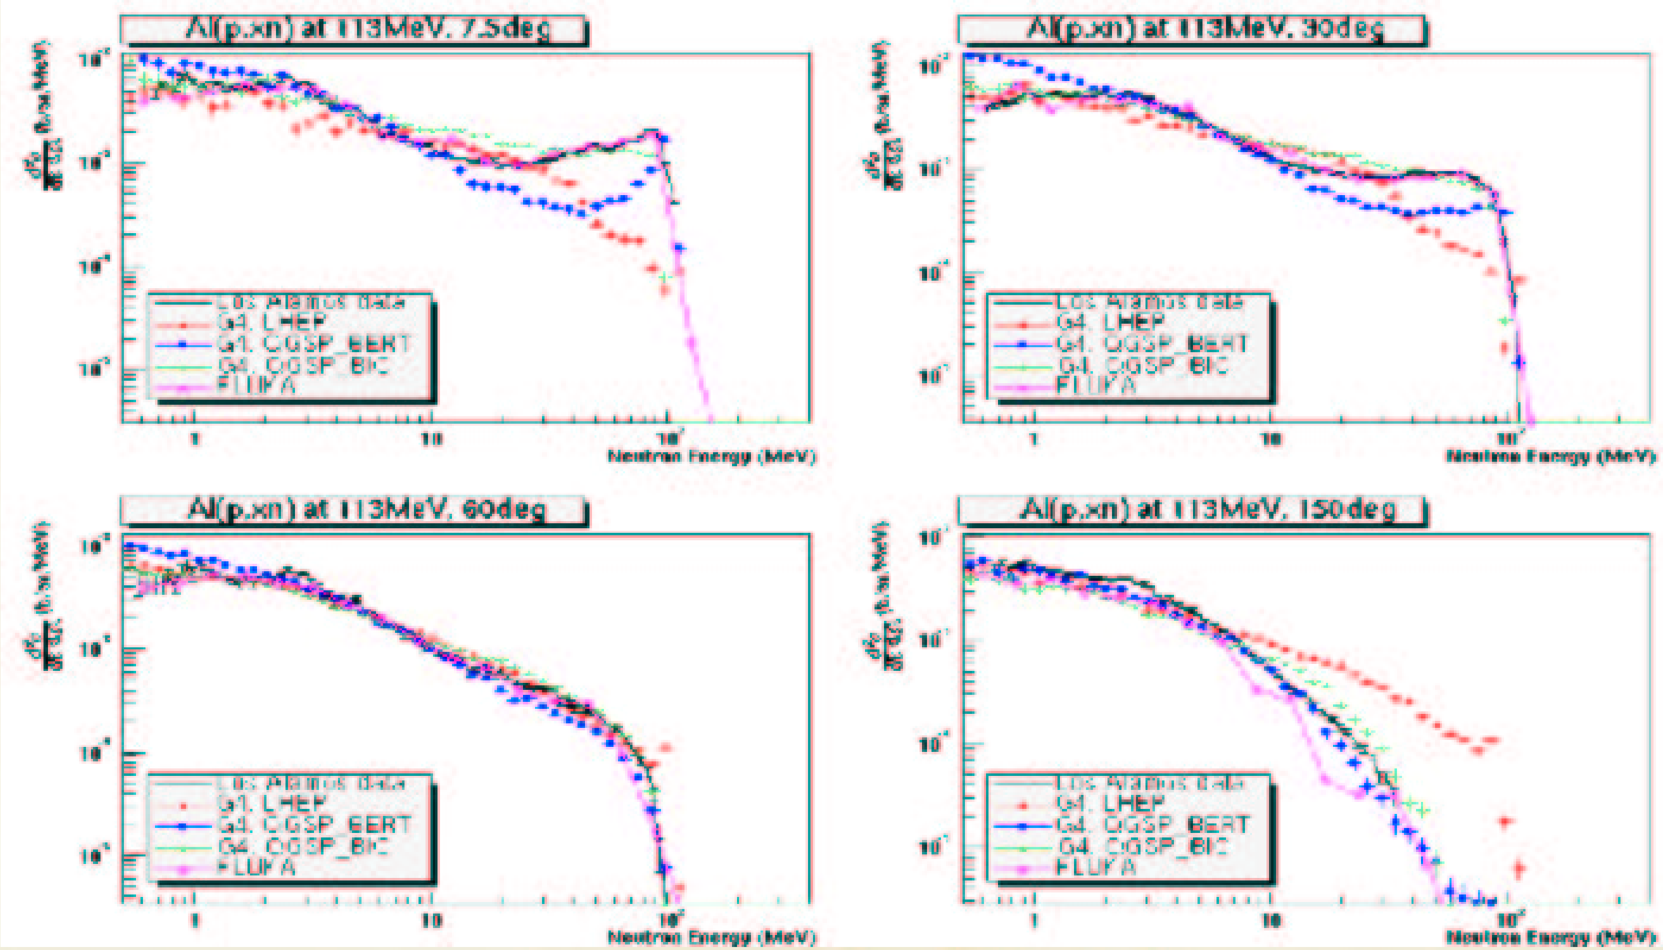

Low energy neutrons

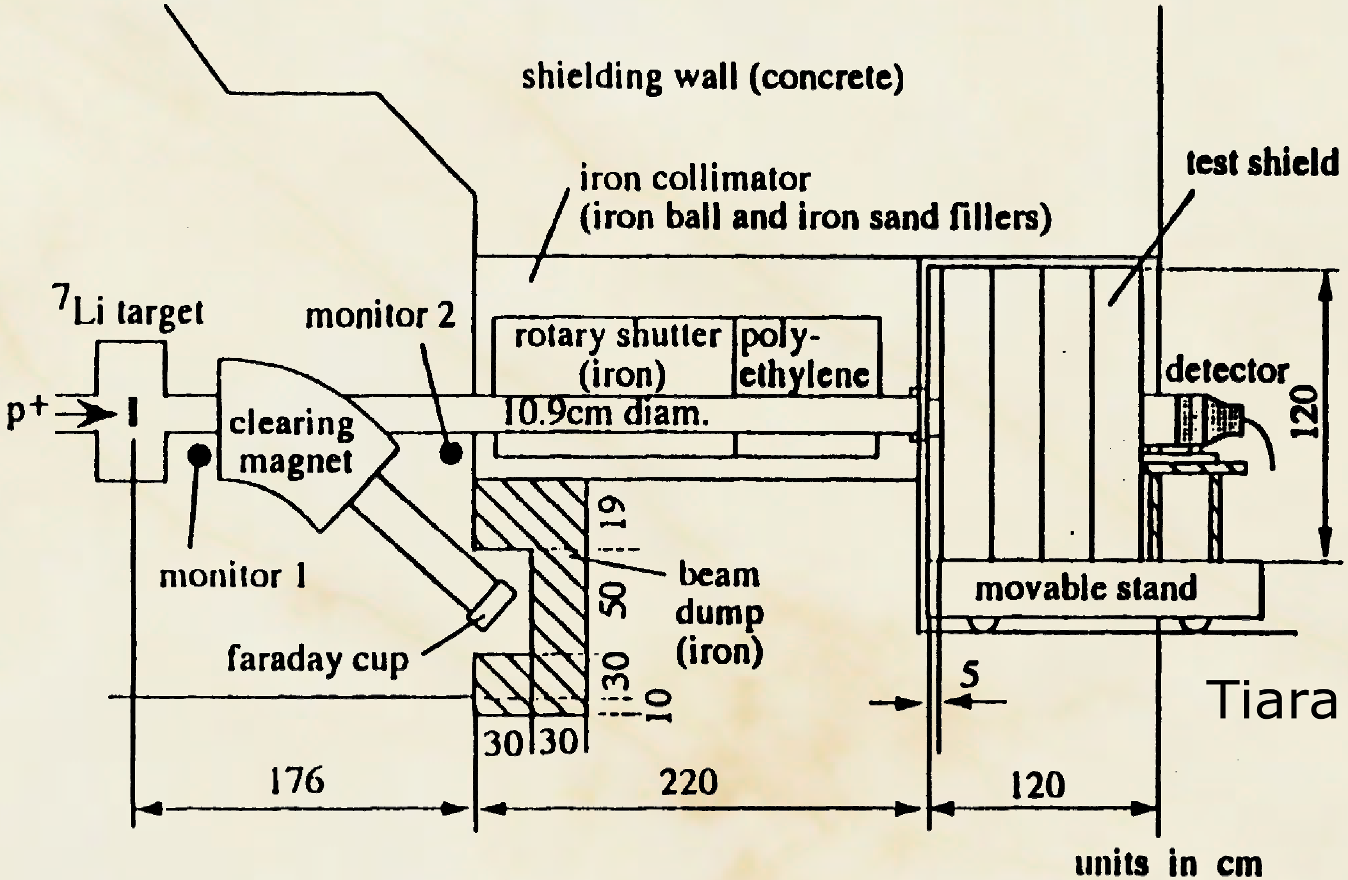

Tiara facility

Preliminary Results: 43 *MeV*

Test Shield: Iron – Thickness: 20 *cm*

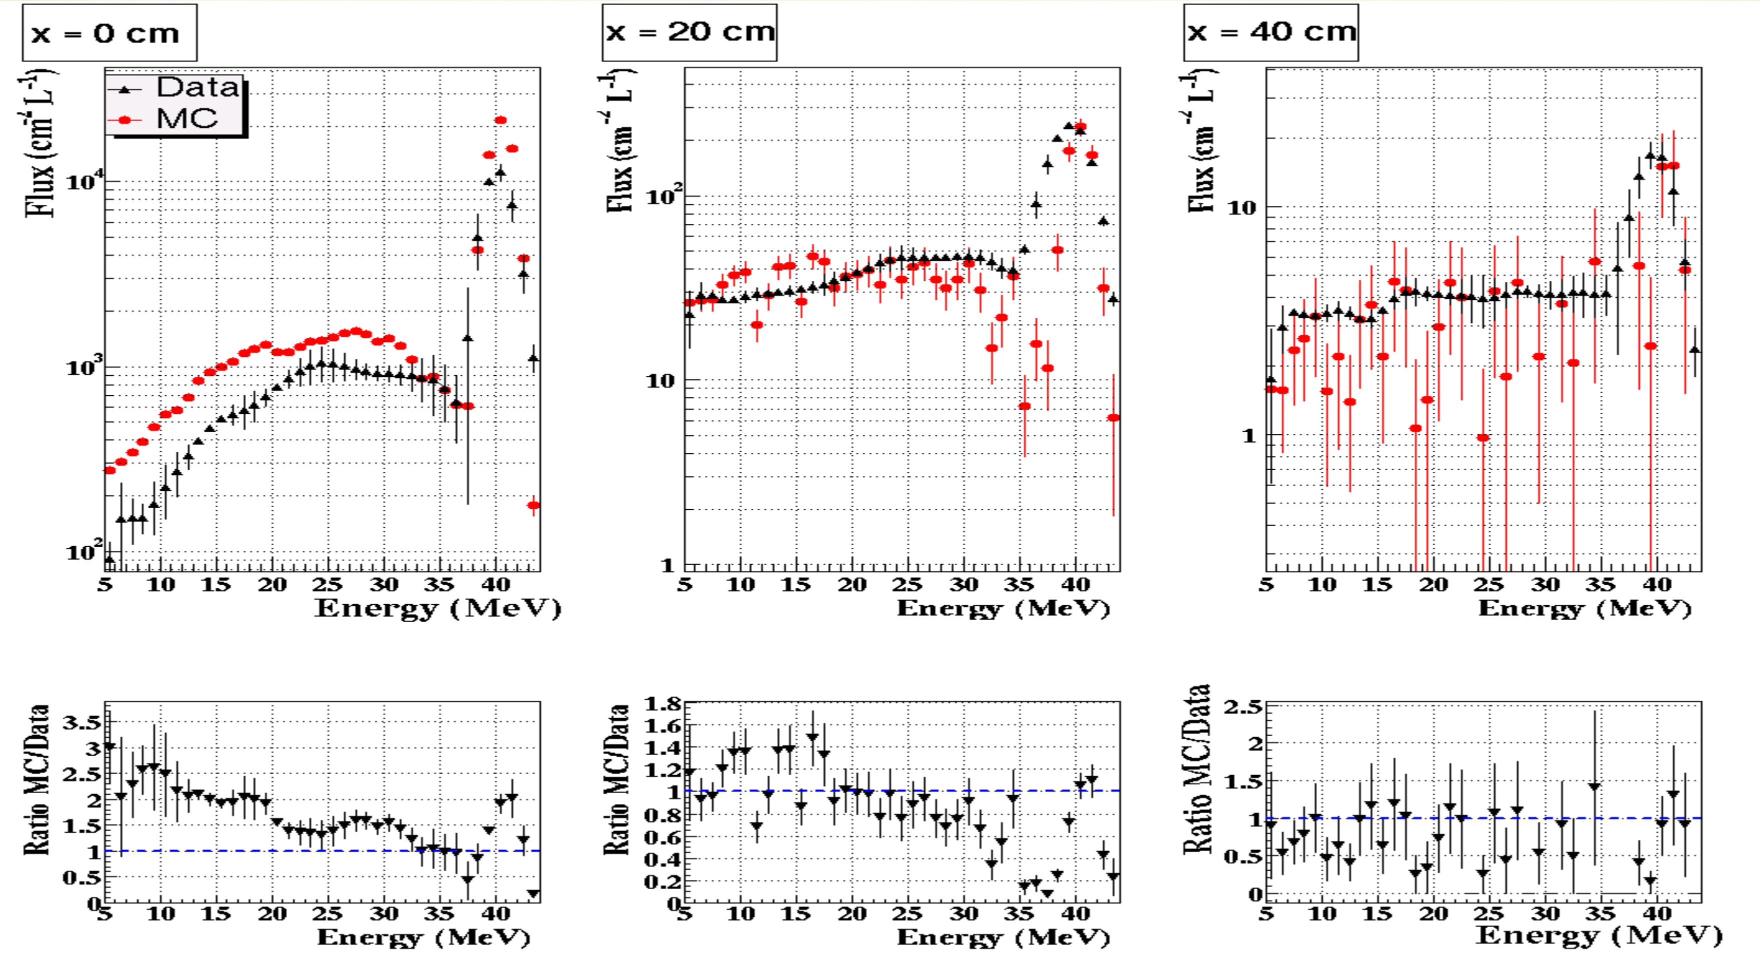

Simulation and LCG

- The LCG project is taking very seriously simulation activity
  - (At last!) a common project includes FLUKA and GEANT4
  - A principle agreement has been found for the distribution of the FLUKA source
- Single process and test beam validations are programmed for FLUKA and GEANT4
  - Should have been started 10 years ago!
- The project has decided to
  - Adopt the ALICE Virtual MonteCarlo as generic framework
  - Adopt the ALICE geometrical modeller

But where to find the necessary computing power?

