GROUPE ESEO
4, rue Merlet de la Boulaye BP 30926
49009 Angers cedex 01 - France
Tél. : +33 (0)2.41.86.67.67
http://www.eseo.fr

European Organization
for Nuclear Research
CERN CH-1211 Genève 23
Switzerland

RAPPORT DE STAGE

(*)TYPE DE STAGE : Stage de fin d'etudes (I3)

(*) AUTEUR: Axel Voitier

(*) DATES : 01/01/2009 au 28/02/2010

NIVEAU DE CONFIDENTIALITE

Aucune	Niveau I	Niveau II	Niveau III
X			

TITRE DU STAGE

Test suite for the archiver of a SCADA system

Responsable de stage : M. Manuel Gonzales Berges
……………………………………………

CERN-THESIS-2009-205
17/12/2009

Summary sheet

Test suite for the archiver of a SCADA system

Topic:

The group responsible for providing the main control system applications for all machines at CERN has to validate that every piece of the control systems used will be reliable and fully functional when the LHC and its experiments will do collisions of particles.

CERN use PVSS from ETM/Siemens for the SCADA part of its control systems. This software has a component dedicated to archive into a centralised Oracle database values and commands of tenth of thousands hardware devices.

This component, named RDB, has to be tested and validated in terms of functionality and performance. The need is high for that because archiving is a critical part of the control systems. In case of an incident on one of the machine, it will be unacceptable to not benefit of archiving the machine context at this moment just because of a bug in RDB.

Bugs have to be spotted and reported to ETM.

Results:

The proposed solution is an extensible automatic tester able to evaluate currently around 160 cases of potential bugs.

Since the beginning of this project last January, the RDB tester has been able to discover more than 10 problems in RDB.

Most of them have since been fixed by ETM. Nowadays CERN is confident that the machine exploitation will not encounter any problem related to RDB.

Still, RDB needs more effort and attention to ensure that future usage of PVSS will not hit other problems RDB may cause. The proposed solution has taken in account that CERN will need to extend the tests to newer cases.

Acknowledgment

To those who make this place so amazingly enjoyable

Manuel Gonzales Berges, my section leader and supervisor.
Piotr Golonka, my secondary supervisor.
Roberto Saban, my group leader.

All the EN-ICE group and especially:
Frederic Bernard, Paul Burkimsher, Daniel Davids, Oliver Holme, Herve Milcent (Alumni ESEO), Fernando Varela Rodriguez, Jonas Arroyo Garcia, Alexander Egorov, Alvar Cuevas, Pawel Macuda, Clair Minna Temper, Enrique Blanco-Vinuela, Jeronimo Ortola Vidal, Jacques Rochez, Guilio Morpurgo, Brice Copy, Mathias Dutour, Geraldine Thomas, Artem Burmyakov, Noora Tuovinen.

All the other students I met at CERN and especially:
Ada Alexandre Metola, Alexander Stadler, Alexandros Attikis, Aurelien Marsili, Benjamin Frish, Florent Bouvier, Juha Kemppinen, Maja Olvegard, Marika Alm, Marta Ruiz, Michael Karsch, Michael Monteil, Michele Quinto, Montse Montero Pineiro, Petri Lehtonen, Samir Arfaoui, Seda , ersembe, Serena Oggero, Tiia Uusimaki, Tiina SalmivValerio Rossetti, Yisel Martinez Palenzuela.

The Technical Student Program team at CERN and the internship department of ESEO.

The LHC, SPS, PS and all the others for giving us work.

And the Higgs boson for its strength to gather so much people in such confined space.

Table of content

I. European Organization for Nuclear Research

i. Creation, composition and vocation

CERN, "Conseil européen pour la recherche nucléaire", was founded in 1954 as one of Europe's first joint organisations, bringing specialists from 12 Member States together to pursue a common dream. The motivation behind the creation of CERN was to keep European physicists in Europe after the World War II. Established on the Franco-Swiss border near Geneva, it has become a shining example of successful international scientific collaboration. Today, CERN has 20 Member States from Europe, and additional nations from around the globe also contribute to and participate in its research programme. CERN is now a European laboratory for the world.
CERN is an intergovernmental organisation exactly as UNESCO, ESA, EU, UN, OECD, IMF or many other important organisations in the world are. Such organisations exist by a treaty ratified by their member states. They are legally out of any state, but are highly dependent of their member-states wills.

CERN's business is fundamental physics, finding out what the Universe is made of and how it works. The instruments used at CERN are particle accelerators and detectors. Accelerators boost beams of particles to high energies before they are made to collide with each other or with stationary targets. Detectors observe and record the results of these collisions. By studying what happens when particles collide, physicists learn about the laws underlying the evolution of the Universe.

CERN mandate has been defined as following:
The Organization shall provide for collaboration among European States in nuclear research of a pure scientific and fundamental character, and in research essentially related thereto. The Organization shall have no concern with work for military requirements and the results of its experimental and theoretical work shall be published or otherwise made generally available.
Article II, §1 of CERN's convention.

The aim of CERN is to provide for physicists from all over the globe the needed tools to realise experiments in subatomic fields. CERN pays for all infrastructures and the engineering behind them (from civil engineering to radio-frequency and magnet development, from building an accelerator to building a restaurant or a hostel). It also provides for human needs like safety, security and other living-related conditions (catering, housing, heating, cleaning). CERN host experiments, but experiments' collaborations have to fund there own detector hardware and grant their scientists and students.

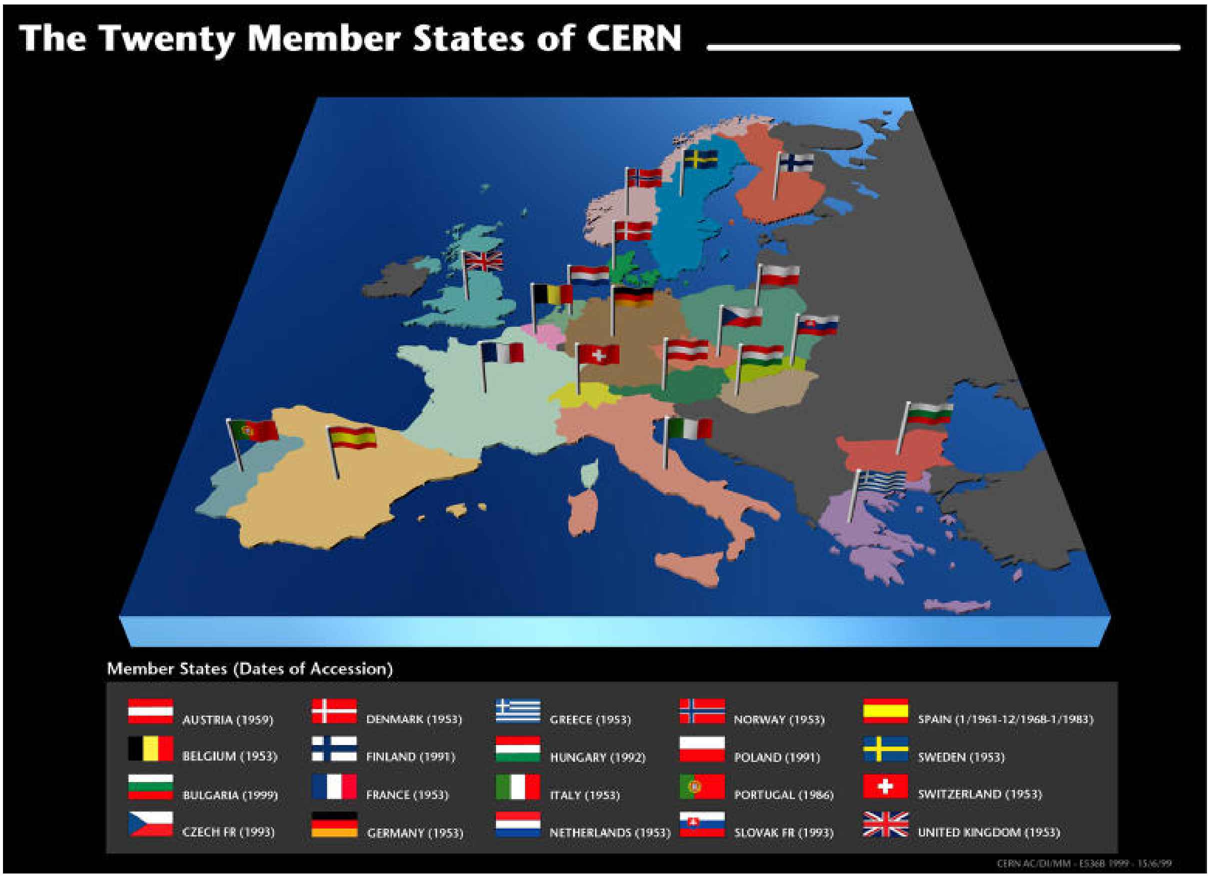

Member states have two representatives at the CERN's council, one politics (usually from their Department of Research) and one scientists.

