6th Summer School on INtelligent signal processing for FrotIEr Research and Industry

Introduction to Big Data, massive parallel and High Performance Computing

[the fundamental sciences case and outcomes]

J. Flix [PIC/CIEMAT]

INFIERI 2020 / UAM (Madrid)
23/8-4/9 2021

PIC port d'informació científica

Institut de Física d'Altes Energies

Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas

Who am I?

PIC port d'informació científica

@JosepFlixMolina

Physics PhD (High Energy Physicist) and Computer Scientist

ATLAS Collaboration / MAGIC Collaboration / CMS Collaboration

Spanish WLCG Tier-1 P.I. and Project Manager @pic_es

LHC Distributed Computing / CMS Computing management / WLCG Grid Deployment Board chair

Associate Professor at Universitat Autònoma de Barcelona (UAB)

Classical mechanics, Special relativity, Numerical methods (Python), High Energy Physics

Consultancy/Private sector experience

C++ programmer @InfoJobs / Linux embedded systems @AnaliticaSl / NeiC consultant @NeICnordic

Contents

PIC port d'informació científica

INTRODUCTION TO BIG DATA, MASSIVE PARALLEL AND HIGH PERFORMANCE COMPUTING: THE FUNDAMENTAL SCIENCES CASE and OUTCOMES

Contents

PIC
port d'informació científica

INTRODUCTION TO **BIG DATA**, MASSIVE PARALLEL AND HIGH PERFORMANCE COMPUTING: THE FUNDAMENTAL SCIENCES CASE and OUTCOMES

Contents

PIC
port d'informació científica

INTRODUCTION TO BIG DATA, **MASSIVE PARALLEL** AND HIGH PERFORMANCE COMPUTING: THE FUNDAMENTAL SCIENCES CASE and OUTCOMES

Contents

PIC
port d'informació científica

INTRODUCTION TO BIG DATA, MASSIVE PARALLEL AND **HIGH PERFORMANCE COMPUTING**: THE FUNDAMENTAL SCIENCES CASE and OUTCOMES

Contents

PIC port d'informació científica

INTRODUCTION TO BIG DATA, MASSIVE PARALLEL AND HIGH PERFORMANCE COMPUTING: THE FUNDAMENTAL **SCIENCES CASE** and OUTCOMES

Big Data and Data Science

PIC port d'informació científica

Data science encompasses all the ways in which information and knowledge is extracted from data

It reflects the ways in which data is discovered, conditioned, extracted, compiled, processed, analyzed, interpreted, modeled, visualized, reported on, and presented regardless of the size of the data being processed

Complex field, which is largely due to the diversity and number of academic disciplines and technologies it draws upon

It incorporates mathematics, statistics, numerical methods, programming, statistical modeling, database technologies, signal processing, data modeling, artificial intelligence and learning, natural language processing, visualization, predictive analytics, and so on...

Data science is highly **applicable to many fields** including social media, medicine, security, health care, social sciences, biological sciences, engineering, defense, business, economics, finance, marketing, geolocation, and many more..

Big Data and Data Science

Big data is a field that treats ways to analyze, systematically extract information from, or otherwise deal with data sets that are too large or complex to be dealt with by traditional data-processing application software

The Vs of big data: they represent the qualities of big data in **volume**, **variety**, **velocity**, **veracity**, and **value**. **Variability** is often included as an additional quality!

Developed economies increasingly use data-intensive technologies. The exchange of information through telecommunication networks has reached **a few thousands of Exabytes per year!**

Dedicated research program to improve facilities and techniques that allows enhancing big data processing, for both industry and science

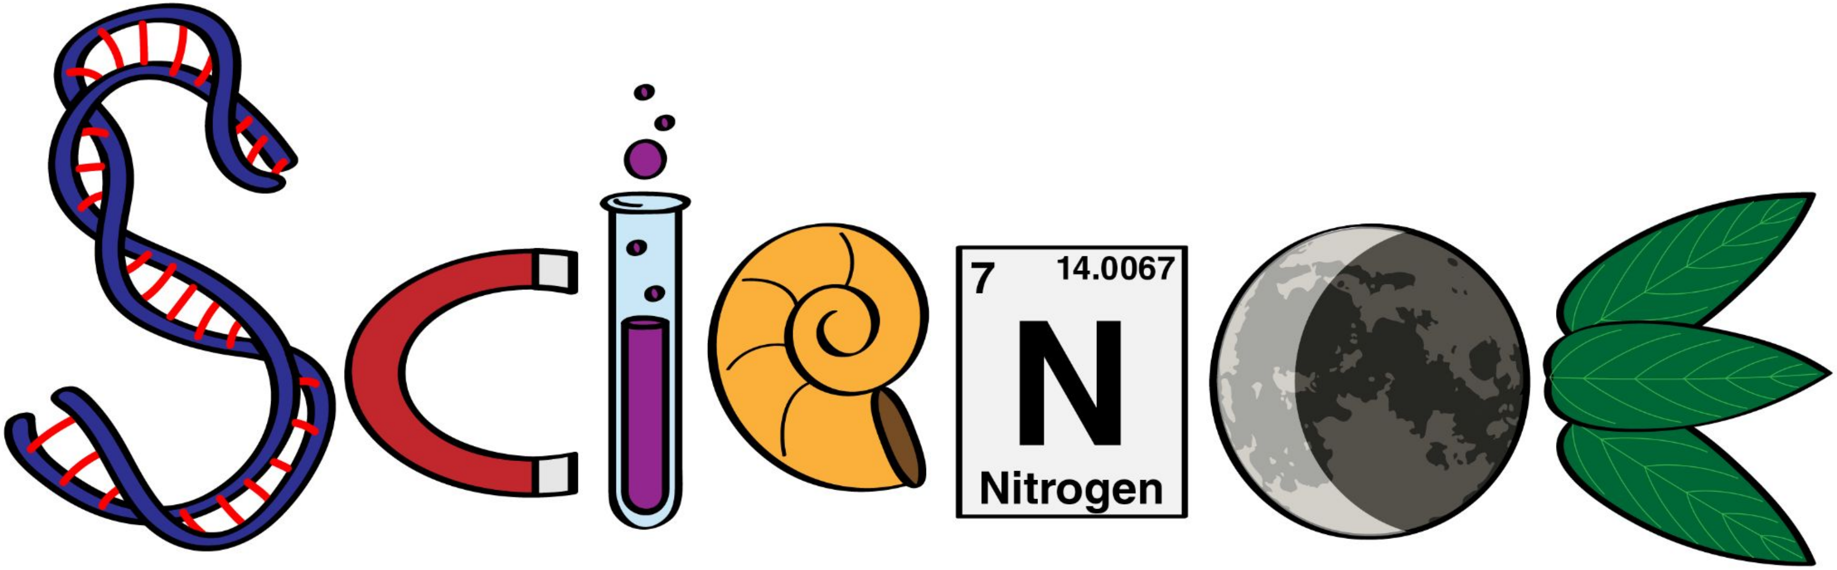

Scientific Research and Big Data - Stanford report

Big Data in Science

Wikipedia
The Free Encyclopedia

Science [edit]

- The Large Hadron Collider experiments represent about 150 million sensors delivering data 40 million times per second. There are nearly 600 million collisions per second. After filtering and refraining from recording more than 99.99995%[114] of these streams, there are 1,000 collisions of interest per second.[115][116][117]
  - As a result, only working with less than 0.001% of the sensor stream data, the data flow from all four LHC experiments represents 25 petabytes annual rate before replication (as of 2012). This becomes nearly 200 petabytes after replication.
  - If all sensor data were recorded in LHC, the data flow would be extremely hard to work with. The data flow would exceed 150 million petabytes annual rate, or nearly 500 exabytes per day, before replication. To put the number in perspective, this is equivalent to 500 quintillion (5×10^{20}) bytes per day, almost 200 times more than all the other sources combined in the world.
- The Square Kilometre Array is a radio telescope built of thousands of antennas. It is expected to be operational by 2024. Collectively, these antennas are expected to gather 14 exabytes and store one petabyte per day.[118][119] It is considered one of the most ambitious scientific projects ever undertaken.[120]
- When the Sloan Digital Sky Survey (SDSS) began to collect astronomical data in 2000, it amassed more in its first few weeks than all data collected in the history of astronomy previously. Continuing at a rate of about 200 GB per night, SDSS has amassed more than 140 terabytes of information.[5] When the Large Synoptic Survey Telescope, successor to SDSS, comes online in 2020, its designers expect it to acquire that amount of data every five days.[5]
- Decoding the human genome originally took 10 years to process; now it can be achieved in less than a day. The DNA sequencers have divided the sequencing cost by 10,000 in the last ten years, which is 100 times cheaper than the reduction in cost predicted by Moore's law.[121]
- The NASA Center for Climate Simulation (NCCS) stores 32 petabytes of climate observations and simulations on the Discover supercomputing cluster.[122][123]
- Google's DNAStack compiles and organizes DNA samples of genetic data from around the world to identify diseases and other medical defects. These fast and exact calculations eliminate any "friction points", or human errors that could be made by one of the numerous science and biology experts working with the DNA. DNAStack, a part of Google Genomics, allows scientists to use the vast sample of resources from Google's search server to scale social experiments that would usually take years, instantly.[124][125]
- 23andme's DNA database contains genetic information of over 1,000,000 people worldwide.[126] The company explores selling the "anonymous aggregated genetic data" to other researchers and pharmaceutical companies for research purposes if patients give their consent.[127][128][129][130][131] Ahmad Hariri, professor of psychology and neuroscience at Duke University who has been using 23andMe in his research since 2009 states that the most important aspect of the company's new service is that it makes genetic research accessible and relatively cheap for scientists.[127] A study that identified 15 genome sites linked to depression in 23andMe's database lead to a surge in demands to access the repository with 23andMe fielding nearly 20 requests to access the depression data in the two weeks after publication of the paper.[132]
- Computational fluid dynamics (CFD) and hydrodynamic turbulence research generate massive data sets. The Johns Hopkins Turbulence Databases (JHTDB) contains over 350 terabytes of spatiotemporal fields from Direct Numerical simulations of various turbulent flows. Such data have been difficult to share using traditional methods such as downloading flat simulation output files. The data within JHTDB can be accessed using "virtual sensors" with various access modes ranging from direct web-browser queries, access through Matlab, Python, Fortran and C programs executing on clients' platforms, to cut out services to download raw data. The data have been used in over 150 scientific publications.

Scientific data: it's a (long) journey

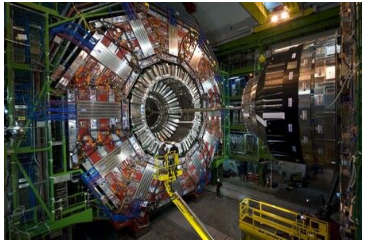

Experiments

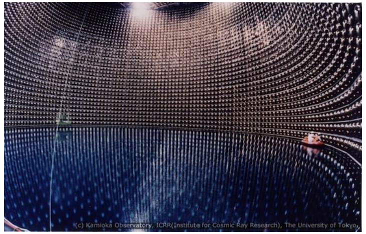

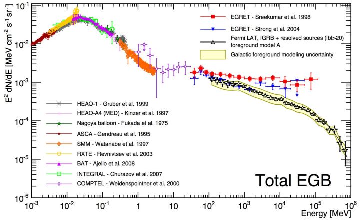

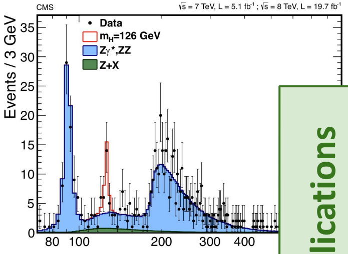

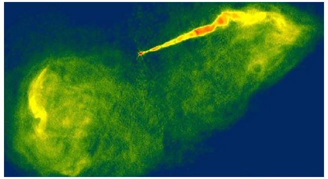

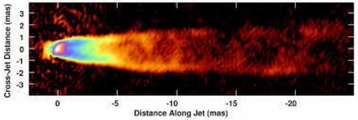

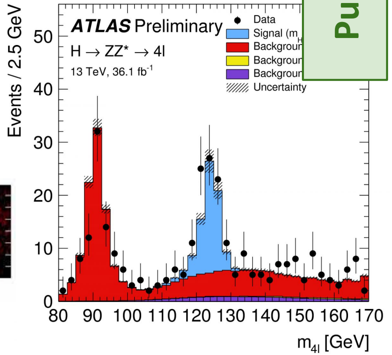

Publications

Simulations: ‘data’ as well

Codes

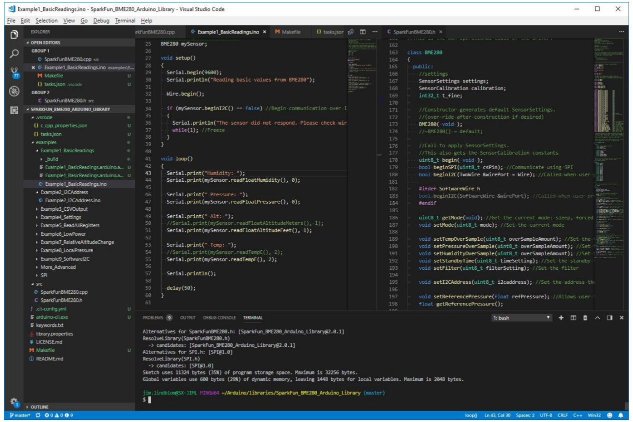

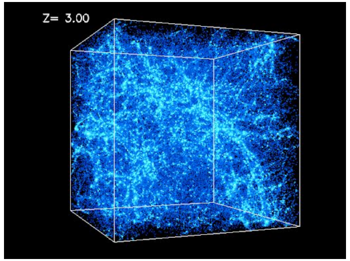

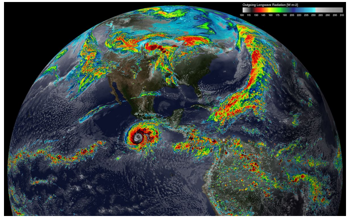

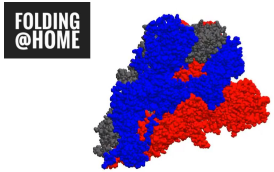

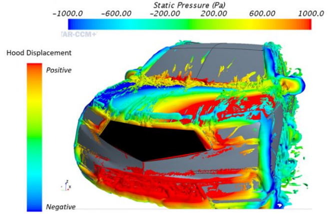

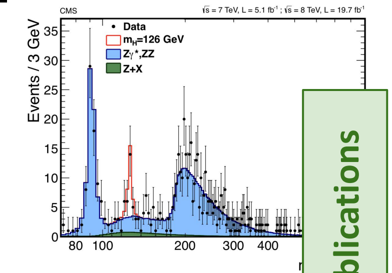

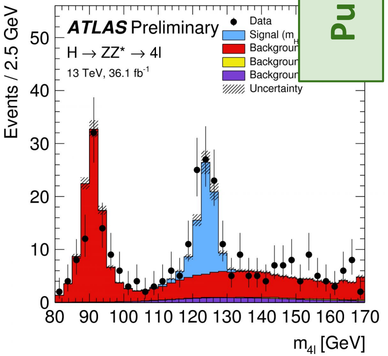

Publications

Design/build of an experiment

Define the scope/need of an experiment → **Technical Design reports**

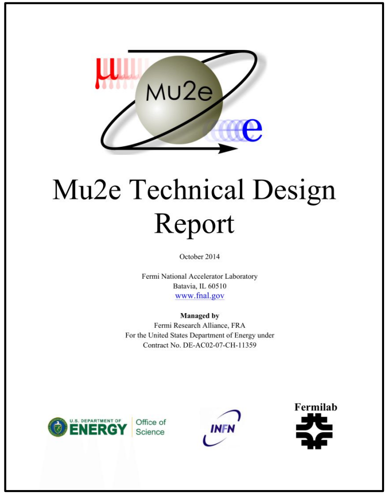

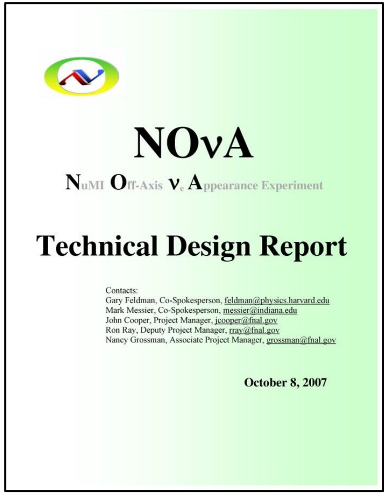

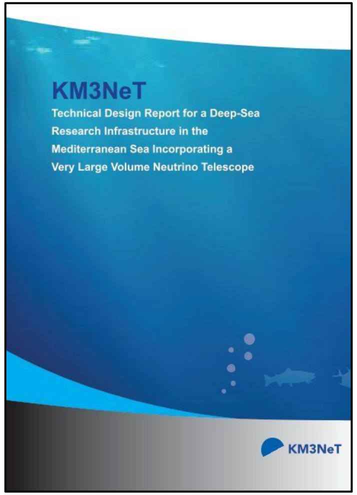

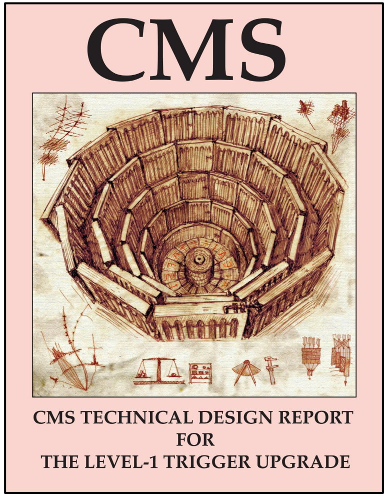

Discovery potentials - measurement(s) window(s) - sensitivity - accuracy - ...

→ Cost driven or limited by R&D efforts

Design/build of an experiment

PIC
port d'informació
científica

Build, in-house tests, install, commissioning and first measurements

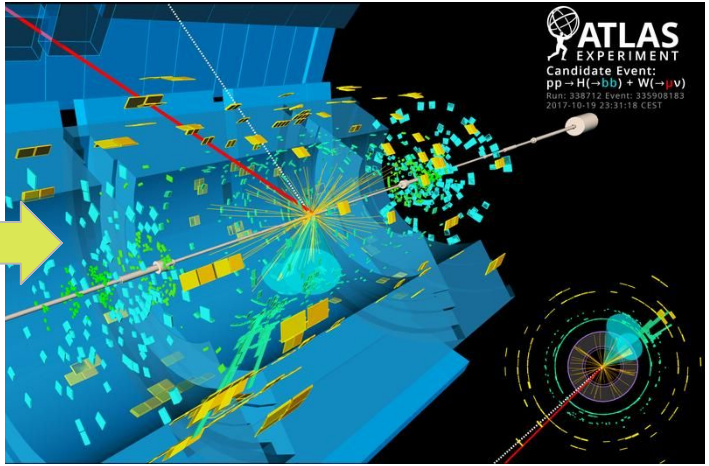

“Big” experiments → **huge worldwide collaboration effort**

Design/build of the needed computing

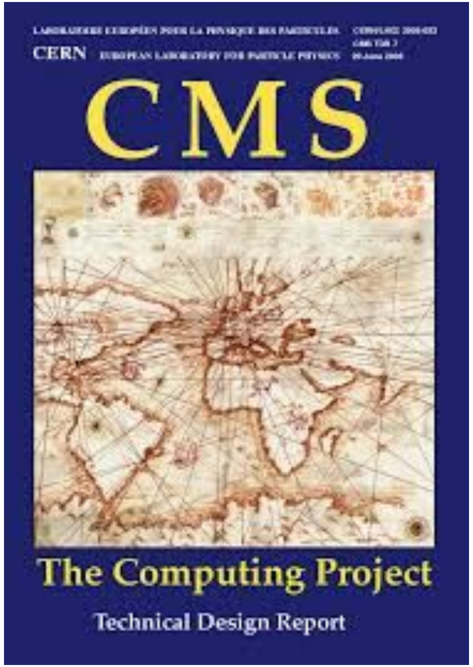

The compute systems of the experiments is also part of the detector!

Design/build of the needed computing

PIC port d'informació científica

Trigger / DAQ system → data volume produced by the experiment

Experience, techniques and available technology allow us to be more sensitive → This can be **large**!

Also, add the simulation needs to get results

Software elements are defined and code architecture is proposed

Data processing characterization - Data formats - Data replication - ...

Compute resources (CPU / storage) needed during the experiment lifetime

Single cluster or distributed

Add data preservation needs

Analog to Digital Converters

Where the experimental data starts!

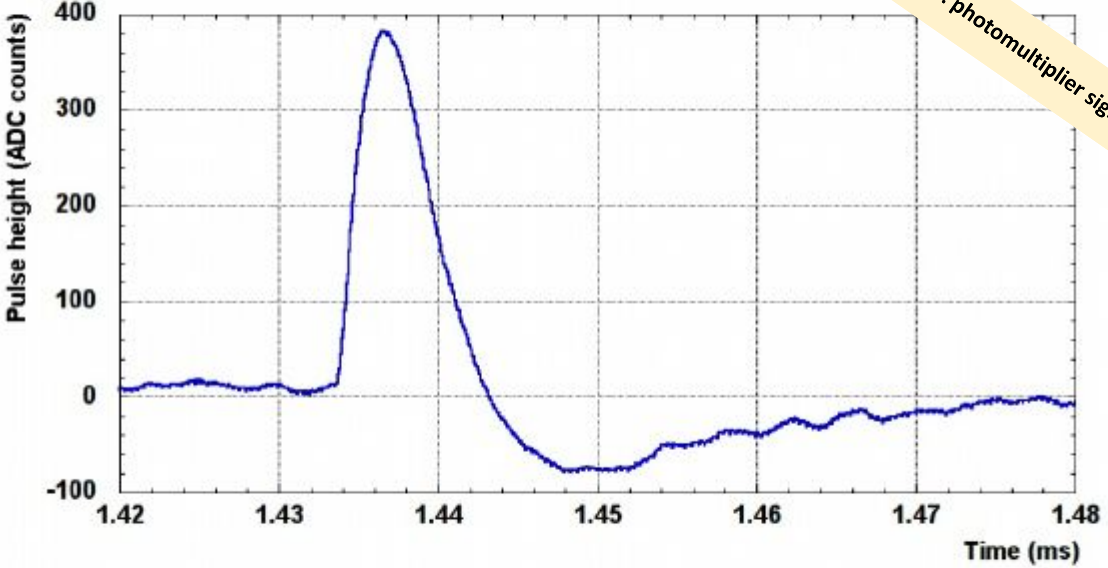

E.g.: photomultiplier signal

Gain adjustments (linearity, saturation), calibrations, pedestal subtraction, signal extraction (integral), ...