- The needs of simulation are very large
  - For the next Data Challenge ALICE will need 1400kSI2K and 300TB for six months
    - Approximately 1400 high-end PCs running continuously!
- But the computing needs to collect, store, reconstruct and analyse LHC data are even larger!

The LHC Detectors

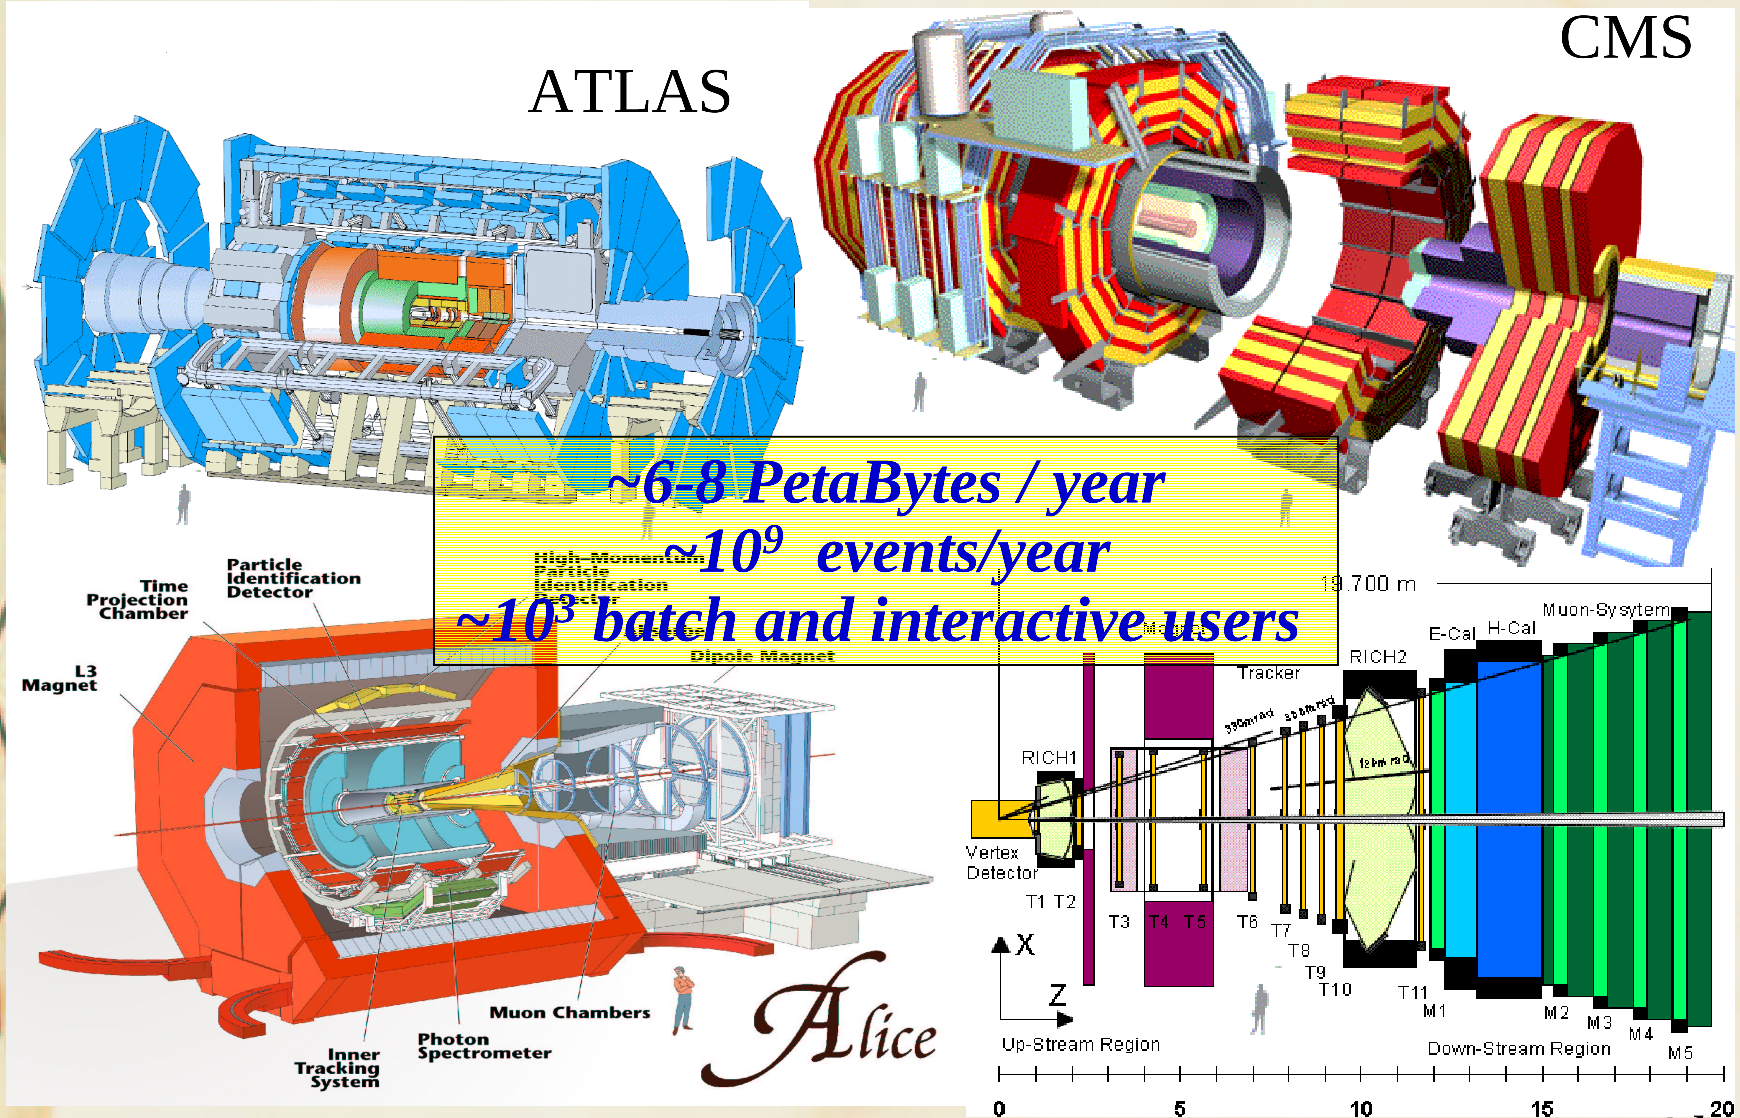

~6-8 PetaBytes / year

~10^9 events/year

~10^3 batch and interactive users

CERN's Network in the World

Europe: 267 institutes, 4942 users
Elsewhere: 208 institutes, 1752 users

Why distributed computing?

- The investment for LHC computing is massive
- For ALICE only
  - 1.25 GB/s in HI mode
  - ~3 PB/y of tape
  - ~1.5 PB of disk
  - ~16,000 kSI2k (~16,000 PC2003)
  - ~ 8M€ of hardware
    - Without personnel + infrastructure and networking
  - Millions lines of code to develop and maintain for 20 years
- Politically, technically and sociologically it cannot be concentrated in a single location
  - Whenever possible countries prefer national investments
  - Competence is naturally distributed
  - A concentrated facility would force people to travel to CERN often

The distributed computing model

- Every physicist should have equal access to data and resources
- Resources for HEP computing (CPUs' and data) will be distributed
  - Co-located in so-called regional centres
- The centres should work as an integrated system to provide
  - Maximisation of the usage of the resources
  - Redundancy and fault tolerance
  - Security
  - Maximum transparency of usage
- The system will be extremely complex
  - Number of sites & components in each site
  - Different tasks performed in parallel: simulation, reconstruction, scheduled and unscheduled analysis
- Physicists have realised the challenge of this since few years
  - Studies started already some years ago (MONARC project)

The Monarc Model

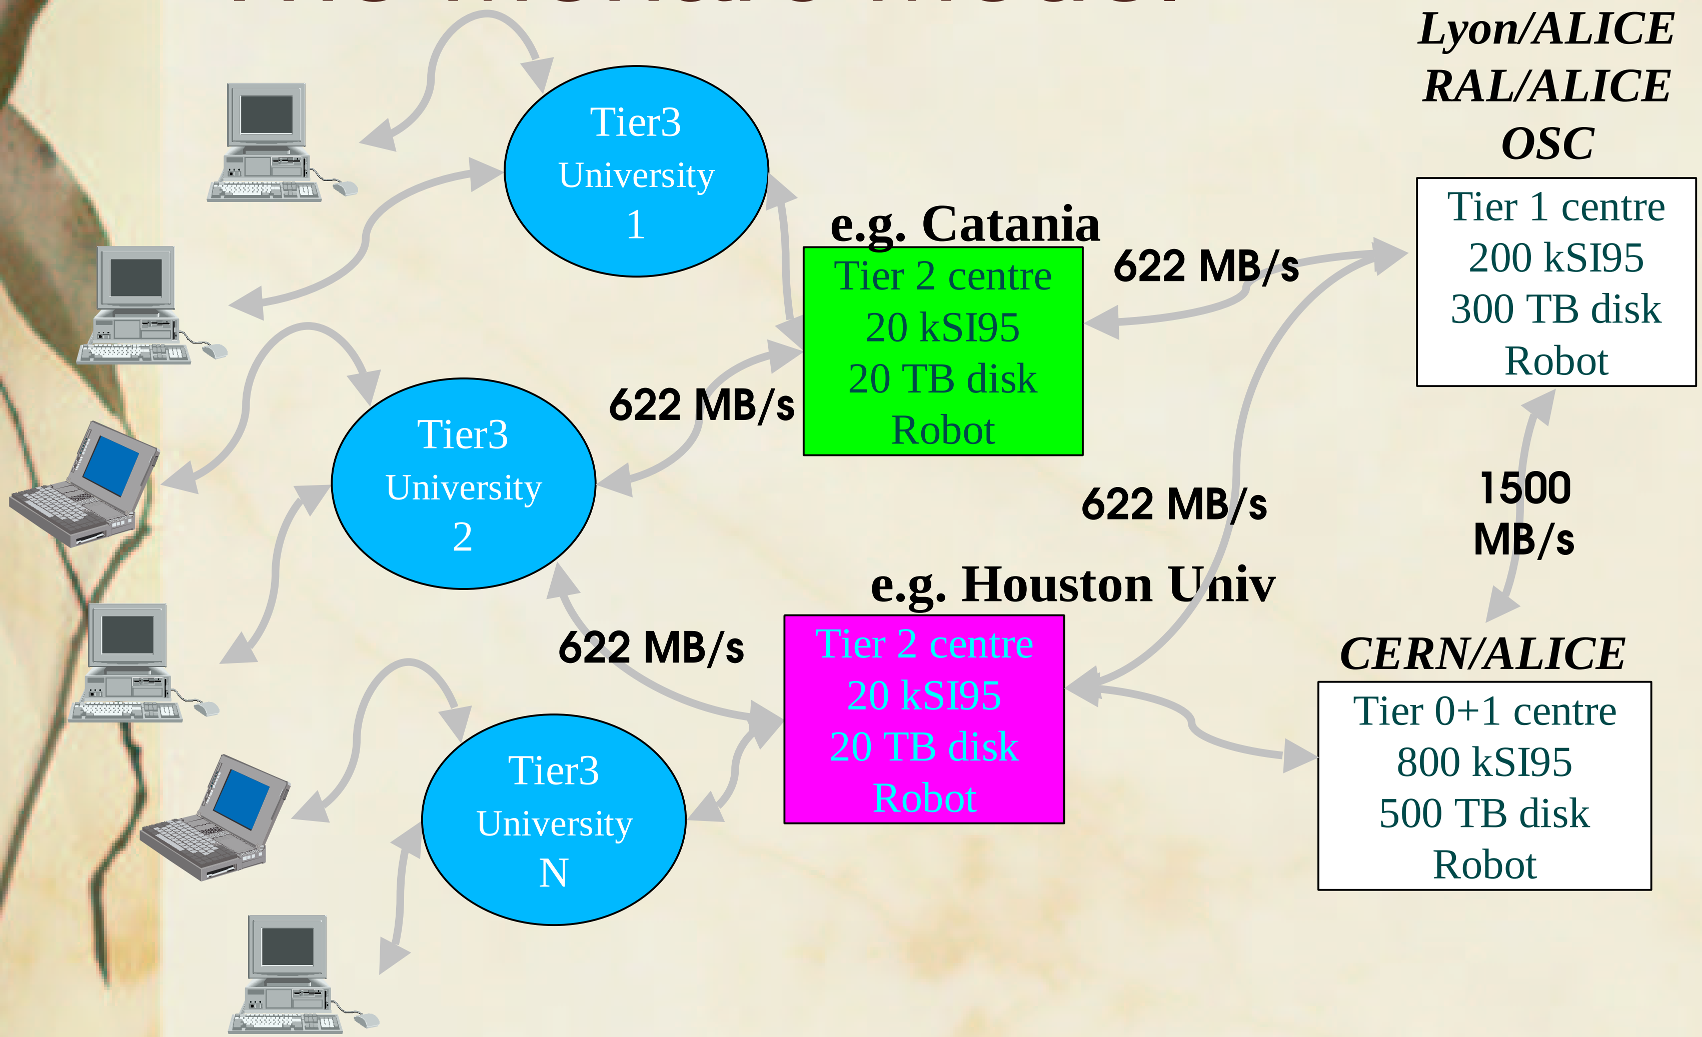

The challenge

- Bad news is that the basic tools are missing (at production quality)
  - Distributed resource management
  - Distributed namespace for files and objects
  - Distributed authentication
  - Local resource management of large clusters
  - Data replication and caching
  - WAN/LAN monitoring and logging
- Good news is that we are not alone
  - All the above issues are central to the new developments going on in the US and in Europe under the collective name of GRID