Observer states: *India (2003), Israel (1991), Japan (1995), Russian Federation (1991), Turkey (1961), U.S.A. (1998)*.
Observer organisations: *European Commission (1985), UNESCO (1955)*.
Observers can assist at the council's meetings and receive documents. But they cannot be involved in the discussions.

Non-member states: *Algeria, Argentina, Armenia, Australia, Azerbaijan, Belarus, Brazil, Canada, Chile, China, Colombia, Croatia, Cuba, Cyprus, Estonia, Georgia, Iceland, Iran, Ireland, Lithuania, Macedonia, Mexico, Montenegro, Morocco, New Zealand, Pakistan, Peru, Romania, Serbia, Slovenia, South Africa, South Korea, Taiwan, Thailand, Ukraine, Vietnam*.
Citizen from a non-member state can work in a CERN project.


The Parking of Flags at CERN main entrance

ii. The accelerator complex and general infrastructures

CERN's flagship is the Large Hadron Collider (LHC), the particle accelerator that will provide the highest particle energies ever achieved in a laboratory. The energy density generated in its 14 TeV proton–proton collisions will be similar to that existing a few instants after the Big Bang. Recreating such conditions is a tremendous way to look back to the birth of the Universe; it provides the only way to do experiments to find answers to very fundamental questions concerning, for example, the origin of mass, the nature of dark matter, and the balance of matter and antimatter in the Universe.

The LHC will also accelerate beams of lead ions to energies never reached before. In the collisions, a new state of matter will be created in which quarks and gluons — otherwise confined in more complex particles, the hadrons — will be set free, giving birth to the 'quark–gluon plasma'.

A physicists once explained me what is done in a particle collider is like a vegetable soup. When one put together a bunch of vegetables in water and heat it up, it becomes a liquid where everything is melted. The same thing is done in a collider. By putting together a bunch of particles, accelerating them and collide with another bunch of particle going in the opposite direction so that their energies add up. The very high energy reached, it forces the original particles to change state, "melt" and let appear particles of higher energy.

It is popularly said that these subatomic particles are what compose matter. But a collider is not actually breaking "atom's shell". Particles appearing in a particle collider only exist here, and also during collisions with cosmic rays in our atmosphere. Otherwise, these particles are not "inside" the matter.

The comparison with a vegetable soup stop here. Because when a soup is cooling down, it does not recreate the original vegetables. In the case of elementary particles, which are of a very high energy, when they are "cooling down", that is when they decay, they turn into a state of lesser energy: a bound state. The composite particle hence created leads progressively, decay by decay, recombination by recombination, to the matter as we know it.

One eV (electron-volt) is the energy of one electron accelerated by a voltage of 1V. A cathodic television has an energy of about 10 keV. Thus, 1 TeV is roughly equivalent to

100 millions television. But seen with another comparison, 1 TeV is also the energy of a flying mosquito. With a collision energy of 14 TeV, the LHC is like colliding a flock of 14 mosquitoes... except that the energy is concentrated here in a space a thousand billion smaller than a mosquito!
Also, to stop a beam generated by the LHC, one need a dumping material capable of resisting the same energy as mostly 80 kilos of TNT would produce.

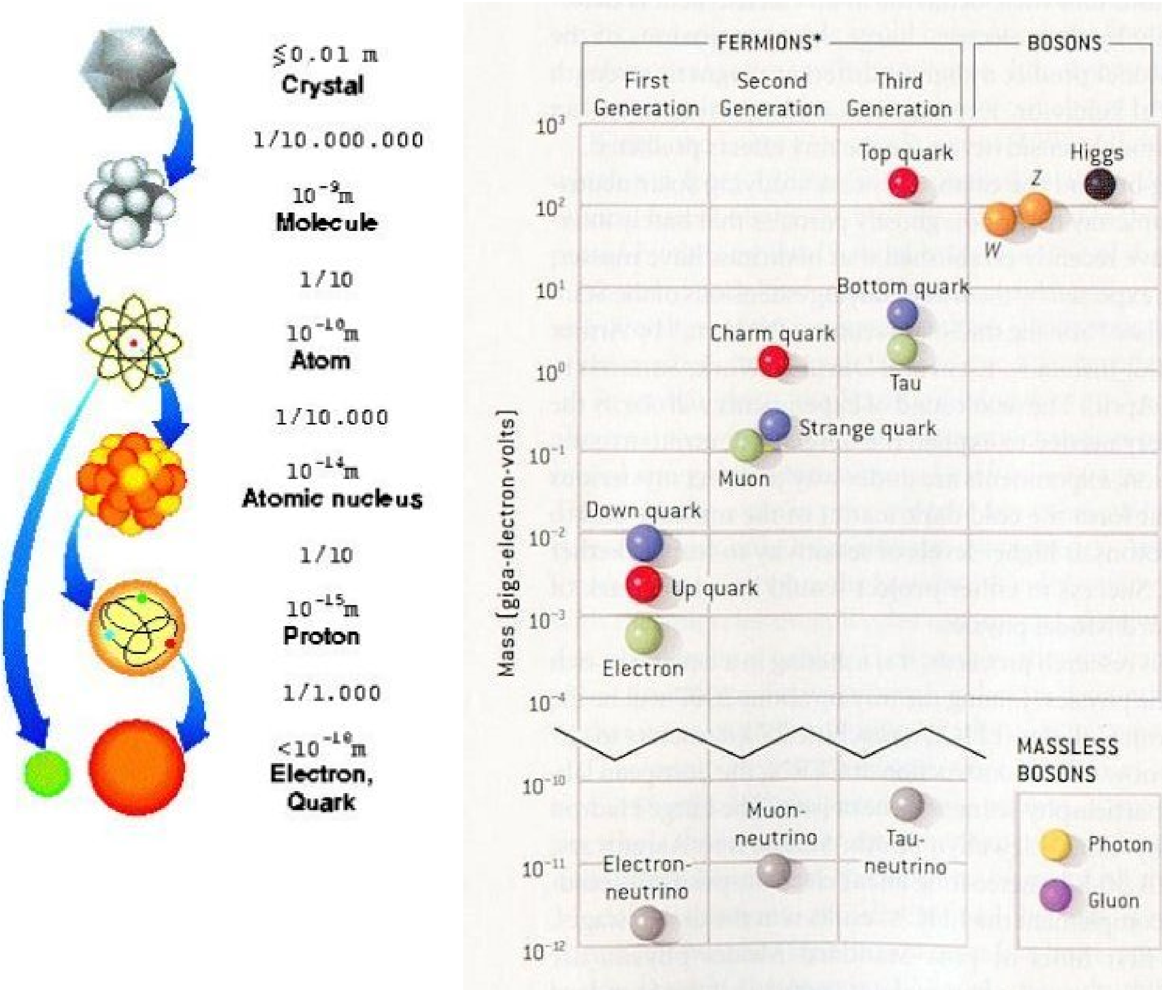

Metric and mass/energy scale of matter and particles, from
http://www.zamandayolculuk.com/cetinbal/elementaryparticles.htm

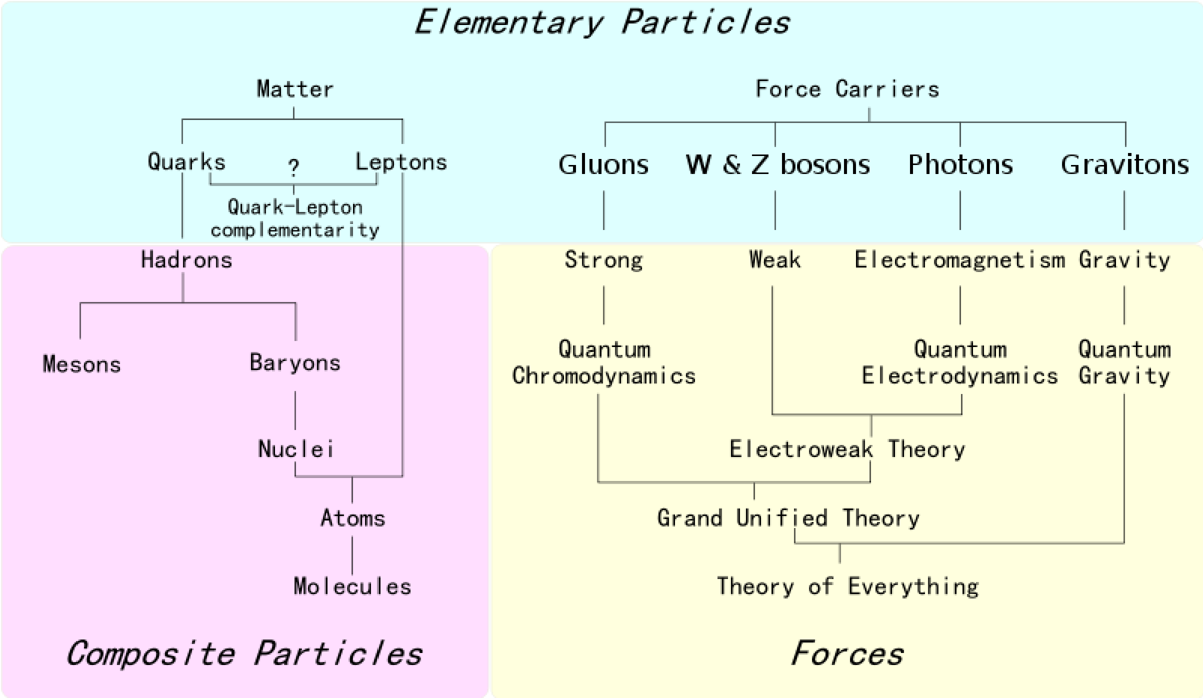

Elementary particles overview, leading to composite particles or forces governing their interactions, according to the Standard Model, from http://en.wikipedia.org/wiki/File:Particle_overview.svg

The LHC is a more than a 20-year-old dream for thousands of scientists around the world that came true in September 2008 when the first two beams of protons circulated along the 27 kilometers circumference beam pipe. It is an incredibly complex machine containing more than 9000 magnets and operates at an extremely low temperature (-271.25°C / 1.9°K).