Limiting the data volumes!

Store only the **"useful" data**

Limiting the data volumes!

PIC
port d'informació científica

Store only the **“useful” data**

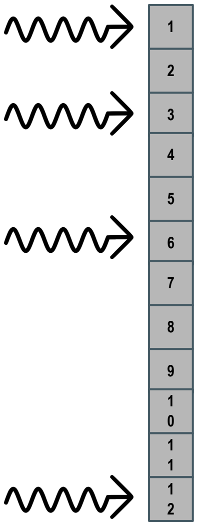

E.g.: 12 channel detector

{114, 0, 90, 0, 0, 73, 0, 0, 0, 0, 0, 85}

[1:114, 3:90, 6:73, 12:85]

Limiting the data volumes!

Store only the **“useful” data**

Use **binary formats** and even **compress the data - structure the data**

Limiting the data volumes!

Store only the **“useful” data**

Use **binary formats** and even **compress the data - structure the data**

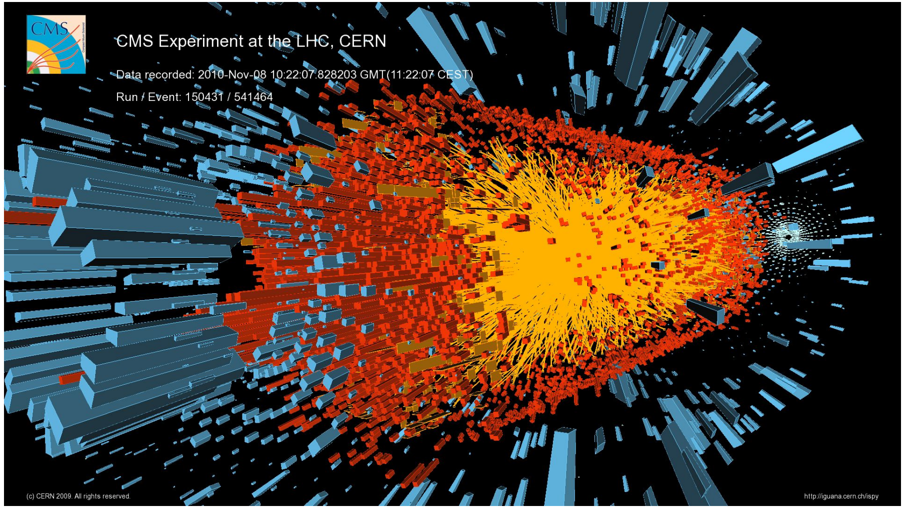

CMS detector has ~300M sensors

A raw event → ~1.5 MB

.root files generated

Limiting the data volumes!

Store only the **“useful” data**

Use **binary formats** and even **compress the data - structure the data**

Define **reduced datasets** for users (no need of continuous data re-reprocess)

Minimize the **data replicas** in a distributed system

Allow your application to **stream data** (if efficient!) from other clusters

Use cold storage (aka Tape systems) to store **unaccessed/old data**

Even delete **intermediate data**, if it can be easily regenerated

Open data (FAIR principles)

Processing (understanding) the data

PIC port d'informació científica

Raw data needs to be reconstructed to get "signals" from sensors

Dead-channels corrections / Calibration constants / ADC signal extraction / … **[dynamic!]**

Then apply sets of algorithms to get the experiment measurement(s) [event]

These algorithms are **refined/optimized** in time

Adding many events to perform statistical analysis and derive results

These can be extremely complex when dealing with **rare physics events**

Processing (understanding) the data

PIC port d'informació científica

Raw data needs to be reconstructed to get “signals” from sensors

Dead-channels corrections / Calibration constants / ADC signal extraction / … **[dynamic!]**

→ **person(s) close to sub-detectors**

Then apply sets of algorithms to get the experiment measurement(s) [event]

These algorithms are **refined/optimized** in time

→ **experienced users / software developers**

Adding many events to perform statistical analysis and derive results

These can be extremely complex when dealing with **rare physics events**

→ **final users / data analyzers**

Writing scientific software

PIC port d'informació científica

During your scientific career you will contribute to the experiment software(s)

You as a **software developer**, rather than a software engineer

Large collaborations have **software gurus**, even software engineers, that help to improve, make more efficient and sustainable, the produced software

But, typically many of the experiments scientific software is not fully optimized

Writing scientific software

PIC port d'informació científica

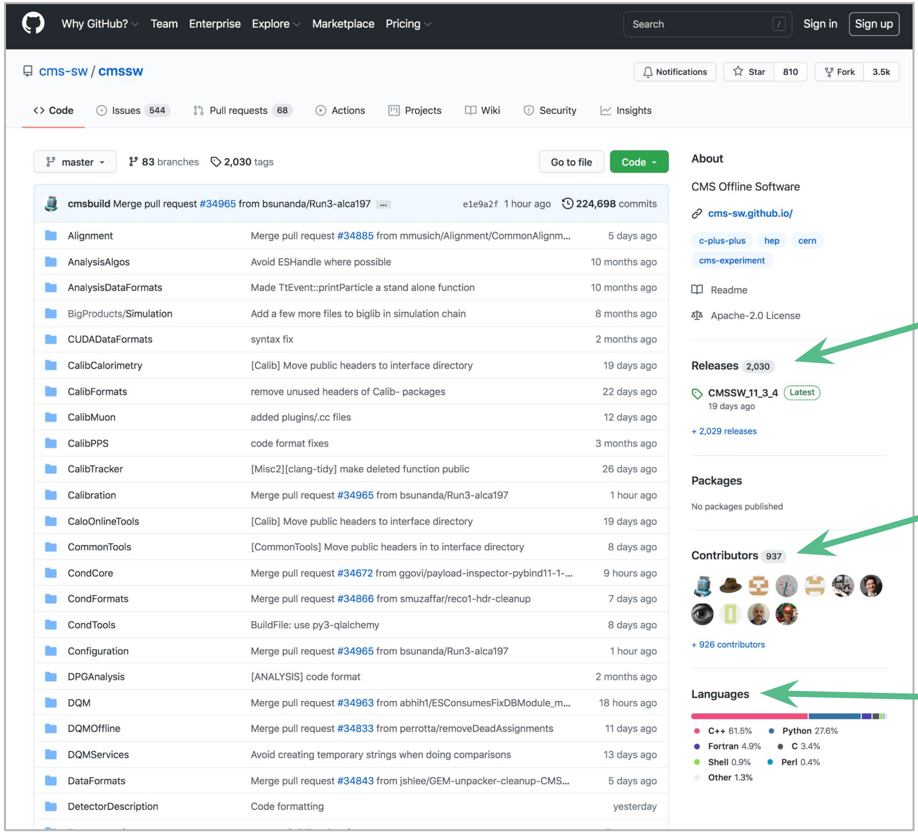

+5 millions code lines

It would take 50 experienced FTEs to re-write completely the code (3 years)

Python becoming popular (for the users)

https://github.com/cms-sw/cmssw

Writing scientific software

PIC port d'informació científica

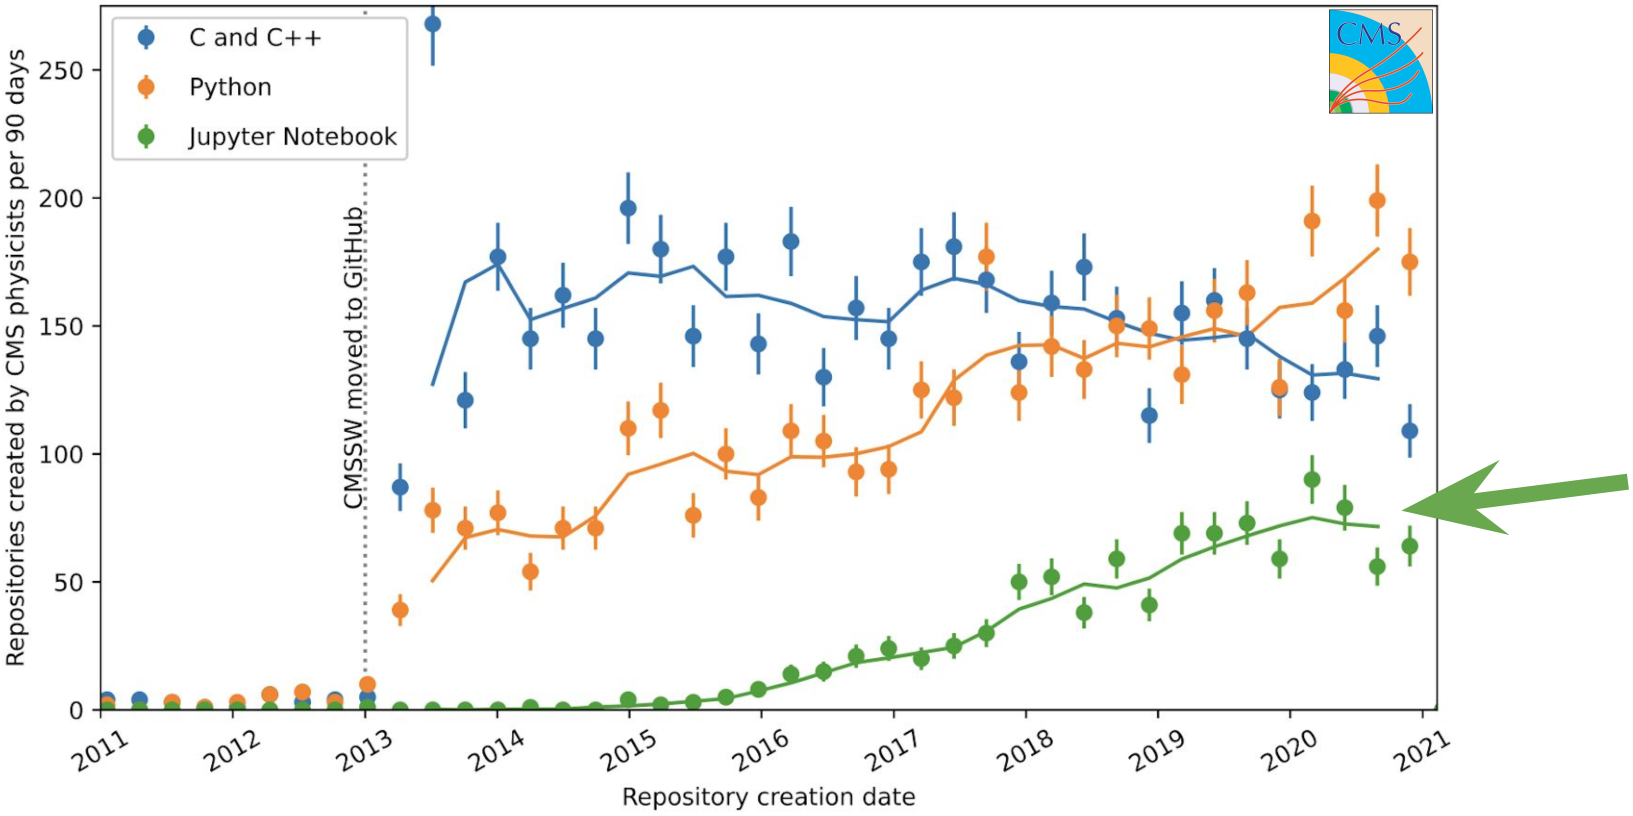

Python: Jupyter Notebooks

Narrative Text

Notebook title and introduction

Description of model parameters

Description of need to profile data

Code and Visualizations

Importing external packages

Implementation of parameters

Profile plotting code

Inline plot

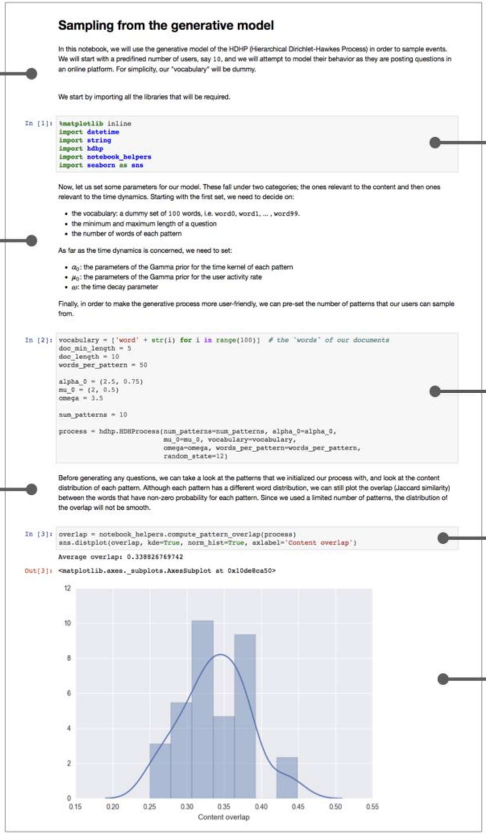

Sampling from the generative model

In this notebook, we will use the generative model of the HDHP (Hierarchical Dirichlet-Hawkes Process) in order to sample events. We will start with a predifined number of users, say 10, and we will attempt to model their behavior as they are posting questions in an online platform. For simplicity, our "vocabulary" will be dummy.

We start by importing all the libraries that will be required.

In [1]:
```
%matplotlib inline
import datetime
import string
import hdhp
import notebook_helpers
import seaborn as sns
```

Now, let us set some parameters for our model. These fall under two categories; the ones relevant to the content and then ones relevant to the time dynamics. Starting with the first set, we need to decide on:

- the vocabulary: a dummy set of 100 words, i.e. word0, word1, ... , word99.
- the minimum and maximum length of a question
- the number of words of each pattern

As far as the time dynamics is concerned, we need to set:

- α_0: the parameters of the Gamma prior for the time kernel of each pattern
- μ_0: the parameters of the Gamma prior for the user activity rate
- ω: the time decay parameter

Finally, in order to make the generative process more user-friendly, we can pre-set the number of patterns that our users can sample from.

In [2]:
```
vocabulary = ['word' + str(i) for i in range(100)]  # the `words` of our documents
doc_min_length = 5
doc_length = 10
words_per_pattern = 50

alpha_0 = (2.5, 0.75)
mu_0 = (2, 0.5)
omega = 3.5

num_patterns = 10

process = hdhp.HDHProcess(num_patterns=num_patterns, alpha_0=alpha_0,
                          mu_0=mu_0, vocabulary=vocabulary,
                          omega=omega, words_per_pattern=words_per_pattern,
                          random_state=12)
```

Before generating any questions, we can take a look at the patterns that we initialized our process with, and look at the content distribution of each pattern. Although each pattern has a different word distribution, we can still plot the overlap (Jaccard similarity) between the words that have non-zero probability for each pattern. Since we used a limited number of patterns, the distribution of the overlap will not be smooth.

In [3]:
```
overlap = notebook_helpers.compute_pattern_overlap(process)
sns.distplot(overlap, kde=True, norm_hist=True, axlabel='Content overlap')
```

```
Average overlap: 0.338826769742
```

Out[3]:
```
<matplotlib.axes._subplots.AxesSubplot at 0x10de8ca50>
```


Popular for final analysis
Running on reduced users datasets

→ **https://github.com/joseflix/DocenciaUAB/tree/master/MN1/2020-2021**

Python: Jupyter Notebooks

PIC
port d'informació
científica

"Why Jupyter is data scientists' computational notebook of choice"

[https://www.nature.com/articles/d41586-018-07196-1]

Programming: good style

PIC
port d'informació
científica

Document your code

Reuse code/routines/packages (be careful to understand what you reuse!)
Import what you need - don’t “ import * “ → memory wastage!

Numerical **accuracy**

Errors / exceptions handling

Good variables **allocation**: caches / pointers / malloc

Beware: **unallocated** mem, **overwriting** mem, **memory** leaks... Use debuggers!

Use **revision control** (CVS, GitHub, GirLab, ...)

Profilers to improve your code efficiency

Programming: good style

Use array programming: operations are applied to whole arrays instead of individual elements

Automatic vectorization: compiler optimization that transforms loops to vector operations

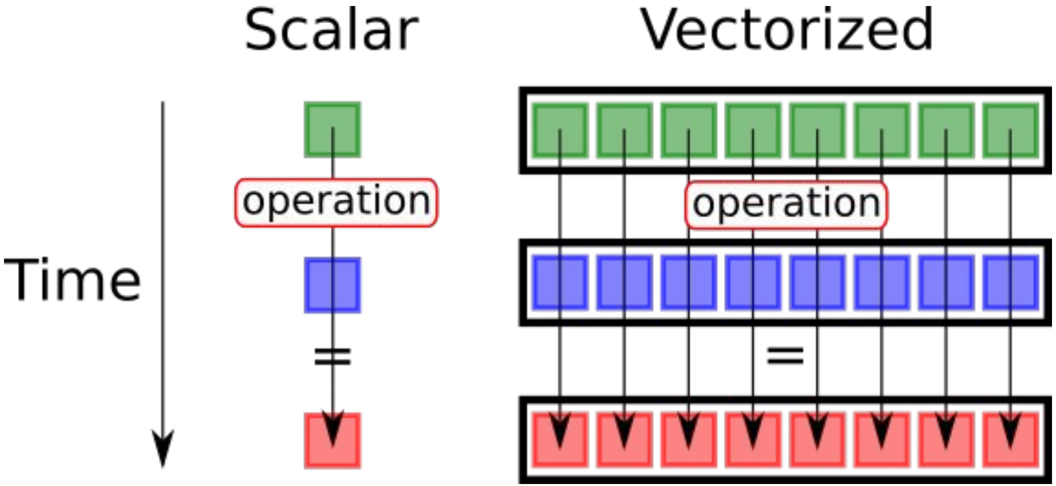

Programming recommendations

CMS Naming, Coding, And Style Rules

By the CMS Offline Software Development Team

Abstract

This document describes the design, naming, coding, and style rules and recommendations for CMS software written in C++, plus a few guidelines for Python configuration.

Outline

1 – Introduction

This document describes the CMS C++ software naming, coding, style, and documentation rules and recommendations. All CMS C++ software is expected to comply with the rules. The asterisk (*) after some rules indicates that there may be exceptional use cases where the rule may be violated with good justification. Intentional violations must be documented in the code next to the violation. Coding rules are meant to prevent serious problems in software function, performance, maintainability, usability, and portability. The Packaging Rules section also has some brief guidelines for Python configuration.

2 – Naming Rules

1. C++ header files use the suffix .h, e.g. `CaloCluster.h`. (*)
2. For C++ source files, the preferred suffix is .cc, e.g. `CaloCluster.cc`. (*)
3. For a header file that contains a class, name that file after the class.
4. Name source files after the class.
5. A geometry XML file should be named as follows:
   base/filename/physical_version_key/implementation_version_key/filename.xml
   E.g.: Geometry/GEMGeometryBuilder/data/GEMSpecs/2019/v1/GEMSpecs.xml
   A revision to a file that has already been in use and that needs to be preserved requires a new version of the file with an incremented version number ("v2" in the example). On the other hand, an upgrade of the detector would require a new `physical_version_key`. For example, an upgraded GEM might become:
   Geometry/GEMGeometryBuilder/data /GEMSpecs/2023/v1/GEMSpecs.xml
6. For class, struct, type, and enumeration names use upper class initials, e.g. `GeometryBuilder`.
7. For namespaces use lower case, e.g. `namespace edm`.
8. Start method names with lowercase, use upper case initials for following words, e.g. `collisionPoint()`.
   Allowed exception: Implementation of virtual methods inherited from external packages, e.g. `ProcessHits()` method required by Geant4.
9. Start data member names with lower case. A trailing "_" is the preferred method to distinguish a data member from the getter method (e.g. `momentum_`).
10. Using "set" for a setter method is preferred, e.g. `setMomentum(double momentum)`.
11. For a getter method, using the value name is preferred, e.g. `momentum()`.

Quick start

- Main page
- FAQ: CMSSW on Github
- CMS Naming, Coding, And Style Rules
- Proposing changes to CMSSW
- Setting up CMS environment using Singularity (new)
- CMSSW pull request approval workflow
- Collaborating with peers
- Working with CMSSW and UserCode
- Resolving conflicts & porting features
- How to use git through a proxy

Release Management

- CMSSW release notes
- Building a CMSSW release
- Integration Builds results
- Historical plots
- Starting a new release cycle
- Forward ports
- Managing users and categories
- List categories and packages
- Latest available IBs
- List of pending PRs
- Circle of pending PRs
- CMS-bot documentation

Infrastructure setup

- Monitoring
- Managing secrets with puppet
- Setting up builder machines
- Setting up builder machines [aarch64]
- Setting up builder machines [power8]
- Setting up cmsdev machines
- Troubleshooting Mesos setup

https://cms-sw.github.io/cms_coding_rules.html

Serial/parallel computing

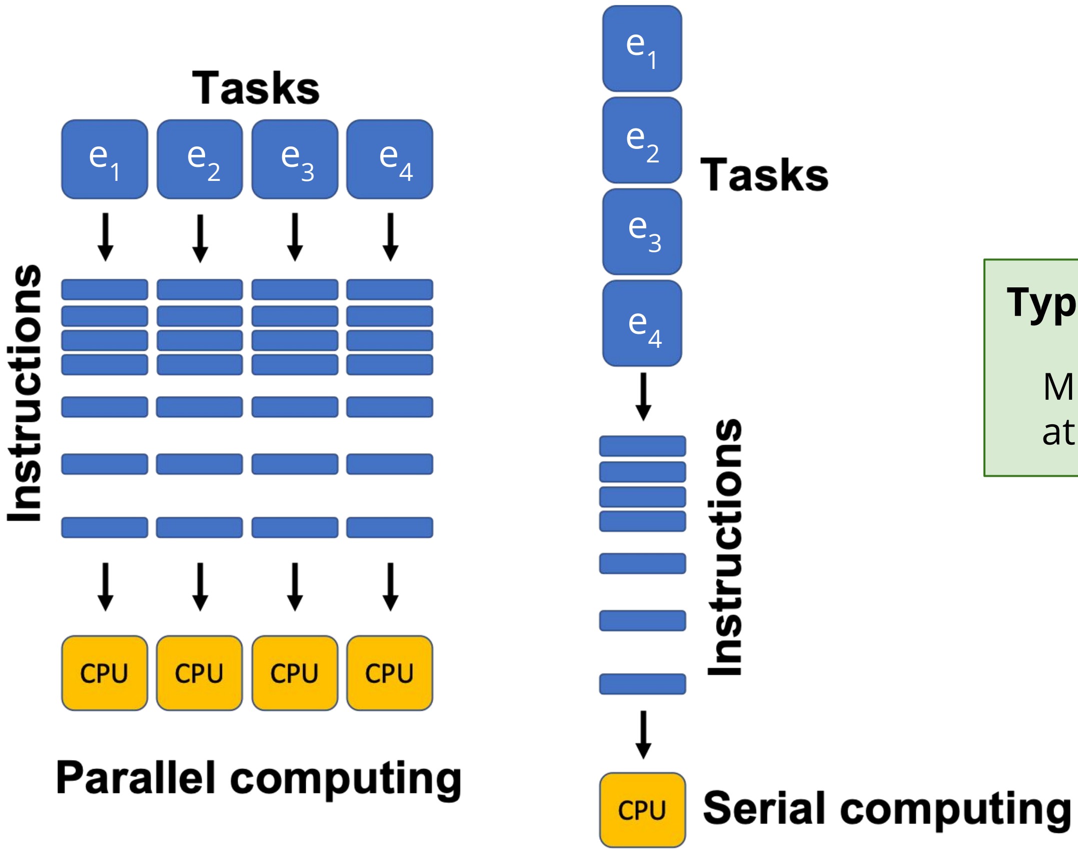

Typically we analyze event by event

Multiple compute nodes can be used at once to speed up the computation

Serial/parallel computing

Since modern CPUs are multi-core, in principle we could use all of the cores of a compute node to run:

- **Multi-process** code
- **Multi-threaded** code (if most of the code can run using threads - efficient)

If you run your application in a HPC cluster, where all of the CPUs talk to each other via Infiniband*, you can also add:

- Processes that **expand to multiple compute nodes**

Typically the **data is read from a local high performance storage** or streamed via network

* InfiniBand (IB) is a computer networking communications standard used in high-performance computing that features very high throughput and very low latency.

