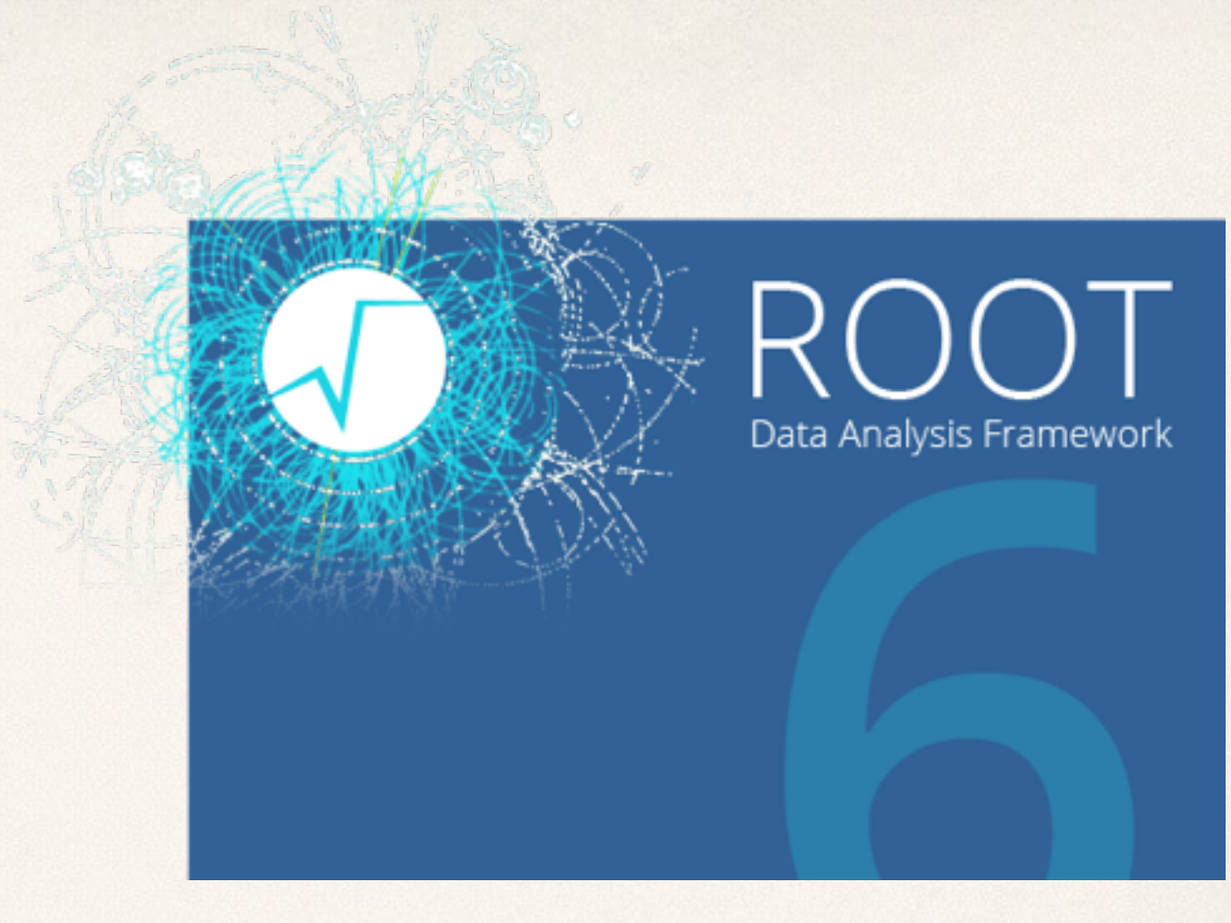

ROOT
A framework for Big data analysis

Pere MATO, CERN

13/06/2014

What do **Events** Look Like?

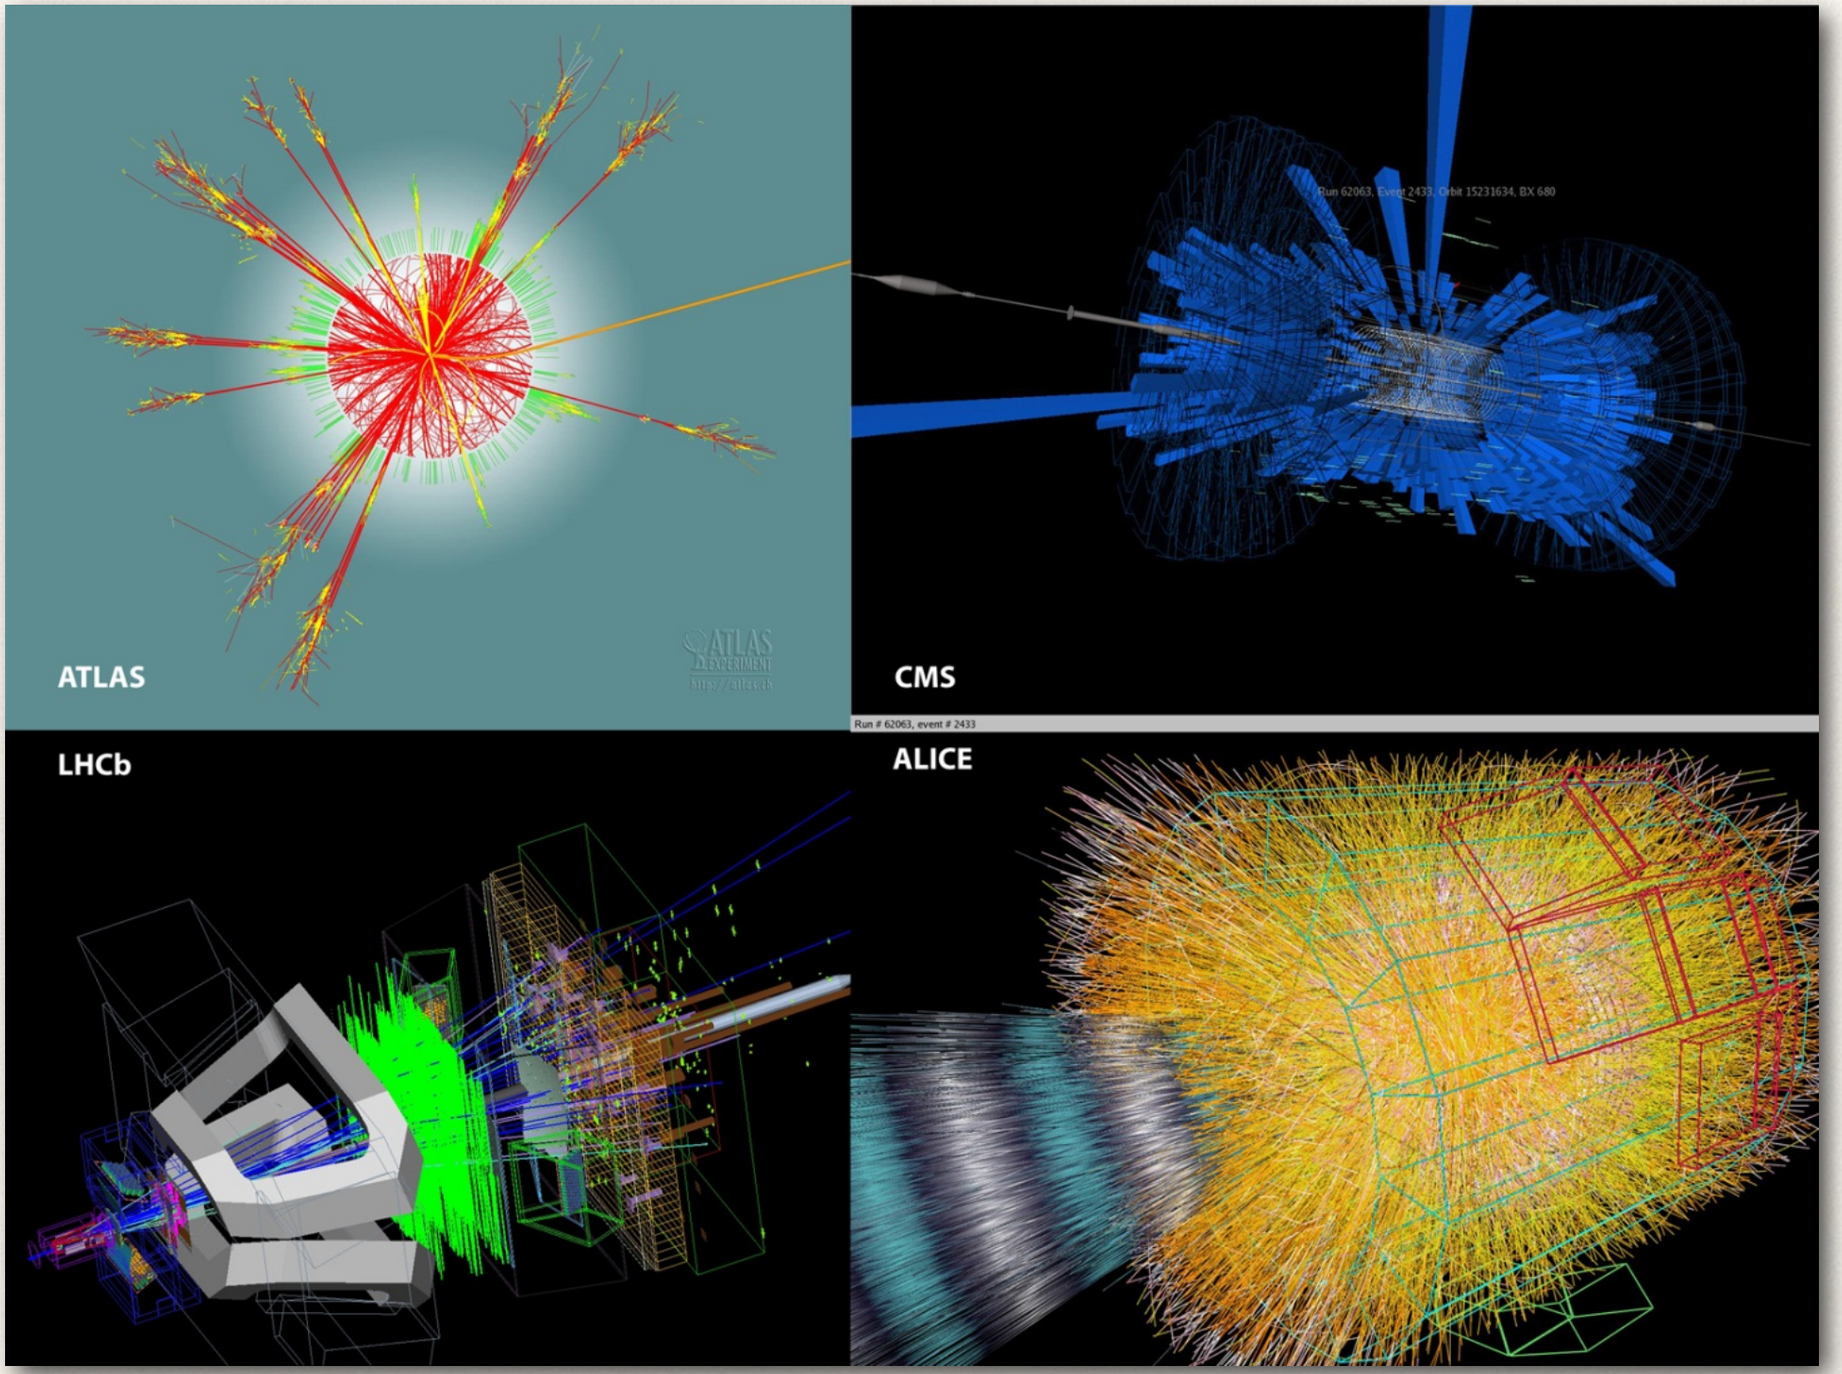

ROOT
Data Analysis Framework

"Event" == data produced in a particle collision (proton-proton)

The needle in the hay-stack

- $\sigma(pp) = 70$ mb $\rightarrow 7 \times 10^8$ / s (!)
- In ATLAS and CMS
  20 – 30 minimum-bias events overlap
- $H \rightarrow ZZ$
  $Z \rightarrow \mu\mu$
- $H \rightarrow$ 4 muons: the cleanest ("golden") signature