To feed the LHC and other experiments in particle, CERN is using a chain of accelerator. The chain start with a proton or heavy ion source. In case of protons it is merely a hydrogen bottle providing a controlled beams of atom at 1 Hz. Hydrogen atoms are firstly stripped of their electrons to leave the hydrogen nucleus: protons. They are then accelerated by an electric field in the LINAC2, a linear accelerator built in 1978, up to one third of the speed of light and at 50 MeV.

Then protons are entering the PSB (Proton Synchrotron Booster). It is a circular 4-beams accelerator with a circumference of 157 meters. There, protons are boosted at 91.6% of the speed of light, and at 1.4 GeV and are injected into the PS (Proton Synchrotron), a master piece of CERN built in 1959. This 628 meters long circular accelerator is used for mostly every particle beam at CERN. It accelerates protons up to 28 GeV and reach the velocity of 99.9% of the speed of light.

Then, protons are injected from the PS to the SPS (Super Proton Synchrotron), another circular accelerator of 6.9 kilometers of circumference. Buried at 40 meters, it accelerates protons at 450 GeV, the injection energy of the LHC. They are interconnected by two transfer lines, one injecting in the LHC clockwise, the other counter-clockwise, going down to the level of the LHC, buried at 100 meters.
The 26.7 kilometers long 2-beams accelerator is currently ramping protons at 1.18 TeV (at the time of writing this thesis), which is already a world record. It should go up to 3.5 TeV per beams of proton by the end of 2010 if everything goes well. Then after an important upgrade, it should be able to reach its designed energy of 7 TeV per beam,

allowing then the 7+7=14 TeV collision energy. Velocity of the protons will then be of 2.7 meters per second less than the speed of light.

It takes between 5.86 seconds and 17.86 seconds for a proton to go from the hydrogen source to the LHC. But it takes 20 minutes to fill the LHC with billions and billions of protons. Then ramping up to 7 TeV per beam take also 20 minutes.

At the injection energy of 450 GeV, LHC's superconducting magnets are at 0.54 tesla. At 7 TeV, they should be at 8.3 tesla, which require to flow a current of nearly 12 kA. That the reason why the LHC's magnet need to be colder than the outer space, at 1.9°K.
To avoid damaging the magnets and other superconducting elements (busbars between magnets) when they quench, a specific device is added to each magnet, the Quench Protection System (QPS). A quench is when a very localised part of a superconducting element is developing a resistive behaviour. With a very high current flowing into, this quickly grow a hot spot on the element, which damage it if the current is not stopped in a few milliseconds and the heat spread.

In September 2008, nine days after the first starting of the LHC, an abnormal resistance of 220 **nano** ohms in an electrical interconnect released heat. An electrical arc developed, puncturing the helium enclosure and leading to a release of helium into the vacuum around the magnets. The rapid pressure rise in the vacuum enclosure in the neighbouring subsectors was responsible for most of the mechanical damage, which extended for some 700 meters.
The reparations lasted 11 months, and cost millions of Euro. The LHC just restarted the 20th of November 2009. The QPS has been largely enhanced to detect also this kind of problem that were not expected before.

But CERN has many other facilities, accelerators, decelerators, and experiment halls. Radioactive isotopes are studied in ISOLDE facilities. Influence of cosmic rays on clouds formation are studied in the CLOUD experiment. Anti-matter is studied in the anti-proton decelerator AD. The time of flight of neutrons are studied in n-ToF. Neutrinos are sent to Gran Sasso, near Rome, with CNGS. And tenth of other experiments are made at CERN.
They are even building AMS, a piece of the international space station (ISS) at CERN.
LHC experiments are explained in details in the next chapter.

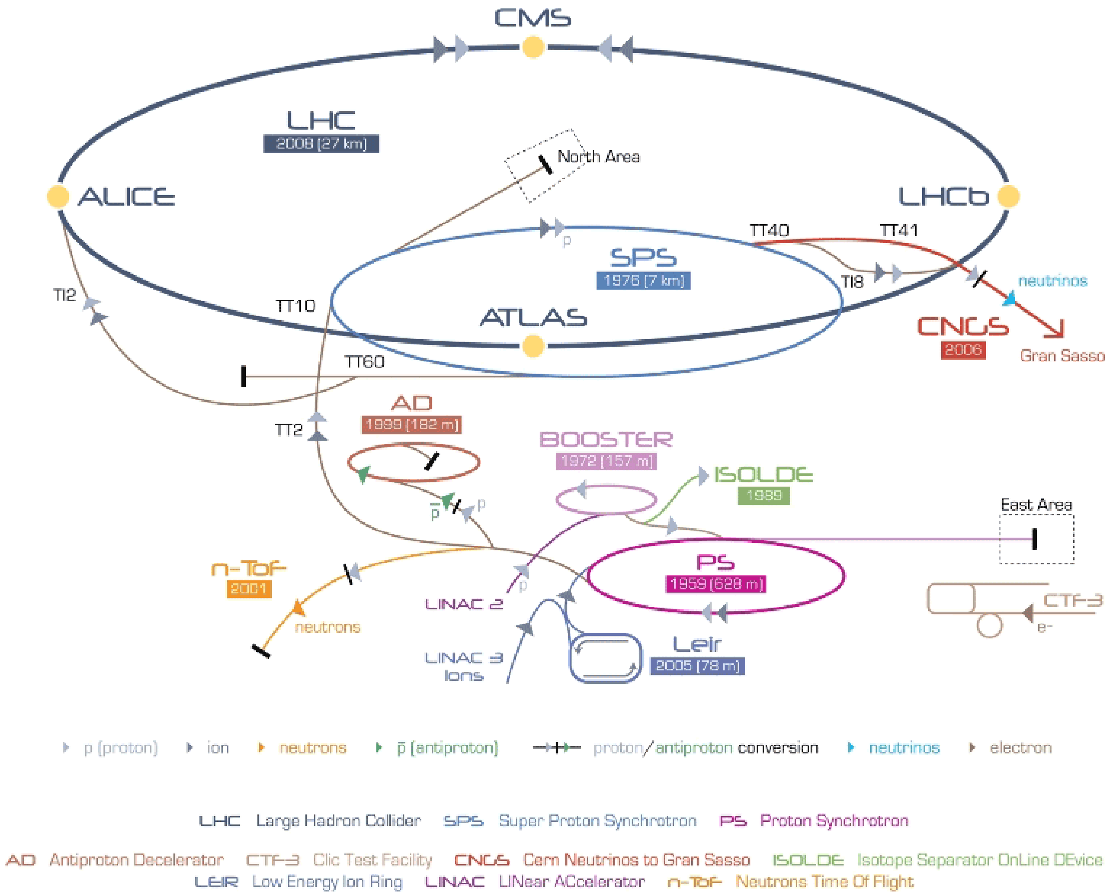

CERN Accelerator Complex

Accelerator	Experiments	Overview
LINAC2		accelerates proton from the source and sends them to the booster (PSB)
LINAC3		accelerates ions from the source and sends them to LEIR
PSB		accelerates protons and sends them to the PS and ISOLDE
PS	EA	East Area - various experiments
	nTOF	neutron time-of-flight
	DIRAC	to observe and then measure lifetimes of Muons and Kaons
	CLOUD	to study possible links between cosmic rays and cloud formation
	SPS	sends beam to SPS (CNGS, fixed target, LHC)
ISOLDE		to produce a range of isotopes for research
AD	ALPHA	to make, capture and study atoms of antihydrogen and compare these with hydrogen atoms
	ASACUSA	to compare anti-protons and protons using antiprotonic helium
	ATRAP	to compare hydrogen atoms with their antimatter equivalents
LEIR		accelerates ions and sends them to the PS
SPS	NA	North Area - various experiments
	CNGS	to send muon neutrinos to to the Gran Sasso National Laboratory in Italy
	COMPASS	to study how elementary quarks and gluons work together to give particles we observe
	LHC	injects beam into the LHC at 450 GeV
LHC	CMS	to search for the Higgs boson, extra dimensions, and particles that could make up dark matter
	ATLAS	to search for the Higgs boson, extra dimensions, and particles that could make up dark matter
	LHCb	to understand why we live in a Universe composed almost entirely of matter, but no antimatter
	ALICE	to study a state of matter known as quark-gluon plasma
	TOTEM	to measure the size of the proton and also monitor the LHC's luminosity
	LHCf	to simulate cosmic rays to interpret and calibrate large-scale cosmic-ray experiments

Description of CERN's accelerator and there attached experiments

CERN infrastructures are spread over two main sites and seventeen smaller. The biggest and oldest one is legally in the swiss city of Meyrin. But actually it is astride on the French-Swiss border. I used to work there during my first 6 months at CERN. Then, because of a recent reorganisation of CERN's structure our section moved to the second site located in the French town of Prevessin.