You can use GPUs as well!

PIC
port d'informació científica

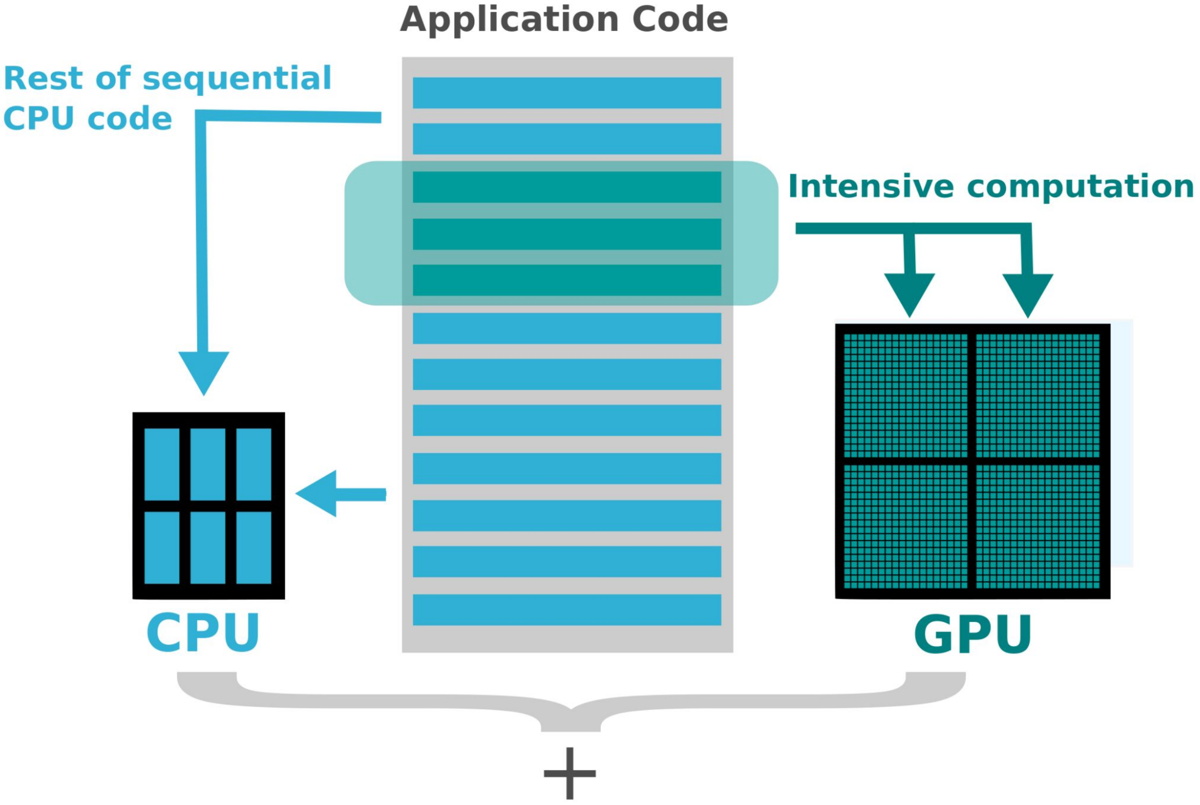

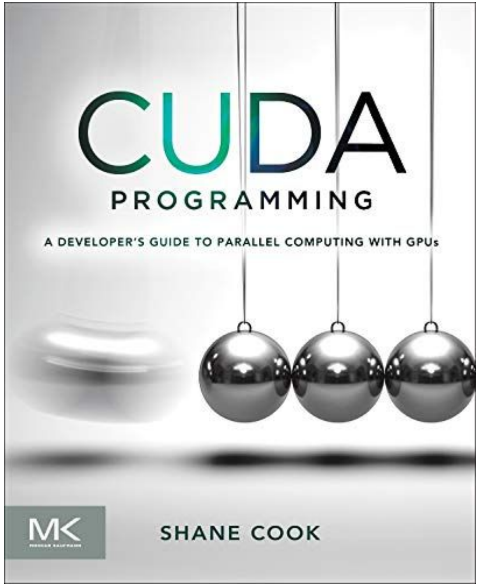

You can even expand to Clouds!

Private clouds

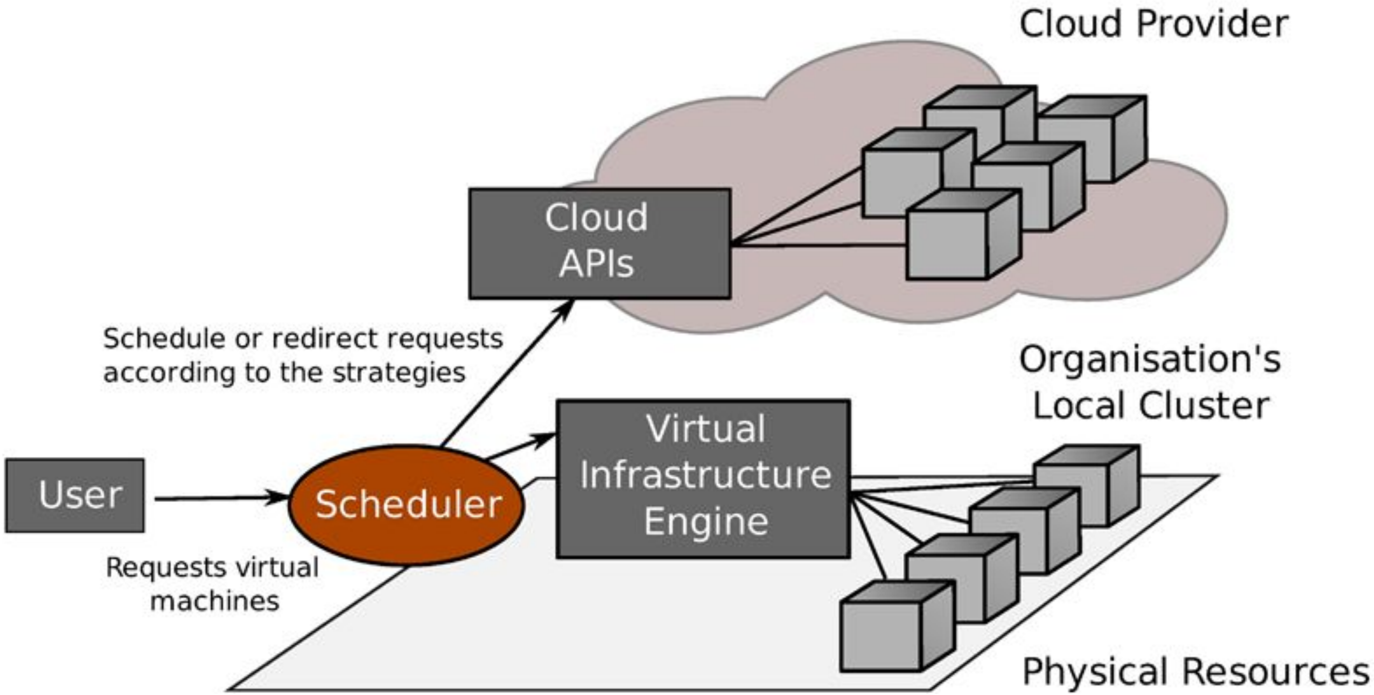

What is High Performance Computing?

High Performance Computing is an aggregation of computing power to solve problems which are either **too large** for standard computers or would **take too long** ← aka **Supercomputing**

It enables **simulation or analysis of huge volumes of data** that would otherwise not be possible with standard computers

This type of computing **speeds up workload execution times** which helps innovation and research, and scientific progress overall

How is HPC achieved?

Use a **compute cluster** to send independent tasks across many compute nodes (using a Job scheduler)

Inside a compute node, they can run **multi-process** or **multi-threaded**

If this is an **"HPC facility"**, then the nodes can talk each other, suitable for applications that need for **large parallelization**

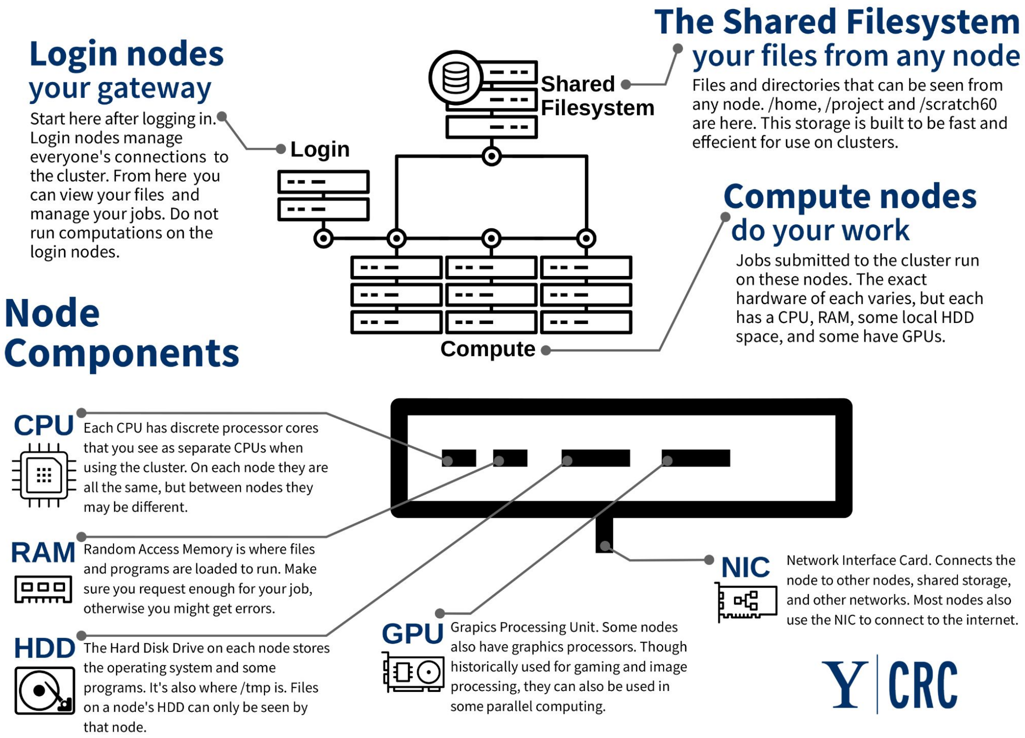

How is HPC achieved?