p-p Collisions at 14 TeV at 10^{34} $cm^{-2}s^{-1}$

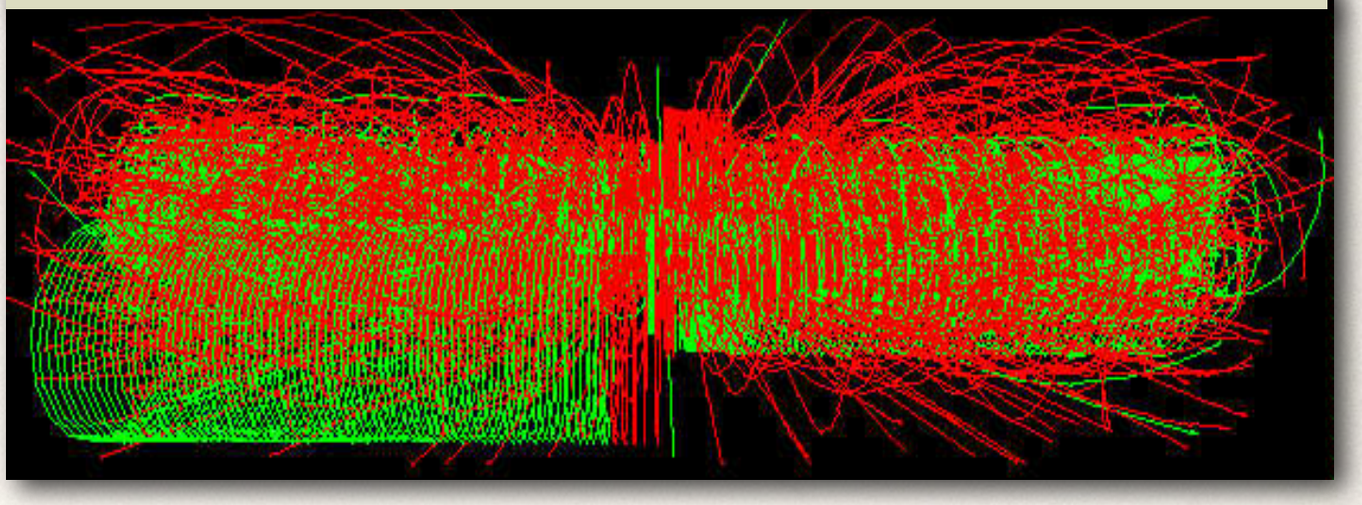

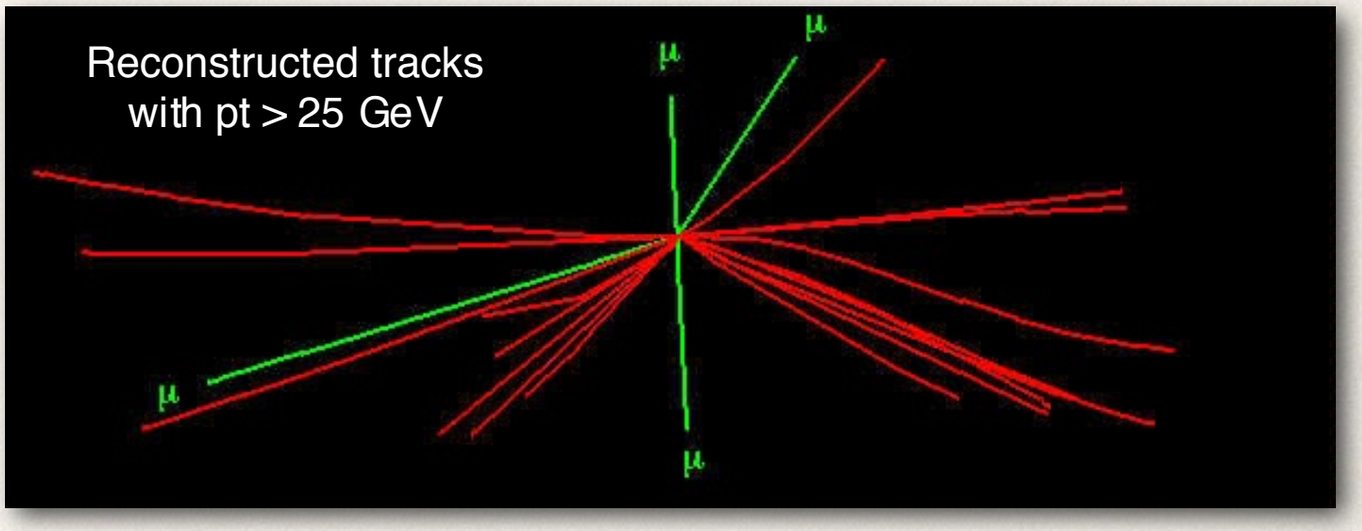

... and this repeats every 25 ns

ROOT
Data Analysis Framework

Physics Selection at LHC

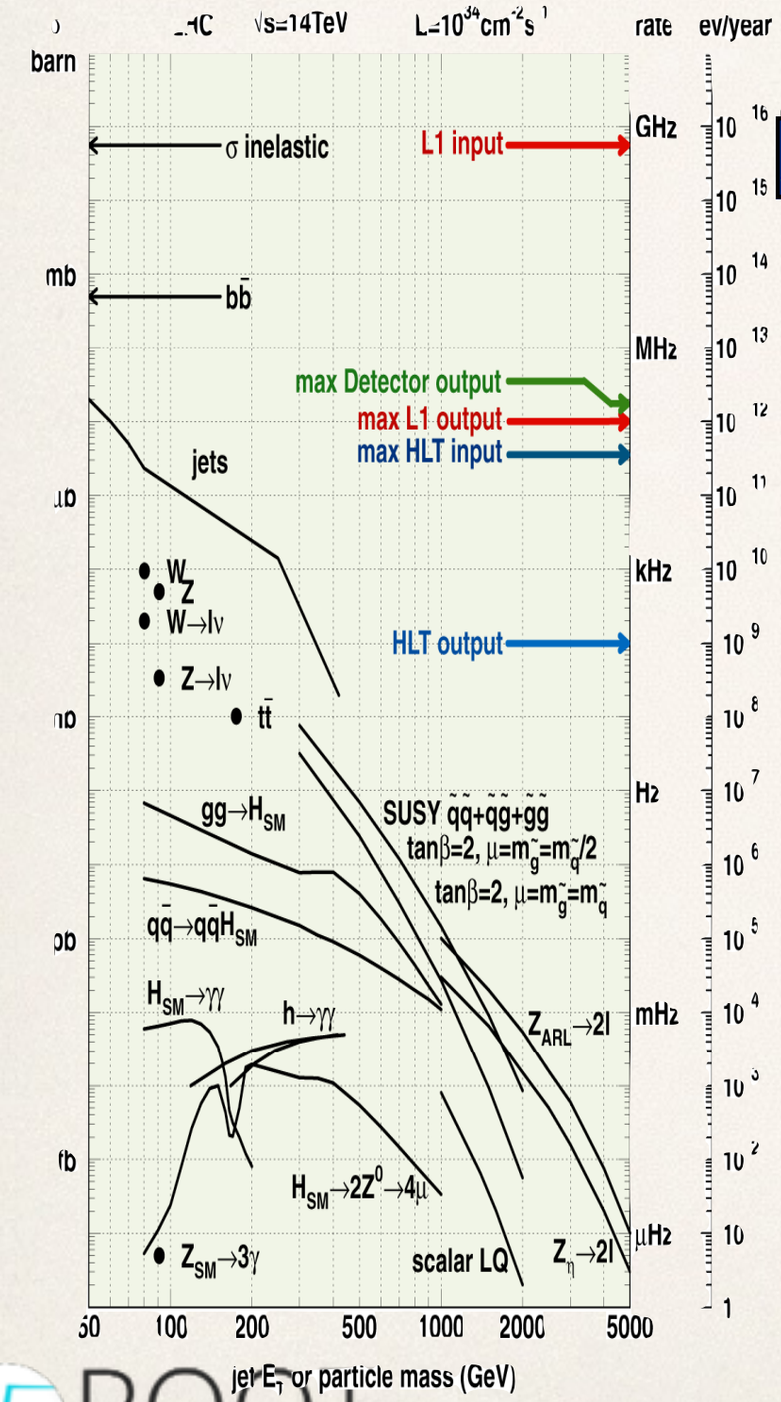

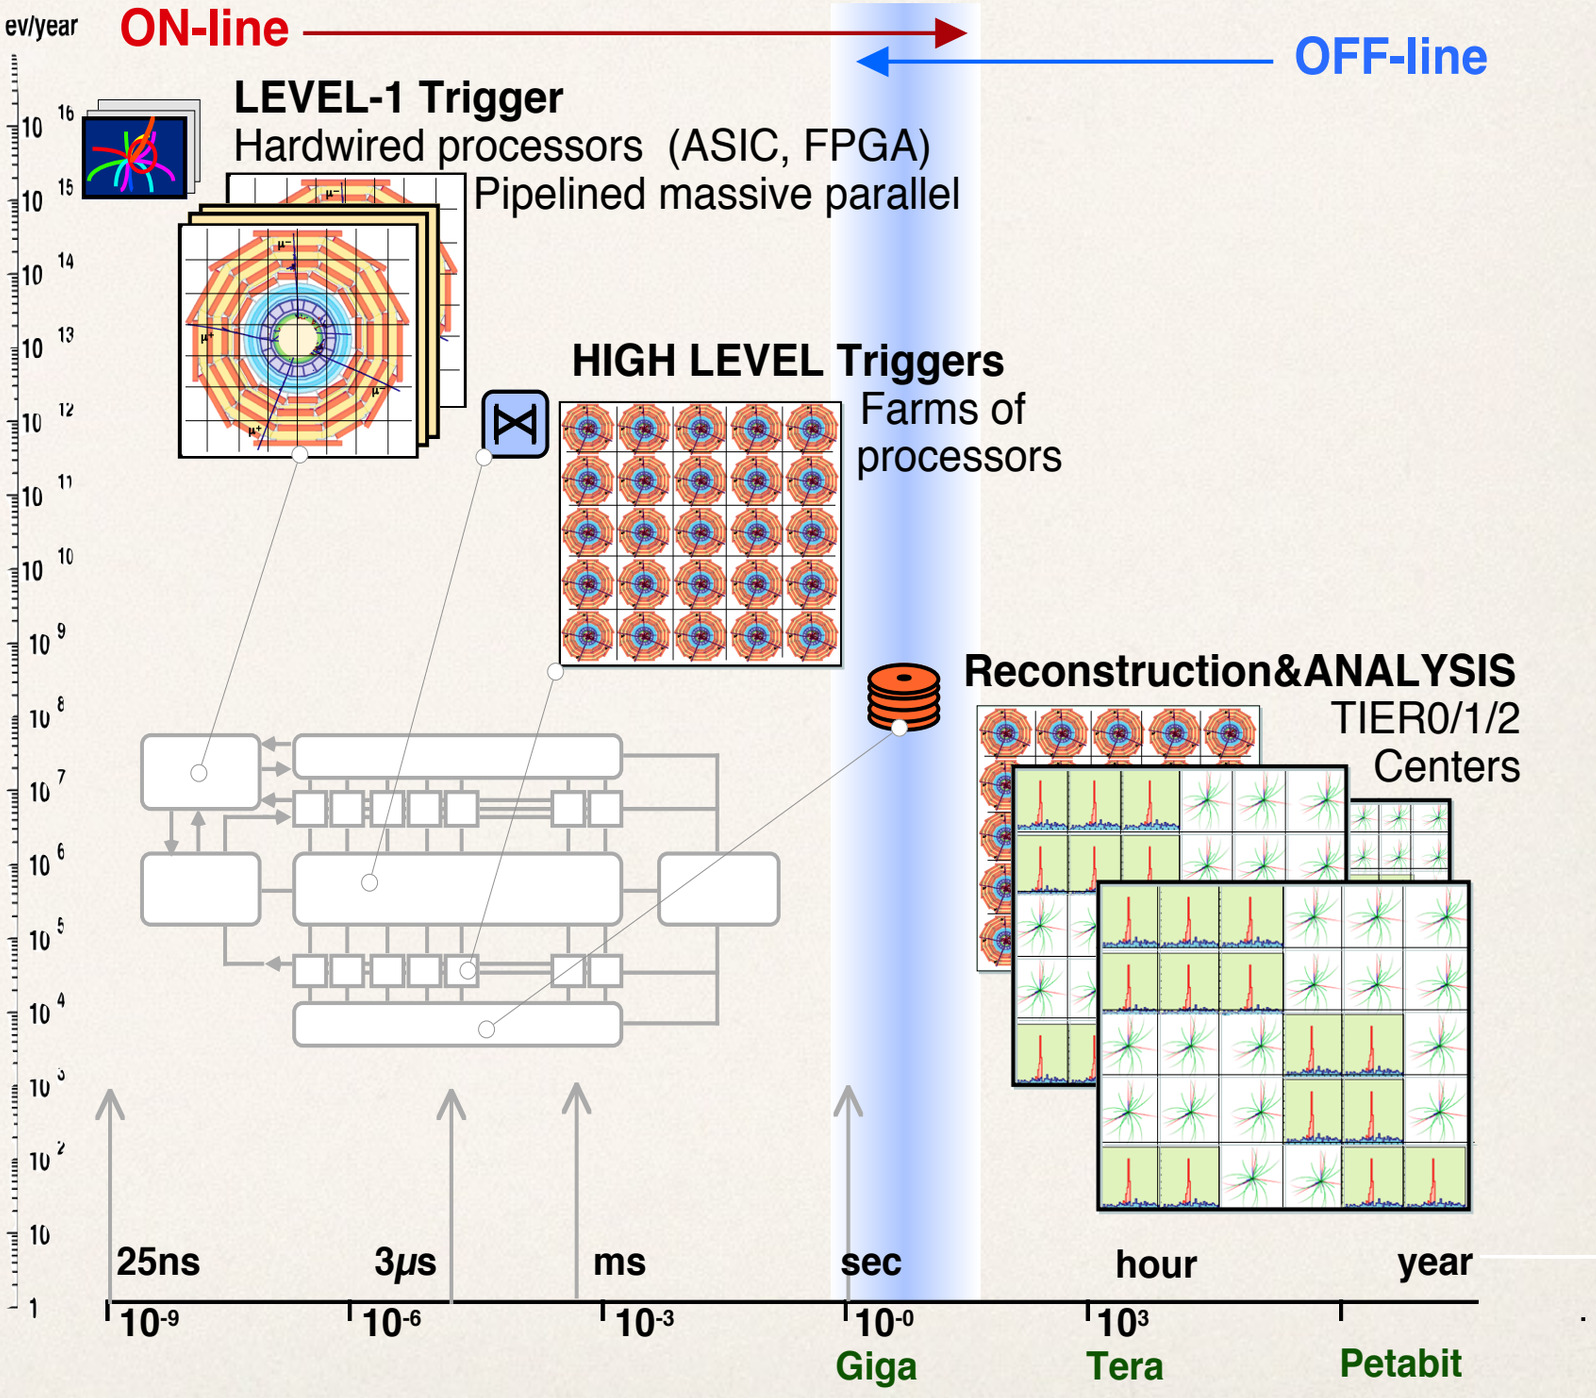

ROOT
Data Analysis Framework

Data Rates

- Particle beams cross every 25 ns (40 MHz)
  - Up to 25 particle collisions per beam crossing
  - Up to 10^9 collisions per second
- Basically 2 event filter / trigger levels
  - Hardware trigger (e.g. FPGA)
  - Software trigger (PC farm)
  - Data processing starts at readout
  - Reducing 10^9 p-p collisions per second to O(1000)
- Raw data to be stored permanently: >15 PB / year

Physics Process	Events/s
Inelastic p-p scattering	10^8
b	10^6
$W \rightarrow eu$; $W \rightarrow \mu u$; $W \rightarrow \tau u$	20
$Z \rightarrow ee$; $Z \rightarrow \mu\mu$; $Z \rightarrow \tau\tau$	2
t	1
Higgs boson (all; m_H = 120GeV)	0.04
Higgs boson (simple signatures)	0.0003

This is our Big Data problem!!

ROOT Data Analysis Framework

Big Data requires Big Computing

- The LHC experiments rely on distributed computing resources:
  - WLCG - a global solution, based on the Grid technologies/middleware.
    - distributing the data for processing, user access, local analysis facilities etc.
    - at time of inception envisaged as the seed for global adoption of the technologies
- Tiered structure
  - Tier-0 at CERN: the central facility for data processing and archival
  - 11 Tier-1s: big computing centers with high quality of service used for most complex/intensive processing operations and archival
  - ~140 Tier-2s: computing centers across the world used primarily for data analysis and simulation.