The Meyrin site host most of the buildings, machines, experiment halls, the administrative departments, the Tier-0 supercomputer centre, as well as offices for all experiment collaborations. There are also two hostels for short duration stay, and a water tower.
The Prevessin site is hosting the CERN Control Center (CCC), some experiments hall (and AMS) and a lot of engineering sections. Both sites have their restaurants (two in Meyrin, one in Prevessin), banks (UBS in Meyrin, Credit Agricole in Prevessin) and national post offices.

In total, CERN sites cover 195 ha. There is 25 kilometers of private road, 574 buildings, 23 pits and 750 green way.

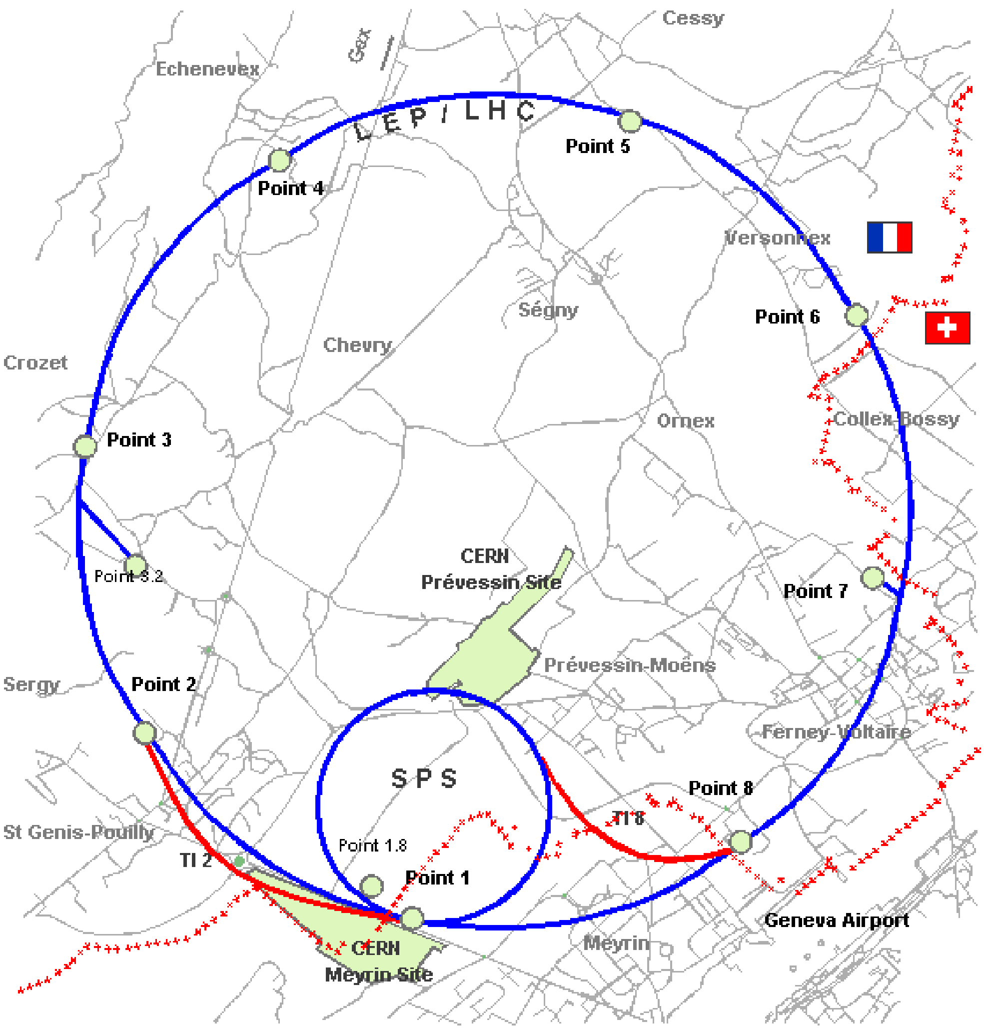

Map of CERN sites and LHC access points

iii. Main experiments and their detectors

Four large experiments, together with two smaller ones, will record the hundreds of particles produced in the LHC's collisions. The basis of particle detections are to measure:

- a particle charge (+1, 0 or -1)
- a particle momentum
- a particle energy as deposited in an electromagnetic calorimeter
- a particle energy as deposited in a hadronic calorimeter

The ratio of the two energies measured in the calorimeters is telling how much a particle looks like an electron/photon, or looks like a quark. To detect the particle charge and momentum, a particle-tracking sensor system is enclosed within a magnetic field. More a particle got a high energy, less the magnetic field will be able to bend its trajectory when it get ejected from a collision. The direction of bending give the charge. The degree of bending give the momentum.
With very strong magnetic field, low energy particles do a spiral. Very low energy particles are even not detected because they don't fly through the tracking system. They are trapped.

Today, it is mostly impossible to detect directly the very high energy particles that the collisions are intended to produce because they decay too quickly. But when these particles decay and recombine with other particles, they produce new particles of lesser energy. And these recombinations follow defined equations according to the original particles they come from. In addition, resulting particles are also flying away by following predictable trajectories.

Hence, when detecting a particle leaving a particular trajectory and a particular energy, along with another different particle following the same kind of predictable behaviour, it could be deduced that the decay of a particle of higher energy has created these particles, which then prove the existence of the "invisible" high energy particle.
Sometimes, theoretical physicists can say before the experiment which high energy particle could be found by predicting the equations of combination and trajectory. Some other times, it is an experiment that lead to discover a new high energy particle.

In both case, the scientific process requires to have these proofs happening a lot of time, and have a high detection probability. To validate the existence of a particle, a stochastic proof is then needed. And, if possible, at least with two different experimental setups.

The LHC is designed to produce beam of bunch of protons in each beam pipe. One beam pipe will contain 2 808 bunches. Each bunch will contain 10^{11} protons. But when the two beams are meeting up, about only 20 collisions should occurs per bunch because of very strong repulsing forces between particles.
The interval between each bunch is 7.5 meters, which is 25 nanoseconds, or 40 MHz.
So, even though a few particles are interacting each time the beams cross, it should produce a number of collision in the order of a billion per second. It will take ten hours to "consume" a whole beam fill. And the LHC and their experiments should be running during several months every year. The high stochasticity needed to prove a new particle will be reached with mostly no doubt on the future scientific discoveries.

1. Compact Muon Solenoid (CMS), the heavy-weight detector

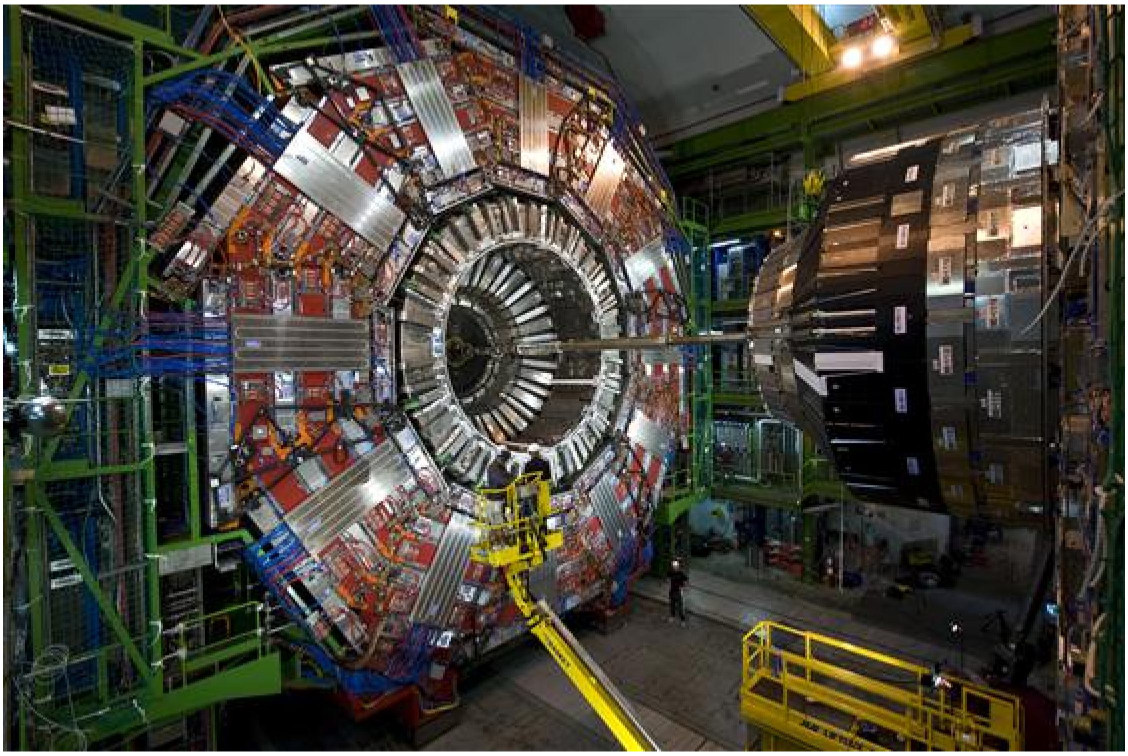
View of the CMS detector before closure, http://cdsweb.cern.ch/record/1133594

CMS is a general purpose detector used to study a large range of physical phenomena produced by particle collisions. In a unique strategy, the detector was assembled above ground concurrently with the excavation of the underground cavern.