PIC port d'informació científica

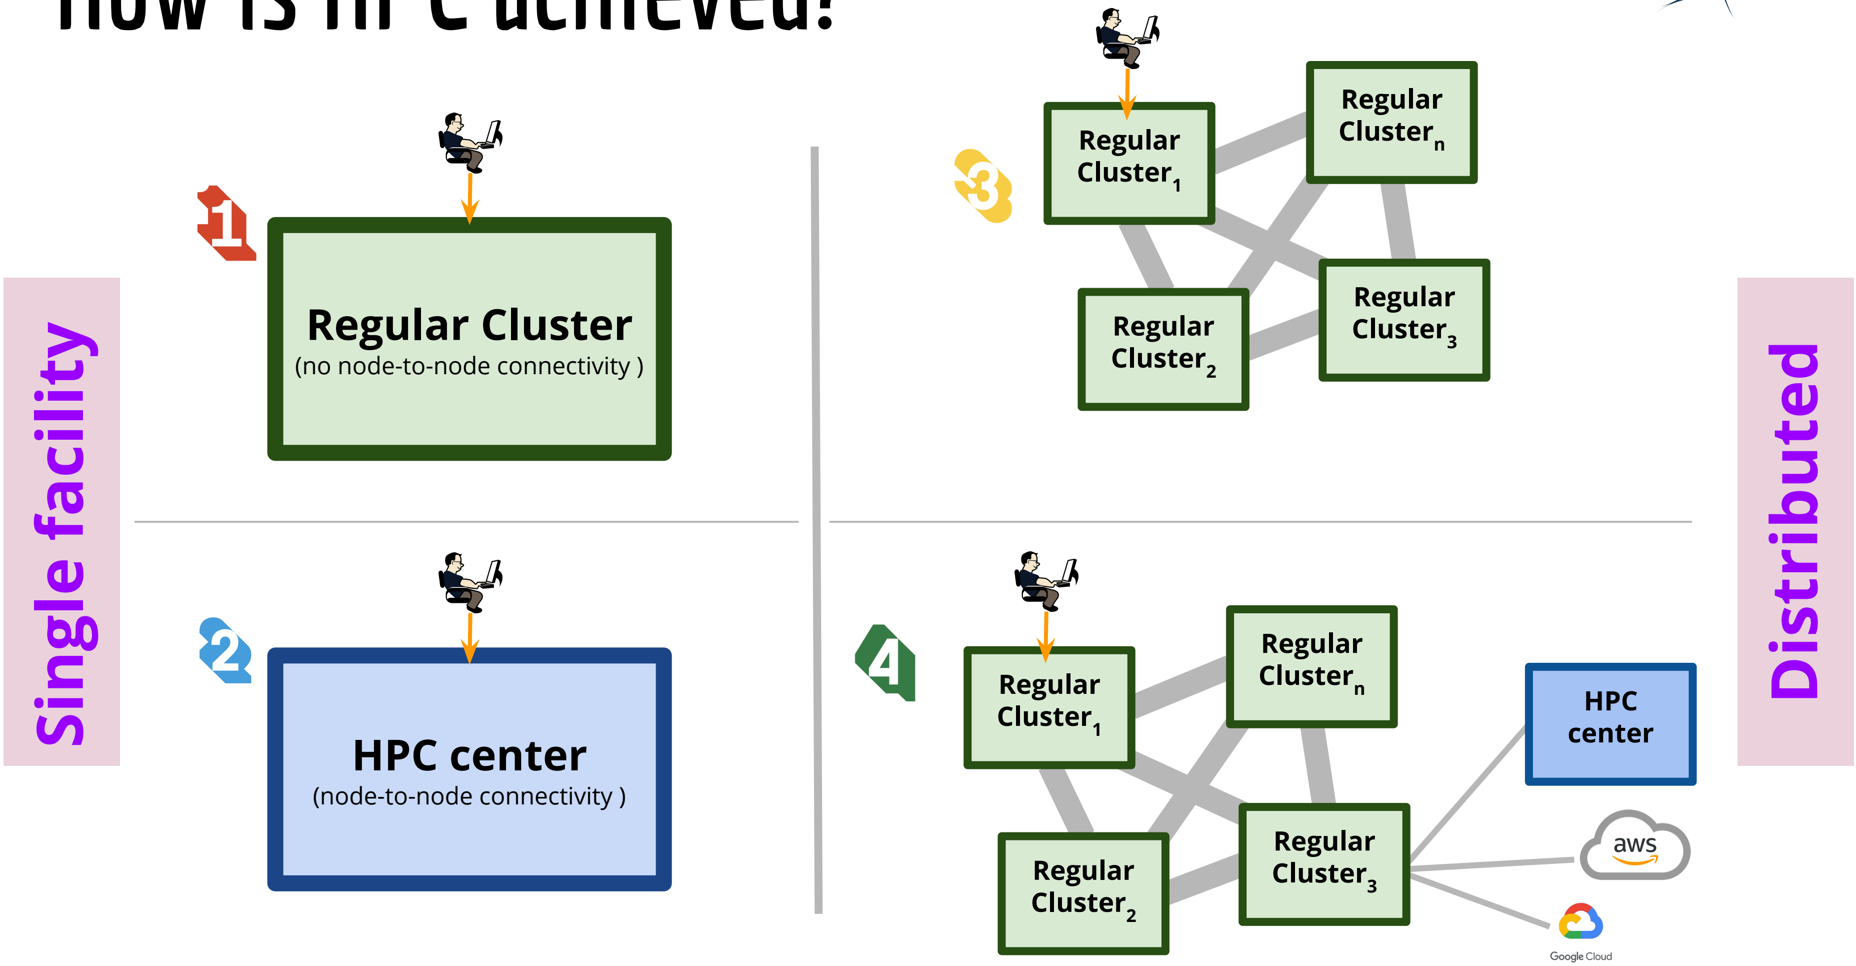

Inside a cluster

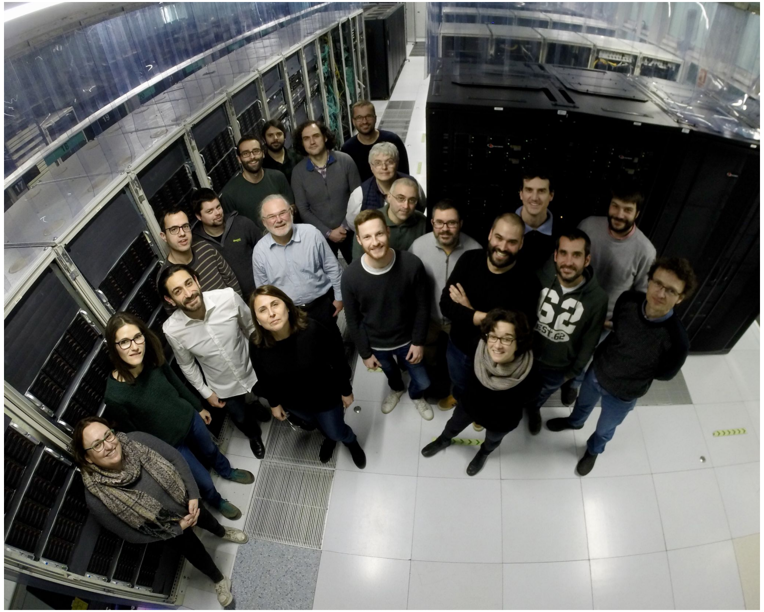

CERN

PIC
port d'informació
científica

Inside a cluster

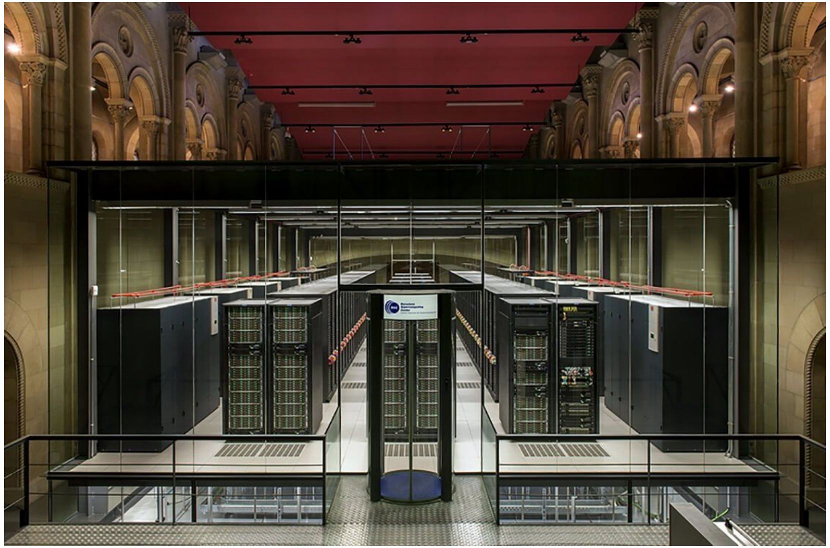

BSC
Barcelona
Supercomputing
Center
Centro Nacional de Supercomputación

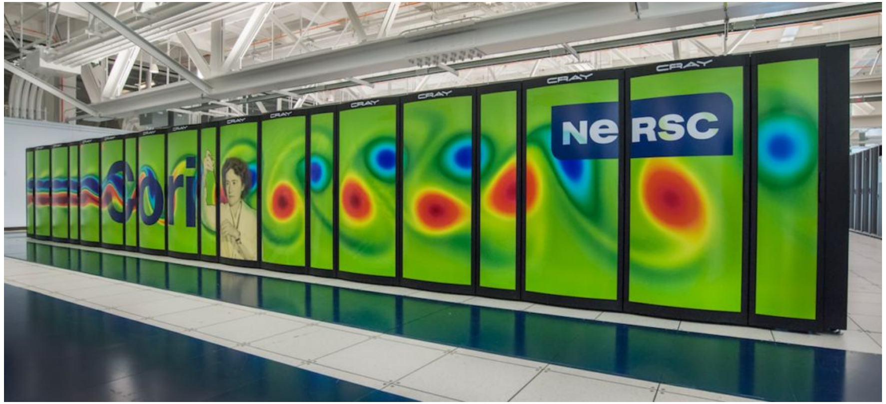

NERSC
National Energy Research
Scientific Computing Center

Use case: LHC computing

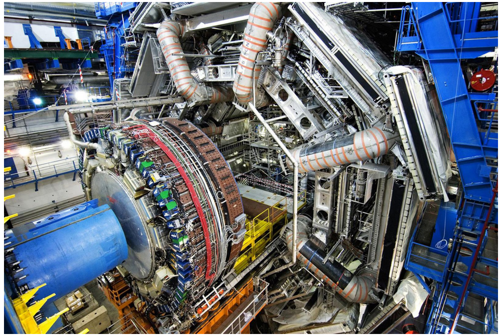

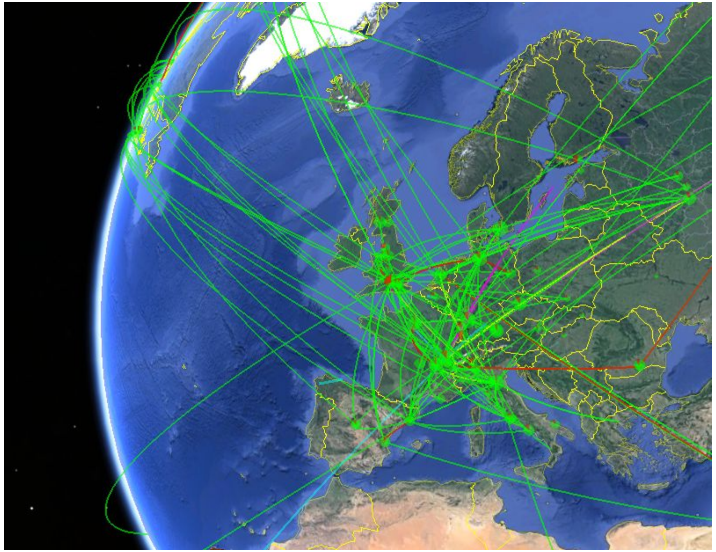

→ https://www.youtube.com/watch?v=jDC3-QSiLB4

→ https://www.youtube.com/watch?v=UBJrNq4rucg

The LHC Accelerator and detectors

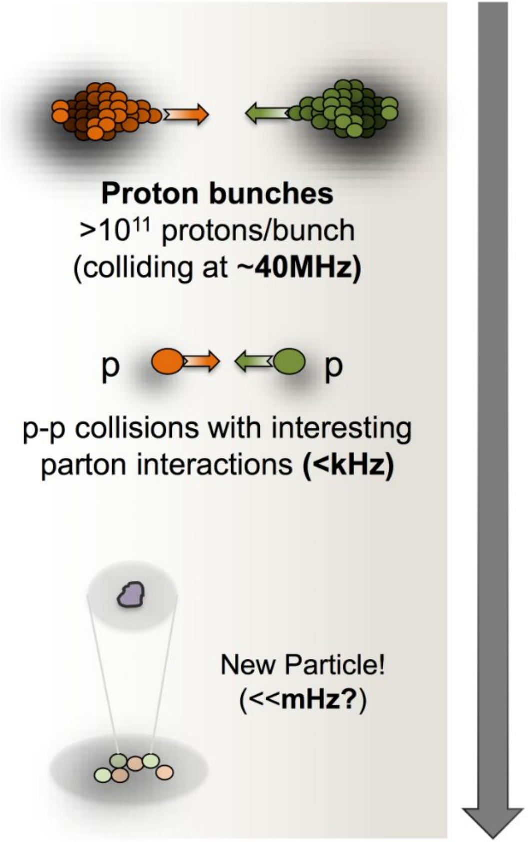

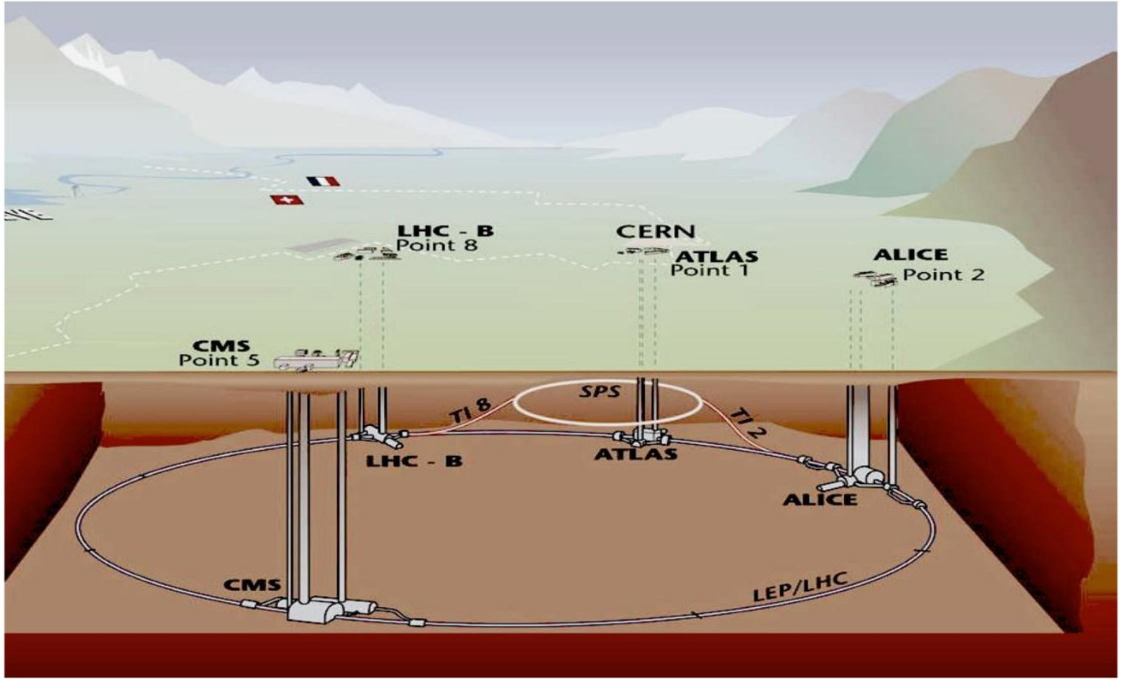

27 km ring of superconducting magnets and accelerating cavities

S. Campana - EIROFORUM 2020

The LHC generates large amounts of data!

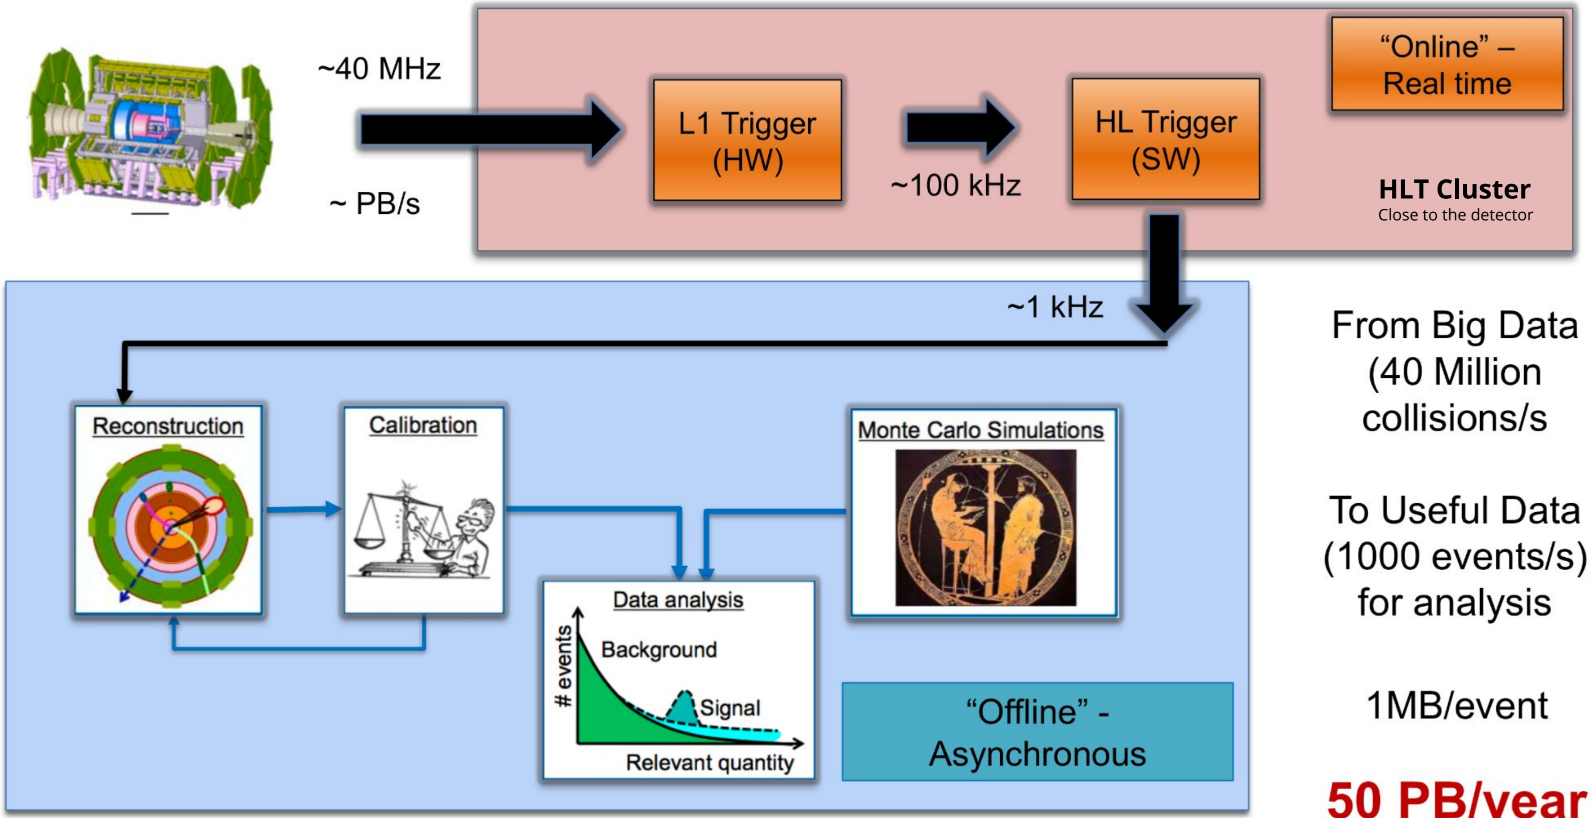

From Big Data
(40 Million
collisions/s

To Useful Data
(1000 events/s)
for analysis

1MB/event

50 PB/year

S. Campana - EIROFORUM 2020

The LHC Computing Grid

Large scale HEP computing embraced a distributed model (**the Grid**) since early 2000s, based on network service technologies, federating national and international initiatives

Worldwide LHC Computing Grid (WLCG): an international collab. to distribute and analyse LHC data

Integrates computer centres **worldwide** that provide computing and storage resource into a single infrastructure accessible by all LHC physicists - dedicated ***middleware*** to transparently access to them

Tier-0 (CERN)
Data recording, reconstruction and distribution

Tier-1s
Permanent storage, re-processing, analysis

Tier-2s
Simulation, end-user analysis

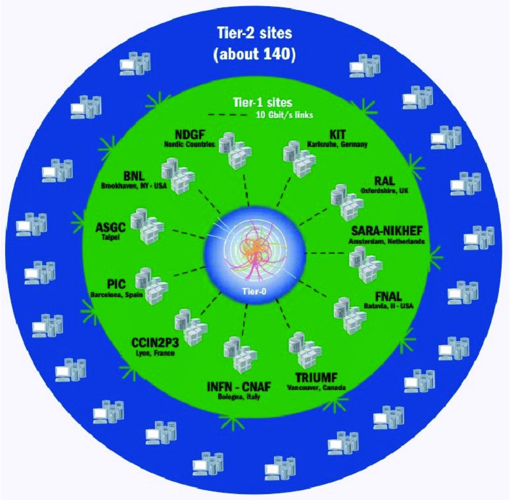

161 sites, 42 countries

~1M CPU cores

~1 EB of storage

> 2 million jobs/day

10-400 Gbps links

The LHC computing Grid

PIC
port d'informació
científica

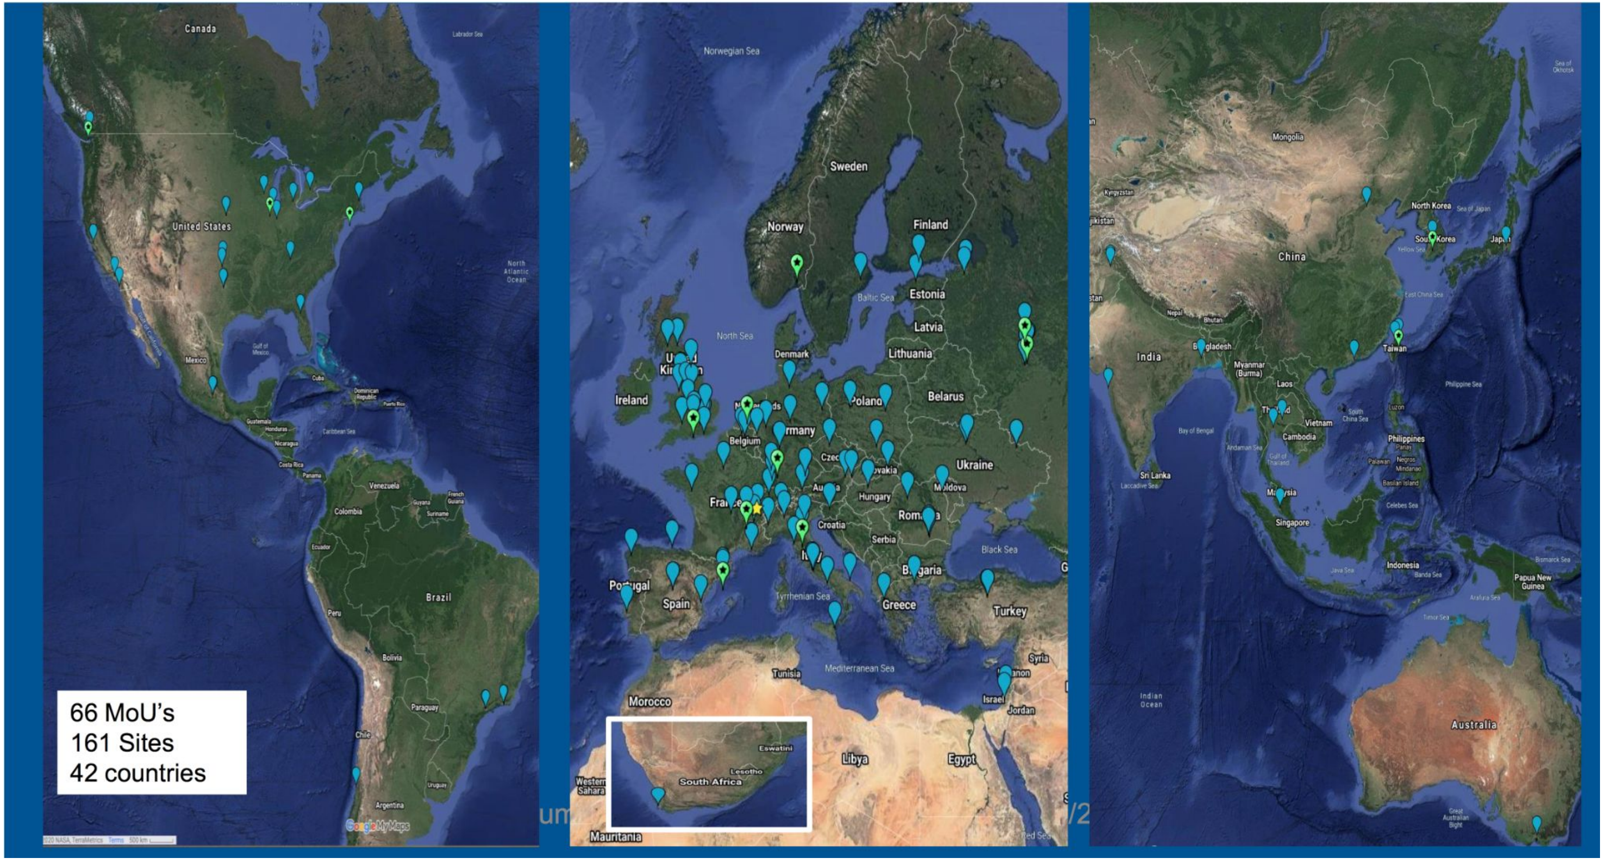

S. Campana - EIROFORUM 2020

LHC Optical Private Network

Private network connecting Tier0 and Tier1s:
- Dedicated to LHC data transfers and analysis
- Secured: only declared IP prefixes can exchange traffic
- Advanced routing: communities for traffic engineering, load balancing
- Star topology around the Tier-0 (CERN)

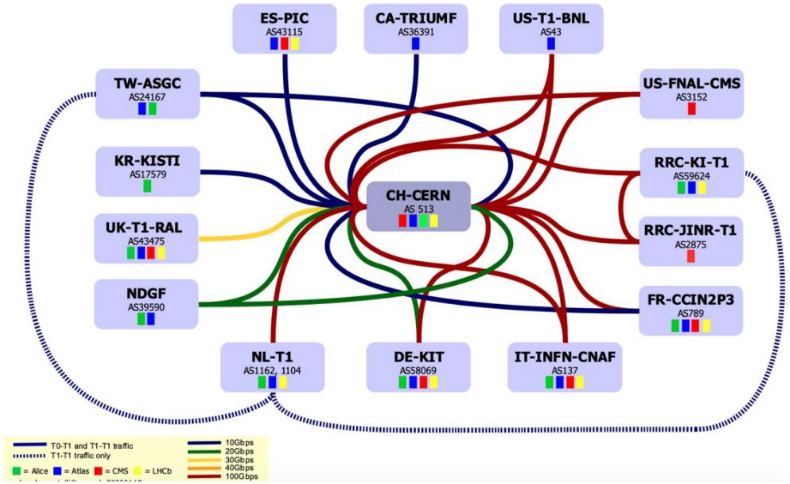

Numbers LHCOPN

- 14 Tier1s + 1 Tier0
- 12 countries in 3 continents
- Dual stack IPv4-IPv6
- 1.1Tbps to the Tier0
- Moved ~288 PB in 12 months

S. Campana - EIROFORUM 2020

LHC Open Network Environment

PIC
port d'informació
científica

Layer3 (routed) Virtual Private Network:
- Worldwide network backbone connecting Tier1s and Tier2s at high bandwidth
- Provided by Research and Education Network providers
- Bandwidth dedicated to High Energy Physics data transfers
- Trusted traffic that can be allowed to bypass slow perimeter firewalls

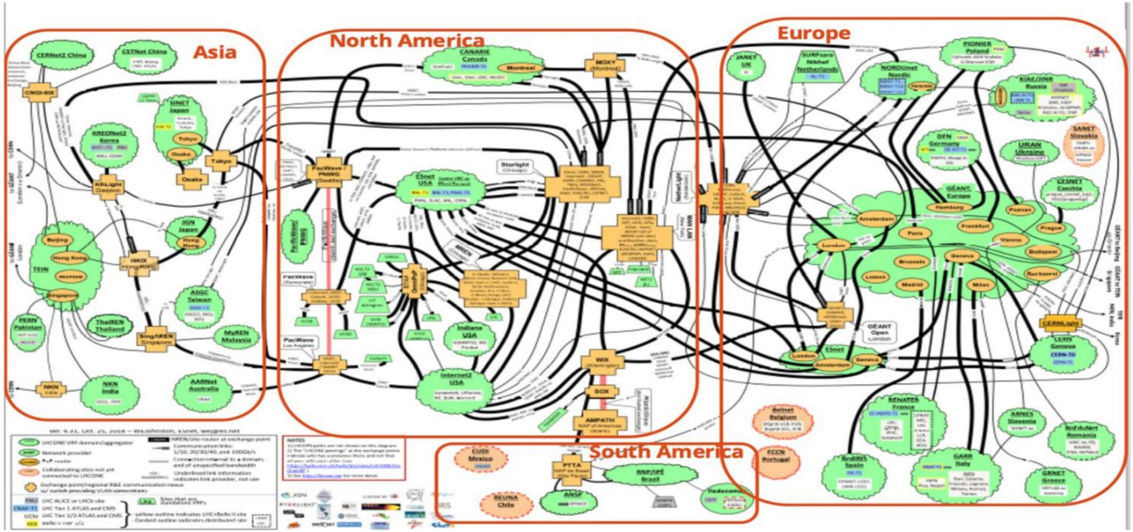

Numbers
- 28 R&E networks
- 14 Tier1s and ~90 Tier2s in 5 continents
- ~250 perfSONAR instances