The Grid: Blueprint for a New Computing Infrastructure

I. Foster, C. Kesselman (Eds), Morgan Kaufmann, 1999

- ✓ Available July 1998;
- ✓ ISBN 1-55860-475-8
- ✓ 22 chapters by expert authors including Andrew Chien, Jack Dongarra, Tom DeFanti, Andrew Grimshaw, Roch Guerin, Ken Kennedy, Paul Messina, Cliff Neuman, Jon Postel, Larry Smarr, Rick Stevens, and many others

"A source book for the history of the future" -- Vint Cerf

http://www.mkp.com/grids

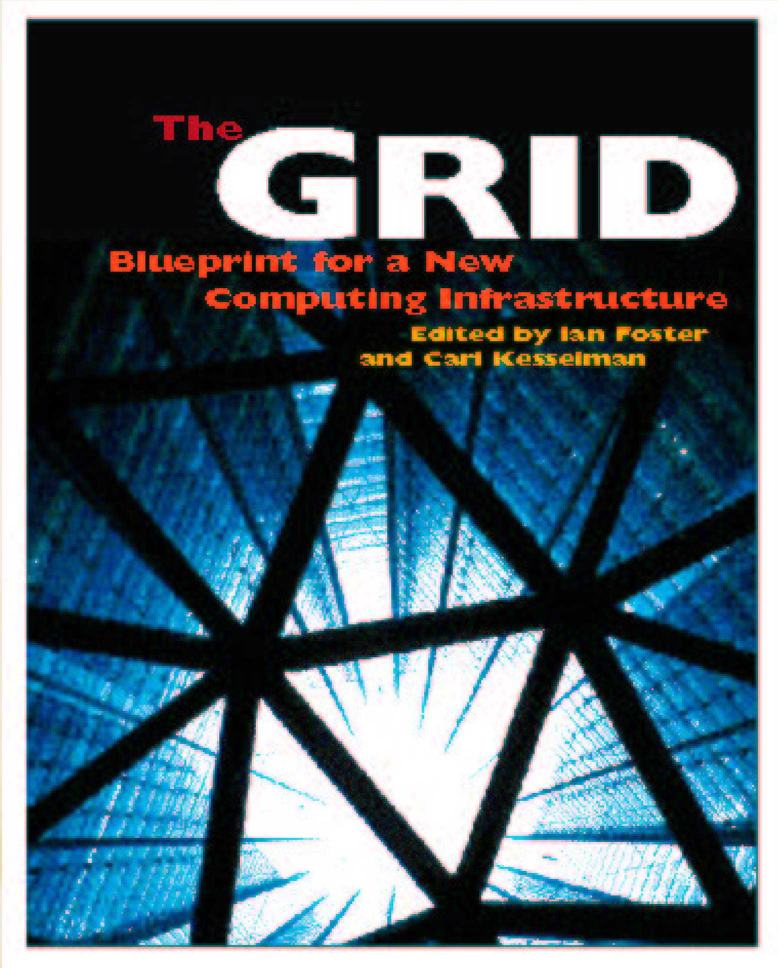

A Brief History

- Early 90s
  - Gigabit testbeds, metacomputing
- Mid to late 90s
  - Early experiments (e.g., I-WAY), academic software projects (e.g., Globus, Legion), application experiments
- 2001 – now
  - Major application communities emerging
  - Major infrastructure deployments
  - Growing technology base
  - Global Grid Forum: ~500 people, >90 orgs, 20 countries
- The “Grid problem” is about *resource sharing & coordinated problem solving in dynamic, multi-institutional virtual organizations*
  - Data is often the focus as opposed to classical numerically intensive simulations
  - Analogy with the power grid

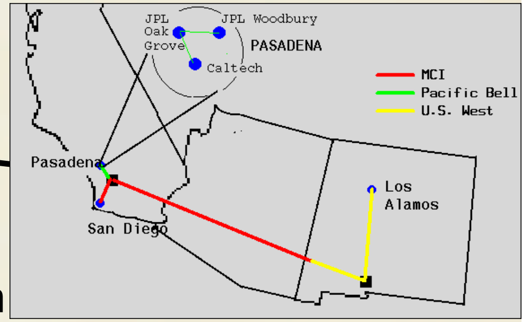

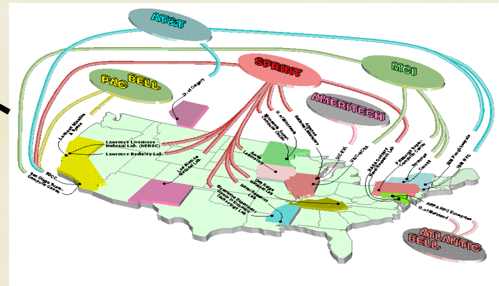

The Grid Vision

Researchers perform their activities regardless geographical location, interact with colleagues, share and access data

The GRID: networked data processing centres and "middleware" software as the "glue" of resources.

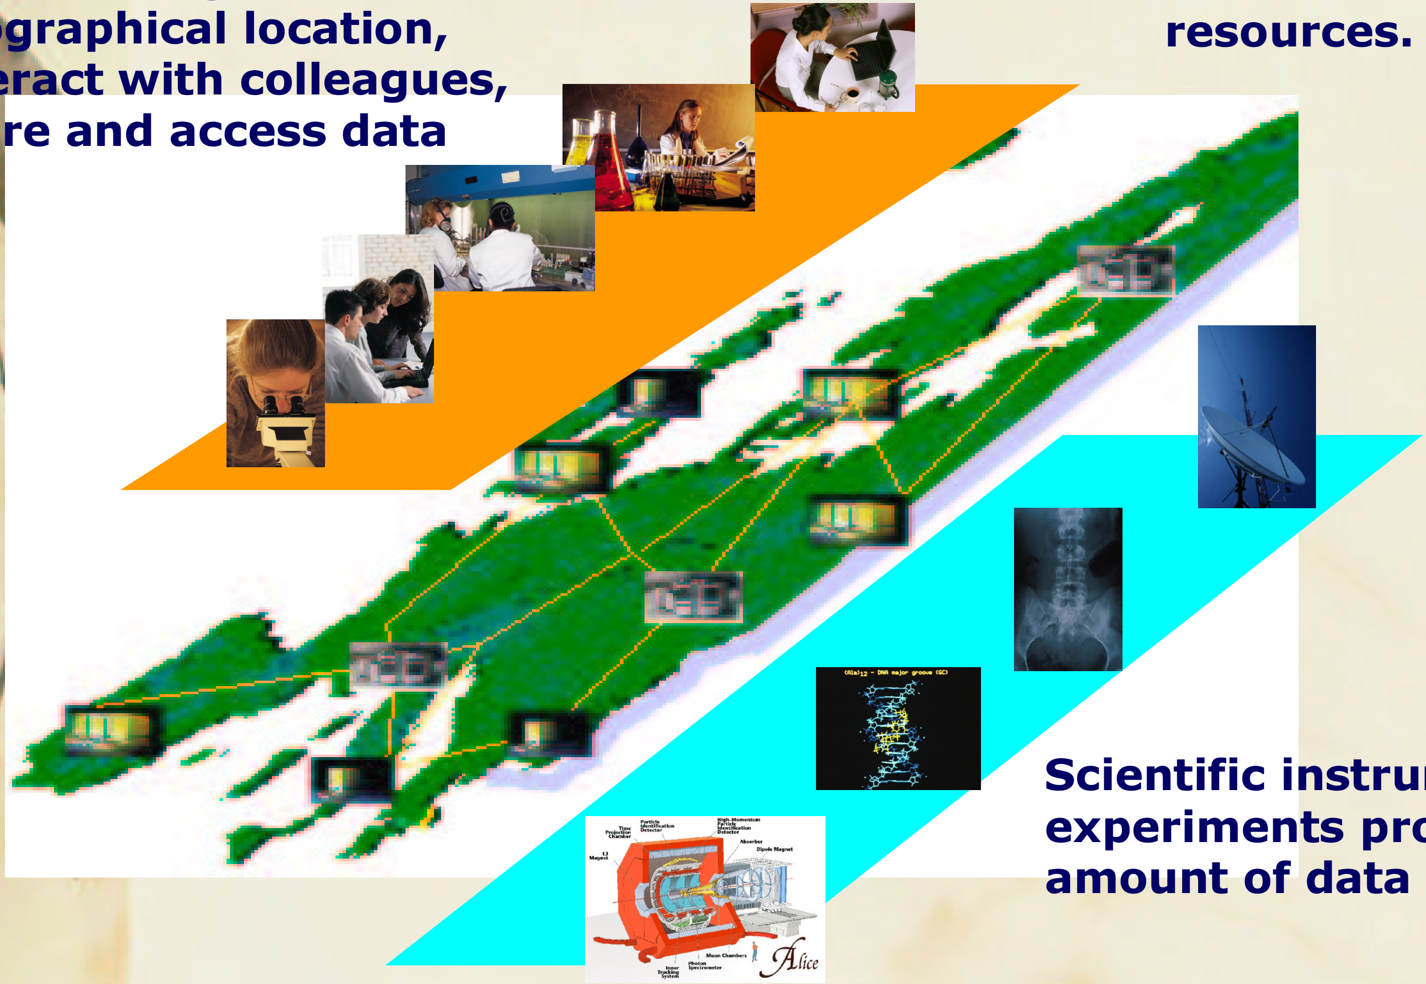

Scientific instruments and experiments provide huge amount of data

GLOBUS hourglass, a model for middleware

- Focus on architecture issues
  - Low participation cost, local control and support for adaptation
  - Use to construct high-level, domain-specific solutions
- A set of toolkit services
  - Security (GSI)
  - Resource management (GRAM)
  - Information services (MDS)
  - Remote file management (GASS)
  - Communication (I/O, Nexus)
  - Process monitoring (HBM)