The CMS Collaboration consists of over 2500 scientists and engineers from over 180 institutes in 38 countries. The main volume of the CMS detector is a cylinder, 21 meters long and 16 meters in diameter, weighing in total 12 500 tonnes. The tracking volume is defined by a cylinder of length 6 meters and a diameter of 2.6 meters.

About 210 m^2 of silicon microstrip detectors (around 10 million channels) provide the required granularity and precision in the bulk of the tracking volume; pixel detectors placed close to the interaction region improve measurements of the track impact parameters and allow accurate reconstruction of secondary vertices. The tracking system is placed inside the huge superconducting magnet, 13 meters long and 6 meters in diameter, which will operate at 3.8 Tesla. The cooldown of the CMS solenoid to the nominal temperature of 4.5 K was achieved at the beginning of August 2008.

The magnet return yoke acts as the principal support structure for all the detector elements. Muons are identified and measured in four identical muon stations inserted in the return yoke. Each muon station consists of many planes of aluminium drift tubes in the barrel region and cathode-strip chambers in the endcap region.

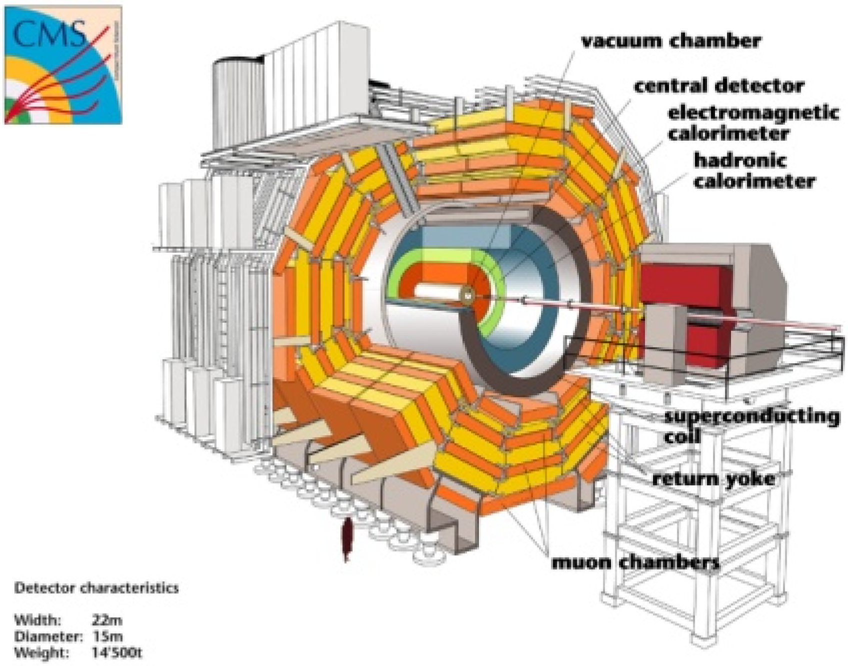

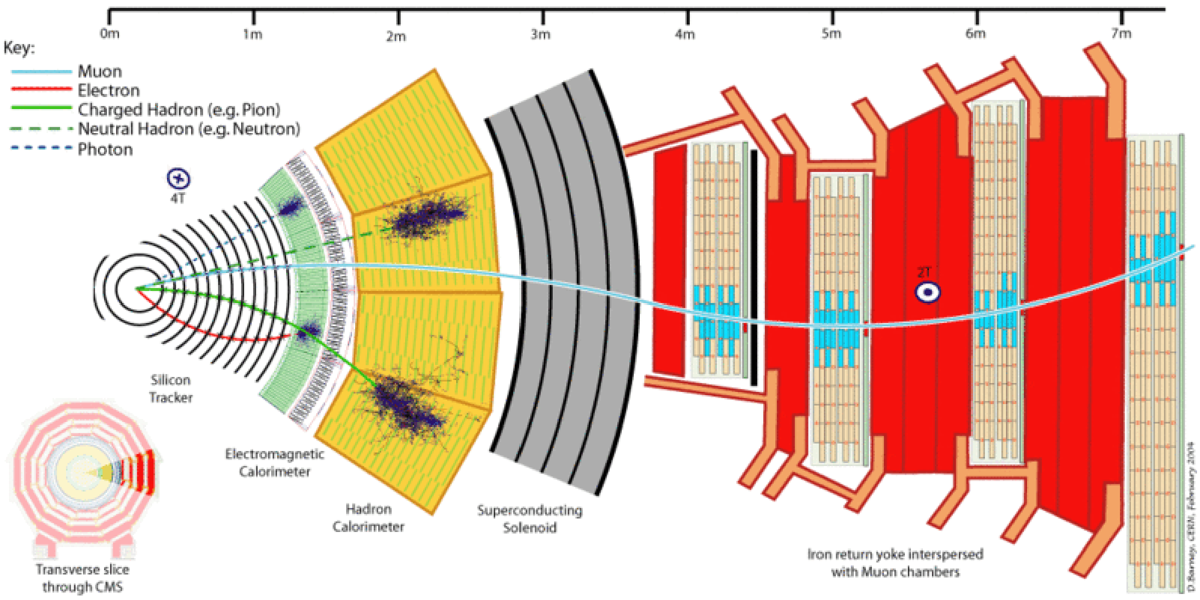

CMS schematic structure

2. ATLAS, the largest volume particle detector

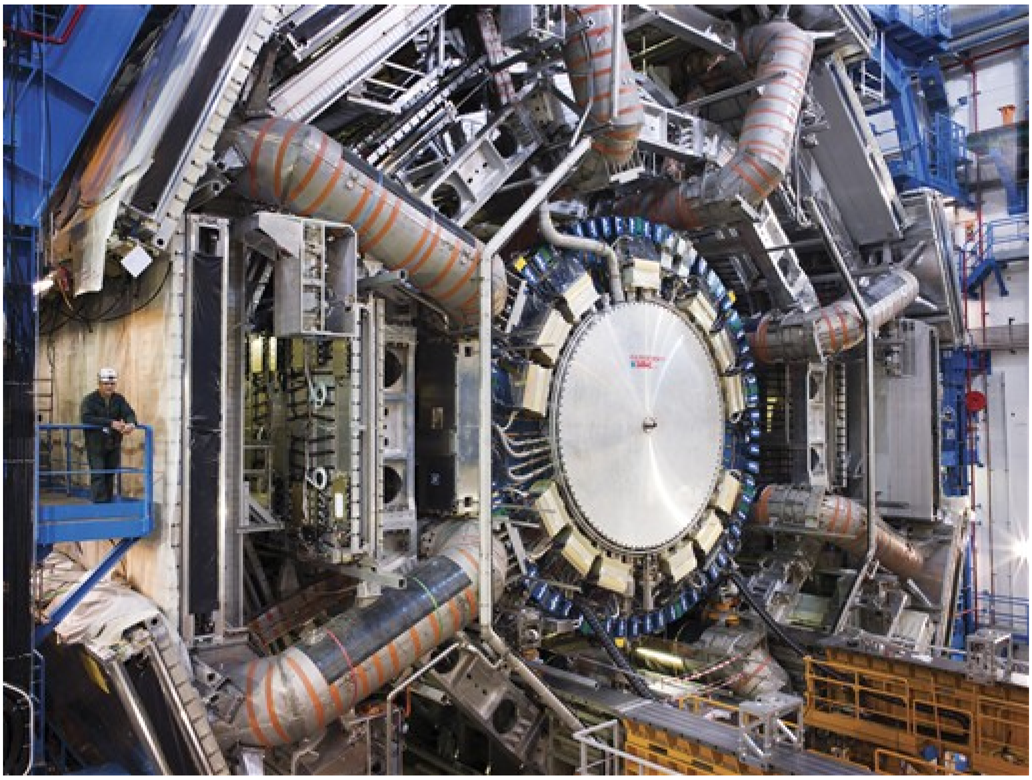

View of the ATLAS detector open, from http://cdsweb.cern.ch/record/1206860

ATLAS is a general-purpose experiment for recording proton–proton collisions. The detector design has been optimized to cover the largest possible range of LHC physics. This includes searches for Higgs bosons or alternative schemes to answer the puzzling question about the origin of mass, and searches for supersymmetric particles, and other new physics beyond the Standard Model. The ATLAS Collaboration consists of 169 institutions from 37 countries with roughly 2800 scientific participants.