S. Campana - EIROFORUM 2020

Last 5 years WLCG traffic GB/s

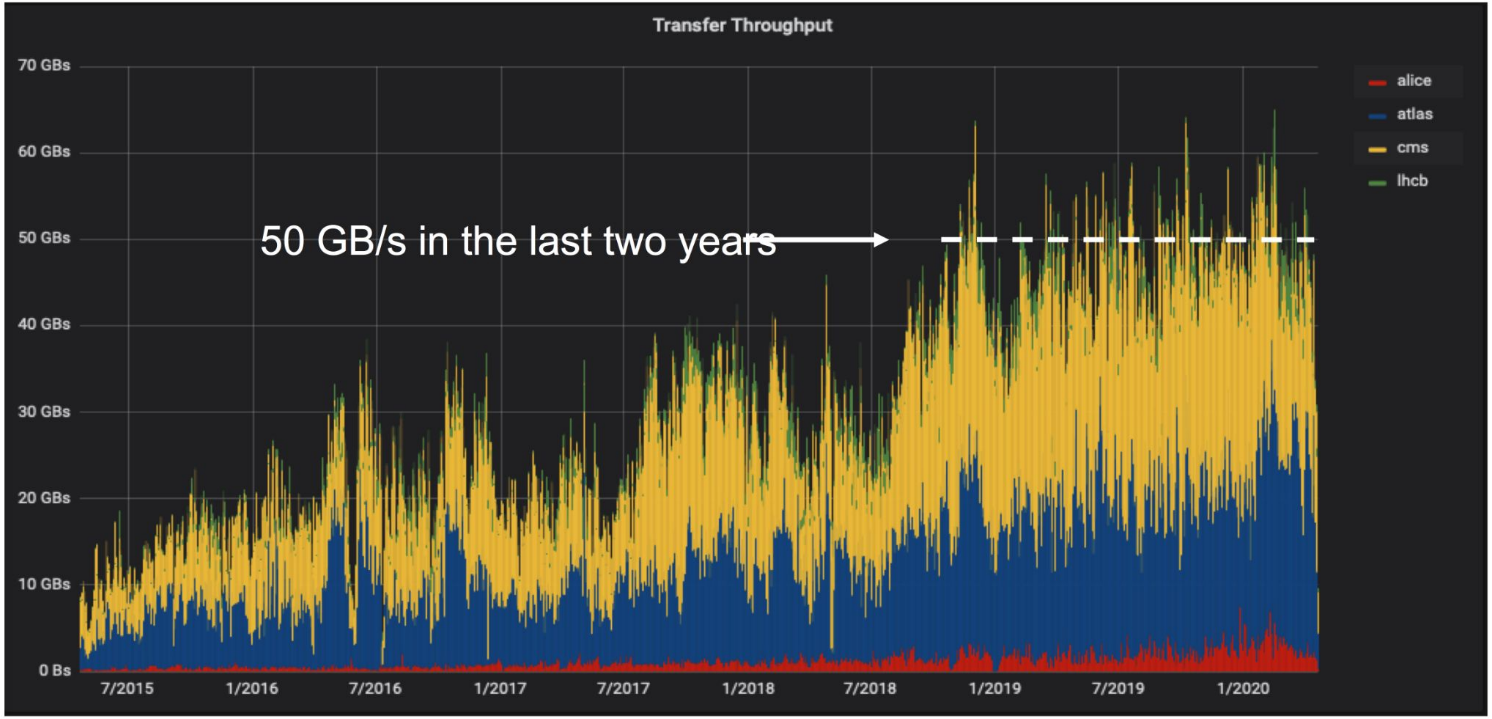

S. Campana - EIROFORUM 2020

Evolution towards a more flexible model

PIC
port d'informació científica

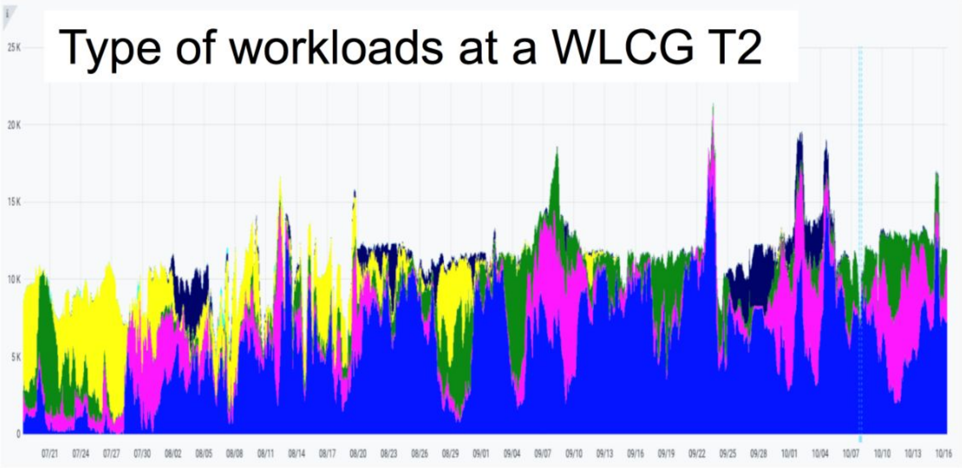

T1s and T2s run flexibly a mixture of workloads depending on their capabilities and the needs of the experiments

Data is transferred in a full mesh rather than hierarchically

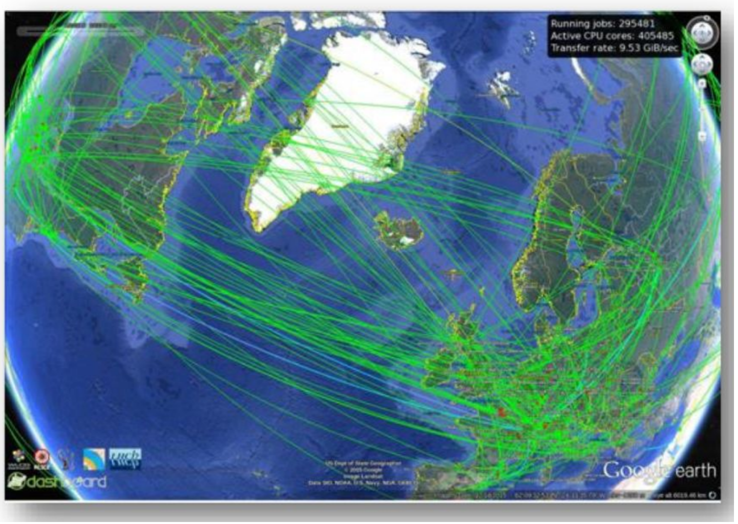

S. Campana - EIROFORUM 2020

WLCG usage and performance

CPU use in WLCG continuously increases despite shutdown periods

We operate 24/7/365

1M cores concurrently in use

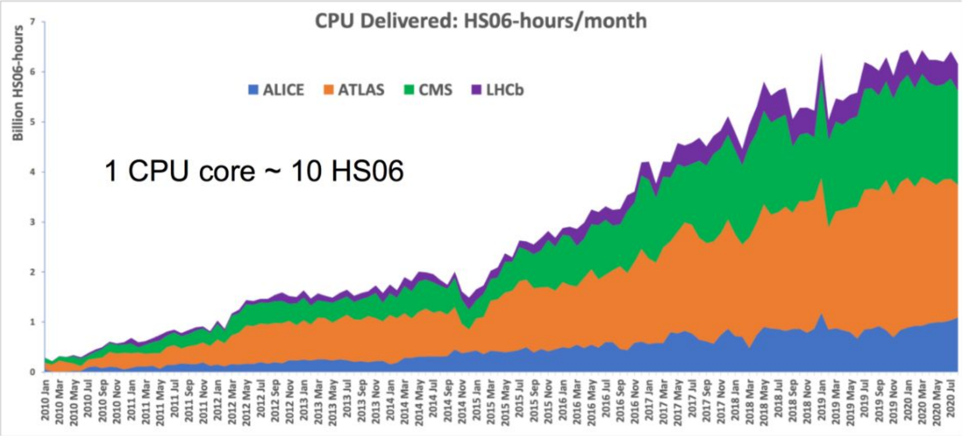

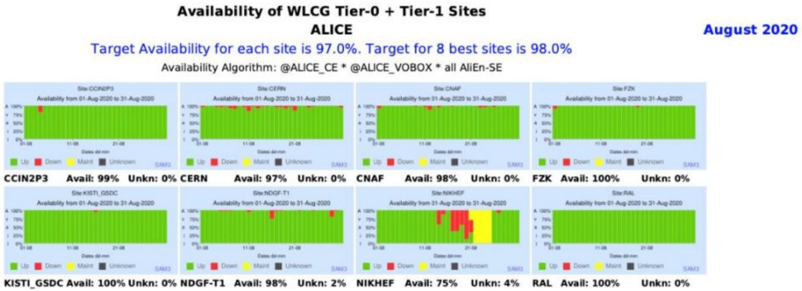

We monitor continuously and and review regularly the performance of the sites against the MoU targets

S. Campana - EIROFORUM 2020

The next challenge: HL-LHC

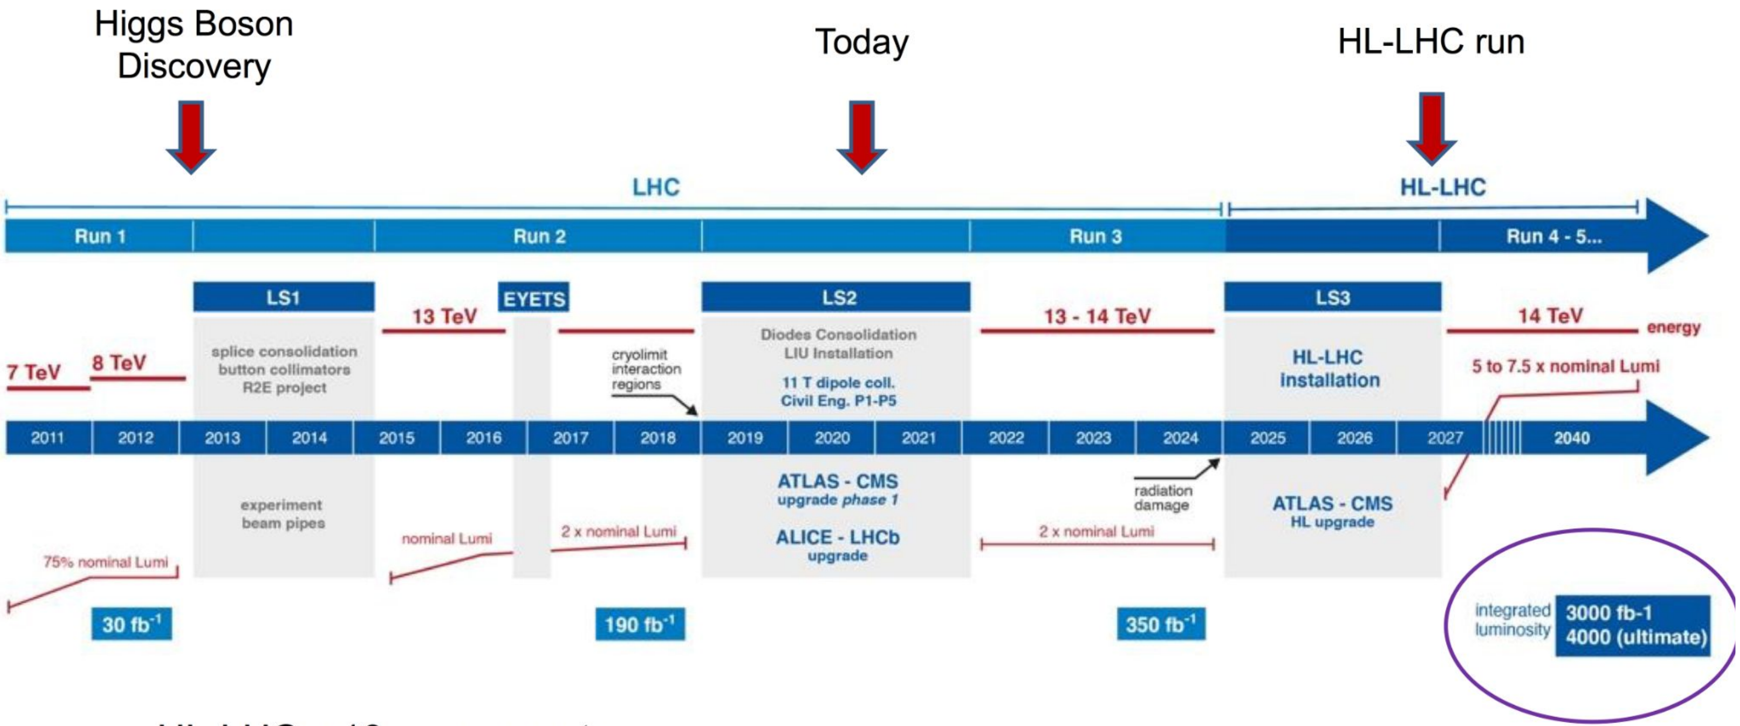

HL-LHC: x10 more events.
Bigger and more complex events (x5)

An unprecedented challenge!

S. Campana - EIROFORUM 2020

Events at the HL-LHC

CMS: a very large event from 2017 with 78 reconstructed vertices

ATLAS simulation for HL-LHC with 200 reconstructed vertices

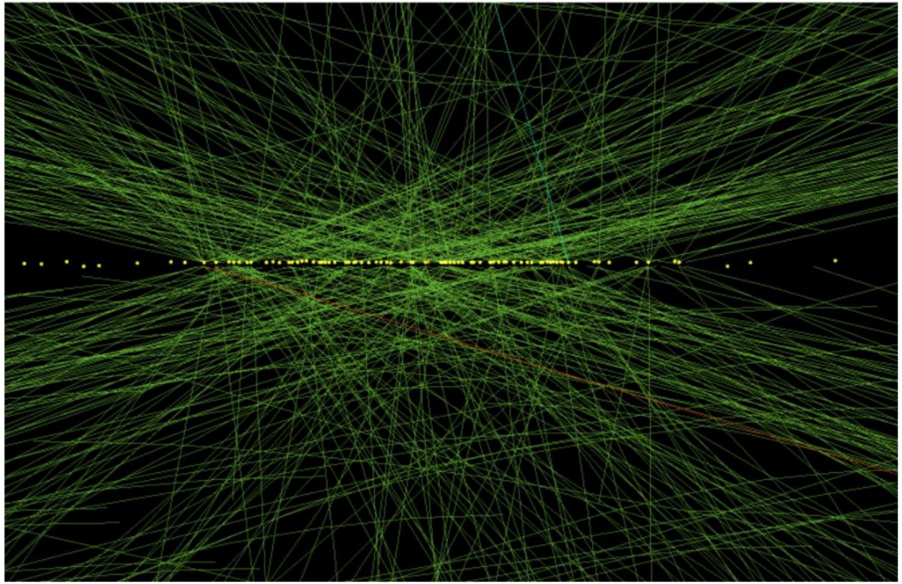

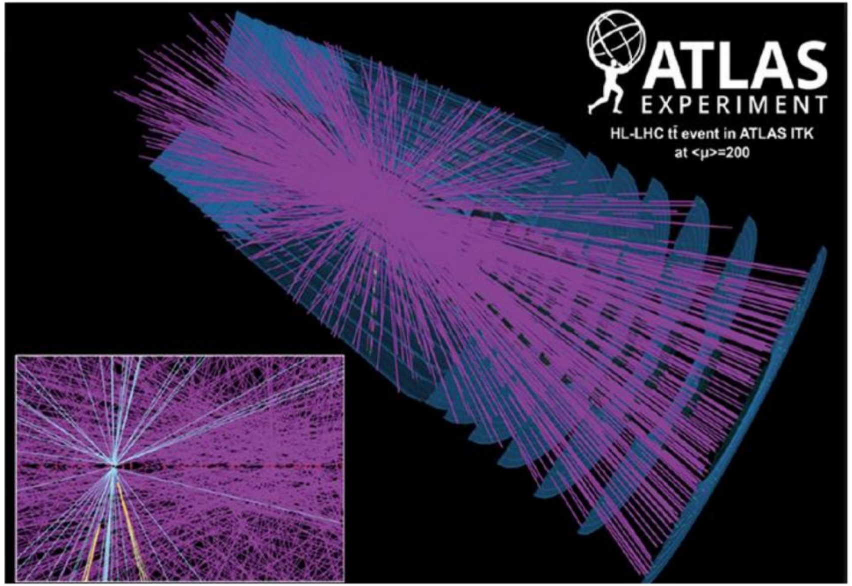

S. Campana - EIROFORUM 2020

Hardware trends

PIC port d'informació científica

- Cost of hardware decreasing exponentially vs time 😄 … but not as steeply as before 😳
- In general, trends driven by market (revenues) rather than technology

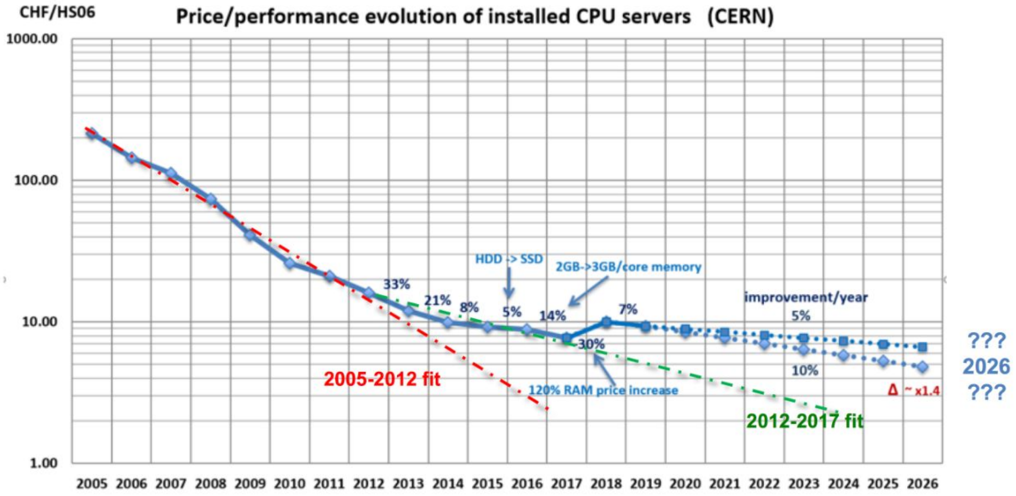

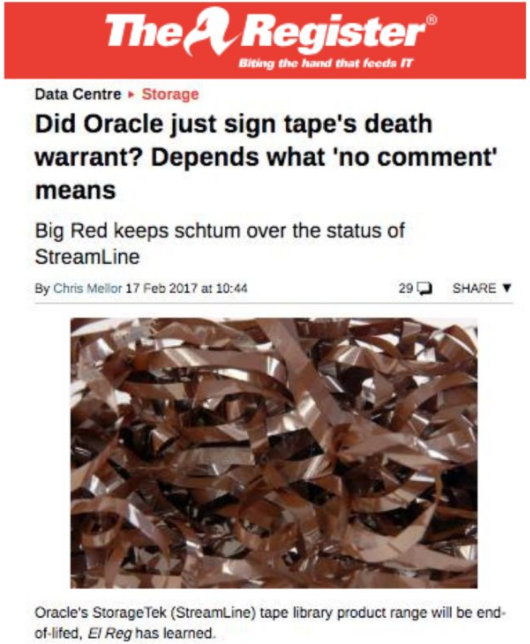
The Register®
Biting the hand that feeds IT

Data Centre ▸ Storage

Did Oracle just sign tape's death warrant? Depends what 'no comment' means

Big Red keeps schtum over the status of StreamLine

By Chris Mellor 17 Feb 2017 at 10:44 29 SHARE ▼

Oracle's StorageTek (StreamLine) tape library product range will be end-of-lifed, *El Reg* has learned.

General loss of long term predictability

S. Campana - EIROFORUM 2020

The HL-LHC Computing Challenge

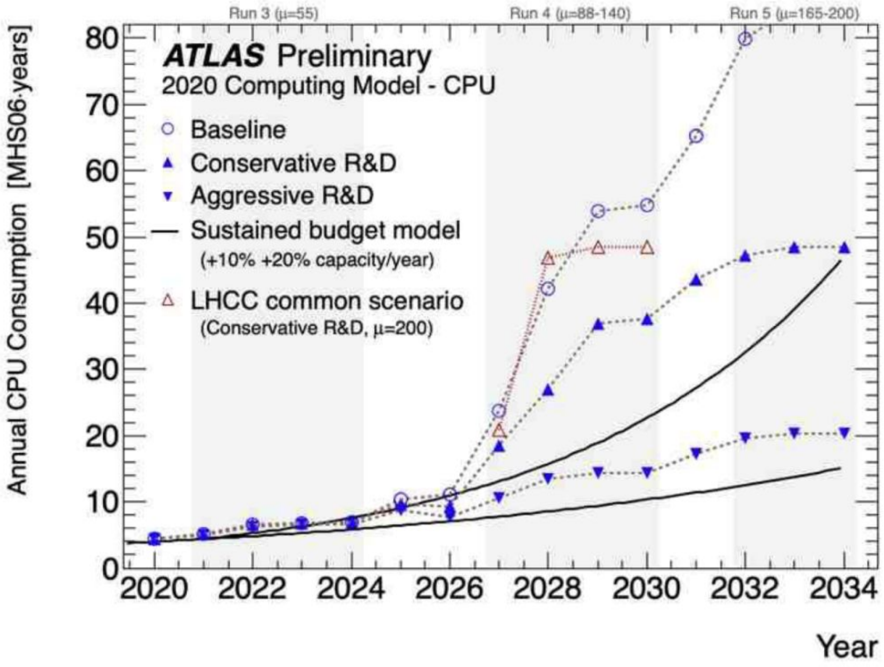

We do not foresee an increase in funding for LHC Computing in the coming years

An aggressive investment in R&D is needed to address the resource challenge

The infrastructure will need to evolve as well, not only in terms of technology but also policies

Long term sustainability of the infrastructure is at least as important as filling the gap in the resource need

The Funding Agencies want to see their investments to benefit multiple sciences

S. Campana - EIROFORUM 2020

Heterogeneous computing

Today we get opportunistic use of many type of compute in particular HPC systems and the experiment online farms presented as Cloud resources

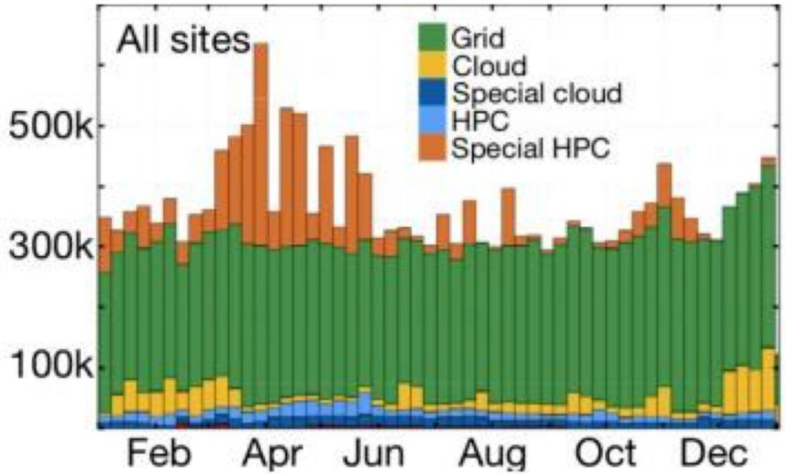

In the future this heterogeneity will expand and we need to be able to make use of the resources provided to us

A lot of work going on at the level of policies, infrastructure/services and software

New Summit HPC @ ORNL:
9.96 Pflops in CPU (**non X86_64**)
215.7 Pflops in GPUs

S. Campana - EIROFORUM 2020

Moving to a network-centric model

Datalake model:

Fewer number of facilities operating storage services, less data replication

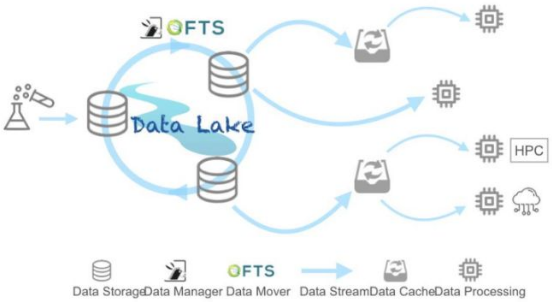

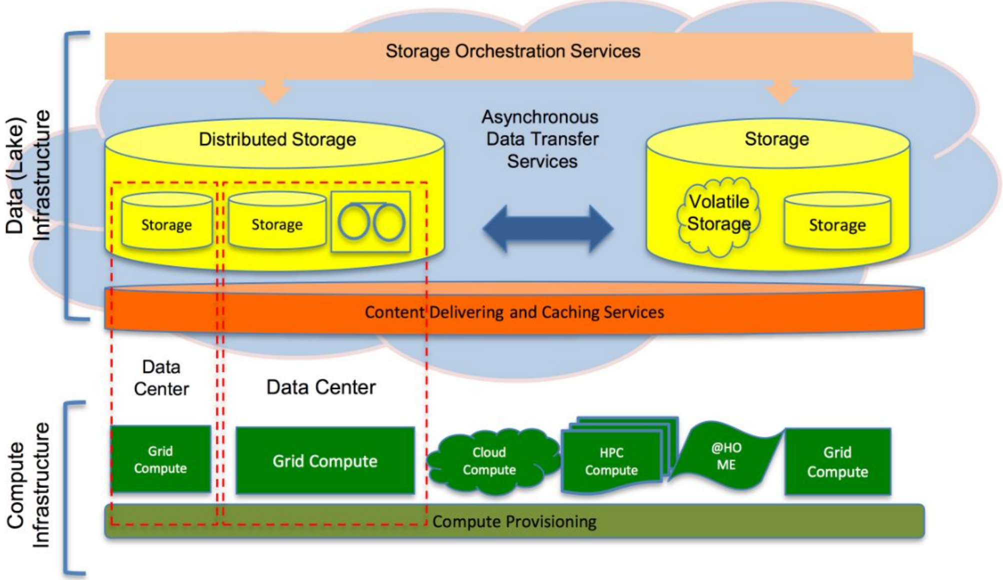

CPUs and storage not necessarily co-located: need to deliver the content over the WAN and/or cache it

Toward a sustainable, open and shared infrastructure

WLCG launched a set of R&D projects to prototype such a data management infrastructure – and associated tools

Aims:

- Reduce the global cost of storage (hw and operations)
- Enable a more effective use of existing storage
- Efficiently and scalably deliver data to large, remote, heterogenous, compute resources

Build a common set of DM tools based on open and standard protocols that can be used by a broad set of scientific experiments

Evolution of the Authentication/Authorisation towards federated identities and token-based systems in line with most modern network services

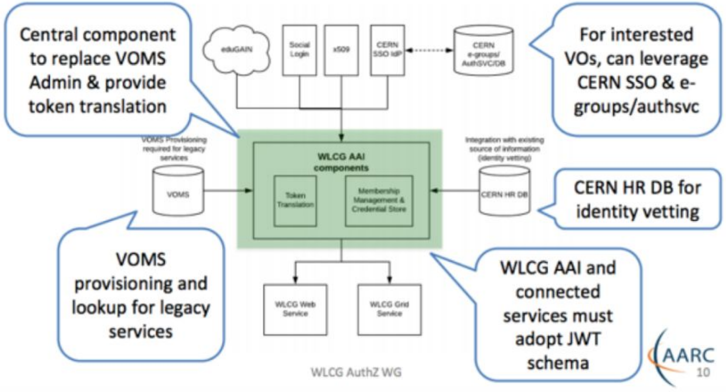

S. Campana - EIROFORUM 2020

Collaboration with data intensive sciences

PIC port d'informació científica


CERN COURIER

Aug 11, 2017

SKA and CERN co-operate on extreme computing

Big-data co-operation agreement

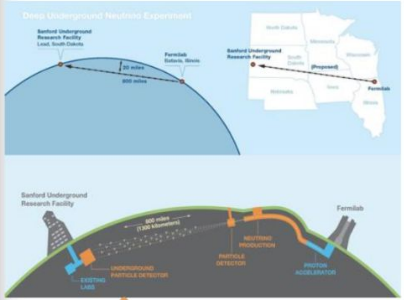

DUNE DEEP UNDERGROUND NEUTRINO EXPERIMENT

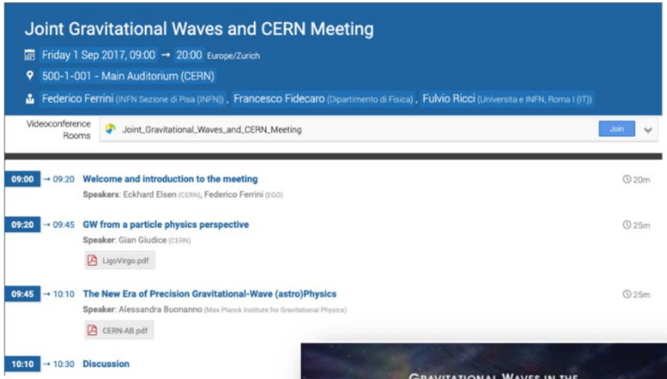
Joint Gravitational Waves and CERN Meeting

Friday 1 Sep 2017, 09:00 → 20:00 Europe/Zurich

500-1-001 - Main Auditorium (CERN)

Federico Ferrini (INFN Sezione di Pisa (INFN)), Francesco Fidecaro (Dipartimento di Fisica), Fulvio Ricci (Universita e INFN, Roma I (IT))

Videoconference Rooms — Joint_Gravitational_Waves_and_CERN_Meeting — Join

09:00 → 09:20	**Welcome and introduction to the meeting** Speakers: Eckhard Elsen (CERN), Federico Ferrini (EGO)	20m
09:20 → 09:45	**GW from a particle physics perspective** Speaker: Gian Giudice (CERN) LigoVirgo.pdf	25m
09:45 → 10:10	**The New Era of Precision Gravitational-Wave (astro)Physics** Speaker: Alessandra Buonanno (Max Planck Institute for Gravitational Physics) CERN-AB.pdf	25m
10:10 → 10:30	**Discussion**	

Belle II

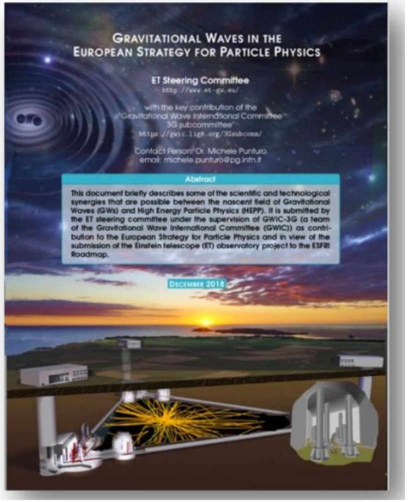

Astroparticle Physics European Consortium (APPEC)

APPEC

APPEC Contribution to the European Particle Physics Strategy

December 17, 2018

Editorial Board:
S. Katsanevas, A. Masiero, T. Montaruli, J. de Kleuver, A. Haungs

Contact Person:
T. Montaruli (APPEC Chair from Jan. 1, 2019)
Email: teresa.montaruli@unige.ch

Website : http://www.appec.org

Belle-2 and DUNE leverage the same infrastructure as WLCG

First formal non-LHC "associate" members of WLCG

S. Campana - EIROFORUM 2020

Collaboration with data intensive sciences

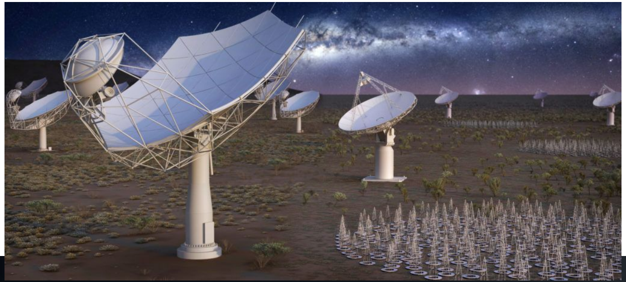

The Square Kilometre Array (SKA) project is an international effort to build the world’s largest radio telescope, with eventually over a square kilometre (one million square metres) of collecting area. The scale of the SKA represents a huge leap forward in both engineering and research & development towards building and delivering a unique instrument, with the detailed design and preparation now well under way. As one of the largest scientific endeavours in history, the SKA will bring together a wealth of the world’s finest scientists, engineers and policy makers to bring the project to fruition.

Software And Computing

Processing the vast quantities of data produced by the SKA will require two (one in South Africa, one in Australia) very high-performance central supercomputers capable of in excess of 100 petaflops (one hundred thousand million million floating point operations per second of raw processing power), which is equivalent to the fastest supercomputer on Earth at the present time (Source: Top500; November 2018).

In total, these two supercomputers will archive 600 Petabytes of data per year. To store this data on an average 500 GB laptop, you would need more than a million of them every year.

WLCG: a story of success, so far!

The **WLCG** organization is providing a shared distributed infrastructure for the LHC experiments **since 15 years** - collaboration with other sciences: similar use cases which will use mostly the same clusters

The computing models together with the infrastructure and services **successfully changed with time** adapting to the evolving landscape (experience and funding)

The HL-LHC will be an unprecedented challenge for us both in terms of scale and sustainability

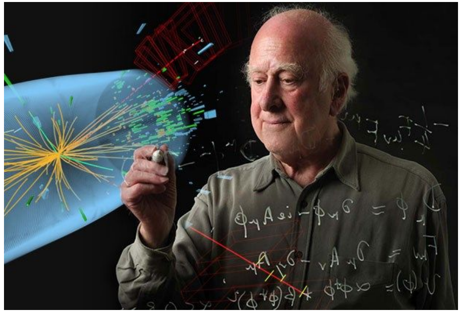

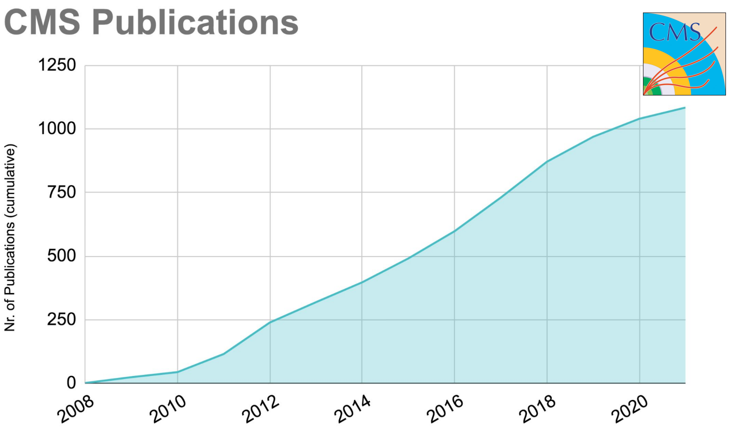

http://cms-results.web.cern.ch/cms-results/public-results/publications/CMS/index.html

Adoption of Grid Computing

PIC
port d'informació científica

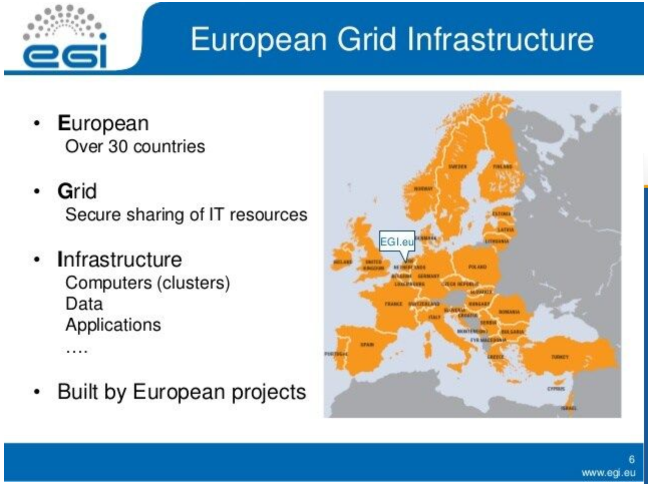

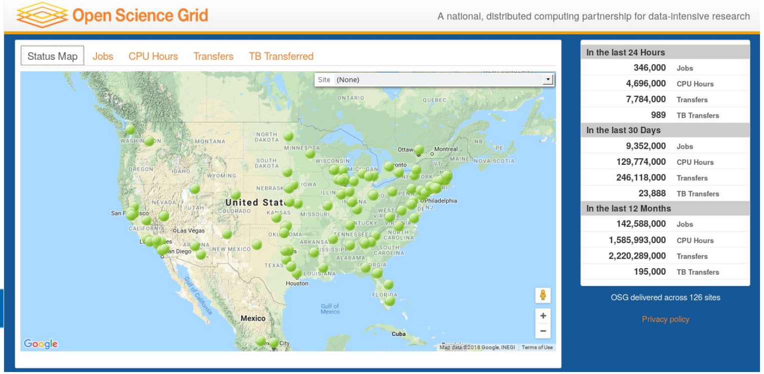

EGI Accounting

https://accounting.egi.eu/

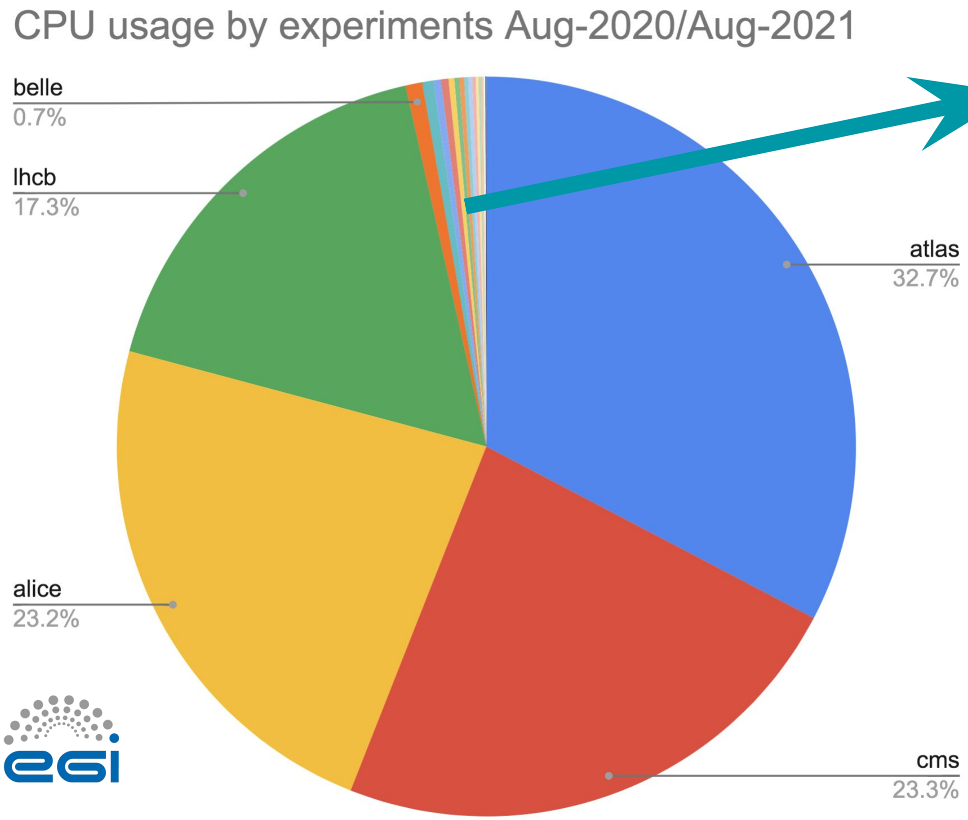

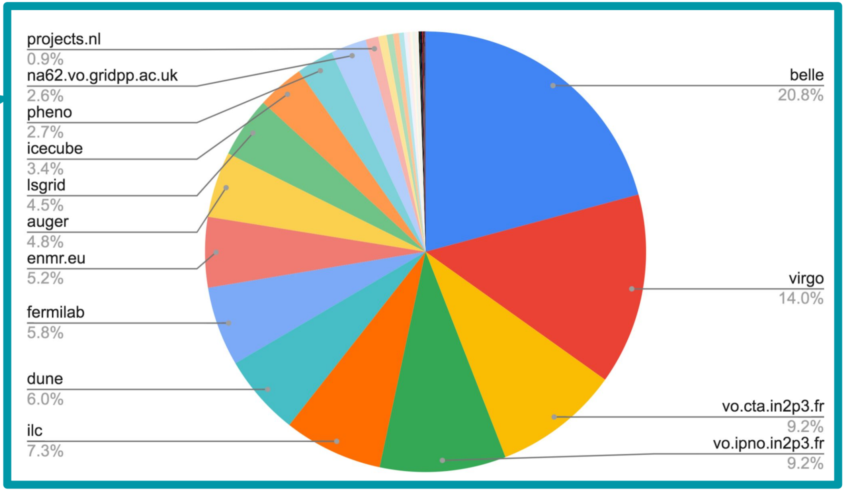

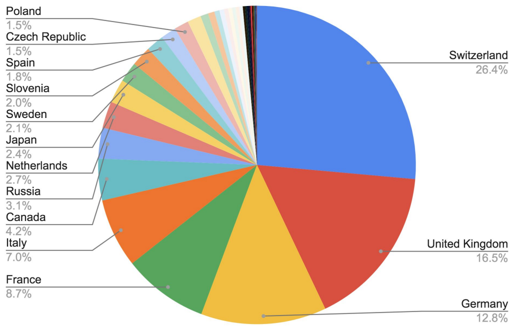

Use case: Folding@Home

Folding@Home (F@H) is a Citizen Science project (implemented as volunteer computing) dedicated to running simulations in order to predict protein folding into their 3D structure

The project **started in 2000**, and by 2005 they were already involved in multiple fields such as antibiotics, Alzheimer's disease and cancer

Folding@Home has been involved in multiple studies on **infectious** diseases (Ebola, Chagas), **neurodegenerative** diseases (Alzheimer's, Parkinson's), **diabetes**, and several forms of **cancer** (breast, kidneys)

Protein **misfolding** and aggregation often associated with **diseases**

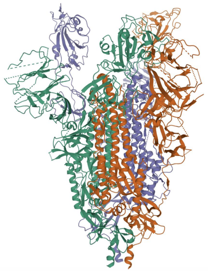

Use case: Folding@Home

PIC port d'informació científica

Folding@Home predicts protein structure by running simulations on **random paths** of structure evolution, starting from an experimental structure of the protein, passing through intermediate **local configuration energy minima (Markov state model)**

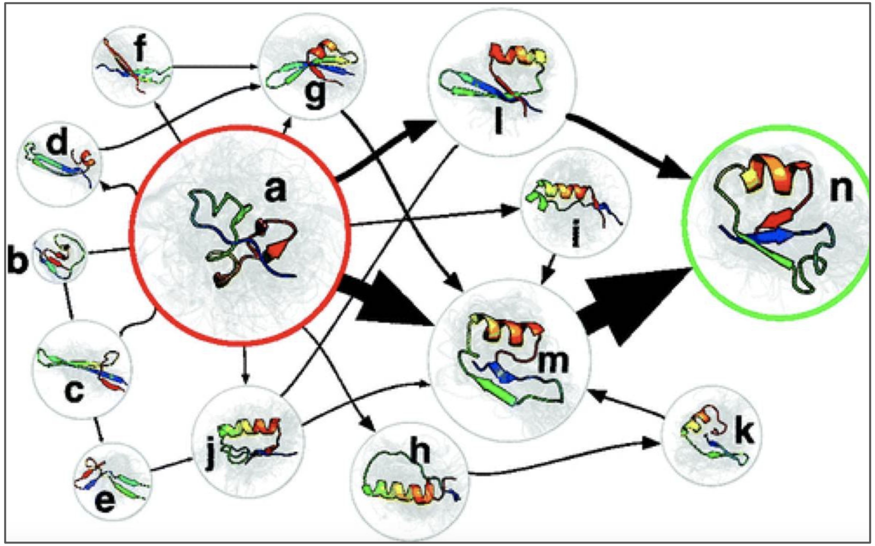

Folding@Home against Covid19

COVID19: coronavirus protein folding boosted to major project soon after the start of the current crisis

Folding@home Turns Its Massive Crowdsourced Computer Network Against COVID-19

By Oliver Peckham

March 16, 2020

For gamers, fighting against a global crisis is usually pure fantasy – but now, it's looking more like a reality. As supercomputers around the world spin up to combat the coronavirus pandemic, the crowdsourced distributed computing platform Folding@home is setting its sights on coronavirus research, spurring a global movement to commit powerful home computers and gaming consoles to the cause.

Folding@home @foldingathome · Mar 25

Thanks to our AMAZING community, we've crossed the exaFLOP barrier! That's over a 1,000,000,000,000,000,000 operations per second, making us ~10x faster than the IBM Summit!