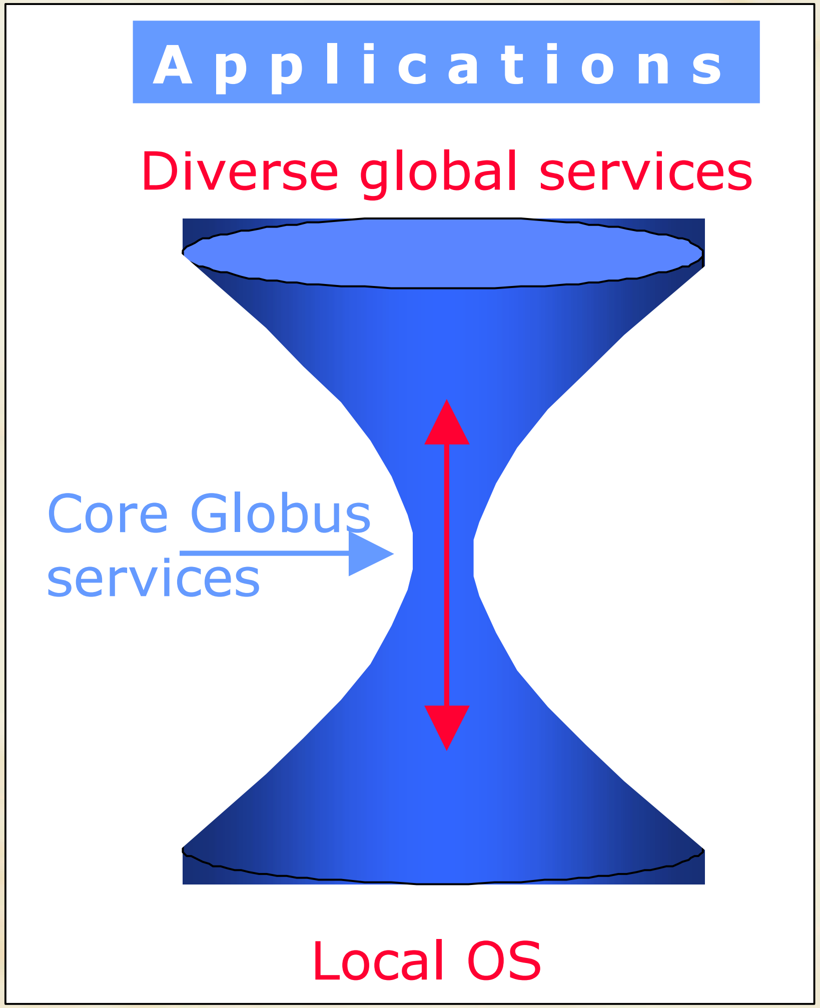

Instruments are expensive

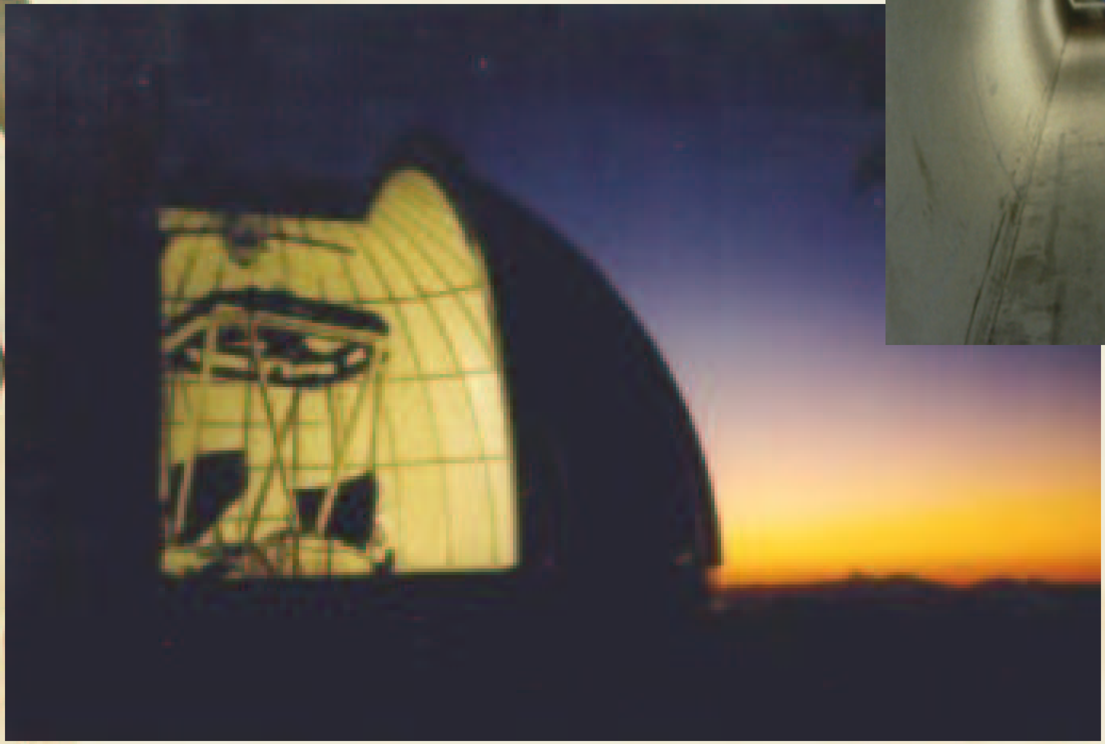

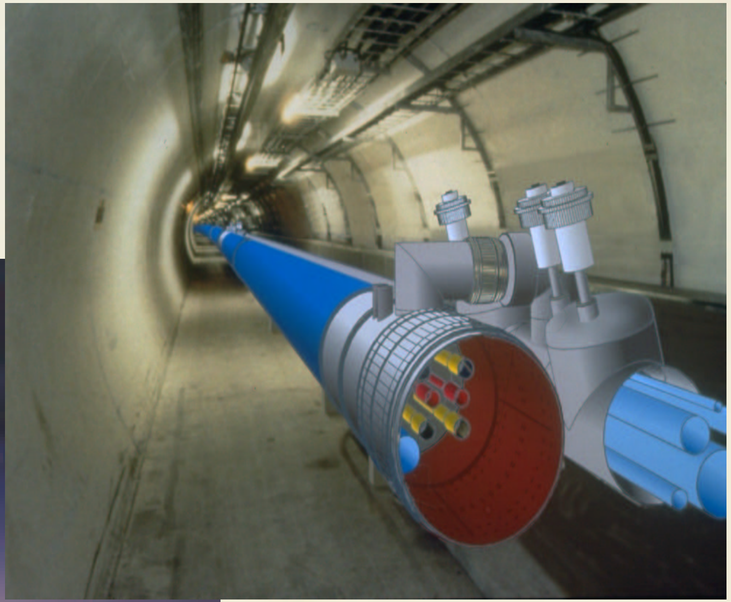

Virtual Observatory: an example of the scientific opportunities

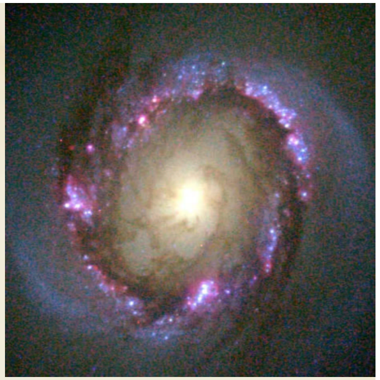

- The planned Large Synoptic Survey Telescope will produce over 10PB/y by 2008!
  - All-sky survey every few days: fine-grain time series for the first time
- The **National Academy of Sciences** recommends, as a first priority, the establishment of a **National Virtual Observatory**
  - http://www.nap.edu/books/0309070317/html/