WLCG
Worldwide LHC Computing Grid

Capacity:
~350,000 CPU cores
~200 PB of disk space
~200 PB of tape space

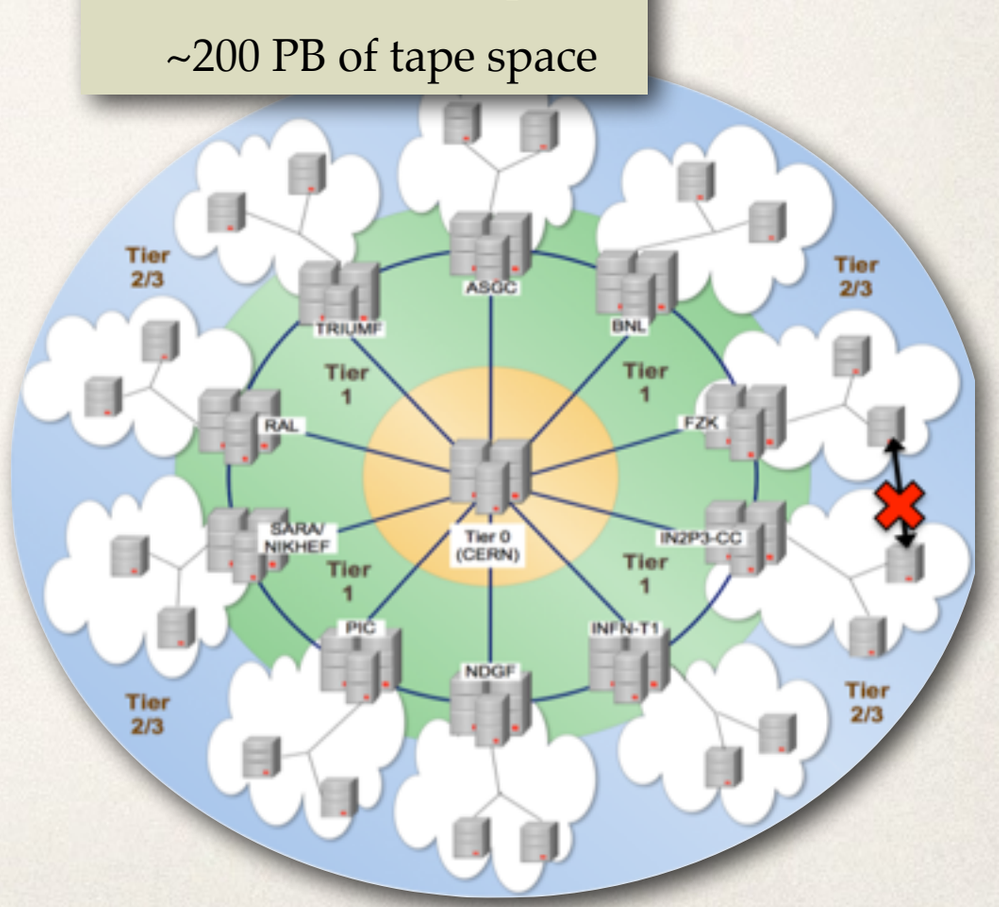

ROOT
Data Analysis Framework

The ROOT Data Analysis

- ROOT is a large Object-Oriented data handling and analysis framework
  - Efficient object data store scaling from KB's to PB's
  - C++ interpreter
  - Extensive 2D+3D scientific data visualization capabilities
  - Extensive set of data fitting, modeling and analysis methods
  - Complete set of GUI widgets
  - Classes for threading, shared memory, networking, etc.
  - Parallel version of analysis engine runs on clusters and multi-core
  - Fully cross platform, Unix/Linux, Mac OS X and Windows
  - 1.7 million lines of C++
  - Licensed under the LGPL
- Used by all HEP experiments in the world
- Used in many other scientific fields and in commercial world

ROOT
Data Analysis Framework

ROOT in Numbers

- Ever increasing number of users
  - 6800 forum members, 68750 posts, 1300 on mailing list
  - Used by basically all HEP experiments and beyond
- Binaries have been downloaded more than 620000 times since 1997

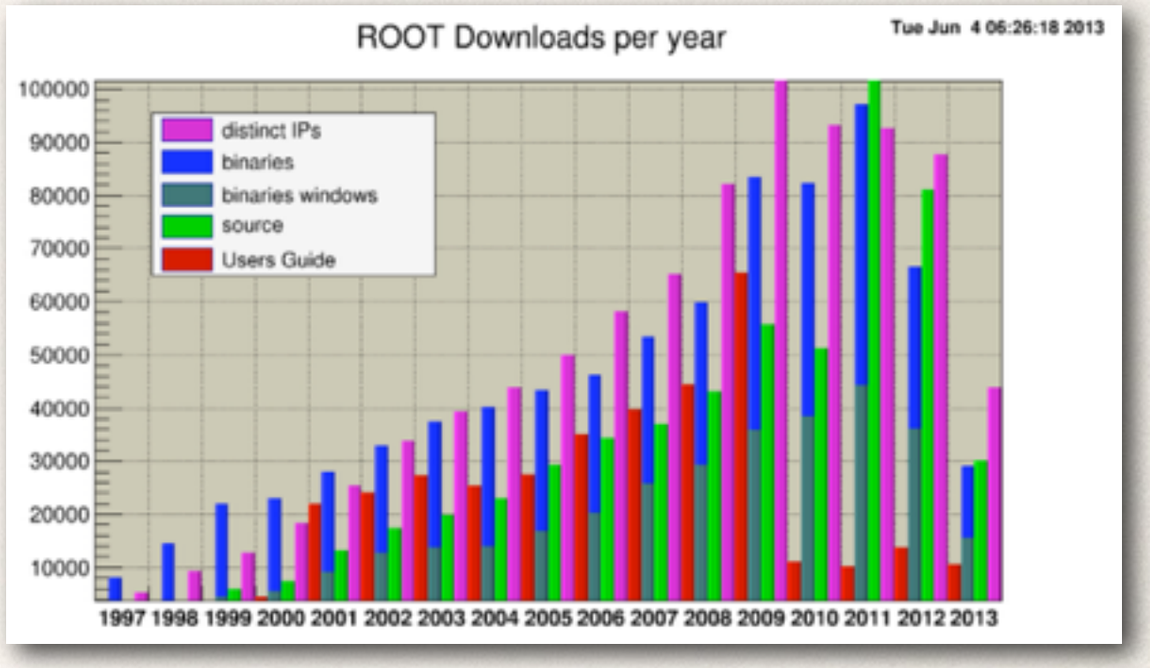

As of today 177 PB of LHC data stored in ROOT format

ALICE: 30PB, ATLAS: 55PB, CMS: 85PB, LHCb: 7PB

ROOT
Data Analysis Framework

ROOT Object Persistency

- Scalable, efficient, machine independent format
- Based on object serialization to a buffer
- Automatic schema evolution (backward and forward compatibility)
- Object versioning
- Compression
- Easily tunable granularity and clustering
- Remote access
  - HTTP, HDFS, Amazon S3, CloudFront and Google Storage
- Self describing file format (stores reflection information)
- ROOT I/O is used to store all LHC data (actually all HEP data)

ROOT
Data Analysis Framework

Object Containers - TTree

- Special container for very large number of objects of the same type (events)
  - Minimum amount of overhead per entry
- Objects can be clustered per sub object or even per single attribute (clusters are called branches)
- Each branch can be read individually
  - A branch is a column

Physicists perform final data analysis processing large TTrees

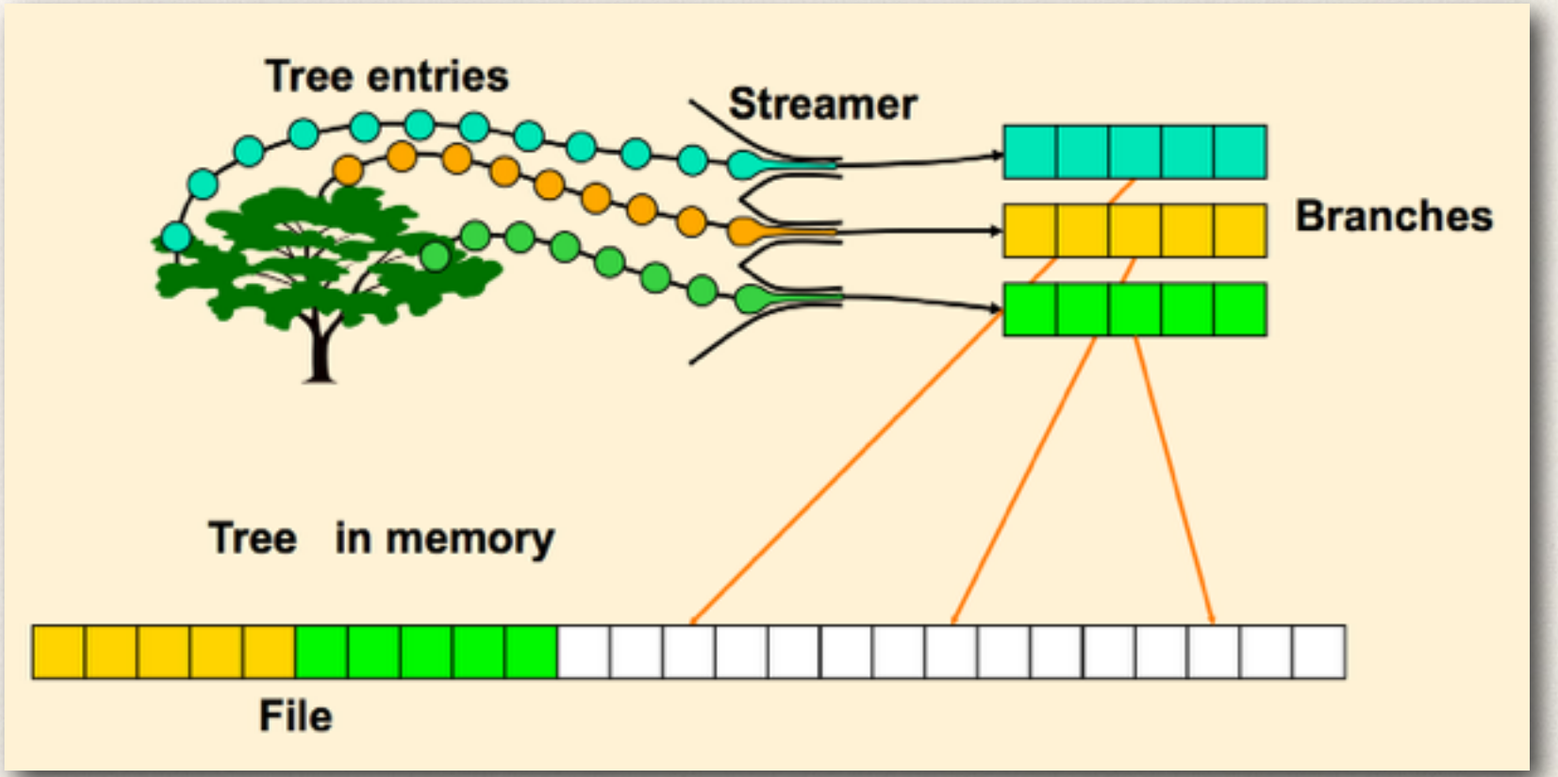

ROOT
Data Analysis Framework

ROOT Interpreter

- ROOT is shipped with an C/C++ interpreter, CINT
  - C++ not trivial to interpret and not foreseen in the language standard!
- Provides interactive shell
- Can interpret "macros" (not compiled programs)
  - Rapid prototyping possible

```
CINT/ROOT C/C++ Interpreter version 5.18.00, July 2, 2010
Type ? for help. Commands must be C++ statements.
Enclose multiple statements between { }.

root [0] TH1D histo("normal","Normal histogram", 100, -10., +10);
root [1] for(int i = 0; i < 10000; i++) {
end with '}', '@':abort >  histo.Fill(gRandom->Gaus());
end with '}', '@':abort > }
root [2] histo.Draw();
```

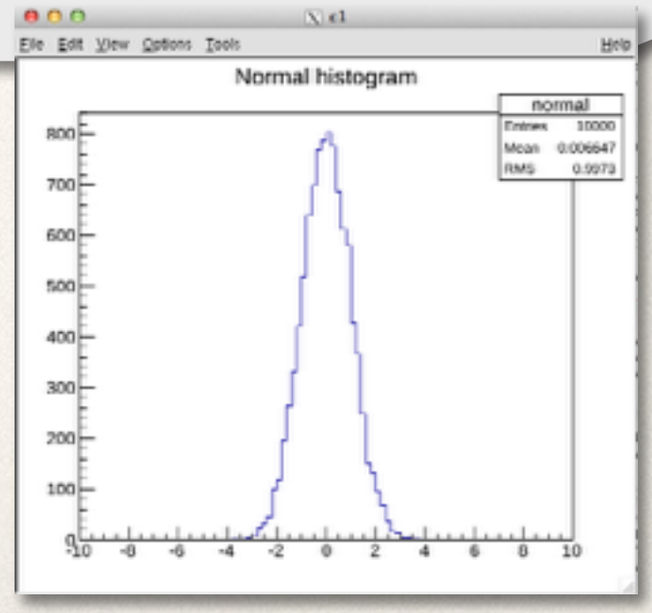

- ROOT provides also Python bindings (PyROOT), which are very popular among physicists
- Starting from ROOT 6, there is the new interpreter Cling (based on LLVM/Clang)