Computational power from donors equivalent to 10 PFlops (2013), 100 PFlops (2016) and 1 EFlop by March 2020

CERN against Covid19

F@H tasks, known as Work Units (**WUs**), are downloaded from **workload servers** and processed by **F@H clients**, which usually run in the background of **donors** computers, using idle CPU cycles

Multiple donors may join into **teams** (e.g. CMS as part of the CERN & LHC Computing team) → WLCG contributed with 5% of its global resources

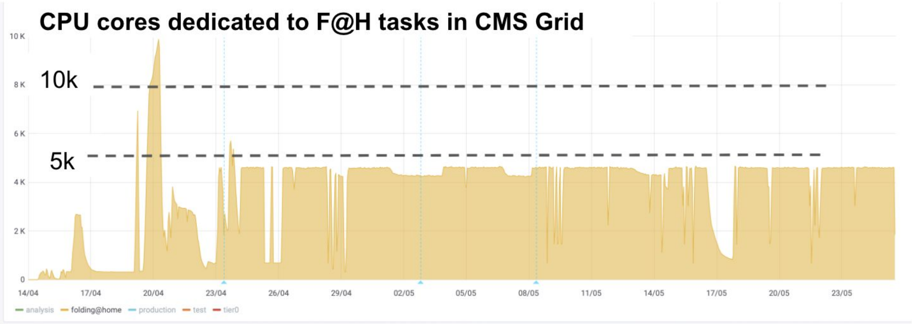

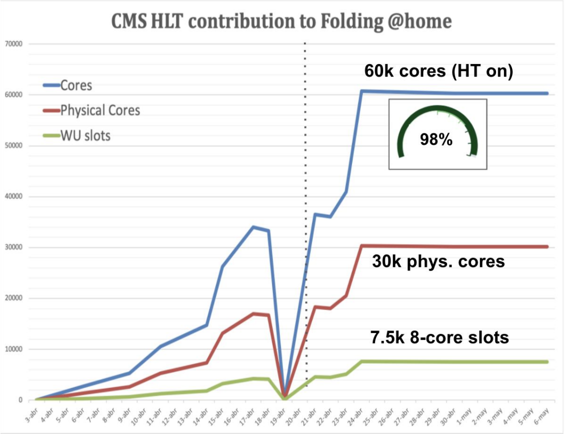

CERN against Covid19

CERN AGAINST COVID-19

PIC port d'informació científica

CERN & LHC Computing quickly rising to **top** contributors to F@H for this activity (April-Oct. 2021)

CERN & LHC Computing team accumulated F@H score

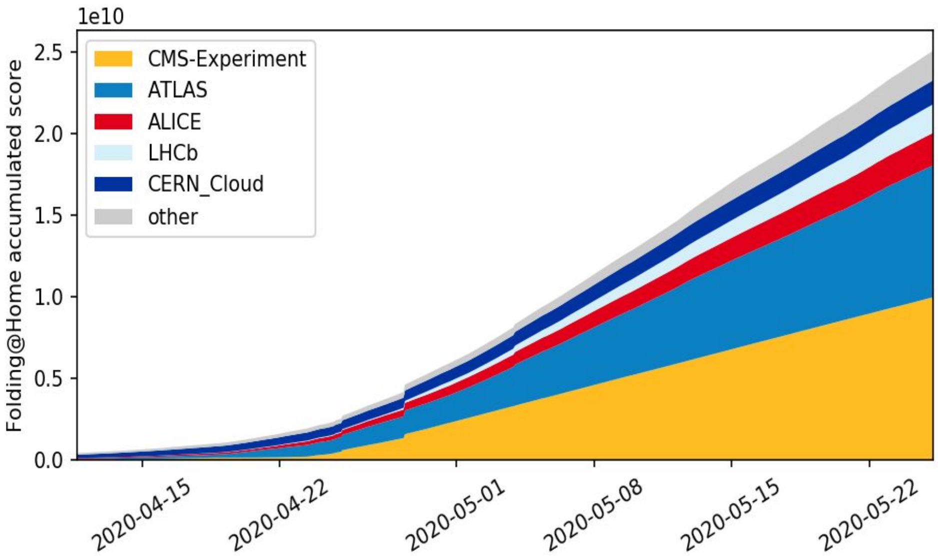

Team: CERN & LHC Computing

Date of last work unit	2020-06-01 21:27:19
Active CPUs within 50 days	1,359,701
Team Id	38188
Grand Score	31,137,092,103
Work Unit Count	8,232,997
Team Ranking	25 of 253746
Homepage	http://public.web.cern.ch/public/
Fast Teampage URL	https://apps.foldingathome.org/teamstats/team38188.html

Team members

Rank	Name	Credit	WUs
29	CMS-Experiment	12,280,780,906	2,402,473
46	ATLAS_CPU	9,848,417,162	2,313,970
293	LHCbHLT	2,281,250,909	339,286
336	ALICE-FLP	2,014,228,852	177,221
426	CERN_Cloud	1,642,060,669	728,133
589	DESY-ZN_GPU	1,265,355,845	9,082
2,827	UC_ATLAS-ML	263,471,027	154,058
3,405	CMSDCS	211,903,168	23,351

F@H against Covid19 highlights

PIC port d'informació científica

CAPTURING THE COVID-19 DEMOGORGON (AKA SPIKE) IN ACTION

April 3, 2020
by Greg Bowman

The spike of the SARS-CoV-2 virus (shown below) is one particularly appealing target for designing therapeutics to combat the COVID-19 disease. It is actually comprised of three identical proteins arranged in a circle. Many copies of the spike protrude from the surface of the virus, where they wait to encounter a protein on the surface of many human cells, called ACE2. Binding of a spike to ACE2 initiates a series of events that ultimately allow the virus to enter the human cell. Therefore, therapeutics that bind the spike in a manner that blocks its interaction with ACE2 could provide a valuable means to prevent infection.

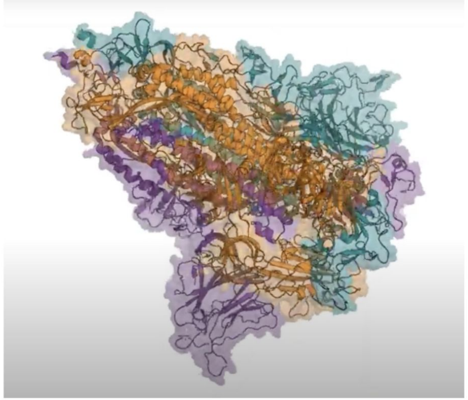

Prevent coronavirus **infection** to human cells

RNA Dependent RNA Polymerase (Nsp12)

A region important for interacting with the genome and Remdesivir

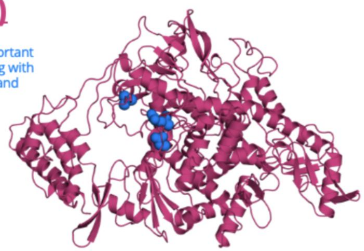

What does Nsp12 do?

Nsp12 copies the SARS-CoV-2 genome to make new viral particles.

How is F@h helpful here?

Nsp12's essential nature make it a great drug target. Drug binding could inhibit the production of new viral genomes and thus viral spread in the body. We are looking for pockets that open, due to protein motions, where a drug could bind. One such example, Remdesivir is in clinal trials. Mutations to Nsp12 are known to make Remdesivir ineffective. F@h can help understand how Nsp12 interacts with this drug and how mutations disrupt that interaction.

With what project IDs is Nsp12 associated?

CPU: p14412, p16424, p16432, p16500, p16501, p16402
GPU: p14436, p14437

Inhibit effective **viral RNA** replication in infected cells

3CL-Protease (3CL-PRO)

What does 3CL-Pro do?

3CL-Pro cuts the immature SARS-CoV2 protein from one piece into many mature viral proteins. It is more active after it assembles into a two-component unit (dimer).

How is F@h helpful here?

We are looking for pockets in the single unit and the dimer (blue and orange) that open, due to protein motions, where a drug could bind. We are also calculating how strongly certain molecules bind to the active site. This would prevent the virus from efficiently replicating. 3CL-Pro is one of the major drug targets across many ongoing studies.

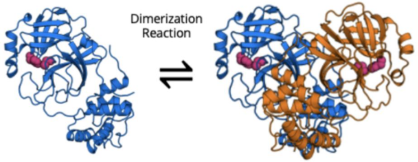

With what project IDs is 3CL-Pro associated?

Compound binding simulations p14350-14399 and p14600-14699
Pocket opening simulations p14582, p14584, p14542, 14592 and p14543

6Y2E

Zhang et al., *Science* 2020
DOI: 10.1126/science.abb3405

Inhibit effective **viral protein** replication in infected cells

A. Pérez-Calero Yzquierdo, IFAE Seminar

Use case: CosmoHub

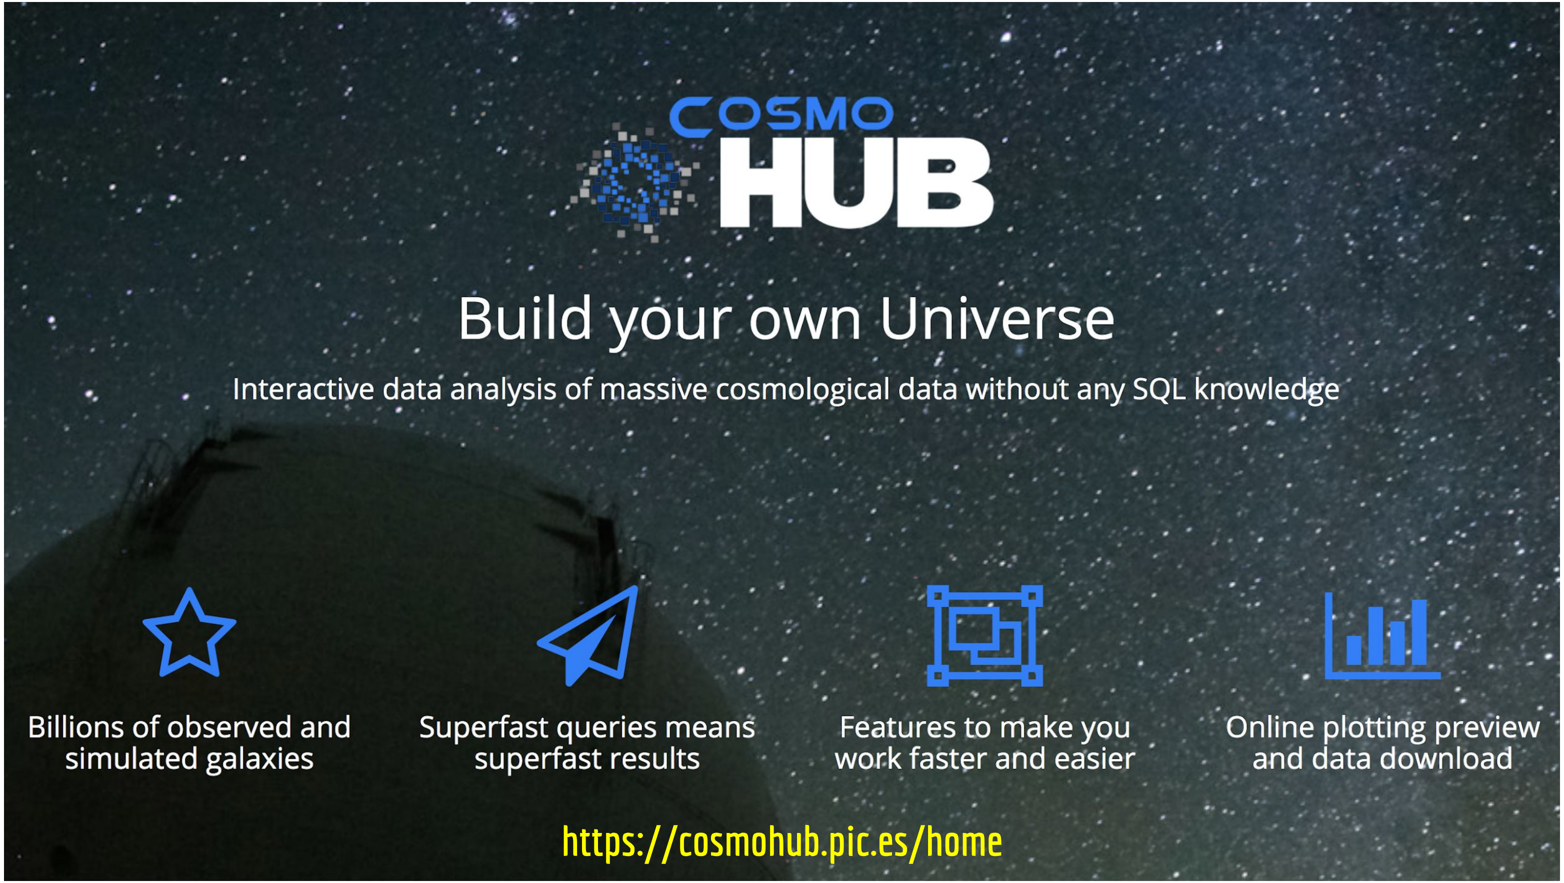

CosmoHub on a Hadoop Cluster

Hadoop + **Hive** + **Spark** cluster + **Jupyter** notebooks @ PIC

Fault tolerant open source Big Data Platform that offers scalable **distributed storage and processing** on commodity computer nodes (no SQL knowledge required)

Multiple cosmological datasets available: **fast interactive exploration** (visualization - 85% < 30s) and **data products distribution** (HTTP/WebDav, FITS formats, 75% in < 3 min)

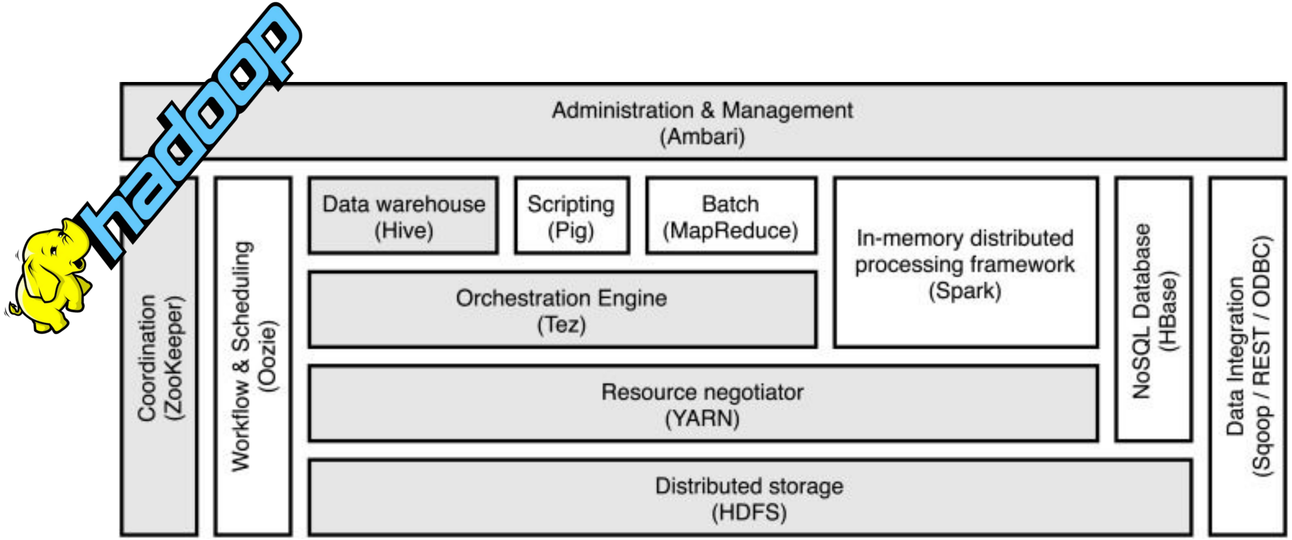

Custom DIY nodes

- 12 nodes AMD Threadripper 1920X
- 128 GB RAM, 12 x 3 TB SATA HDD hot-swap
- 2x1 TB NVMe SSD i 2x10-GBASE-T LAN

Hortonworks HDP 3.1.4

- Hadoop 3.1.0
- Hive 3.1.0
- Spark 2.3.2

J. Carretero - IVOA 2020

CosmoHub queries interface

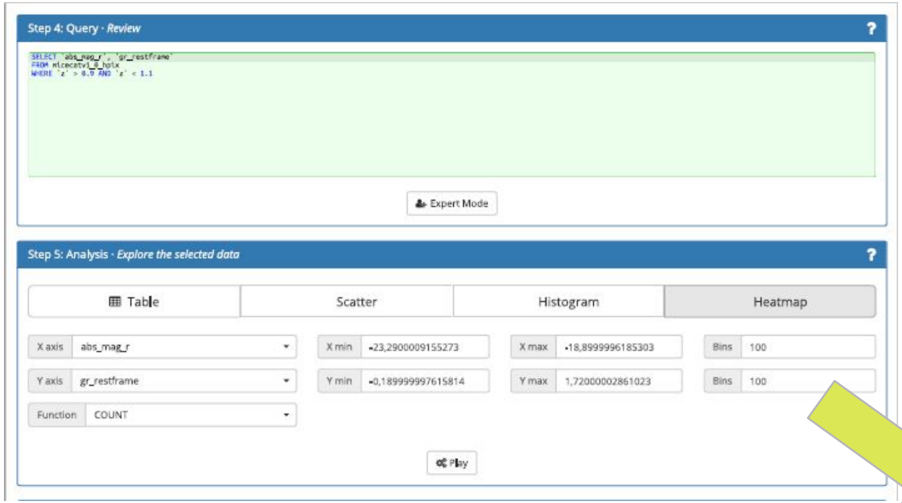

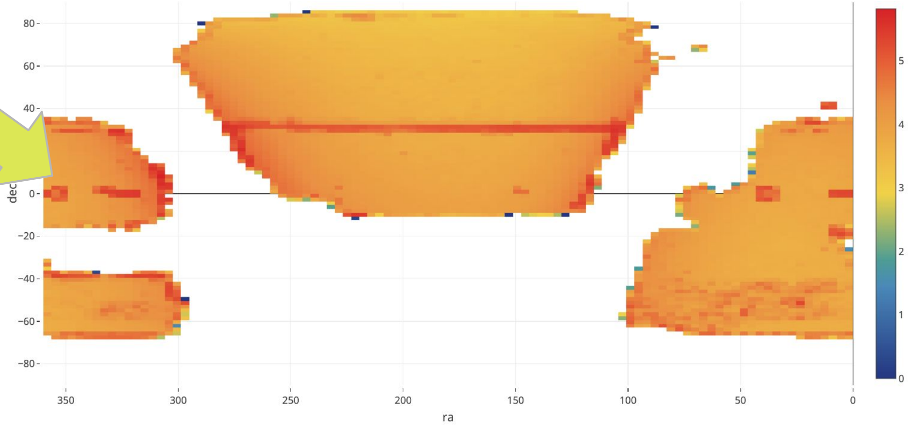

DESI Legacy Survey z_phot < 0.1

J. Carretero - IVOA 2020

CosmoHub catalogues

PIC
port d'informació
científica

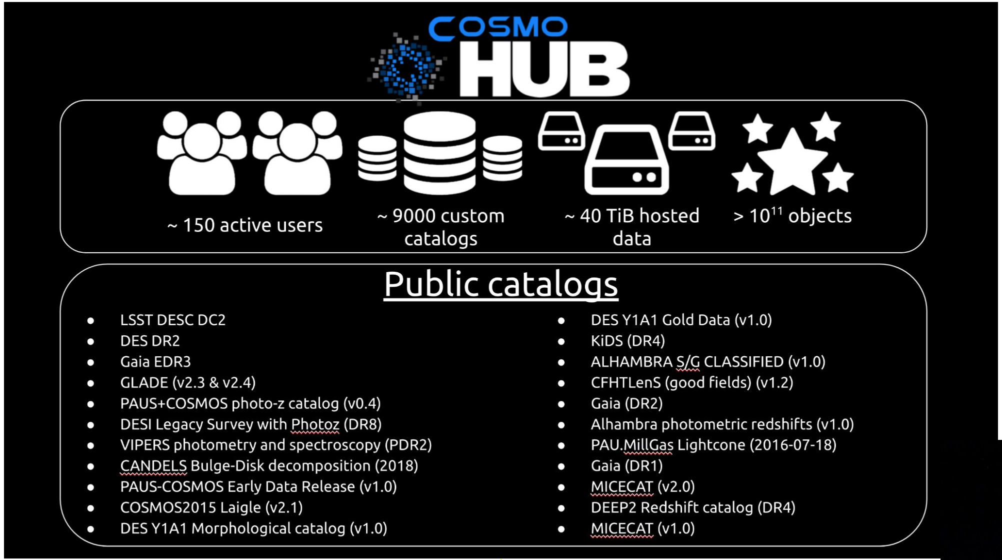

J. Carretero - IVOA 2020

Early Data Release 3 (EDR3) of the Gaia mission. The early access is possible because PIC recently became one of the affiliated sites of Gaia

1.8 billion objects (100x100 pixels < 1 minute!)

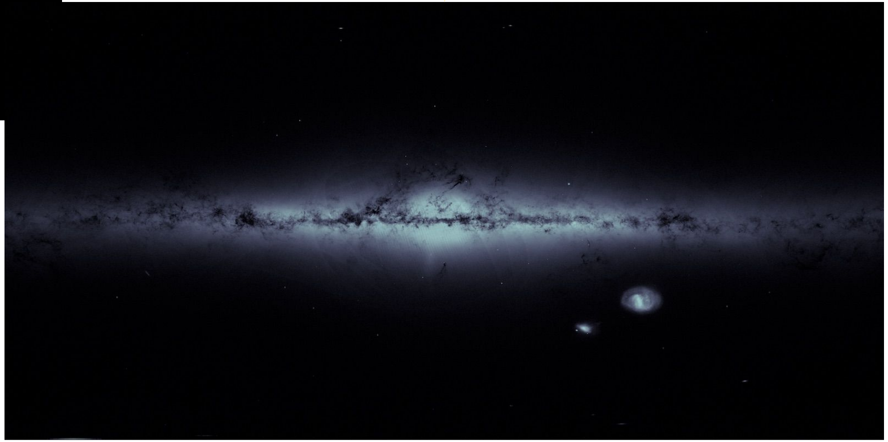

Recently CosmoHub has been accepted by re3data.org as a Research Data Repository

Use case: cosmological N-body simulations

PIC port d'informació científica

Cosmological simulations of galaxy formation are instrumental in advancing our understanding of structure and galaxy formation in the Universe

Nonlinear galaxies evolution / variety of physical processes / enormous range of time/length scales

A better understanding of the relevant physical processes, improved numerical methods and increased computing power have led to simulations that can reproduce a large number of the observed galaxy properties

Modern simulations model **dark matter**, **dark energy** and **ordinary matter** in an expanding space-time starting from well-defined initial conditions

Modelling ordinary matter is the most challenging due to the large array of physical processes

Cosmological simulations have also proven useful to study **alternative cosmological models** and their impact on the galaxy population