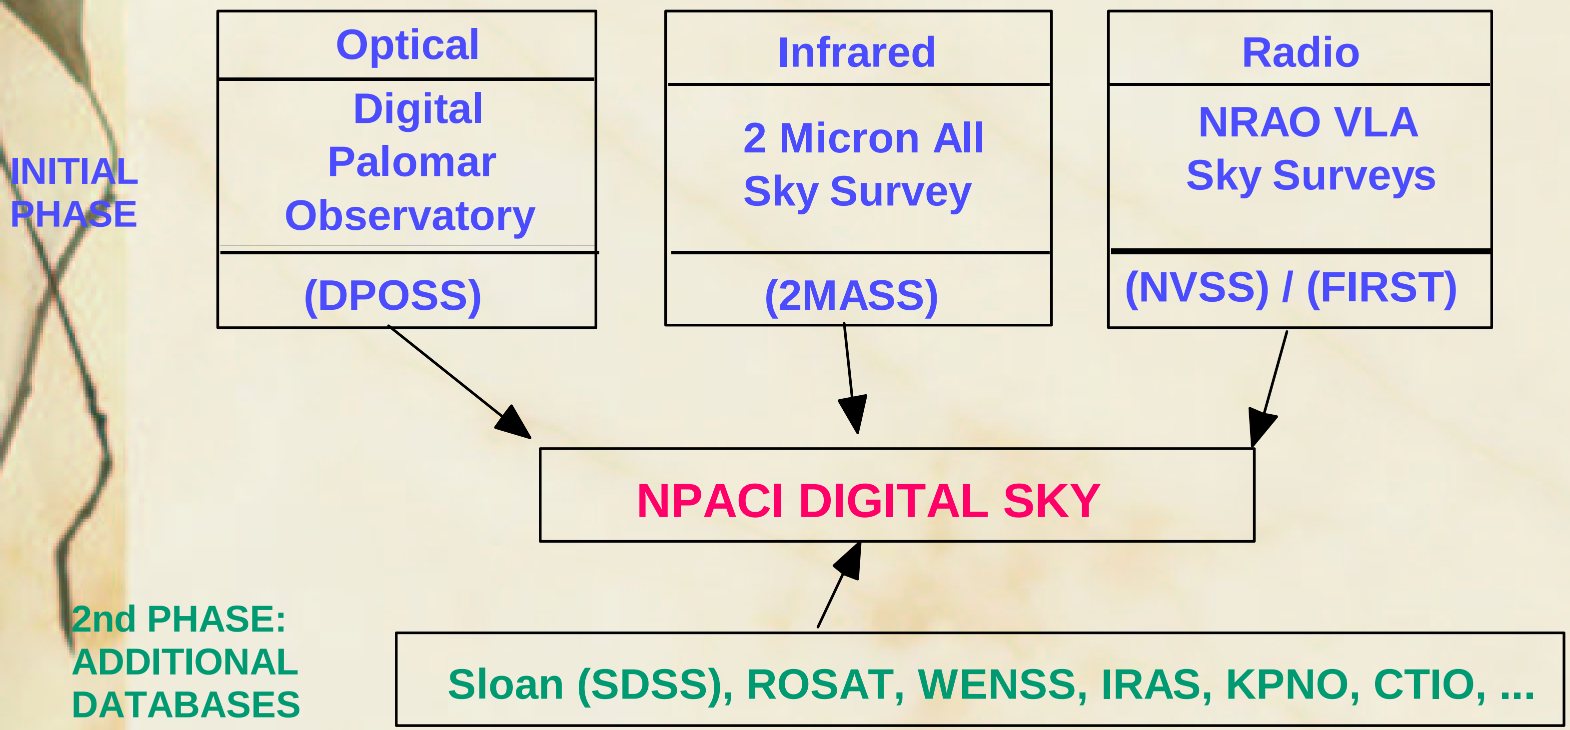

Crab Nebula in 4 spectral regions

X-ray, optical, infrared, radio

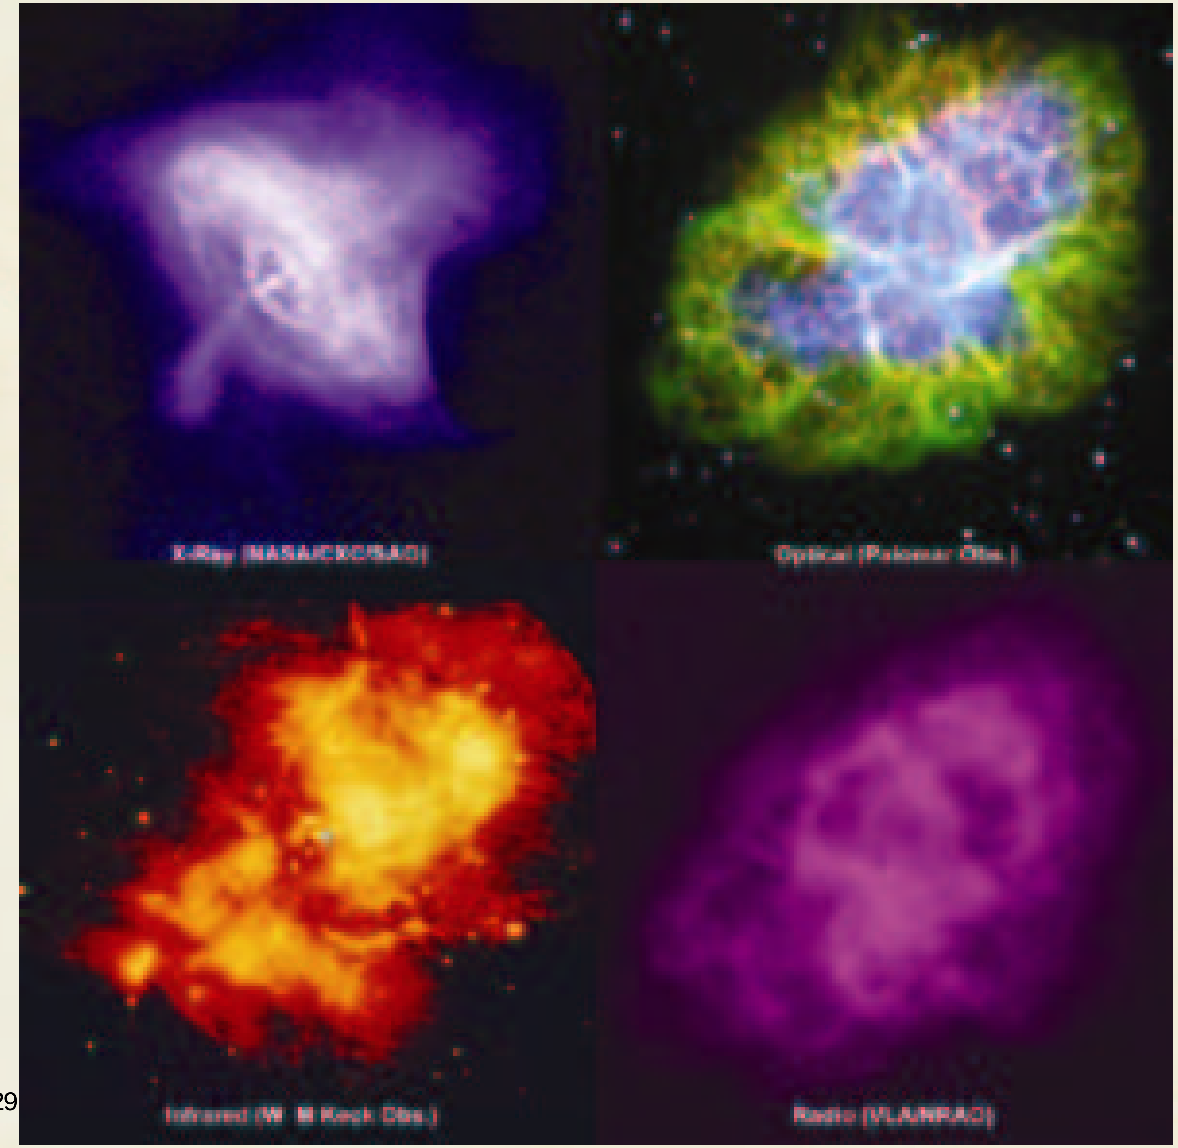

Virtual Sky: Image Federation

http://virtualsky.org/
from
Caltech CACR
Caltech Astronomy
Microsoft Research

Virtual Sky has
140,000,000 tiles
140 Gbyte

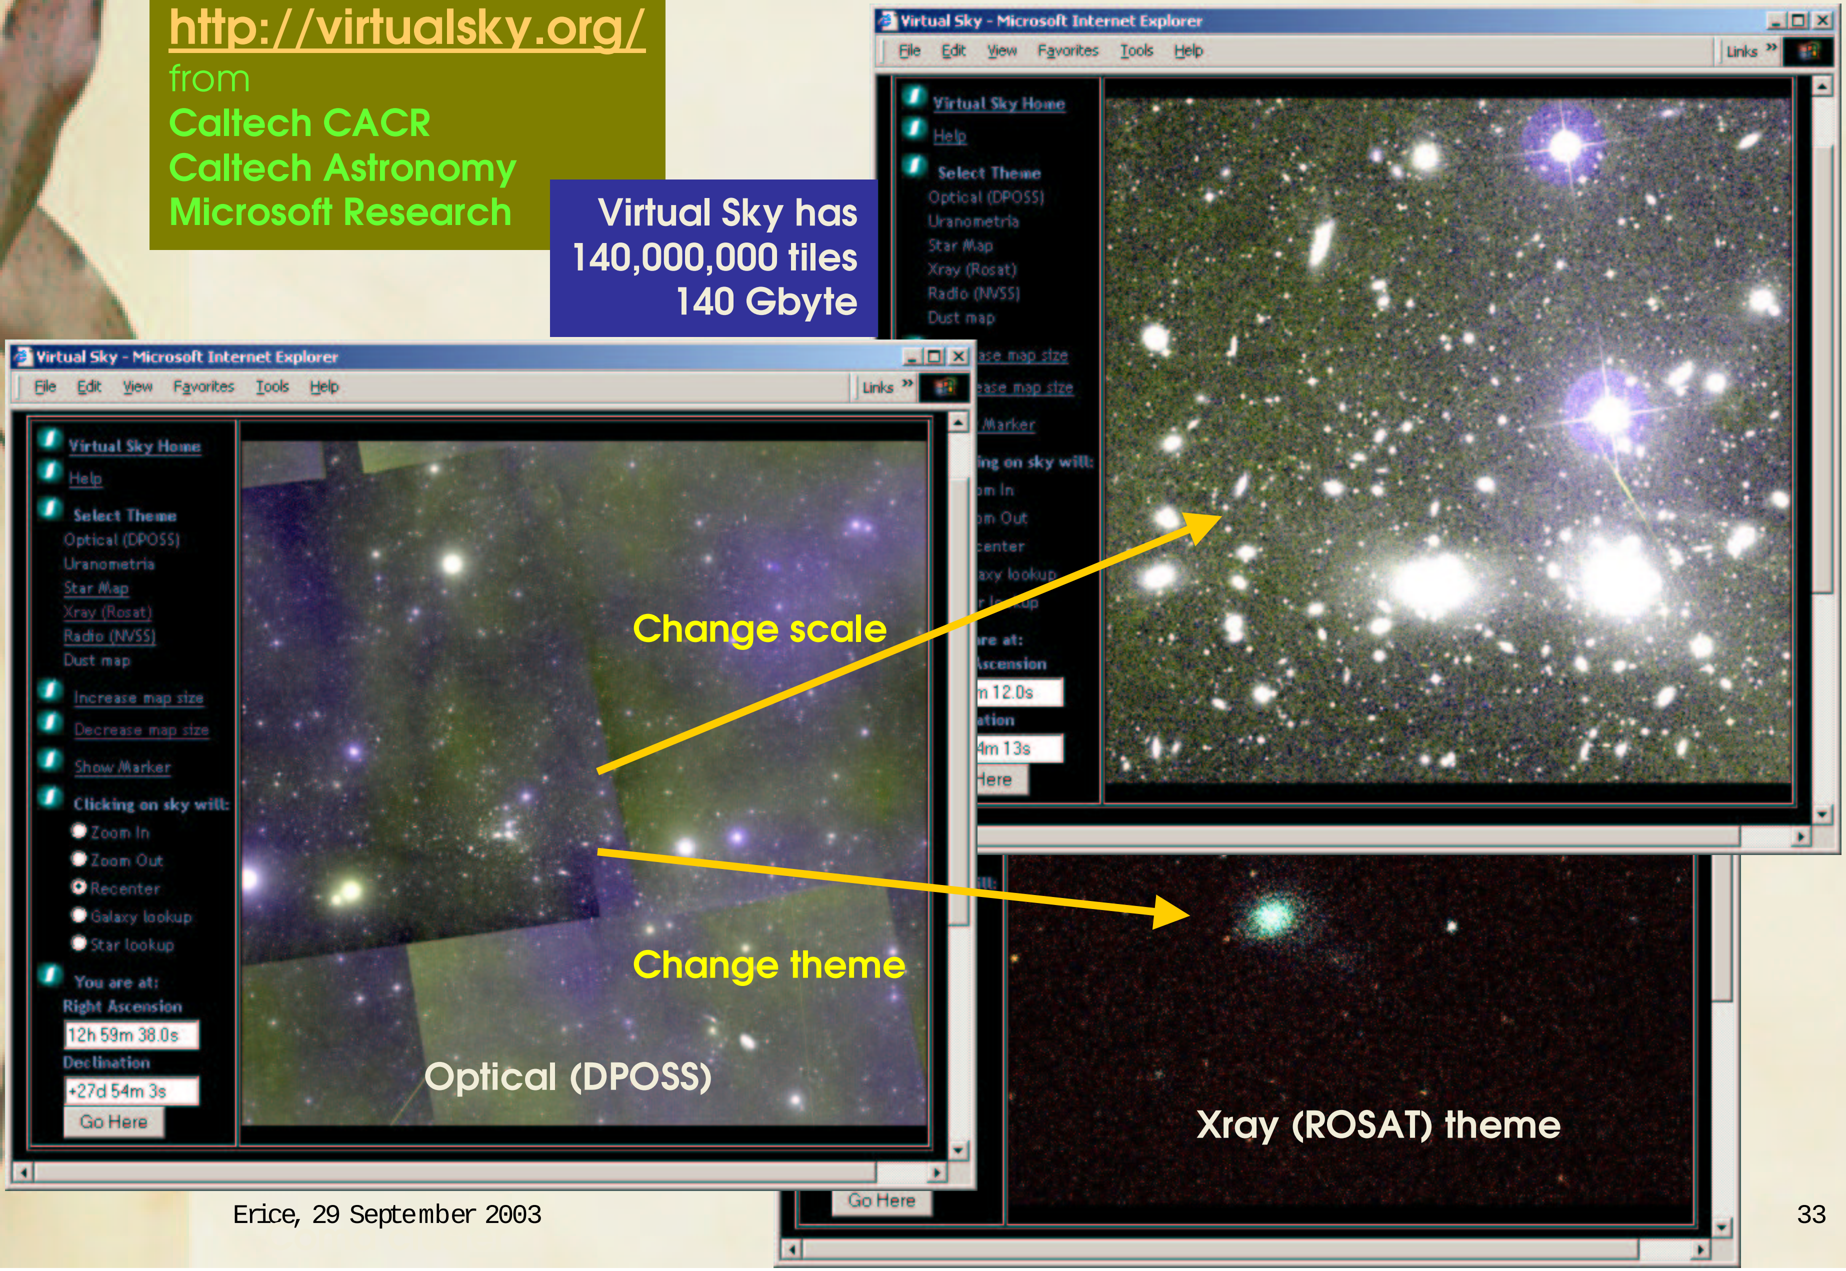

Not a centralized resource

- The catalogs are NOT ingested into the Virtual Sky server, but left in place (maintained by their curators), and accessed by the accompanying web service
- This is an important point: leave the resources and the responsibility for managing them where they would normally be

Results Possible on TeraGrid

Modelling Cell Structures

- **Pre-Blue Horizon:**
  - Possible to model electrostatic forces of a structure with up to **50,000** atoms -- a single protein or small assembly
- **Pre-TeraGrid:**
  - Possible to model **one million** atoms – enough to simulate drawing a drug molecule through a microtubule or tugging RNA into a ribosome
- **TeraGrid:**
  - Models of **10 million** atoms will make it possible to model function, structure movement, and interaction at the cellular level **for drug design and to understand disease**

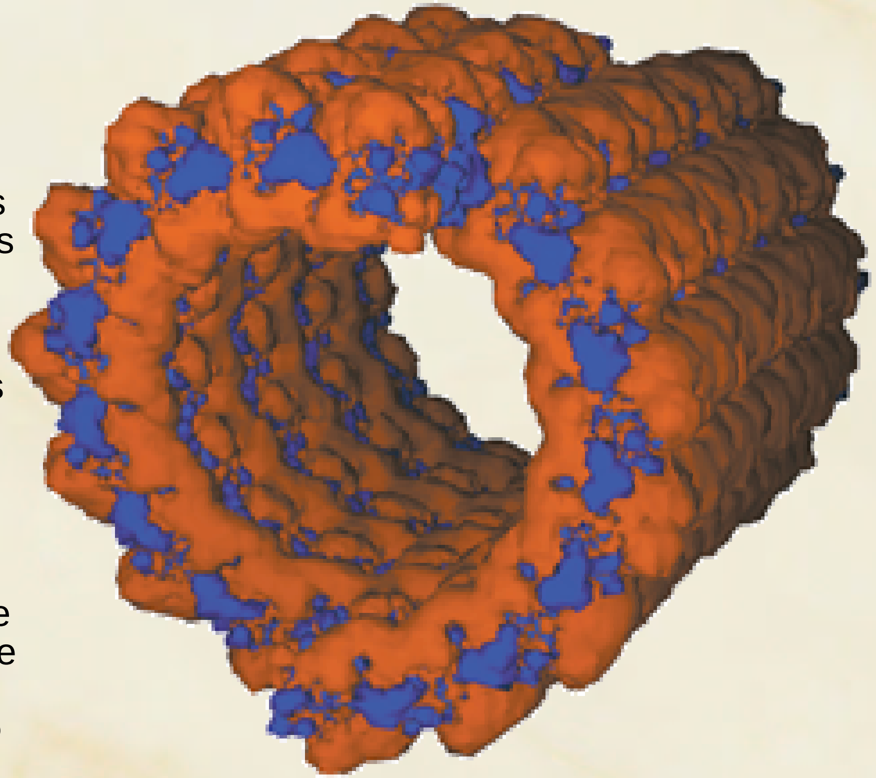

Baker, N., Sept, D., Joseph, S., Holst, M., and McCammon, J. A. *PNAS* **98**: 10037-10040 (2001).

mammograms X-rays
MRI
cat scans
endoscopies, ...

Digital Radiology (Hollebeek, U. Pennsylvania)

- Hospital Digital Data
  - Very large data sources - great clinical value to digital storage and manipulation and significant cost savings
  - 7 Terabytes per hospital per year
  - dominated by digital images
- 2000 Hospitals x 7 TB per year x 2 = 28 PetaBytes per year

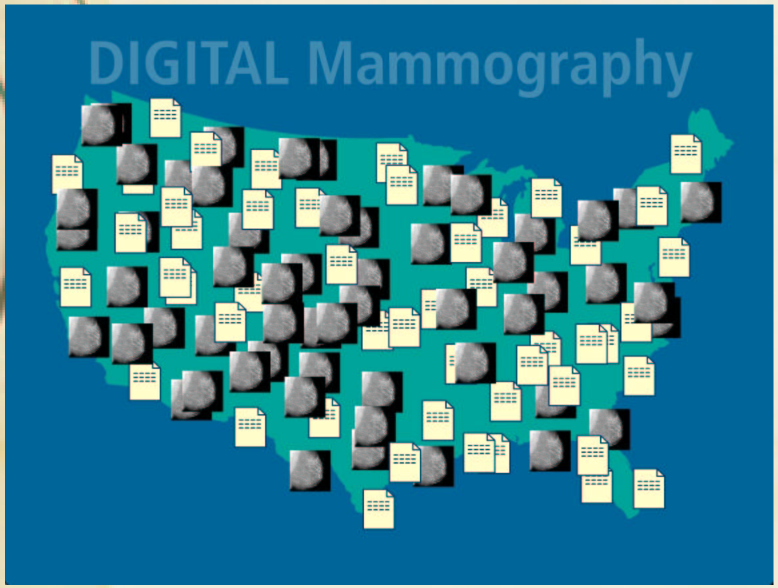

Highly Distributed Source

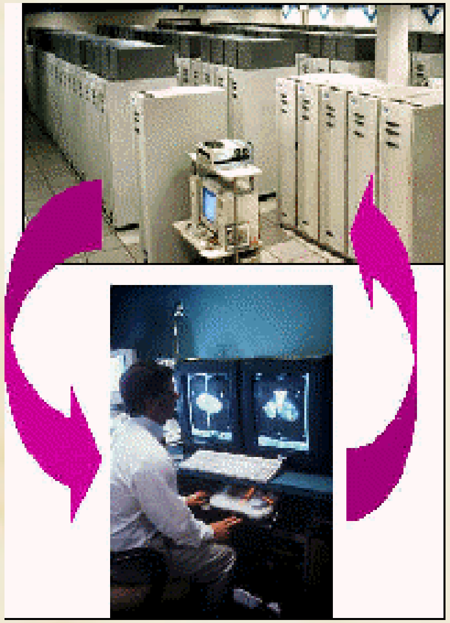

Hierarchical Storage and Indexing

Resource brokering (easy)