ROOT Image Gallery

ROOT Image Gallery

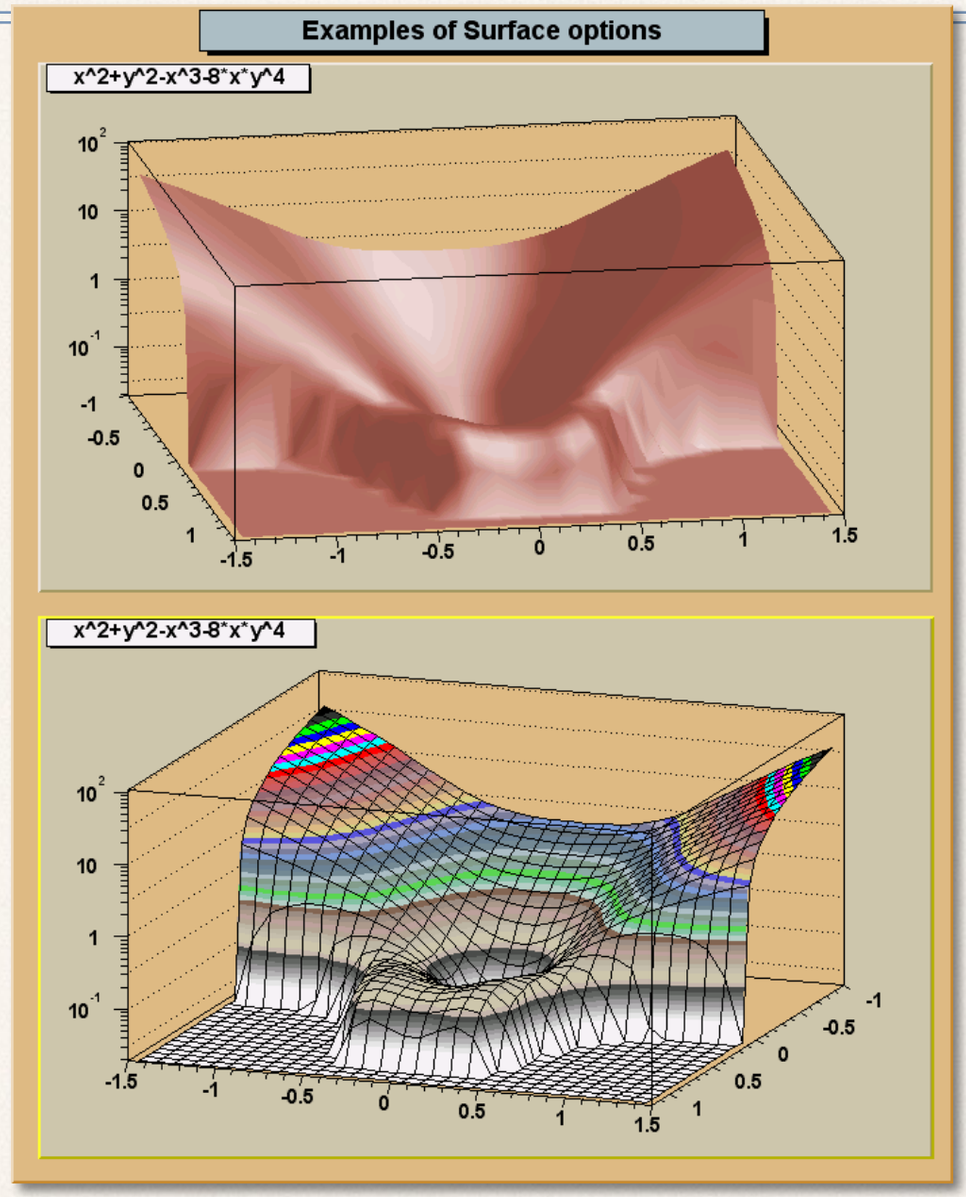

ROOT
Data Analysis Framework

ROOT Image Gallery

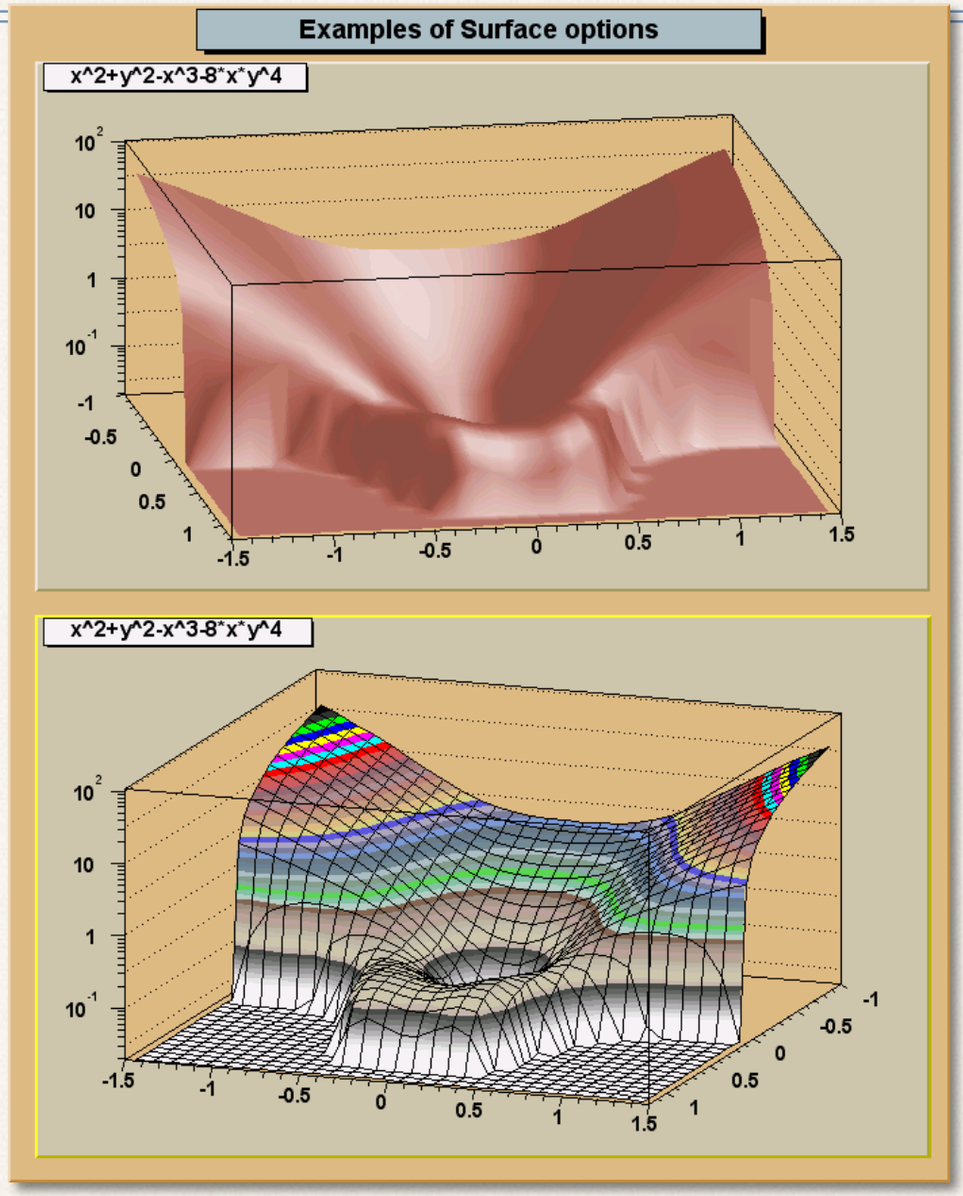

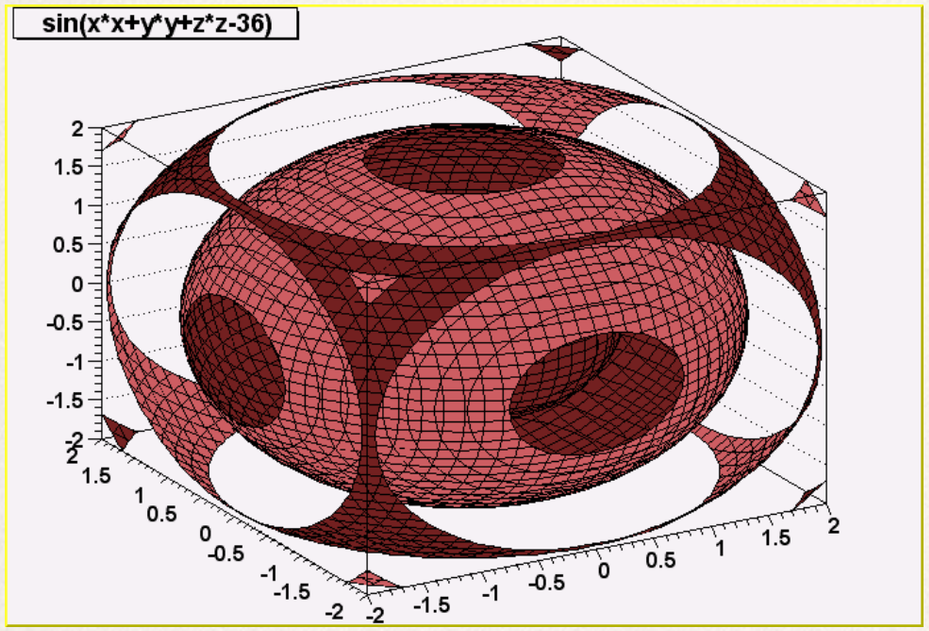

ROOT
Data Analysis Framework

ROOT Image Gallery

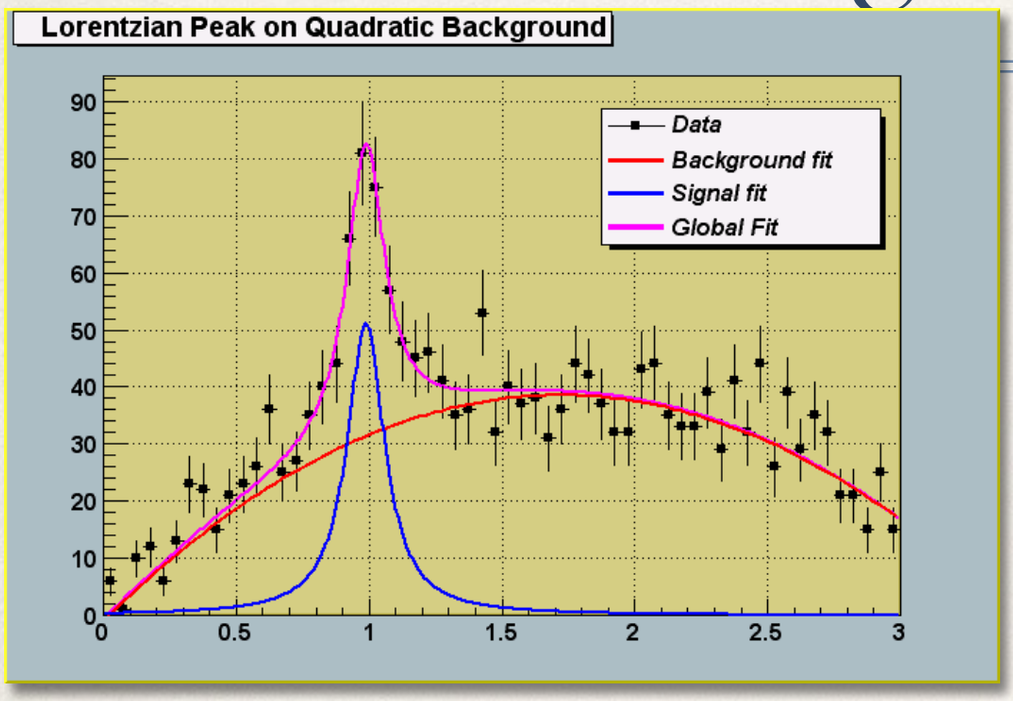

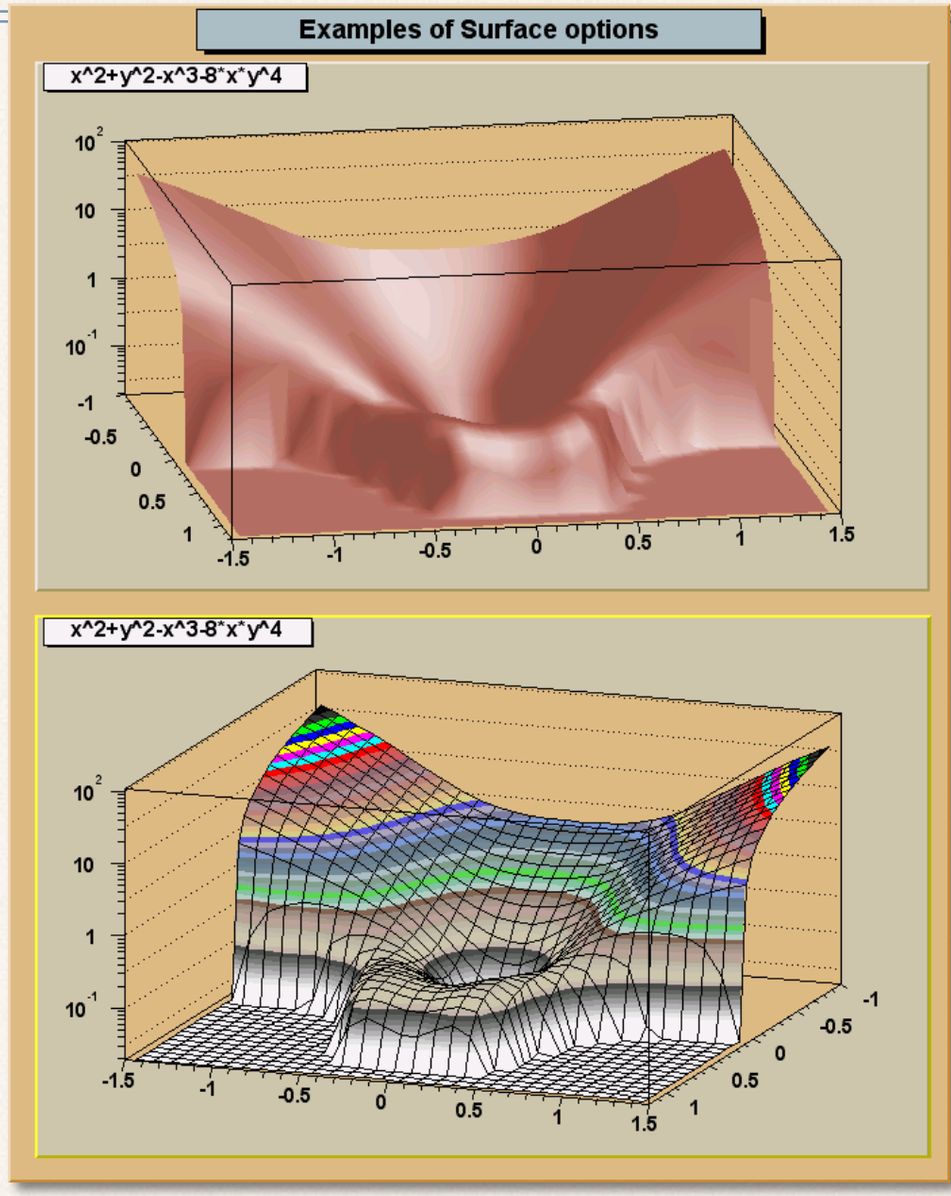