The ATLAS detector has cylindrical symmetry around the beam pipe, with increasingly large layers of subdetectors placed around it and endcaps to ensure hermiticity. The inner detectors — a series of thin silicon and gas detectors immersed in a solenoidal magnetic field — are used for pattern recognition, and for momentum and vertex measurements. In addition to the central solenoid, the magnet system also comprises a barrel toroid and two endcap toroids. The high granularity liquid-argon electromagnetic calorimeters and the hadronic scintillator-tile calorimeter are surrounded by the muon spectrometer, which defines the overall dimensions of the ATLAS detector.

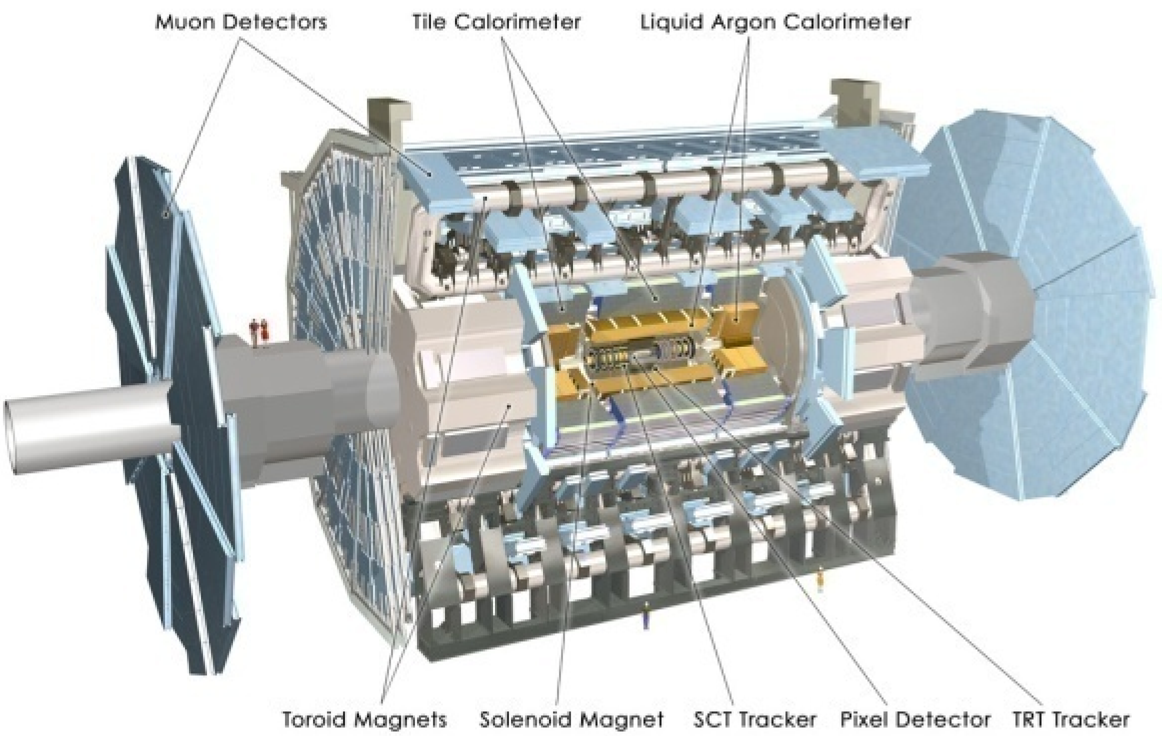

ATLAS schematic structure

3. LHCb, tracking down antimatter

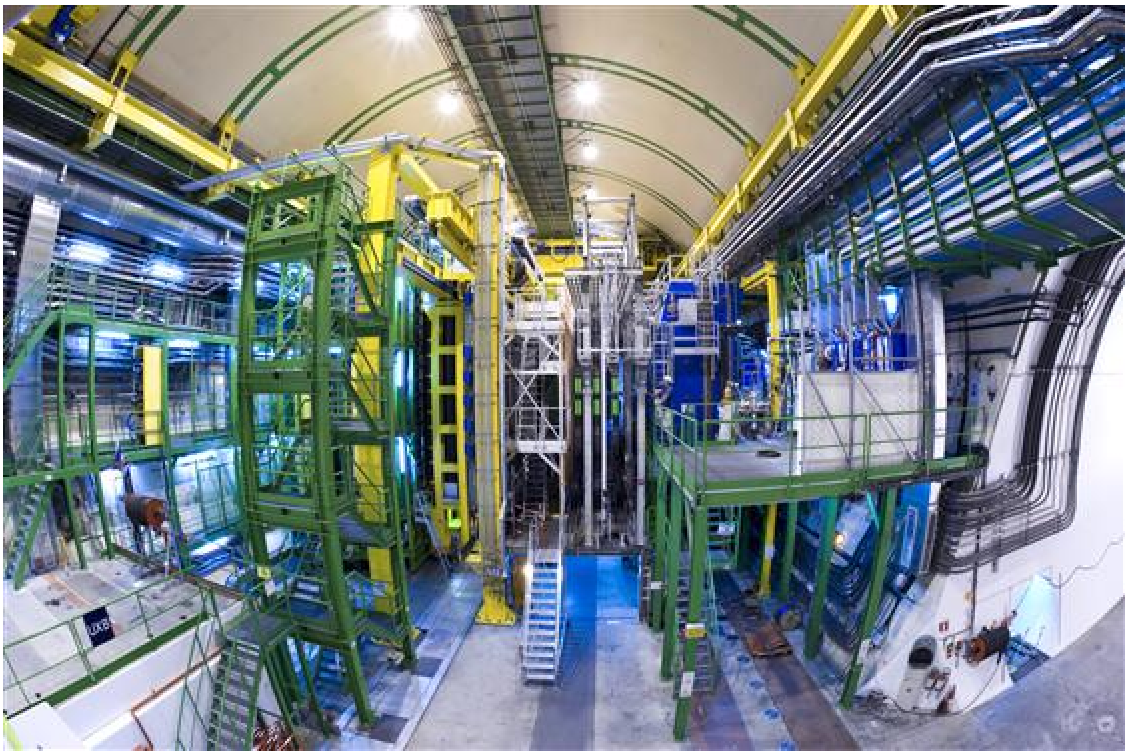

View of the LHCb detector, from http://cdsweb.cern.ch/record/1124325

The main purpose of the Large Hadron Collider beauty (LHCb) experiment is to investigate the phenomenon known as CP violation in the decay of particles containing b and anti-b quarks, collectively known as 'B mesons'. CP violation is a necessary ingredient in explaining the total absence of antimatter in the Universe.

Rather than flying out in all directions, B mesons formed by the colliding proton beams (and the particles they decay into) stay close to the line of the beam pipe. This is reflected in the design of the detector, which stretches for 20 meters along the beam pipe, with its subdetectors stacked behind each other like books on a shelf.

The point where the beams collide, and B mesons are produced, is inside the VErtex LOcator (VELO) subdetector. With its 84 half-moon-shaped silicon sensors, each connected to electronics, the VELO can locate the position of B particles to within 10 micrometers.

Two ring-imaging Cherenkov detectors (RICH), as well as tracking detectors, lie on either side of LHCb's magnet, which consists of two coils, both weighing 27 tonnes, mounted inside a 1 450 tonnes steel frame.

The calorimeter system — a sandwich-like structure, with alternating layers of metal and polystyrene plates — is LHCb's main way of identifying neutral particles, such as photons.

Muons are present in the final states of many B meson decays, so muon detection is vitally important for the LHCb experiment. Located at the far end of the detector, the muon system comprises five rectangular 'stations', gradually increasing in size and covering a combined area of 435 m² — about the same size as a basketball court.