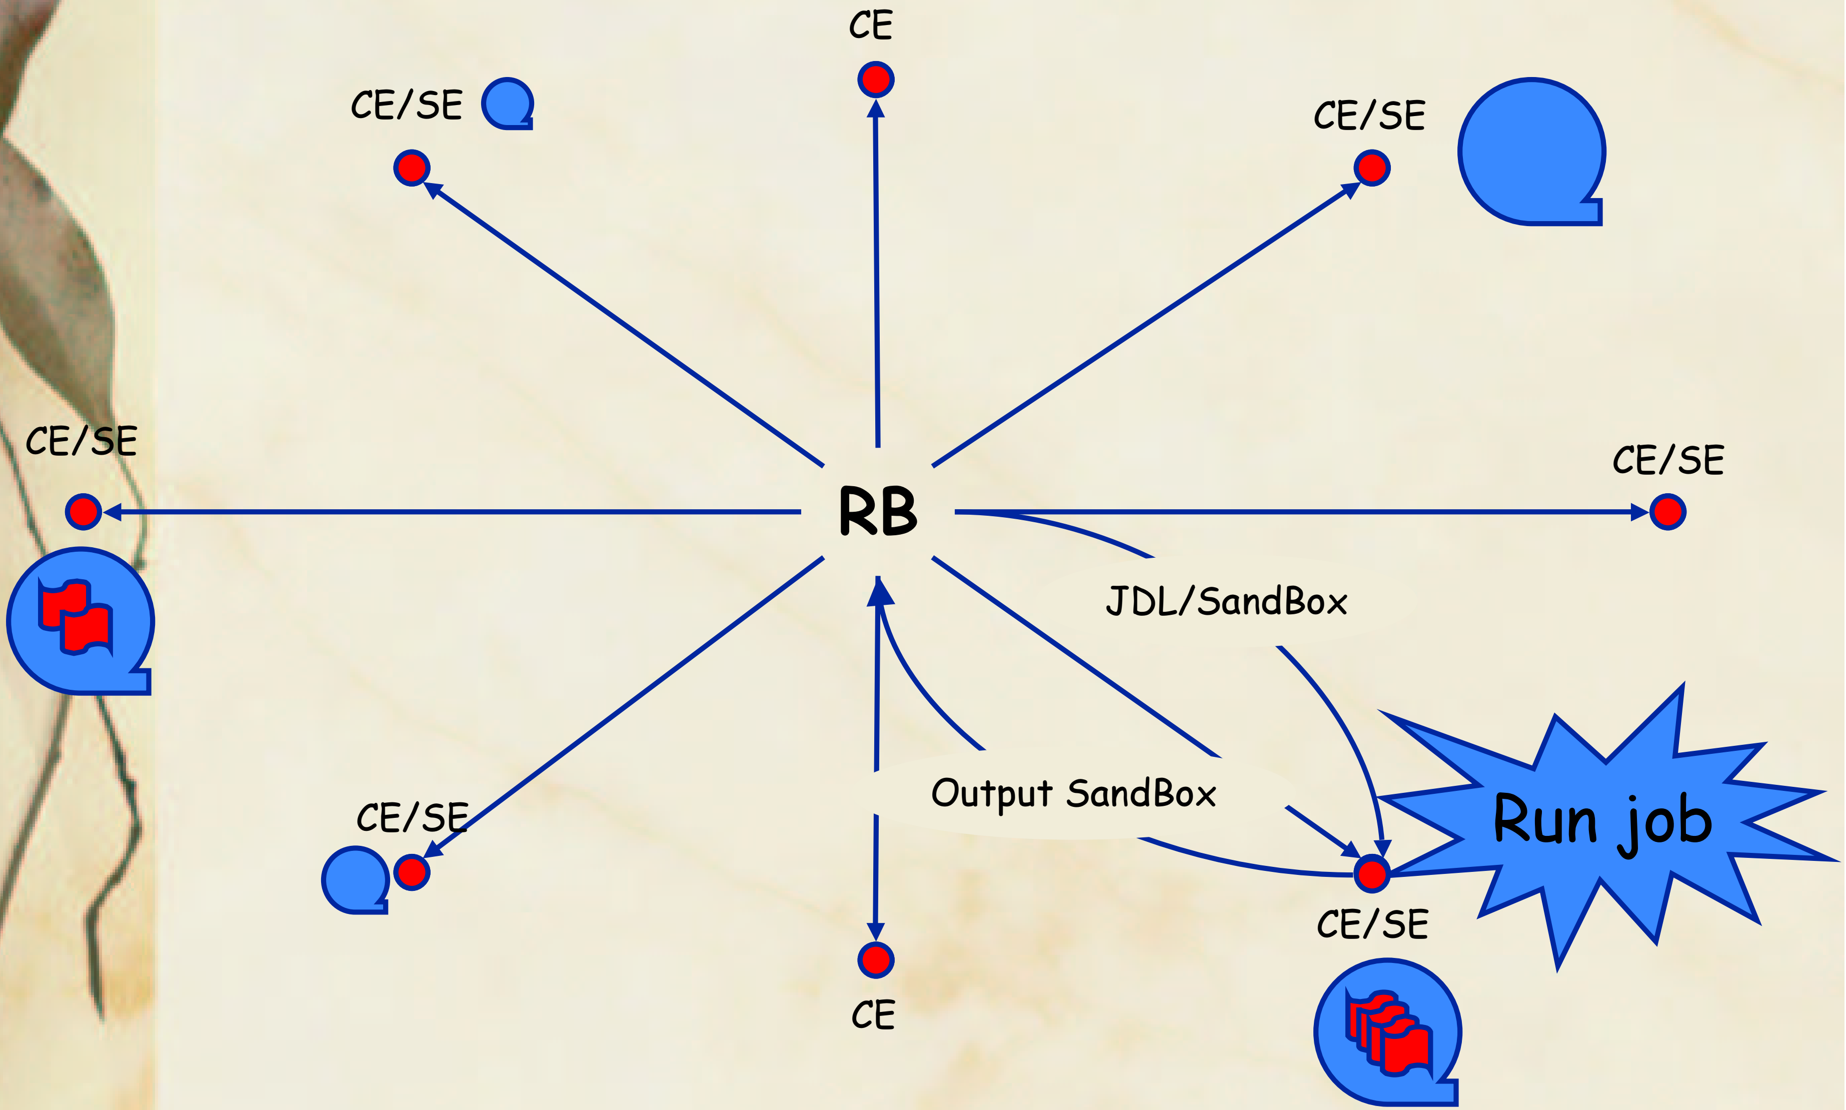

Resource brokering (in reality)

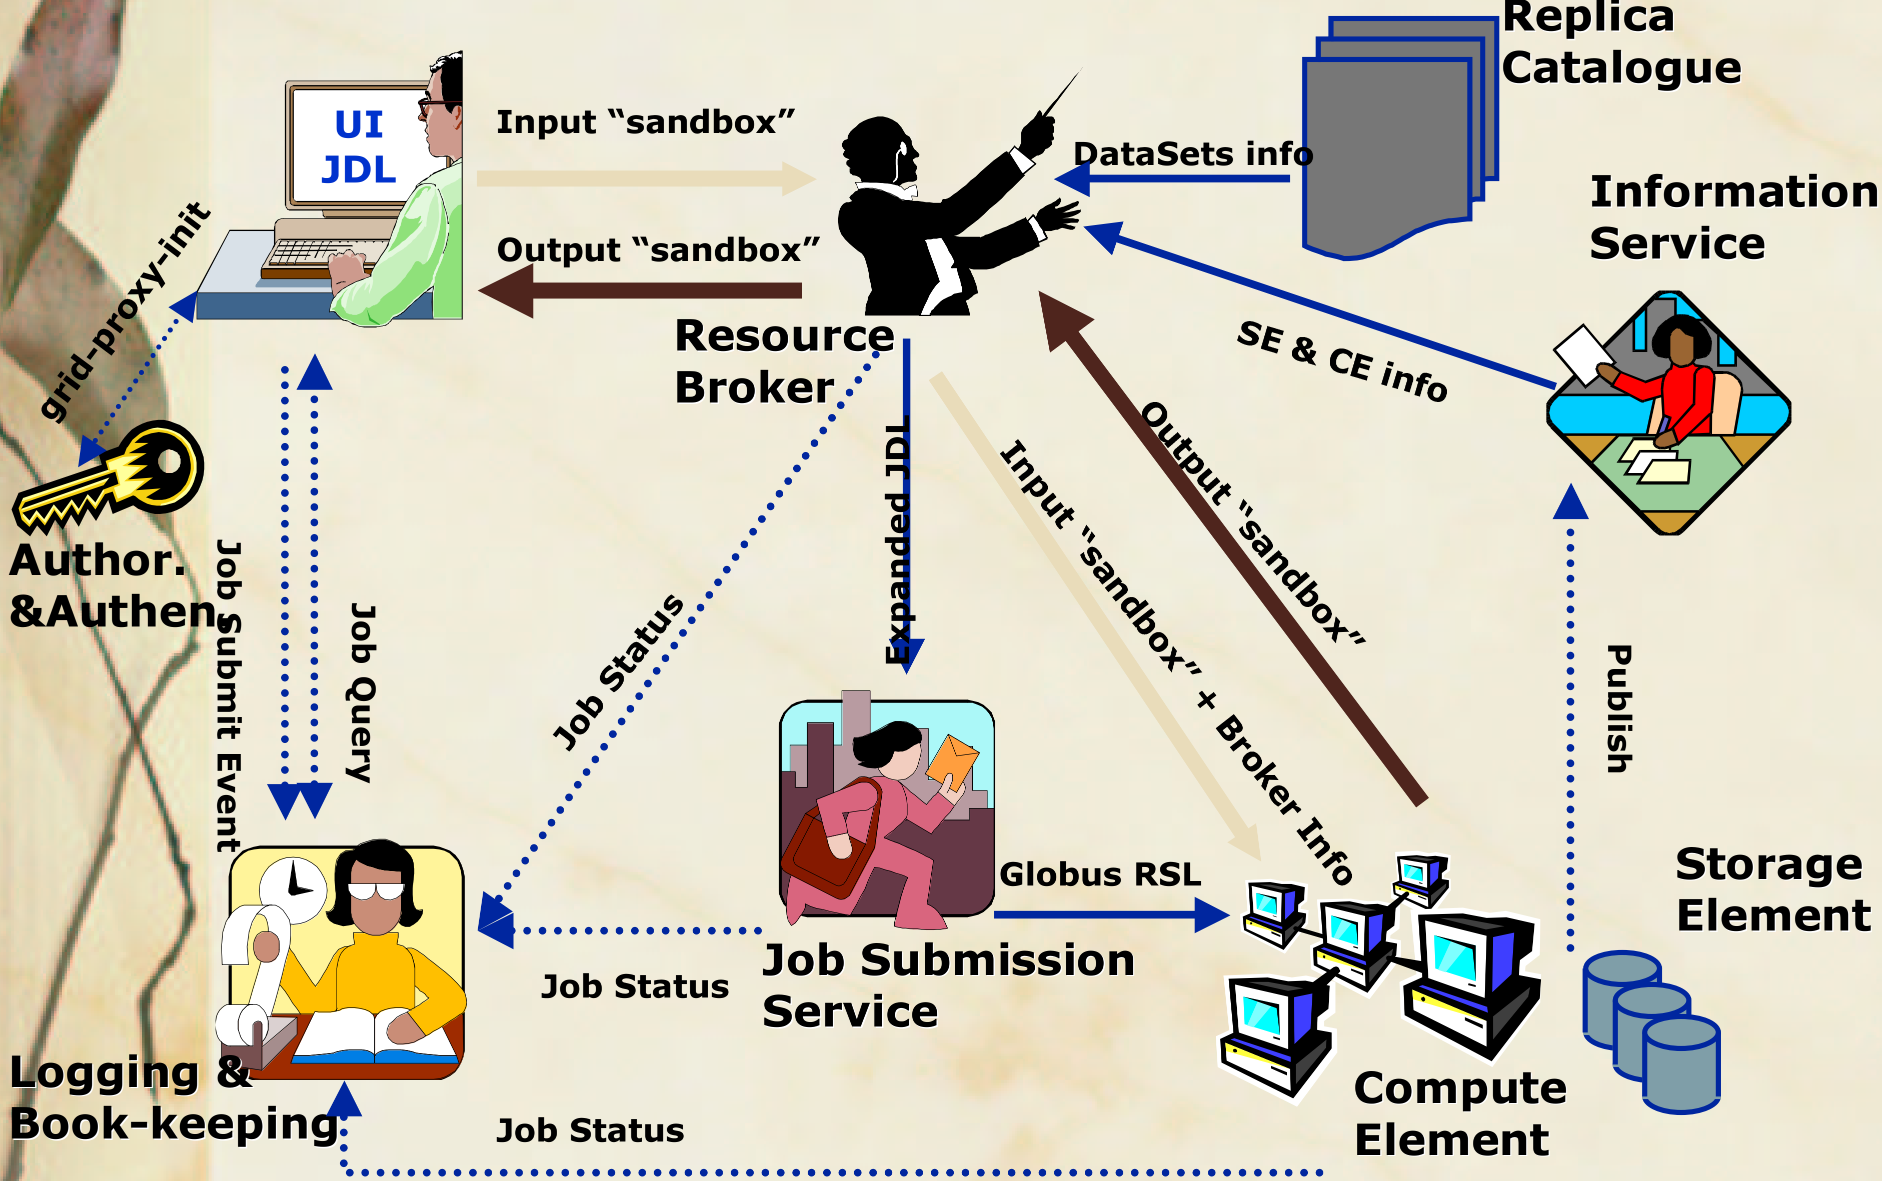

Grid projects

- Several GRID projects have been launched by EU and US funding agencies
- They have all started designing “the GRID”
  - Although based on common components such as GLOBUS
- Tremendous richness of architectures and products
- But worrying lack of stable testbeds where to experiment and provide feedback
- At the moment only friendly and advanced users can use the system
- Which of course creates a vicious circle…

An alternative approach

- Standards are now emerging for the basic building blocks of a GRID
  - There are millions lines of code in the OS domain dealing with these issues
- Why not using these to build the *minimal GRID* that *does the job*?
  - Fast development of a prototype, no problem in exploring new roads, restarting from scratch etc etc
  - Hundreds of users and developers
  - Immediate adoption of emerging standards
- An example, AliEn by ALICE (5% of code developed, 95% imported)