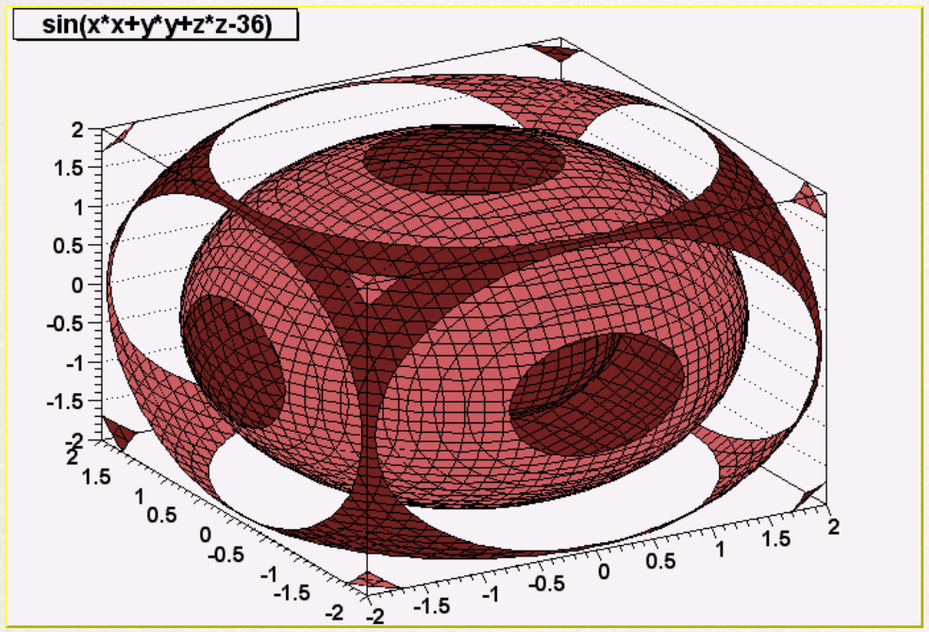

ROOT
Data Analysis Framework

ROOT Image Gallery

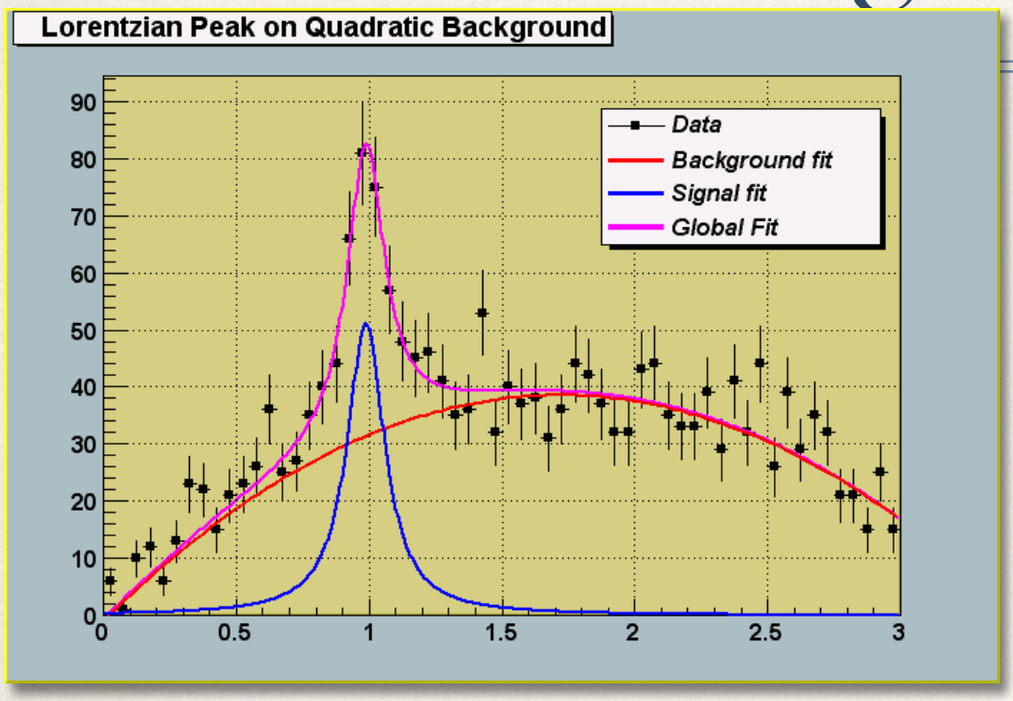

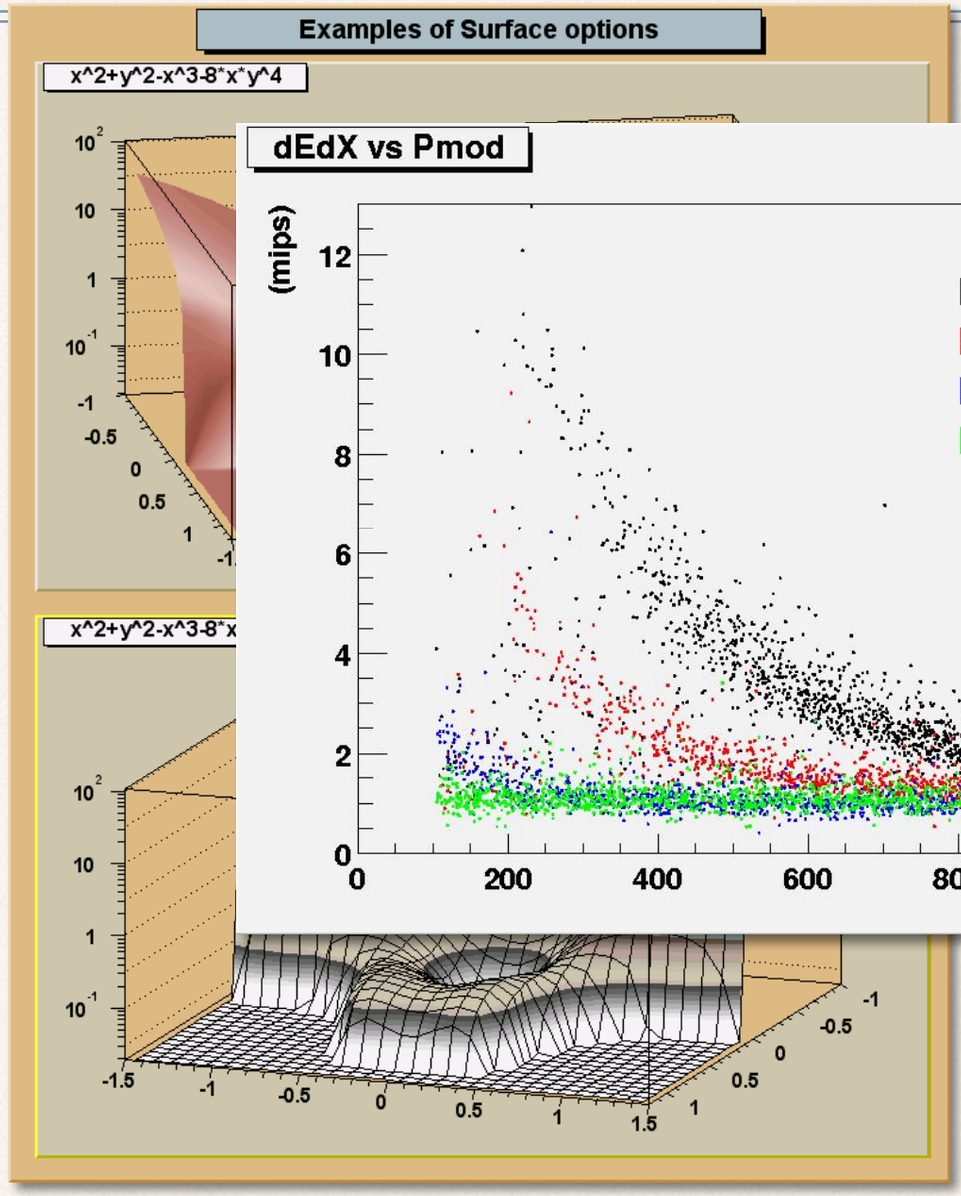

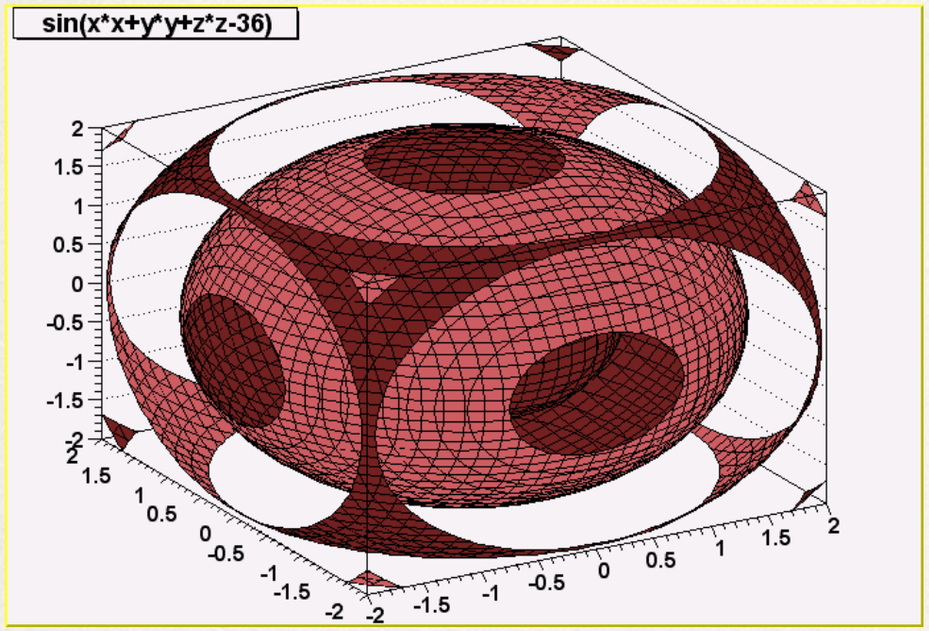

ROOT
Data Analysis Framework

ROOT Image Gallery

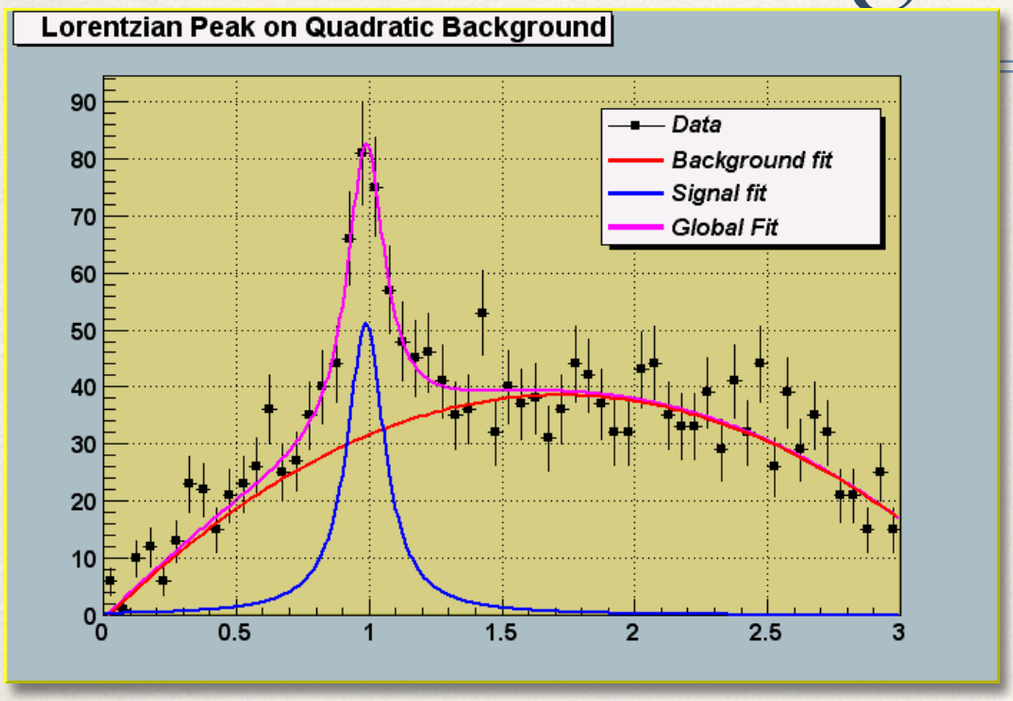

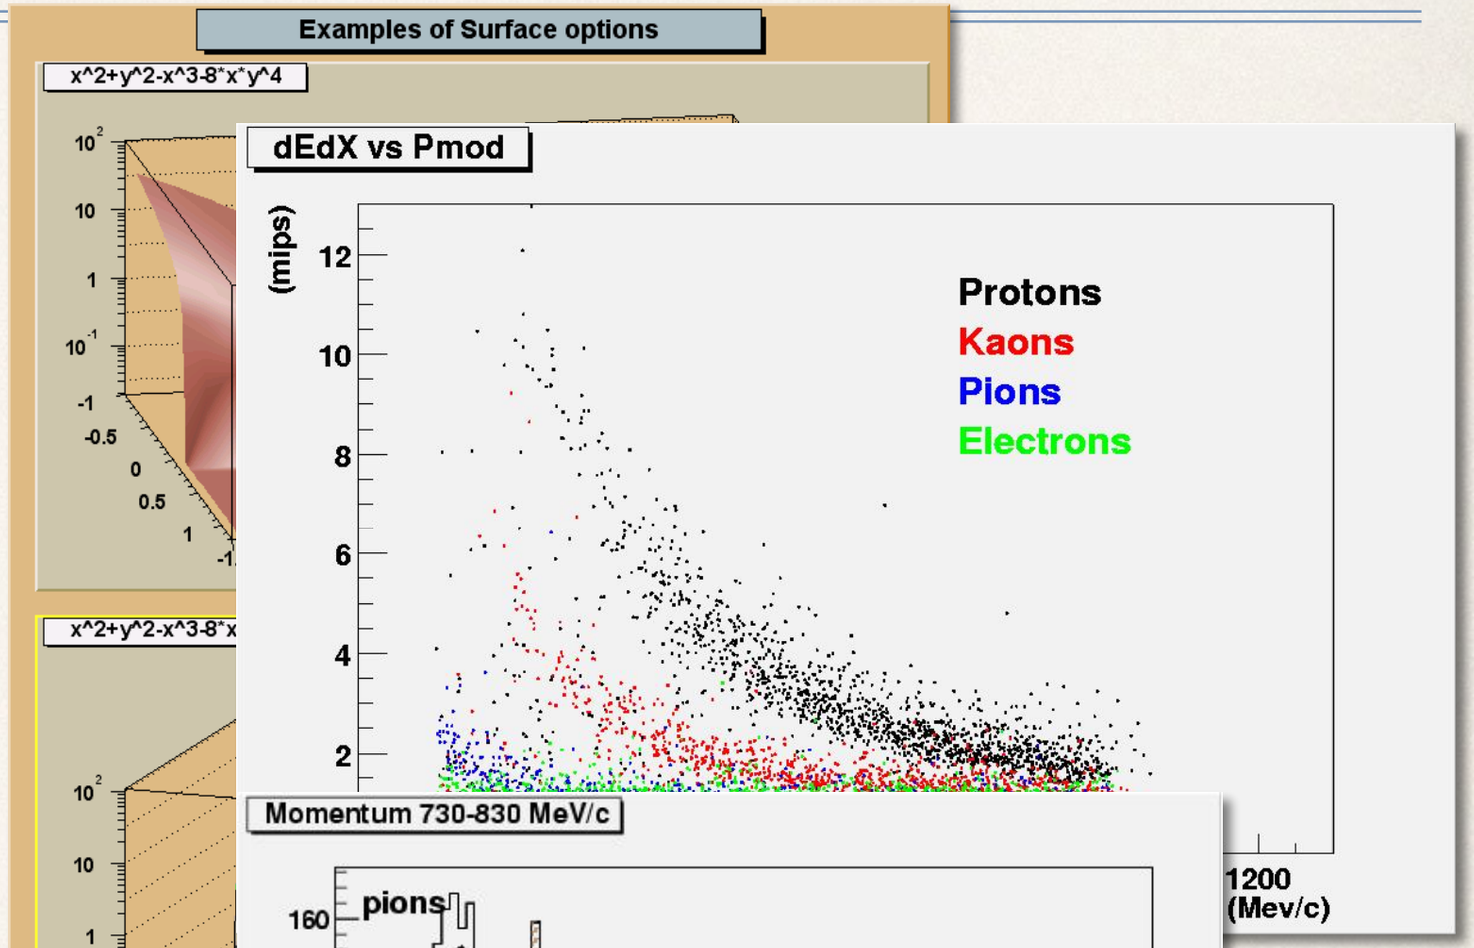

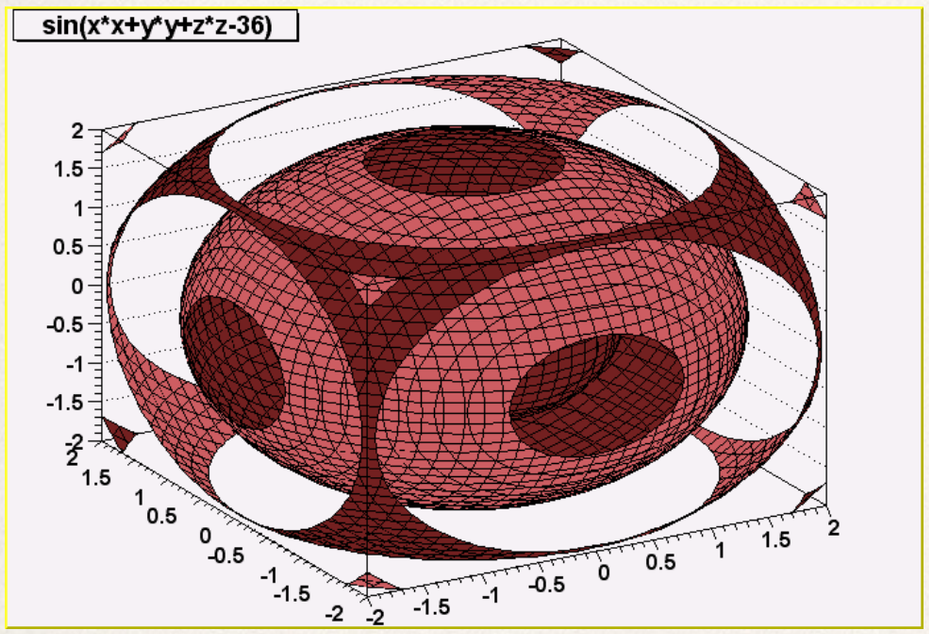