Use case: cosmological N-body simulations

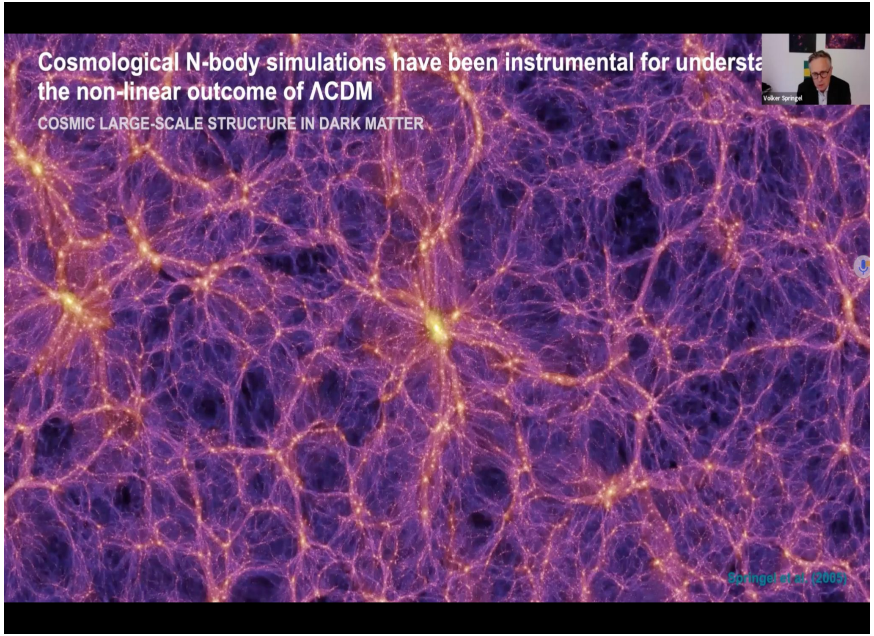

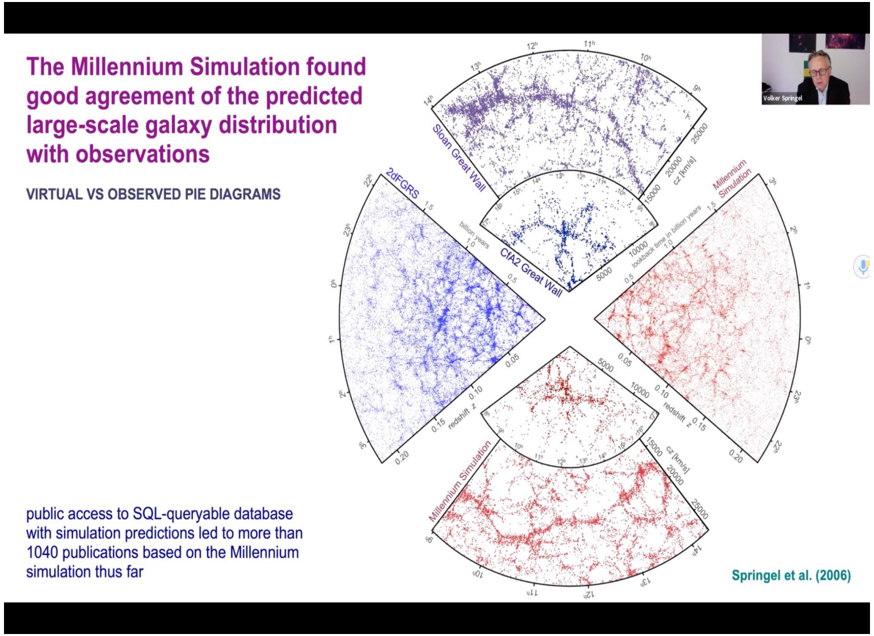

Volker Springel - "Hydrodynamical Simulations of Galaxy Formation"

→ **https://www.youtube.com/watch?v=LlYpsmdBxLM**

Use case: cosmological N-body simulations

PIC
port d'informació
científica

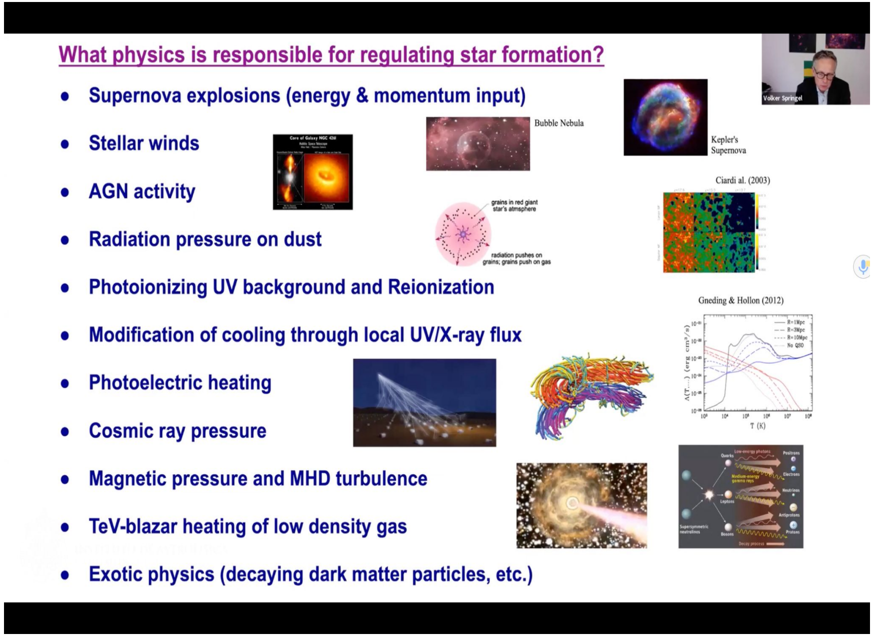

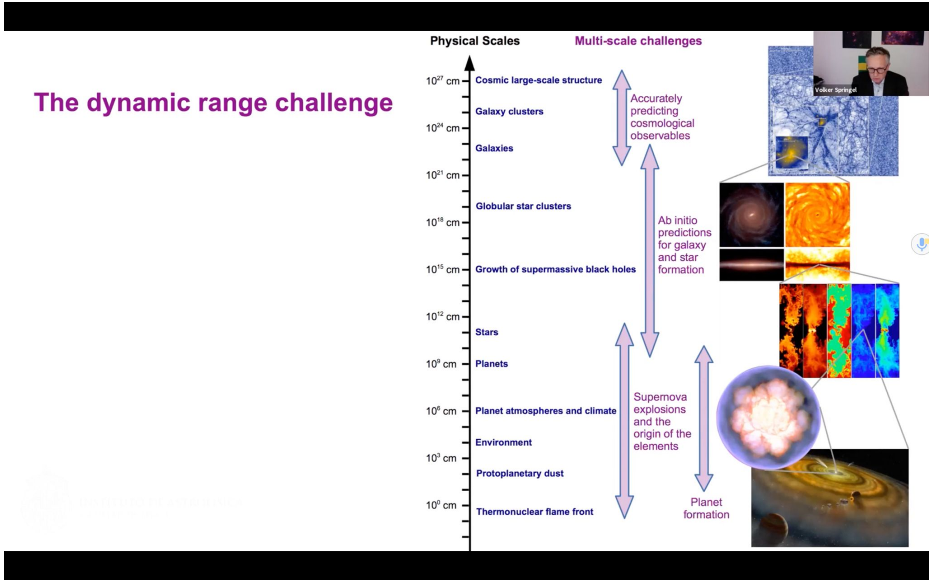

Volker Springel - "Hydrodynamical Simulations of Galaxy Formation"

→ https://www.youtube.com/watch?v=LlYpsmdBxLM

Use case: cosmological N-body simulations

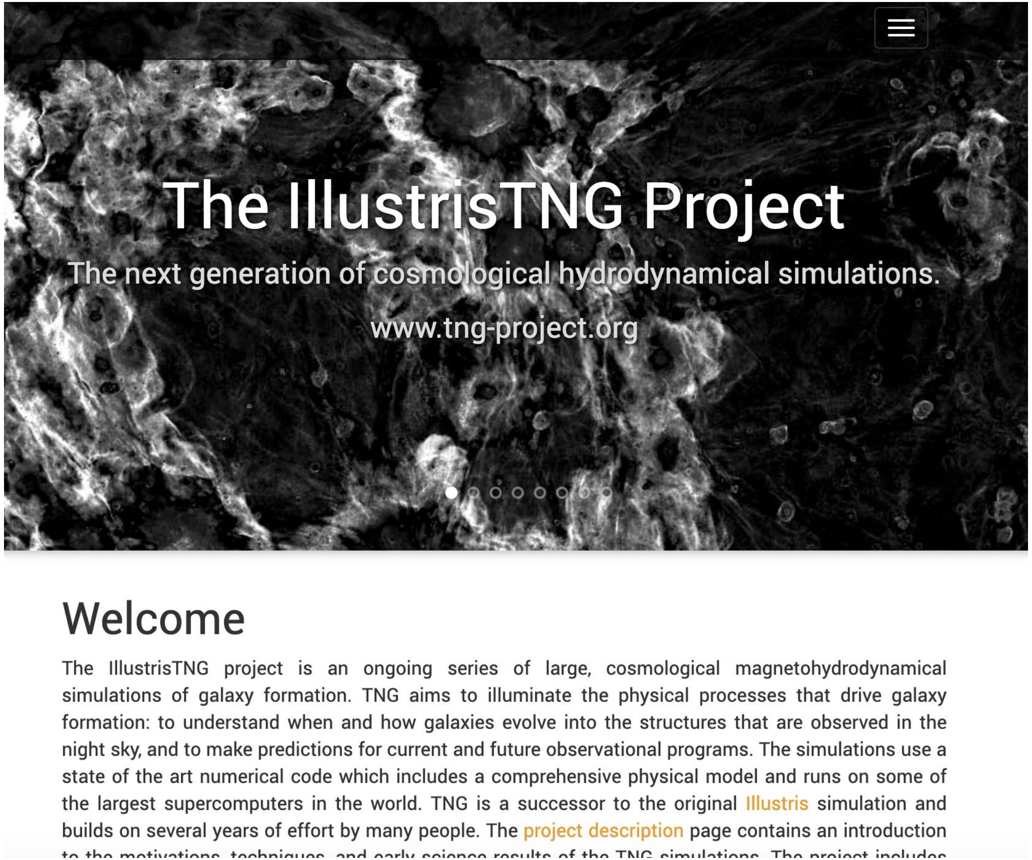

→ **https://www.tng-project.org/media/**

Cosmological simulations: collab. effort!

PIC
port d'informació
científica

Arepo code

moving mesh hydrodynamics

Getting started

Home > Getting started

General

We believe the best way (and the most fun) to familiarise with a new simulation code is to run simulations with it. To make this as easy as possible, Arepo comes with a number of examples that should run with the settings provided.

The only thing you need to do is to make sure the necessary libraries (**mpi**, **gsl**, possibly **fftw**, **hdf5** and **hwloc**) are installed on your system.

Getting the code

Arepo is available in a git repository hosted on MPCDF gitlab. The easiest way to get it is by cloning it from the repository

```
git clone https://gitlab.mpcdf.mpg.de/vrs/arepo.git
```

Search …

News

Code repository online

September 11, 2019

Code release paper submitted

September 10, 2019

Hello world!

September 7, 2019

Other sciences at HPC centers

PIC port d'informació científica

Barcelona Supercomputing Center Centro Nacional de Supercomputación

DISCOVER BSC | RESEARCH & DEVELOPMENT | MARENOSTRUM | TECH TRANSFER | JOIN US | NEWS

Home / Research & Development / **Research Areas**

Example

- < Home
- < Research & Development
- Research Areas
  - Atmospheric Composition
  - Big Data
  - Bioinformatics
  - Biomechanics
  - Climate Prediction
  - Cloud Computing
  - Cognitive Computing
  - Computer Architecture and Codesign
  - Distributed Computing
  - Education
  - Engineering Simulations
  - Fusion
  - Genomics
  - Geophysics

Research Areas

Search for a Research Area ... | Search

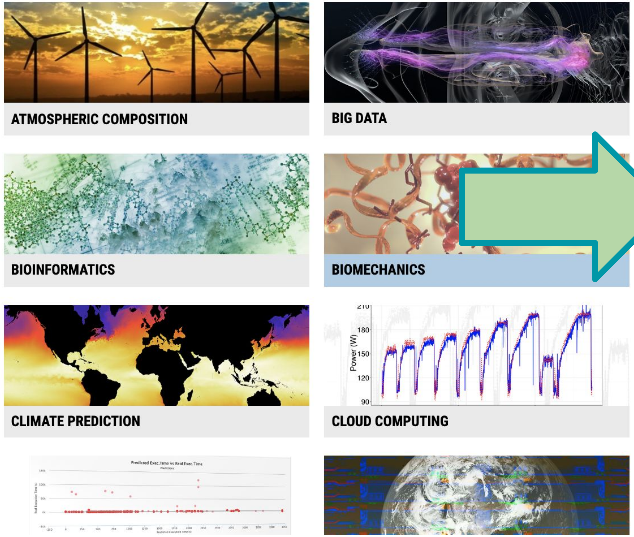

FRIDAY, 3 SEPTEMBER

11:20 **HIGH PERFORMANCE COMPUTING: THE CERN, SKA0, GEANT and PRACE COLLABORATION** 40m

Dr. Maria GIRONE is the Chief Technology Officer (CTO) of the CERN openlab. She leads the development of High Performance Computing (HPC) technologies for particle physics experiments.
CERN openlab was established in 2001 and supports academics at CERN in their collaborations with independent companies.
Dr. Girone has worked on the upgrade of the Large Hadron Collider (the High Luminosity Large Hadron Collider), which will require up to one hundred times more computing capacity than it did originally.This increase in capacity will come through access to commercial cloud computing platforms, data analytics, deep learning and new computing architectures.
(Credit WIKIPEDIA)

Speaker: Dr Maria GIRONE (CERN)

CERN-SKA0-PRACE-...

Next generation exascale HPC centers

EuroHPC Joint Undertaking launches first research and innovation calls

€190 million will fund research and innovation in supercomputing in Europe.

European Commission

The European High Performance Computing Joint Undertaking (EuroHPC) has launched calls for proposals to fund research and innovation activities that will help Europe to remain globally competitive in the field of supercomputing.

EuroHPC began its operations in November 2018, with 29 European countries currently taking part. Its goal is to pool EU and national resources in order to develop an integrated world-class supercomputing and data infrastructure in Europe, and to create a highly competitive and innovative European HPC ecosystem. In June this year, it announced a joint investment, with the participating countries, of around €840 million to acquire and deploy eight world-class supercomputers in the EU before the end of 2020. These machines will multiply Europe's current supercomputing capabilities by a factor of 5-10.

Mariya Gabriel, European Commissioner for Digital Economy and Society, said:

> *These calls complement the substantial investment that is being made by the Joint Undertaking in Europe's supercomputing infrastructure. They will help the Joint Undertaking draw on the skills and knowledge of European SMEs and industry to put its ambitious work plan into action and develop applications and services*

Related topics

HPC

Investing in network and technologies

< Home

News

BSC News

Events

BSC in the Media

Newsletter

BSC has been selected as coordinator of three projects in the first EuroHPC call to promote research and development of supercomputers with European technology

03 March 2021

The center participates in six other projects, out of the 20 selected.

> **The three projects coordinated by BSC have a total budget of 19.6 million euros**
>
> **EuroHPC JU aims to pool resources to develop exascale supercomputers based on European technology to process large amounts of data.**

The Barcelona Supercomputing Center (BSC) has been selected as coordinator of three research projects of the call launched by the EuroHPC Joint Undertaking to promote the research and development of the new supercomputers of the future. In its first call, EuroHPC has selected 20 projects, of which the BSC coordinates three and participates in another six. The projects coordinated by BSC have a global budget of 19.6 million euros, 50% funded by EuroHPC JU and the other 50% by the partner's national governments.

Exascale for large science projects

PIC
port d'informació
científica

ESCAPE
European Science Cluster of Astronomy & Particle physics
ESFRI research infrastructures

Horizon 2020
funded project

EUROPEAN OPEN SCIENCE CLOUD

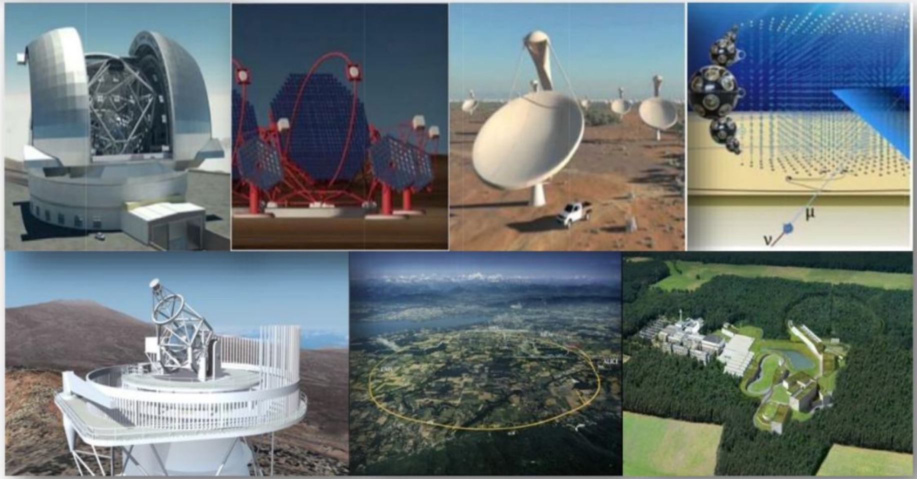

Goals:

Prototype an infrastructure adapted to the Exabyte-scale needs of the large science projects.

Driven by the sciences

Address *FAIR* data management

Science Projects

HL-LHC	SKA
FAIR	CTA
KM3Net	JIVE-ERIC
ELT	EST
EURO-VO	EGO-VIRGO
(LSST)	(CERN,ESO)

Data centres: CERN, INFN, DESY, GSI, Nikhef, SURFSara, RUG, CCIN2P3, PIC, LAPP, INAF

S. Campana - EIROFORUM 2020

Modern analysis / facilities

We are looking for new, scalable tools that have become available to **boost interactive data analysis** and enable scientists **exploiting the maximum scientific potential of the data** - WLCG deploying such facilities!

python jupyter HTCondor High Throughput Computing XRootD

Example

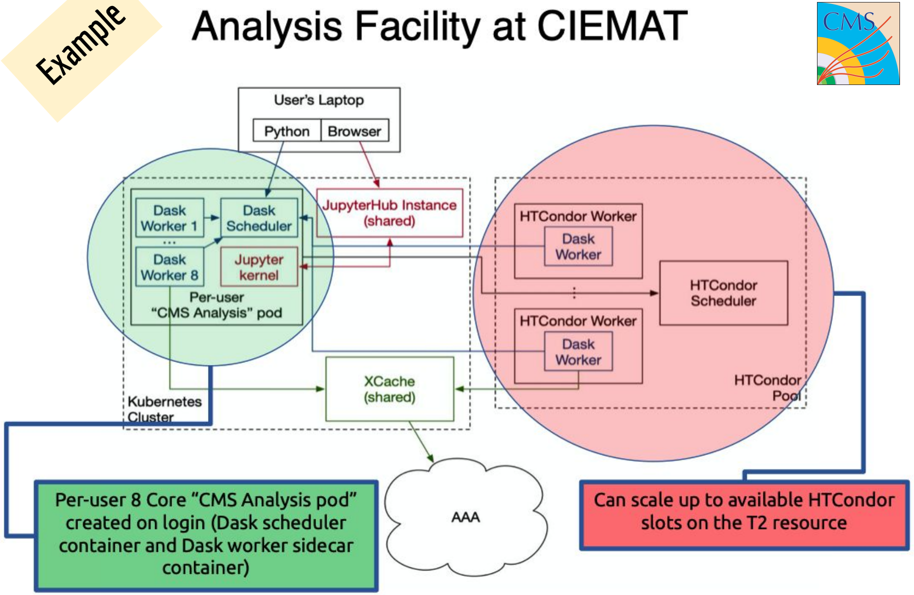

Outlook

PIC port d'informació científica

The methods, skills, technologies and practices involved in **handling data** - and particularly big data - **is crucial** to understanding empirical knowledge-making and science overall

Scientists need to specialize (more) in software engineering and modern computing techniques to boost their scientific efficiency - **learn!**

Compute expert teams and facilities available to boost science

Scientific computing is very dynamic - constantly changing and adopting new techniques - We will soon use **Quantum Computing** on a daily basis?

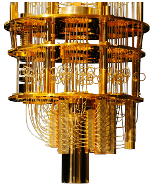

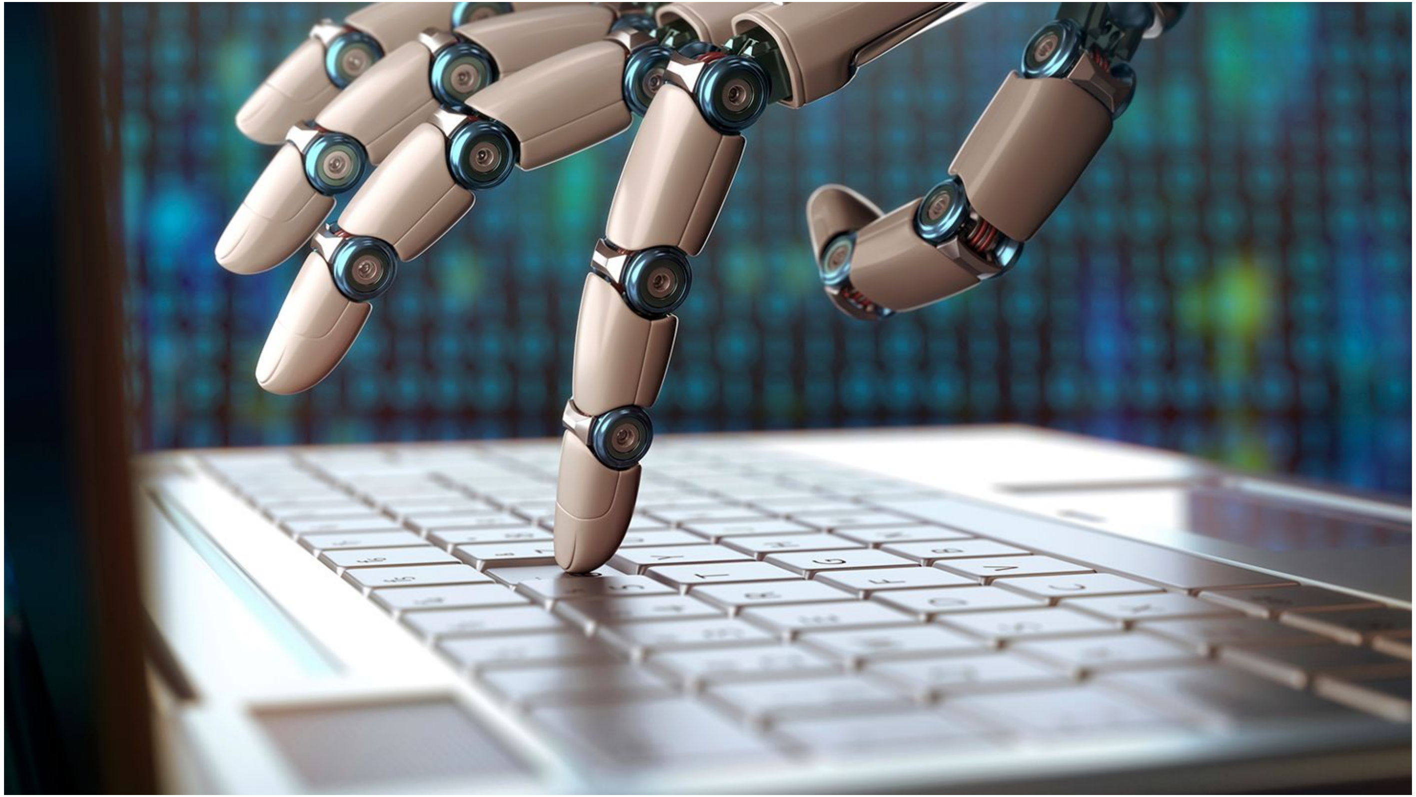

Thanks!
Questions?