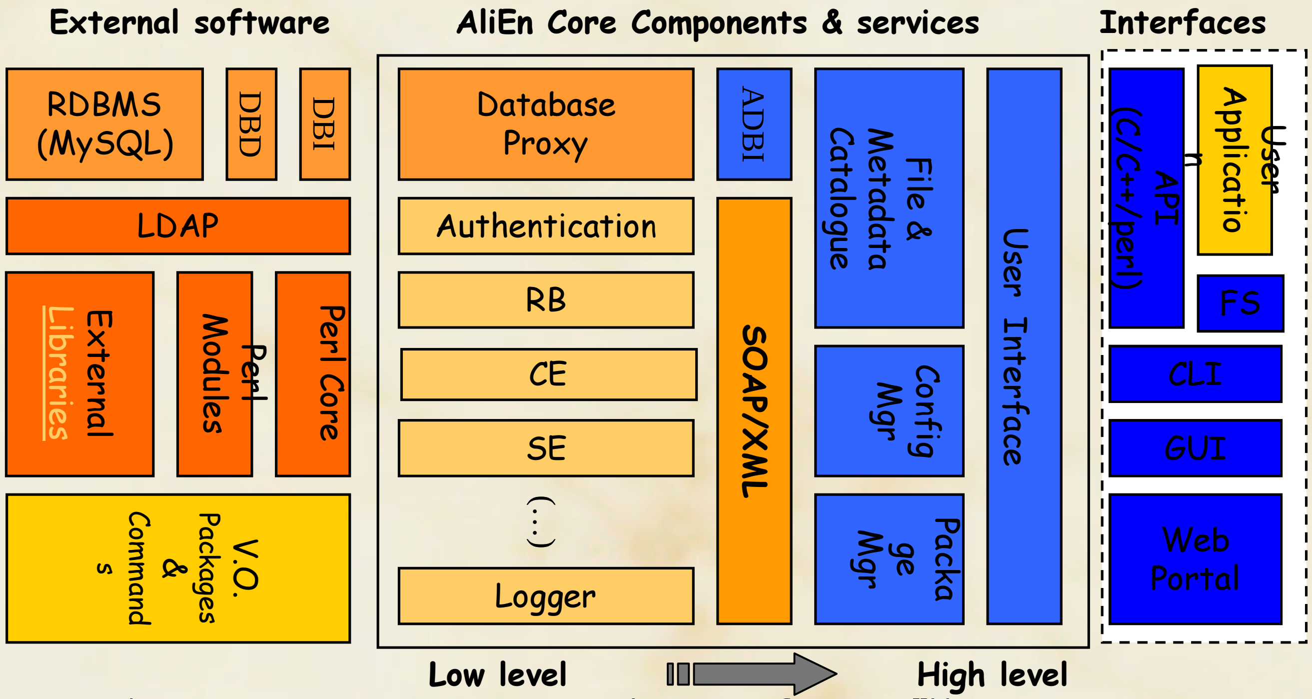

AliEn activity

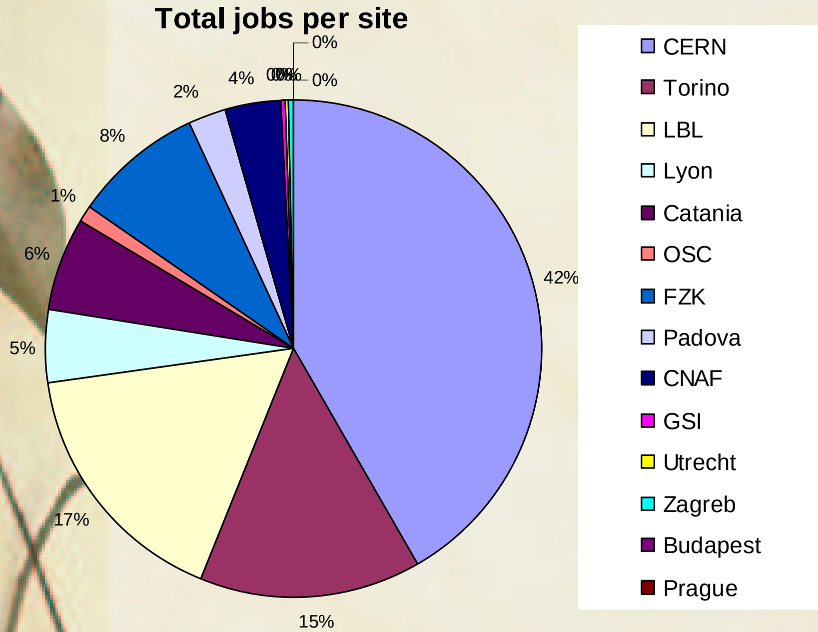

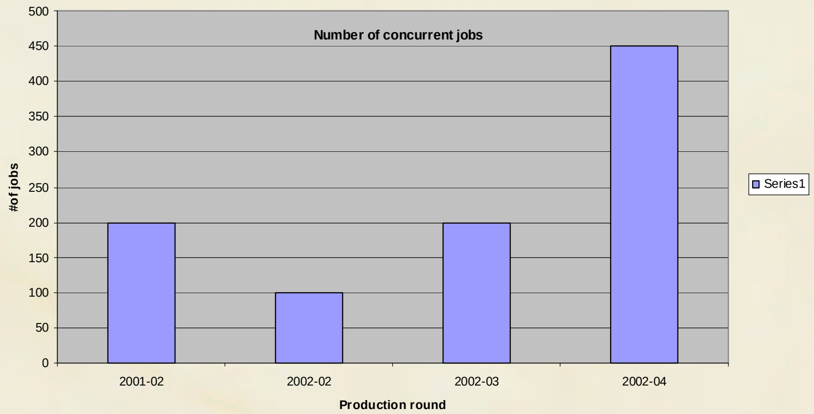

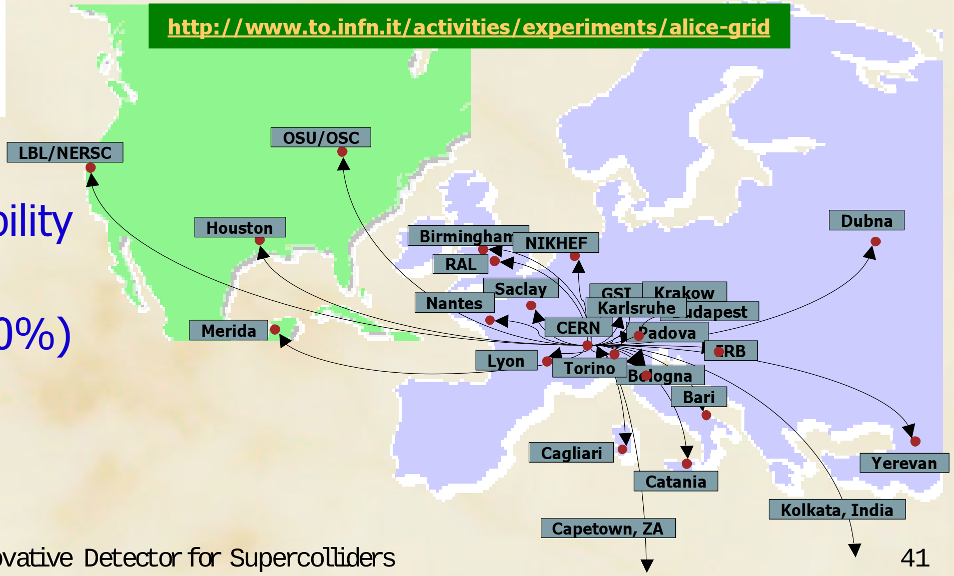

- 32 sites configured
- 5 sites providing mass storage capability
- 12 production rounds
- 22773 jobs validated, 2428 failed (10%)
- Up to 450 concurrent jobs
- 0.5 operators