ROOT
Data Analysis Framework

ROOT Image Gallery

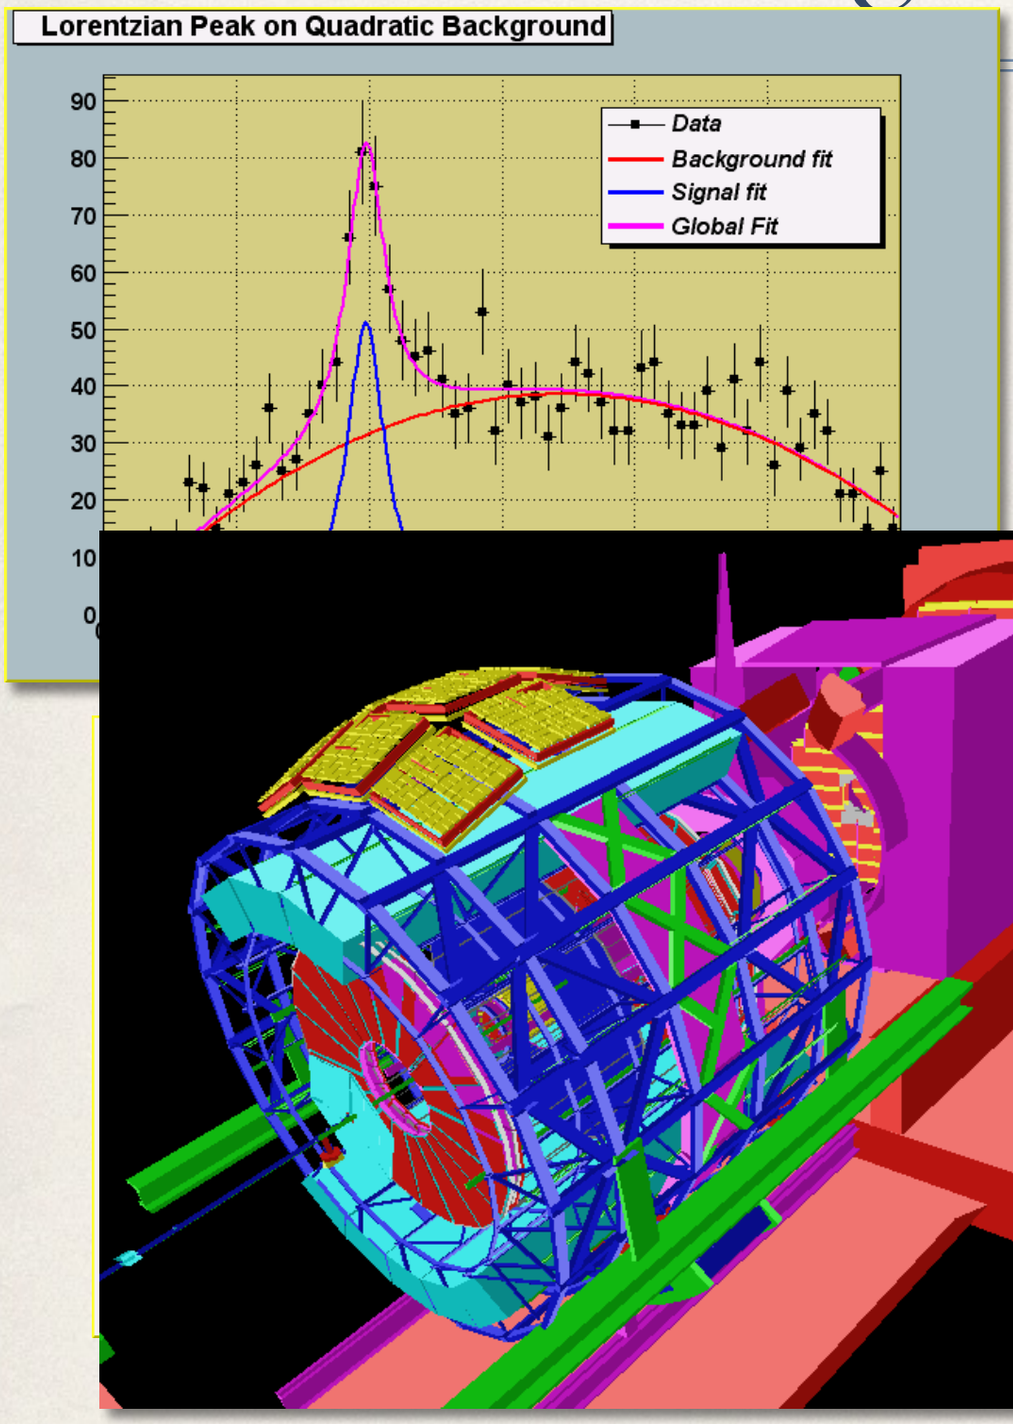

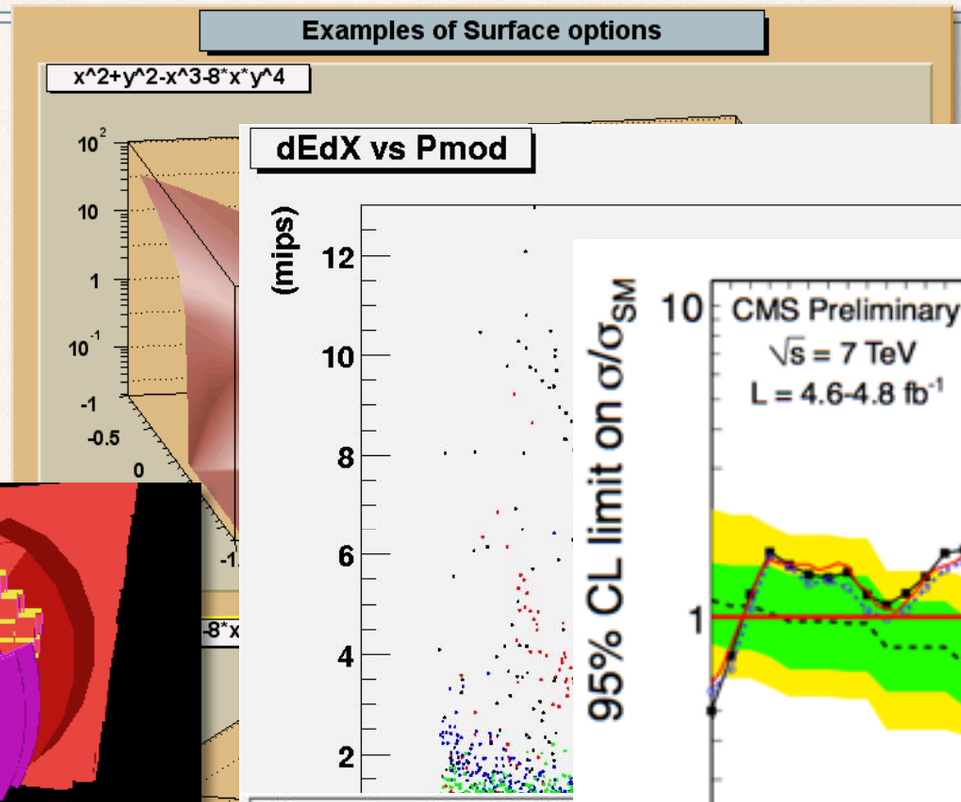

ROOT
Data Analysis Framework

ROOT in Javascript

- Provide ROOT file access entirely locally in a browser without any prior ROOT installation on the server or client
  - ROOT files are self describing ...