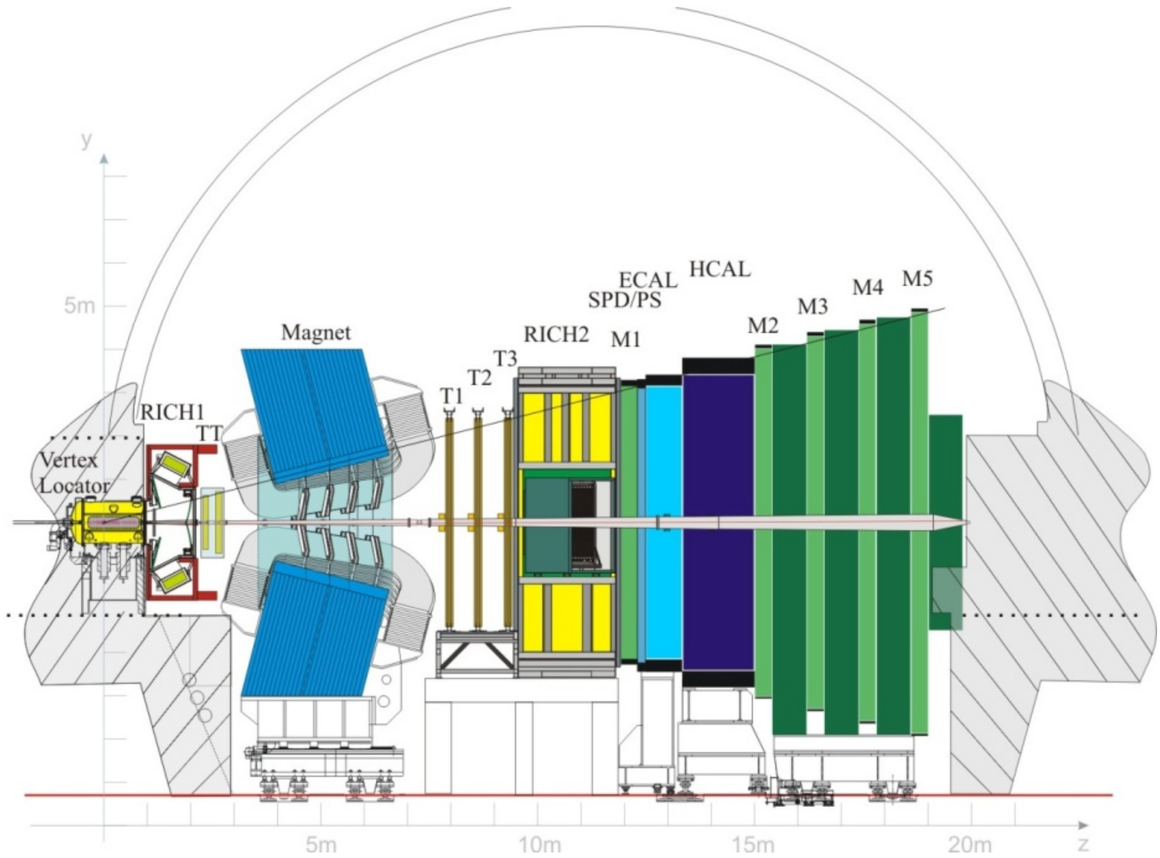

LHCb schematic structure

4. ALICE, probing the quark-gluon plasma

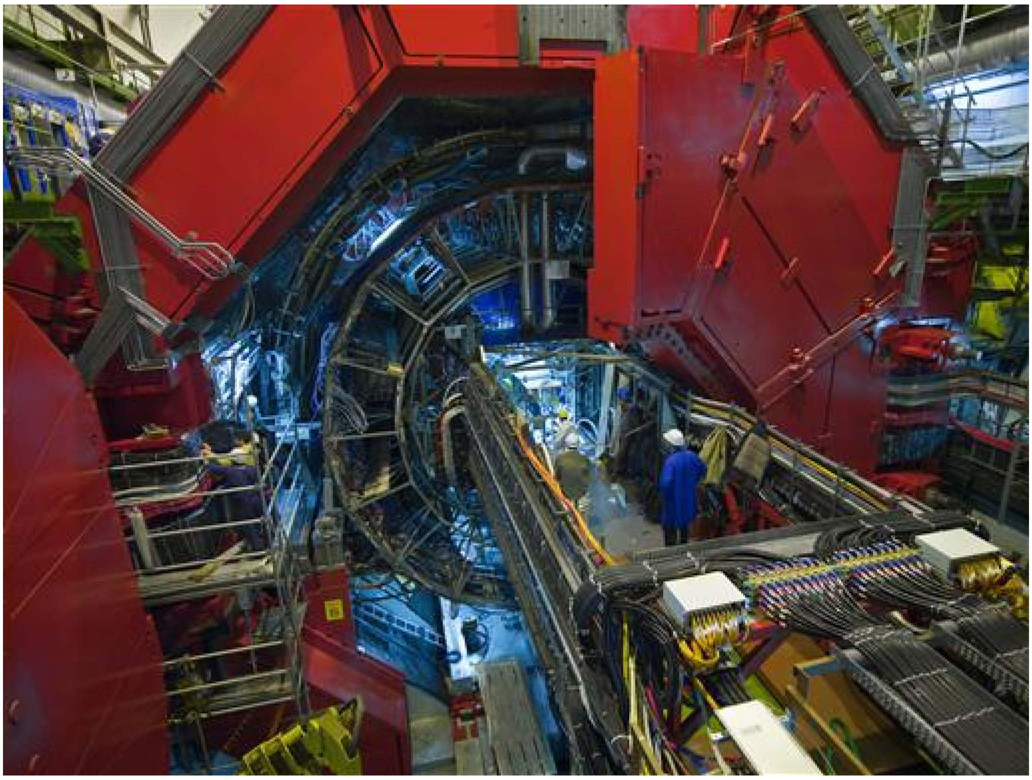

Closing the first door of the L3 magnet in the ALICE experiment, from http://cdsweb.cern.ch/record/1089856

ALICE is a heavy-ion experiment designed to study the physics of strongly interacting matter and the quark–gluon plasma in lead–lead collisions at the LHC. The ALICE Collaboration currently includes more than 1000 physicists and senior engineers — from both nuclear and high-energy physics — from about 100 institutions in some 30 countries.

ALICE consists of a central part, which measures hadrons, electrons, and photons, and a forward spectrometer to measure muons. The central part is embedded in the large L3 solenoid magnet and comprises an inner tracking system (ITS) of highresolution detectors, a cylindrical time projection chamber (TPC), three particle identification arrays of time-of-flight (TOF), ring imaging Cherenkov (HMPID) and transition radiation (TRD) detectors, plus two single-arm electromagnetic calorimeters (the high-resolution photon spectrometer PHOS and the large-acceptance jet calorimeter EMCAL). The forward muon arm consists of a complex arrangement of absorbers, a large dipole magnet, and 14 planes of tracking and triggering chambers. Several smaller detectors (ZDC, PMD, FMD, T0, V0) used for global event characterization and triggering are located at forward angles. An array of scintillators (ACORDE) on top of the L3 magnet is used to trigger on cosmic rays.

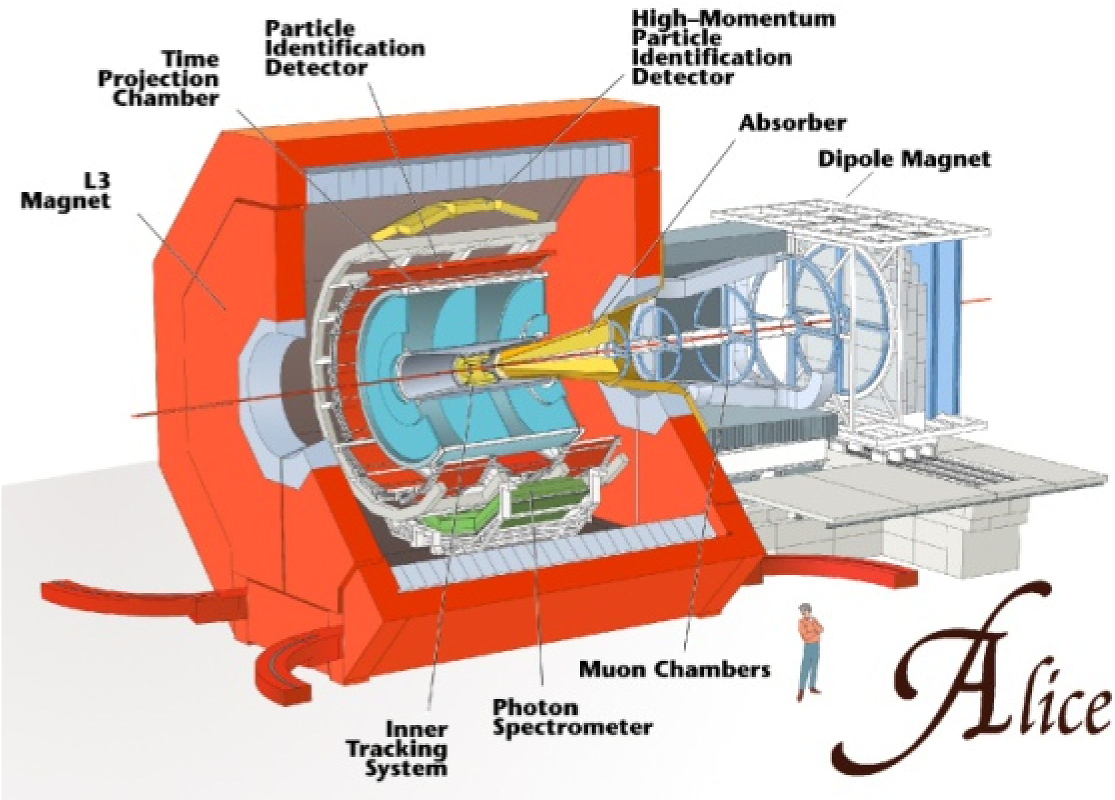

ALICE schematic structure

iv. Data crunching

By detecting around 20 collisions at 40 MHz, the processing and analysis of the data produced by the detectors will present a huge computing challenge. The Worldwide LHC Computing Grid (WLCG) has been developed to harness distributed computing resources for physicists in institutes around the world to make the most of the rich harvest of physics that will be available at the LHC.