Distributed analysis model

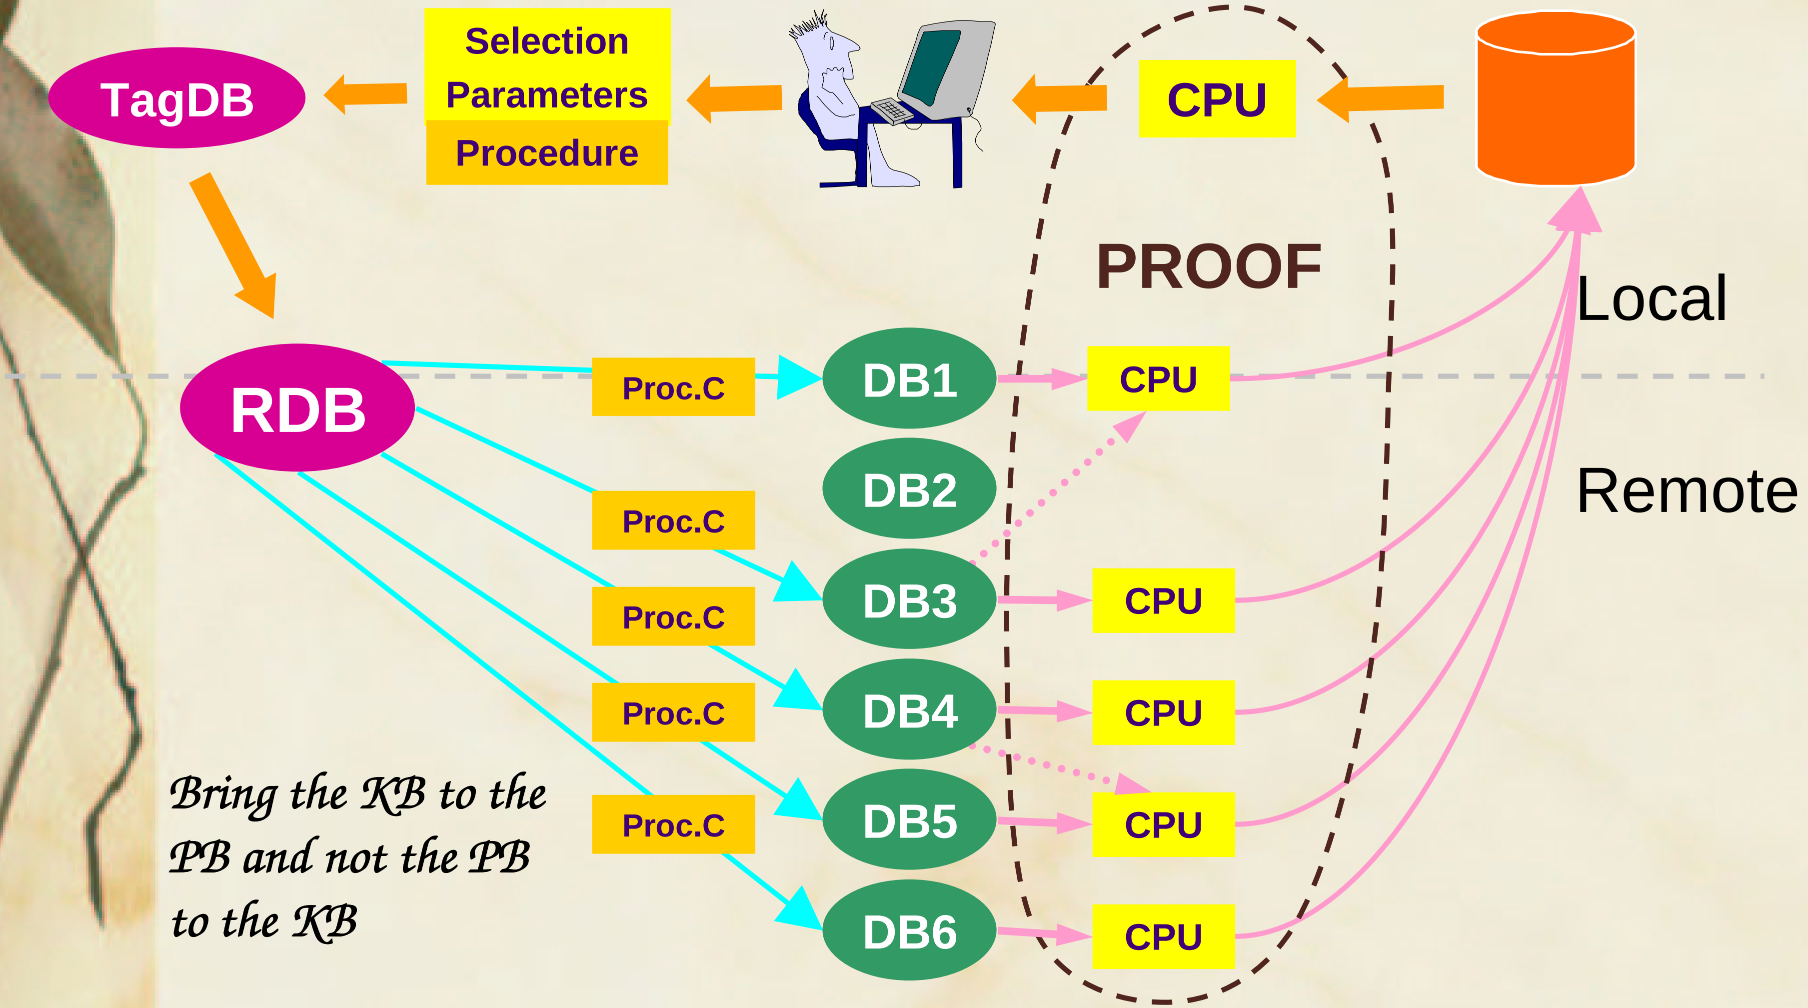

Bring the KB to the PB and not the PB to the KB

The GRID universe

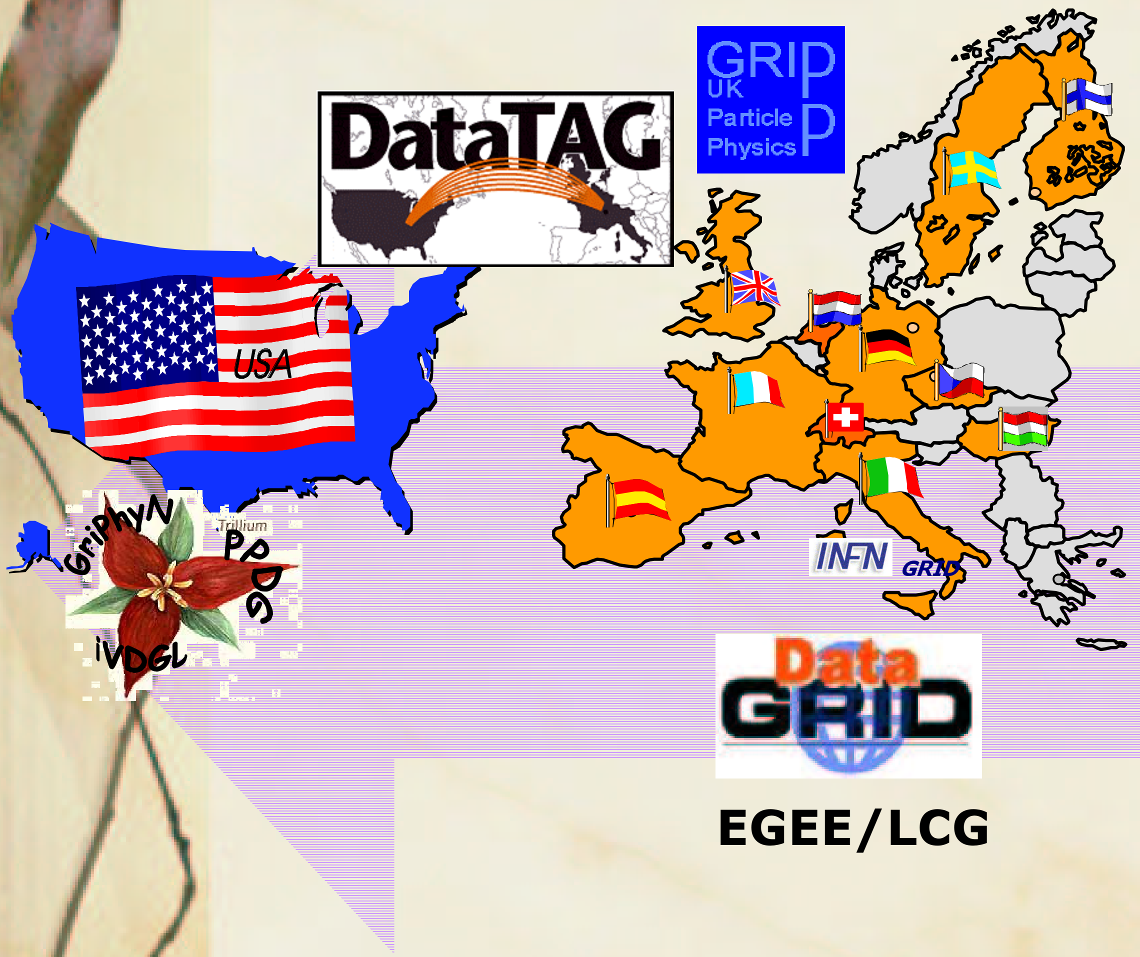

- **The real problem is how to avoid divergence!**

MetaGrid and Grid federations

- It has been realised that there will be many Grids and not a single one
  - However users will not want to learn more than one
- The concept of Grid federation and Meta Grid are now explored
- Unfortunately this sometimes looks like building on sand…
  - As we still do not have a stable base on which to build
- And it does not help an early adoption and response from the users

Meta and Federated grids

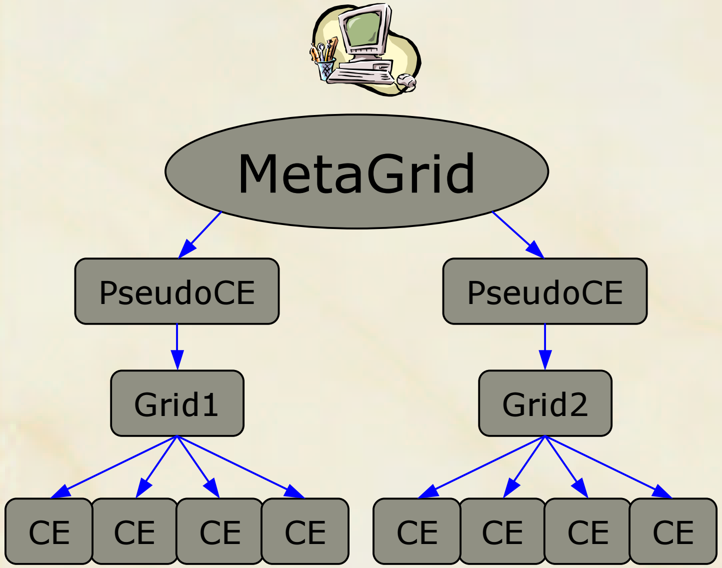

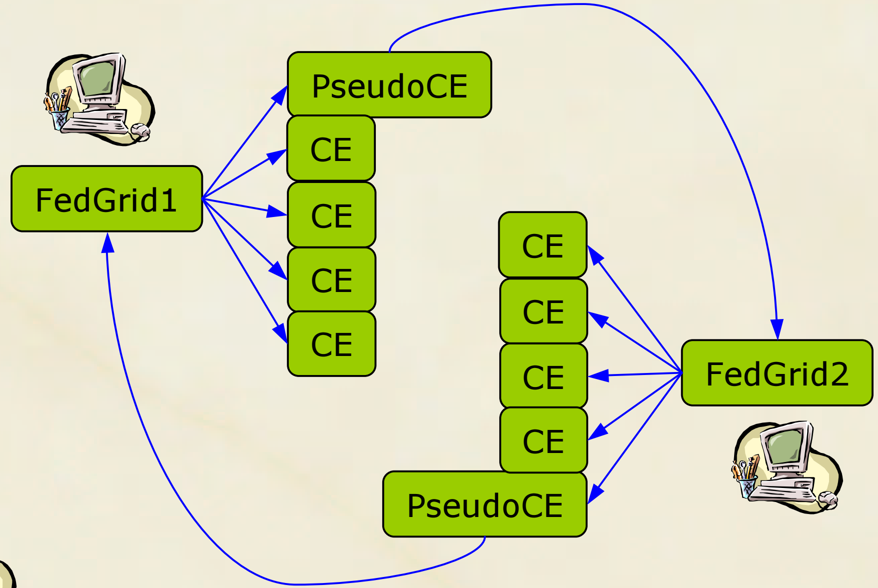

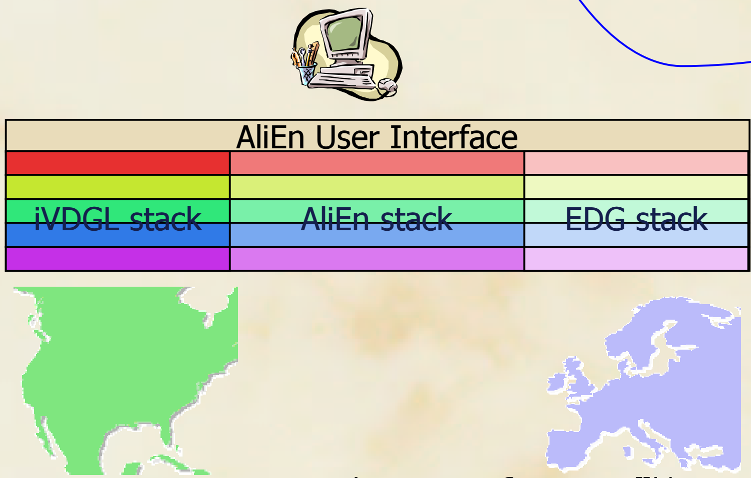

A new standard emerging?

- Grid Services are defined by OGSA: what Grid Services should be capable of, what types of technologies they should be based on
- Grid Services are specified by OGSI, which is a formal and technical specification of the concepts described in OGSA, including Grid Services
- Grid Services extend Web Services
- Suppose you want to build a new house, you need
  - An architecture (OGSA)
  - A detailed engineering plan (OGSI)
  - Workers that build the hous (the real GRID)

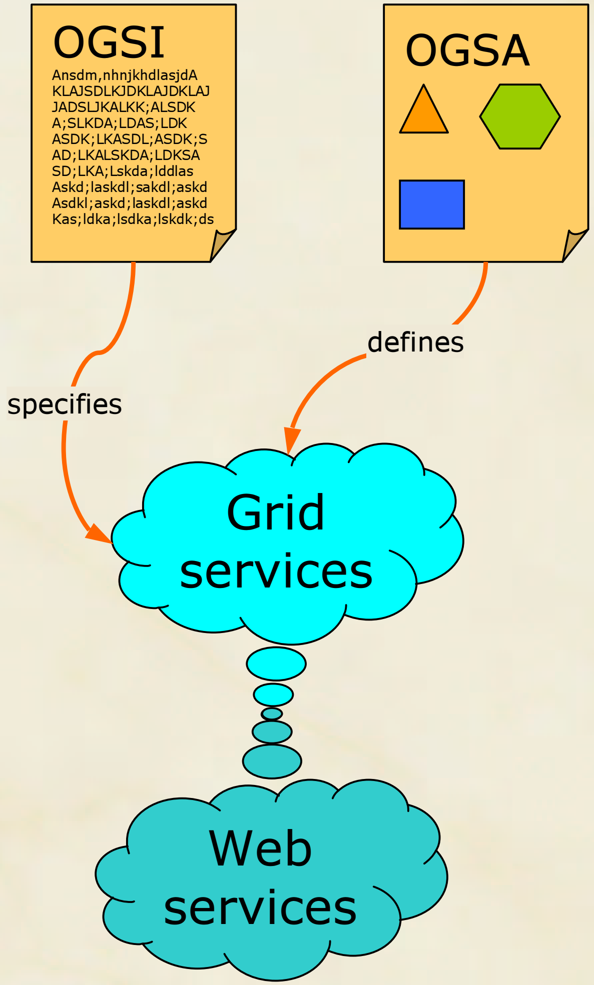

Standard interoperable technologies
XML, WSDL, SOAP

LCG new strategy

- LHC experiments are proposing to LCG to
  - Start from the architecture of a demonstrated working product (strongly inspired from AliEn)
  - Develop components based on standards and/or adopt existing components
  - Deploy very quickly a prototype and expose it to the users
  - Refine iteratively the architecture and the services as needs and standards evolve
  - Don’t write doc, write working code!
- Apparently LCG is receptive to these arguments

Software development

- In the *LEP era* the code was 90% written in FORTRAN
  - ~10 instructions, 50 pages!
- In the *LHC era* the code is in many cooperating languages, mainly C++
  - ~ 50 instructions, 700 pages – nobody understands it completely (B.Stroustrup)
- Users are heterogeneous, sparse and without hierarchical structure
  - From very expert analysts to users, from 5% to 100% of time devoted to computing
- People come and go with a very high rate
  - Programs have to be maintained by people who did not develop them
  - Young physicists need knowledge they can use also outside physics
- Modern SE ("Agile Methodologies") propose to value

Individuals and interactions		*processes and tools*
Working software	**OVER**	*huge documentation*
Customer collaboration		*contract negotiation*
Responding to change		*following a plan*

That is, while there is value in the items on the right, we value the items on the left more.

HEP, LHC & GRID

- Funding for HEP is becoming scarce
- There is a serious personnel deficit (also!) in HEP computing
- The exceptional interest spreading in most countries for GRID resulted in an acute need for GRID-trained CSs
- HEP (which invented the web) looked as an ideal place where to train young scientists in GRID technologies
- This was seen as a unique opportunity to alleviate the personnel problems of LHC computing

HEP, LHC & GRID

- People are sent at CERN to work and train on GRID
- LHC experiments may greatly profit from GRID
- CERN has experience in large distributed collaborations
- This could be a good deal, however
  - Mostly young and non-experienced CSs are sent at CERN
    - No knowledge of HEP habits and needs, little experience
  - CERN has no record in distributed computing research
  - LHC computing needs go beyond GRID middleware
    - There is a pressure to launder personnel into other roles
  - People come at CERN with agendas and constraints
  - GRID developed at CERN is specific for the need of HEP
    - However a working middleware for HEP would go a long way in satisfying most applications

GRID as the environment for science?

- The killer applications are there
  - Biomedical, Environmental, Industrial
- Governments see major potential benefits, economic and societal
- Technologies continue to improve with resources in Peta units
- And then miracle happens…
  - Funding agencies worldwide have MOUs for sharing resources
  - Research projects give timely access to their data
  - GGF etc successfully fostered GRID standards
  - Grid MW is included in the system software

What about the people?

- Opportunity to work on the forefront of development
  - Pushing technology boundary
- Profit from and contribute to GRID research from any location
  - No need of big computers or fast network to be part of the game!
  - Middleware comes from brainware, not hardware!
  - Contributions from “peripheral” areas can have the same impact than large and established institutions
- The old dream of low-inertia hi-tech industry coming true?
- Work now is on basic principles, almost at a philosophical level
  - Prototypes can be assembled quickly and have a large impact
- Agile Technologies are widely used in GRID developments

Summary

- Simulation is a fundamental issue
  - Fundamental work is still needed
  - Finally LHC seems to realise it!
- GRIDs may not be “the solution”, but they will be part of it, helping
  - Make “more real” the *third methodology* for scientific research, alongside experiment and theory
  - Expand the pool of people who can do forefront research
  - Leverage investments in research infrastructure