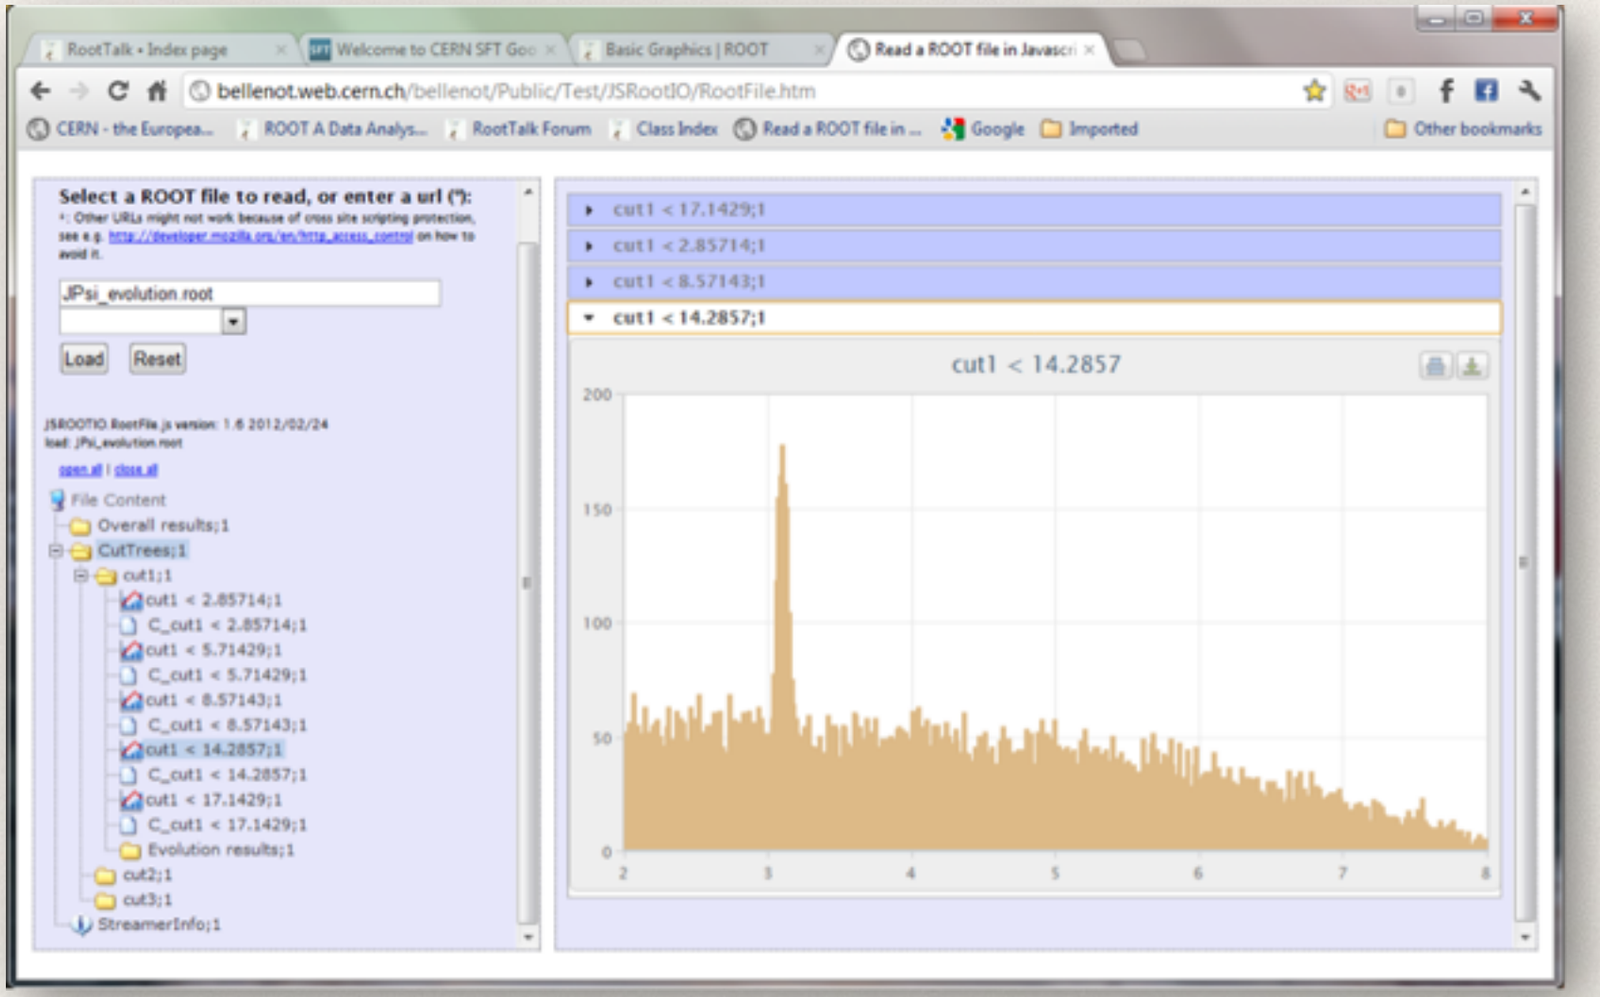

ROOT
Data Analysis Framework

EVE Event Display

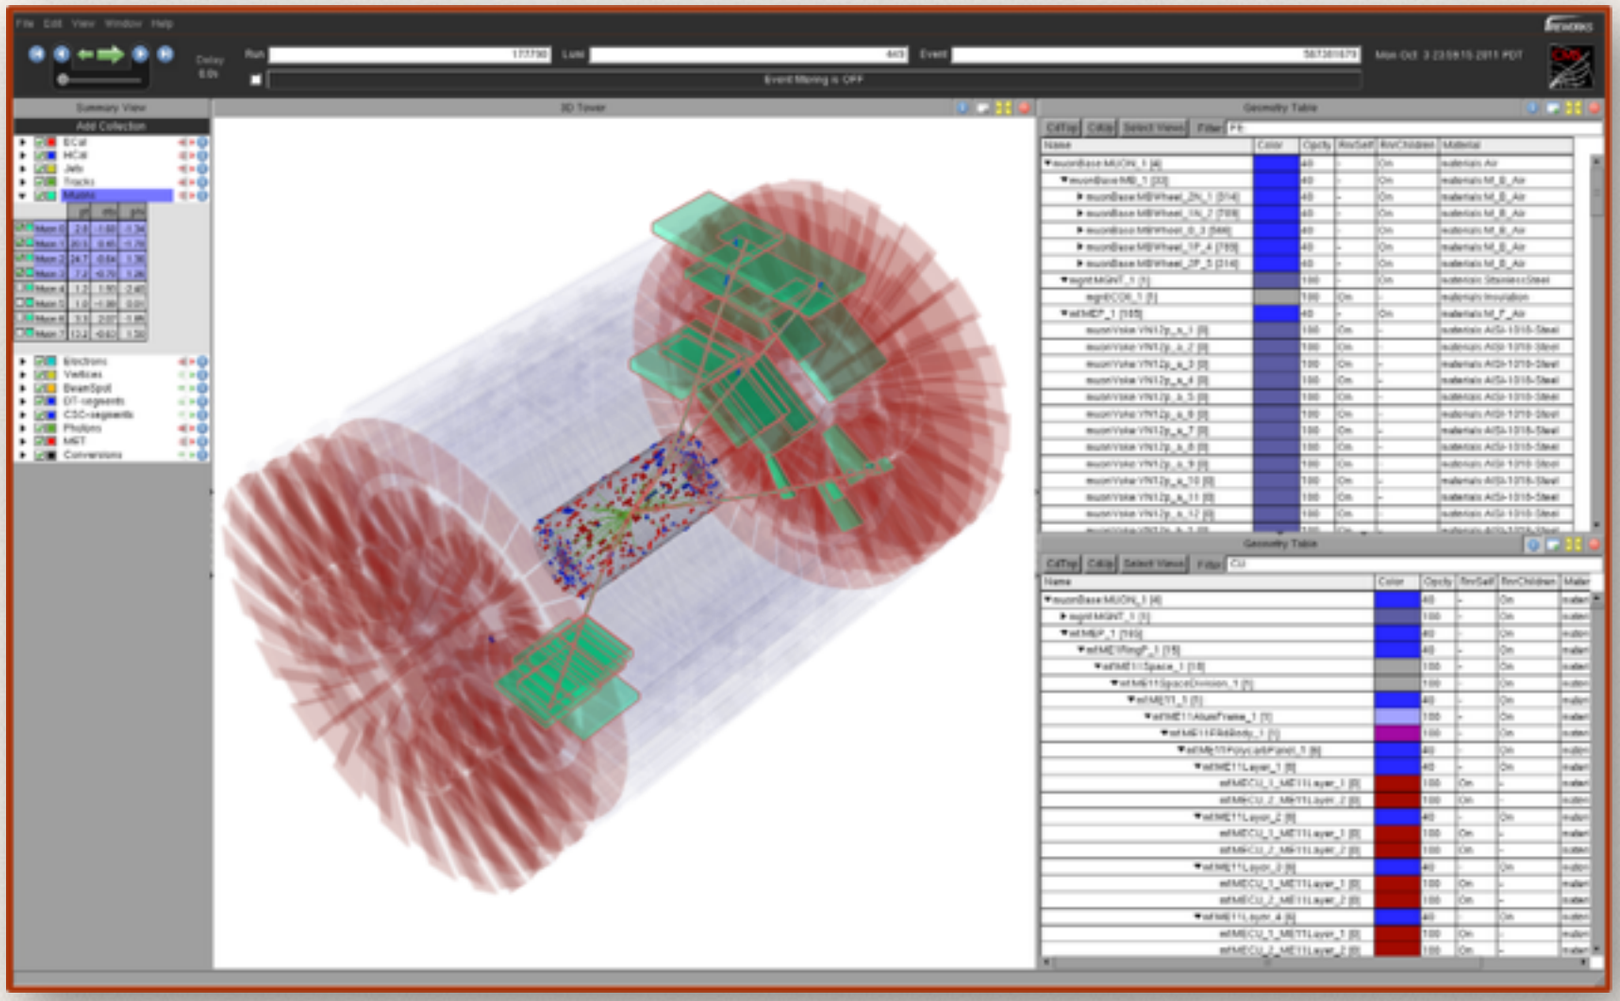

ROOT
Data Analysis Framework

EVE Event Display

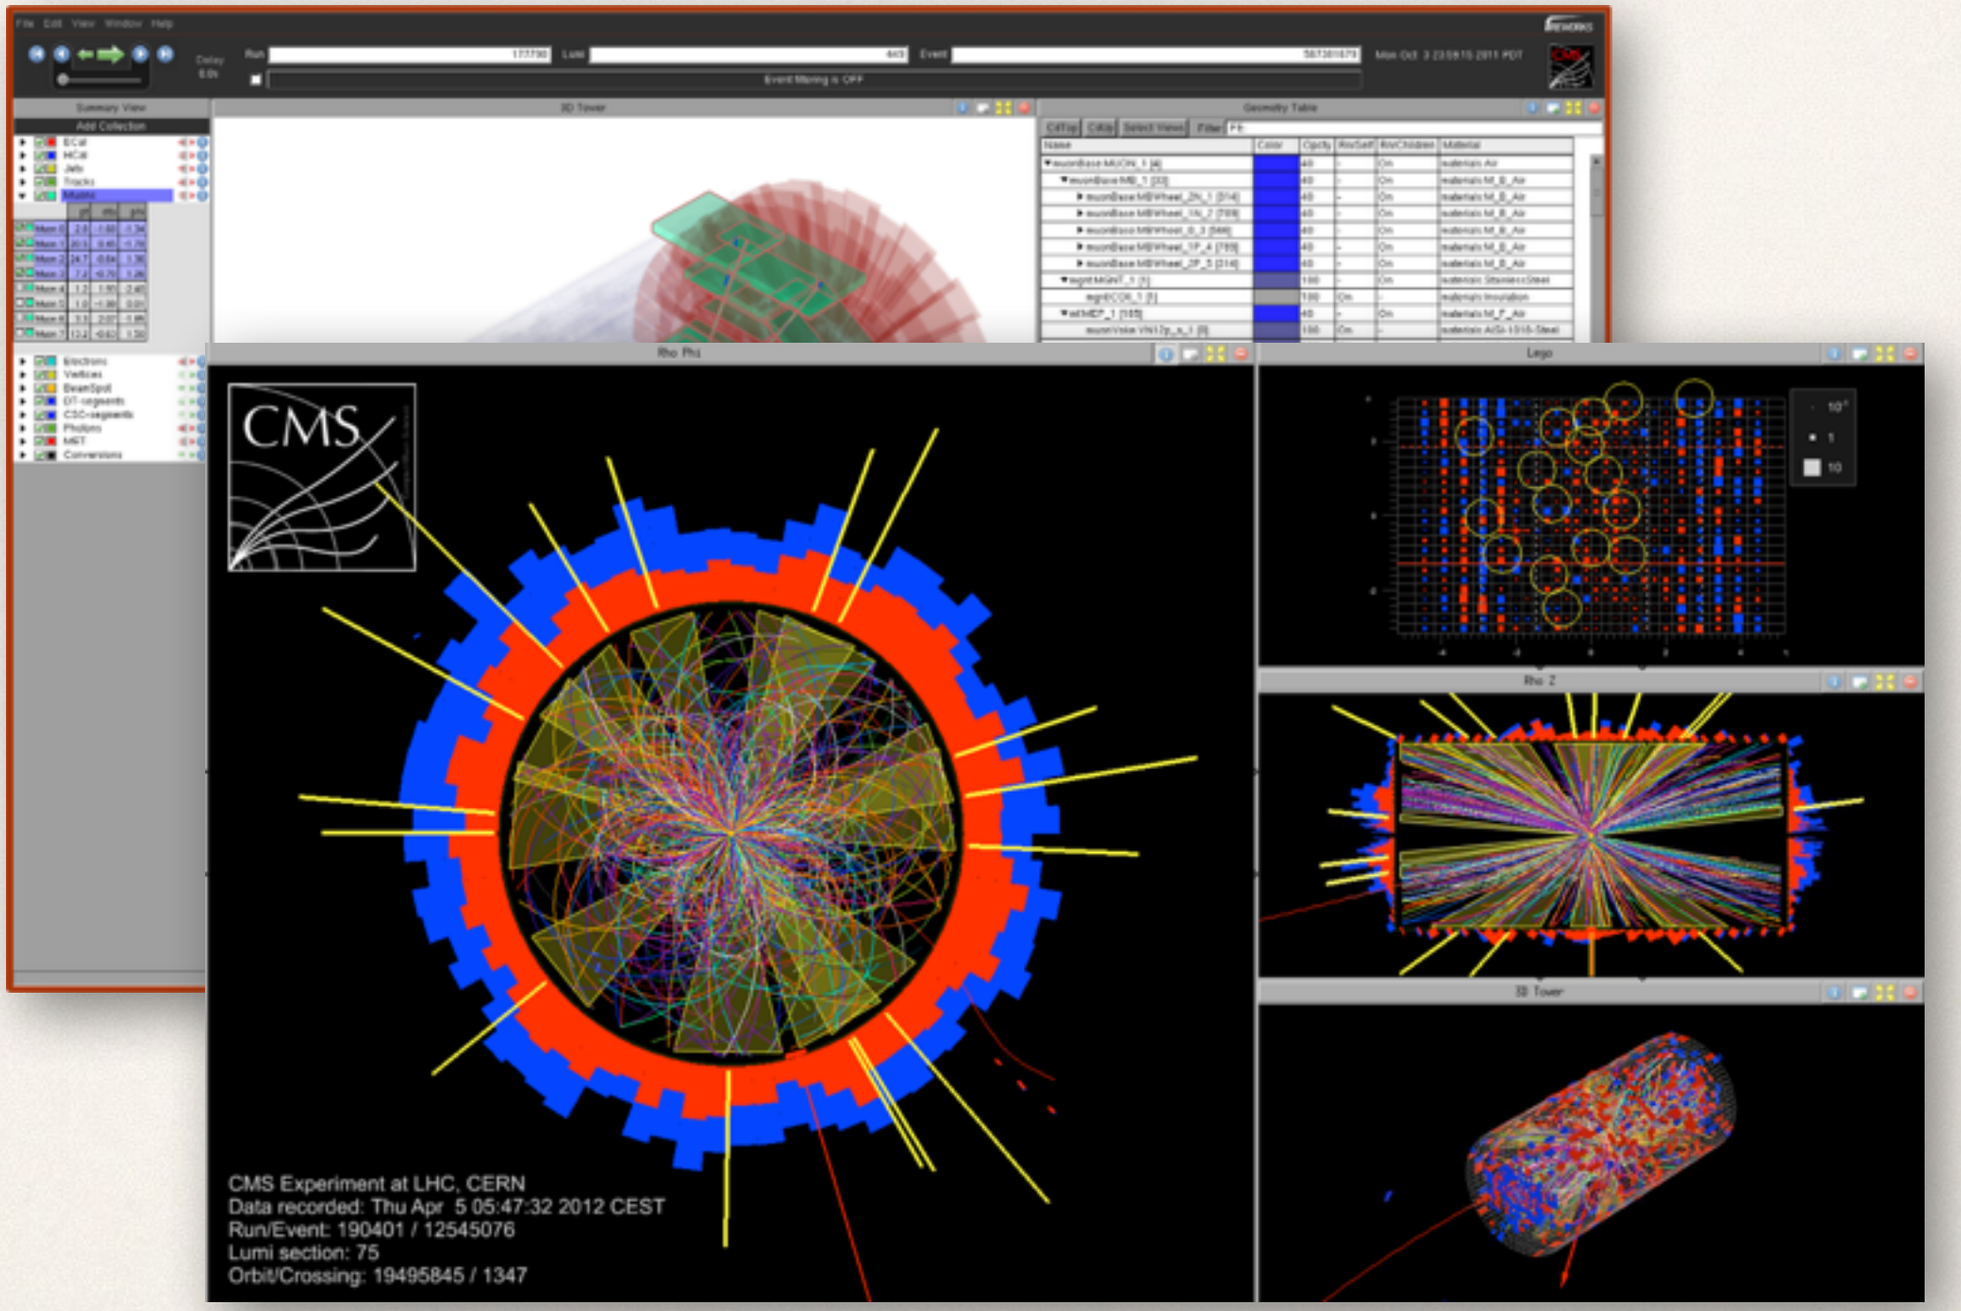

ROOT
Data Analysis Framework

PROOF-The Parallel Query

- A system for running ROOT queries in parallel on a large number of distributed computers or many-core machines
- PROOF is designed to be a transparent, scalable and adaptable extension of the local interactive ROOT analysis session
- For optimal CPU load it needs fast data access (SSD, disk, network) as queries are often I/O bound
- The packetizer is the heart of the system
  - Runs on the client/master and hands out work to the workers
  - Takes data locality and storage type into account
  - Avoids storage device overload
  - Ensures that workers end at the same time

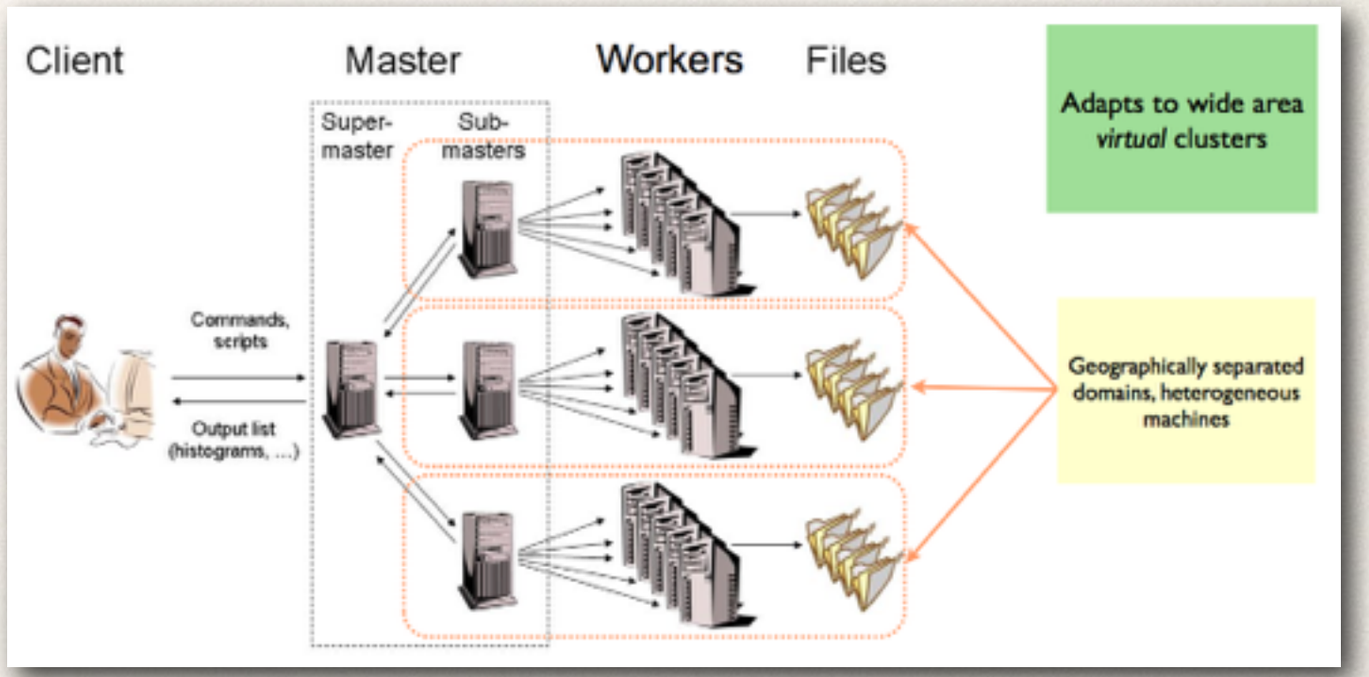

ROOT
Data Analysis Framework

Various Flavors of PROOF

- PROOF-Lite (optimized for single many-core machines)
  - Zero configuration setup (no config files and no daemons)
  - Workers are processes and not threads for added robustness
  - Once your analysis runs on PROOF Lite it will also run on PROOF
- Dedicated PROOF Analysis Facilities (multi-user)
  - Cluster of dedicated physical nodes
  - Some local storage, sandboxing, basic scheduling, basic monitoring
- PROOF on Demand (single-user)
  - Create a temporary dedicated PROOF cluster on batch resources (Grid or Cloud)
  - Uses an resource management system to start daemons
  - Each user gets a private cluster

Usage in Industry

- Overview highly incomplete
  - Very difficult to have an exact picture
- Based on discussions with users
- Based on user registrations
- Based on bug reports

ROOT
Data Analysis Framework

Industries

- Flight planning systems (MITRE)
- Insurance (Nationwide)
- Stock market applications (Merrill Lynch, Renaissance Corp)
- Banking, mortgaging (Countrywide home loan, Landesbank Baden Wurtenberg, Credit Suisse)
- Pharmaceutical research (Merck Frosst)
- Medical imaging, MRI (Philips Medical)
- Telecom (KPN research, Vodafone, Alcatel, RIPE)
- Aerospace research (ELT Rocket Research, Mitsubishi space software, Boeing, DASA)
- Defense (USAF, DoD)

ROOT
Data Analysis Framework

Medical Fraud Detection

- First industrial application, early 1997
- Outsourced to researchers of Los Alamos National Laboratory
- Used to mine and correlate records in:
  - Medical bills database (50 million)
  - Patient data base (3 million)
  - MD data base (30000)
- To discover possible fraudulent billing

Allowed us to improve ROOT for small events (records)

ROOT Data Analysis Framework

Insurance

- Ratemaking
- Modeling
- Simulation

"There are many other reasons why ROOT is an appropriate tool for predictive modeling. But efficiency in storing and accessing the data is where ROOT stands out from any other tool that is in the market today."

Arun Tripathi, at the Casual Actuary Society ratemaking seminar.

Finance

- Used by several hedge fund and Wall Street trading companies (please don't blame ROOT for the credit crunch)
- Renaissance Technologies important user
  - 250 employees, many math, physics and CS PhD's
  - Technical trading: data into computer ➜ trade recommendation
  - They contributed and maintain the TMatrix linear algebra classes
  - They sponsor one developer at CERN

Contributions from industry incorporated into ROOT

ROOT
Data Analysis Framework

Telecom

- KPN Research
  - Mobile network performance monitoring
  - Multi Layer Packet Analysis using ROOT for analysis and plotting

- RIPE
  - Analysis of network monitoring data

ROOT
Data Analysis Framework

Telecom

- KPN Research
  - Mobile network performance monitoring
  - Multi Layer Packet Analysis using ROOT for analysis and plotting
- RIPE
  - Analysis of network monitoring data

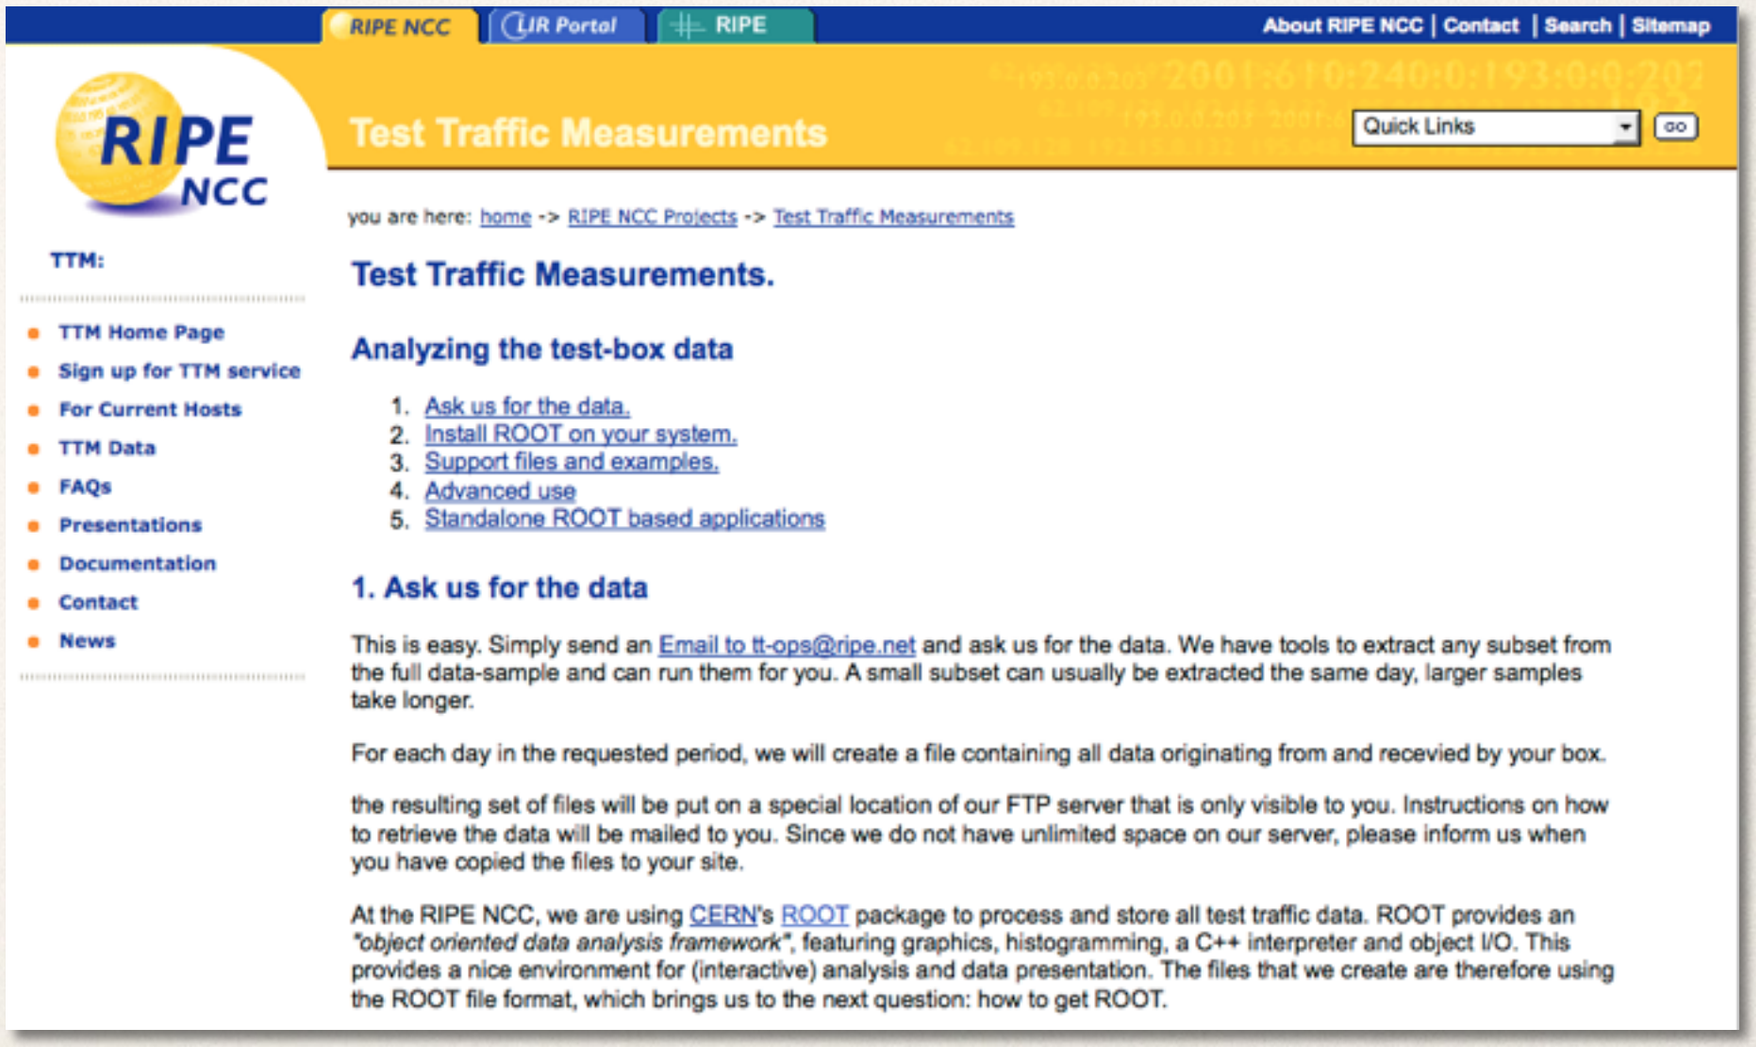

ROOT
Data Analysis Framework

Genetics

xps

Methods for Processing and Analysis of Affymetrix Oligonucleotide Arrays including Exon Arrays and Gene Arrays

The package handles pre-processing, normalization, filtering and analysis of Affymetrix GeneChip expression arrays, including exon array systems (Exon 1.0 ST) and gene array systems (Gene 1.0 ST) on computers with 1 GB RAM only. It imports Affymetrix .CDF, .CLF, .PGF and .CEL as well as annotation files, and computes e.g. RMA, MAS5, DABG-calls. It is an R wrapper to XPS (eXpression Profiling System), which is based on ROOT, an object-oriented framework developed at CERN. Thus, the prior installation of ROOT is a prerequisite for the usage of this package, however, no knowledge of ROOT is required. ROOT is licensed under LGPL and can be downloaded from http://root.cern.ch.

Author	Christian Stratowa, Vienna, Austria
Maintainer	Christian Stratowa

To install this package, start R and enter:

```
source("http://bioconductor.org/biocLite.R")
biocLite("xps")
```

Vignettes (Documentation)

- xps.pdf
- xpsClasses.pdf

Package Downloads

Source	xps_1.0.1.tar.gz
Windows binary	xps_1.0.1.zip
OS X binary	xps_1.0.1.tgz

ROOT
Data Analysis Framework

Astronomical Data Analysis

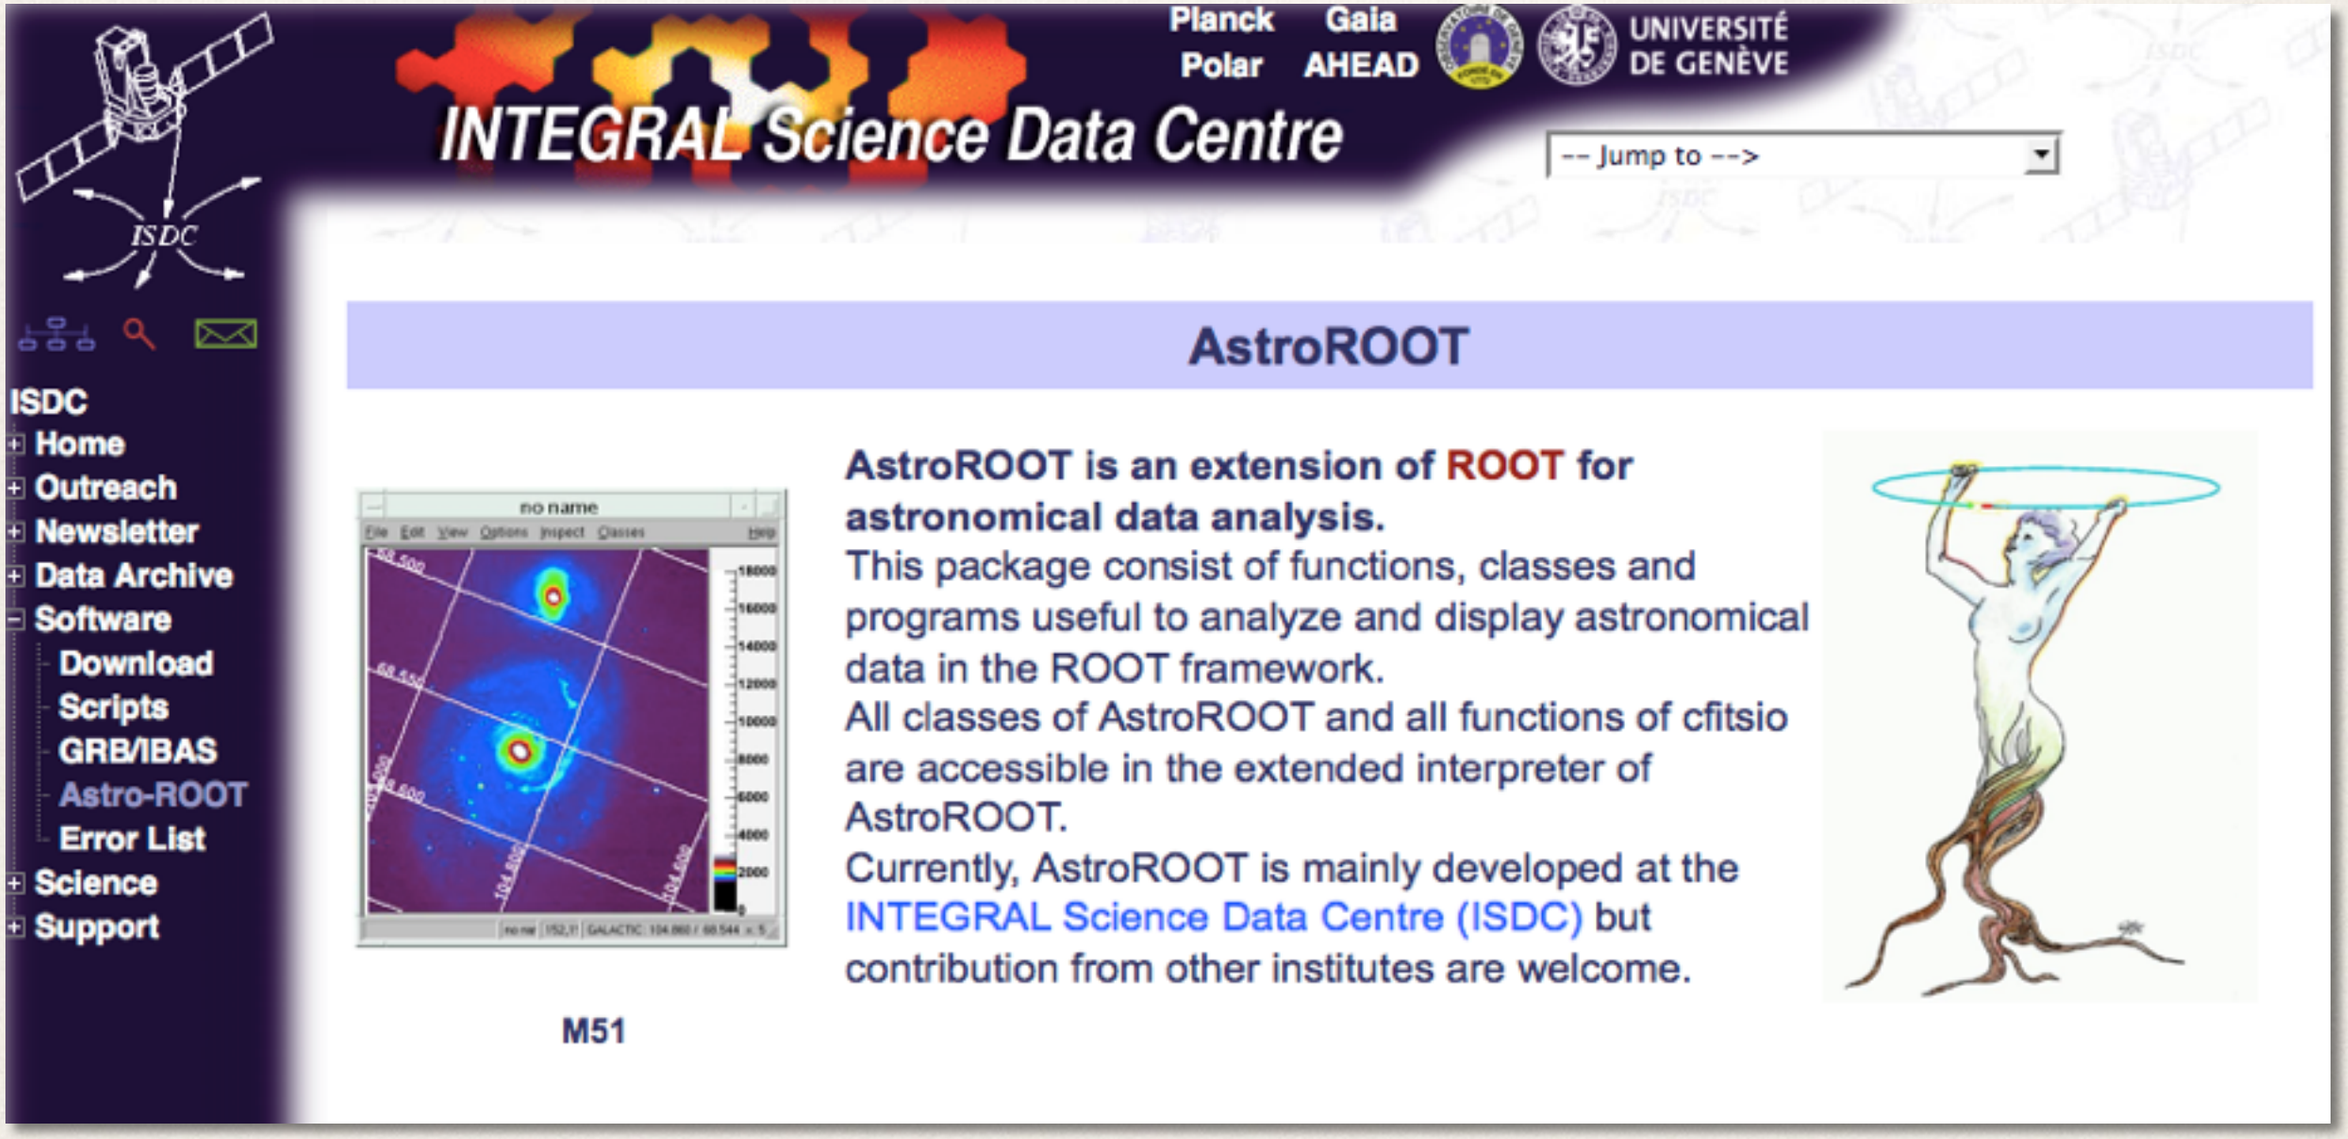

ROOT
Data Analysis Framework

Conclusions

- ROOT is a very successful CERN software spin-off
- It is used everywhere in HEP and widely in science
- It has found good inroads in industry, without explicit advertisement, mainly word of mouth and migrating scientists
- Designed to handle the large quantities of LHC data, it proves to be an attractive application for industry where data quantities are also increasing rapidly
- Being Open Source has been very beneficial for its wide acceptance and has stimulated collaboration

For more see:
http://root.cern.ch

ROOT
Data Analysis Framework