With the amount of sensors in each experiment (about tenth of thousands), detectors are producing tera bytes of data every 25 nano second. A first level of electronic named the trigger is selecting or discarding an event based on its quality. Then a second level of computing is compressing and filtering event's data. But it still produces giga bytes per seconds of data that are sent to the CERN computing centre.

The Tier-0 centre is primarily intended to store all these data. Every year it is 15 peta bytes produced by CERN that have to be processed and analysed. For this purpose "the grid" has been invented. The WLCG spans some 150 computing centres around the world to do so.

In 2008 a test have been made to ensure that the WLCG is able to fulfill its job. Simulated data were used to demonstrate the full chain of data processing, from the detectors to the CERN Tier-0 computer centre, via the Tier-1 centres (11 large facilities around the world) through to the analysis centres at the Tier-2 sites. Data were continually transferred from CERN to the Tier-1 computing centres at a rate of more than 2 gigabytes (GB) per second. This is more than three times the data rate needed to support actual LHC running, and is equivalent to copying one full DVD every three seconds. The WLCG must be able to manage a large number of 'jobs' — runs of a program that analyses a chunk of data and takes some 8–12 hours of computer time. The challenge showed that the system could manage more than half a million such jobs every day — well above the levels required.

To do these physics computing jobs, computer specialists at CERN work in collaboration with physicists from the experiments to develop in-house, general software packages that can be adapted to meet the specific needs of each experiment. ROOT and GEANT4 are the two main projects currently under development. ROOT is a tool for statistical analysis and visualization of data, storage of complex C++ object data structures, and distributed analysis.

GEANT4 is a simulation code, designed to simulate the response of the detectors to traversing particles, once the geometry is specified. It is widely used not only in high-energy physics, but also in medical and space science. In 2008 the code was updated with new models for hadrons, and to respond to the feedback from the cosmic and test-beam runs from the LHC experiments.

But the computing force is also on the other side of the detectors. To have detectors able to run, as well as the LHC, they all need to be controlled. Control systems at CERN have to operate more than hundred of thousands I/O signals. To do so, considerable effort was dedicated to improve the tools available for both deploying the control system software and monitoring the control systems themselves. The close collaboration with ETM, a subsidiary of Siemens who provides PVSS, a supervisory control and data acquisition (SCADA) system used extensively at CERN to protect the experiment equipment, continued to bear fruit with many important enhancements achieved. The two-way benefit of this collaboration was a factor in the decision taken by Siemens this year to join CERN openlab, the partnership between CERN and the IT industry.

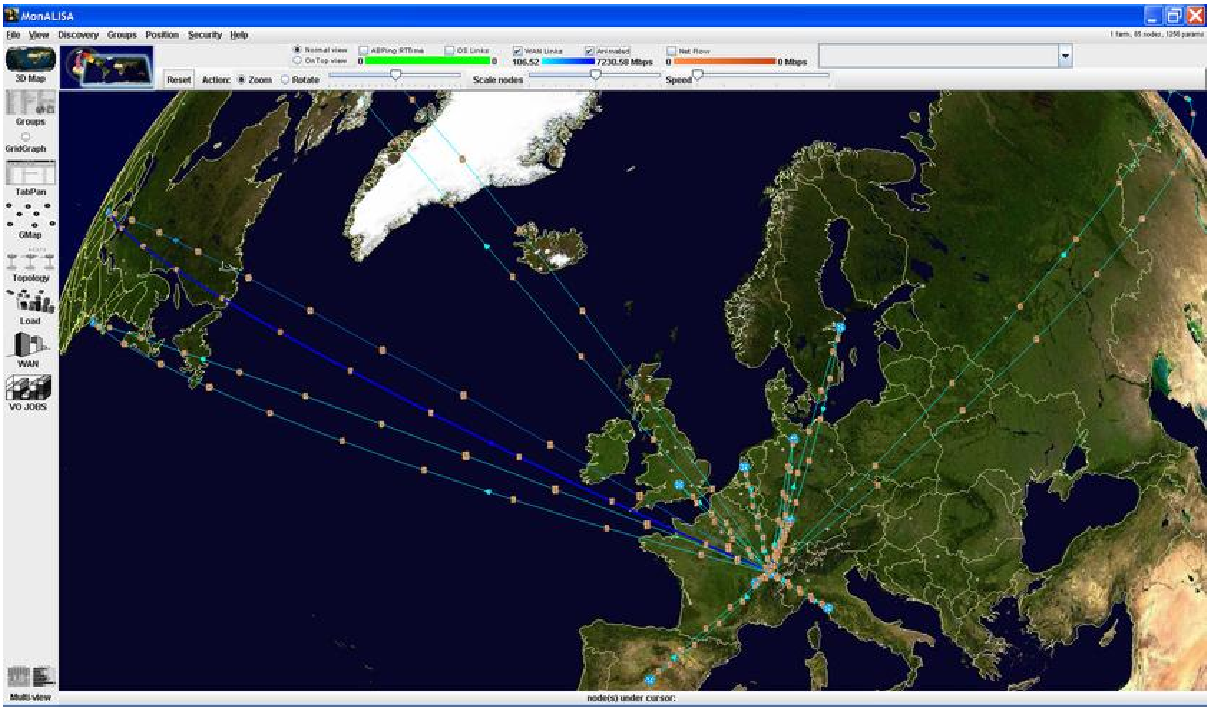

Data flow during a LCG test challenge, from http://cdsweb.cern.ch/record/915161

Left: Inside of a tape storage robot.
Right top: a small part of Tier-0. Right bottom: an IBM storage unit.
http://cdsweb.cern.ch/record/1106092
http://cdsweb.cern.ch/record/1138228
http://cdsweb.cern.ch/record/989403

v. Knowledge and technology transfer

The construction of large fundamental research infrastructures requires a long and intensive R&D phase before the production to industry. This R&D has the following characteristics:

- It is a source of innovation.
- It is a source of new technologies.
- It pushes existing technologies to the limits.
- It is a source of know-how creation.

Although developed for the purpose of fundamental research, many technology developments and know-how made at CERN can have a strong impact on society. Generic technologies developed for accelerators, particles detection and data processing will find applications in domains related to health, information technology, energy and environment, industrial processes and security.

Knowledge and technology transfer is an integral part of CERN's principal mission of fundamental research.

A very well known technology primarily invented at CERN and for CERN needs is the World Wide Web. The principles of hyperlink between documents, of Uniform Resource Identifier (URI), the HyperText Markup Language (HTML) and the HyperText Transfer Protocol (HTTP) have been invented by Tim Berners-Lee at the beginning of the nineties when he was an independent contractor at CERN.

CERN had also an impact on the cancer research by improving particles detectors technologies. The high requirements in superconducting elements for LHC's magnets improved the industrial process of this field. Some of the improvement have actually been possible thanks to previous discoveries made at CERN.
As a last example, the needs in computing power for CERN's experiments leaded to the creation of the Enabling Grids for E-sciencE (EGEE) project which is the world's largest multi-science Grid infrastructure, bringing together more than 140 institutions across the world.

vi. Safety and public concerns

When devising and running its machines, CERN is ensuring a lot of safety and security rules. First of all it prevents as much as possible accidents by reviewing multiple times all engineering work. As CERN as several radioactive zones due to the interaction of particles with matter, each of these zones is controlled by a specific access level that need several security courses before obtaining it. In addition, every person having access to radioactive area has to wear a dosimeter, an electronic device recording radioactivity. Furthermore, the access to the newest machines is now ruled by biometric eye scanners airlock. Ultimately, everybody having access to radioactive area has to do a medical checkup.
And it is worth mentioning that at every car entrance of CERN's sites there is radioactive detector able to block a vehicle having trace of radioactive elements.

CERN rejections of air and water are also controlled by the organisation and by an external and independent entity. Reports on the environmental impact of CERN are regularly published. You can learn for instance that by living one year in the CERN region you receive as much radioactivity as you would receive in a air plane going from Geneva to Athens.

The LHC start-up gave also rise to many fanciful theories about the possible consequences of high-energy collisions. In particular, rumours were rife concerning the hypothetical appearance of black holes. While it is true that certain theories predict the production of mini black holes in LHC collisions, all such theories also predict that they would decay instantaneously and have no macroscopic effect.

The safety of collisions at the LHC was the subject of a detailed study by the LHC Safety Assessment Group (LSAG), comprising scientists from CERN and the Institute for Nuclear Research of the Russian Academy of Sciences. Their in-depth report was presented to the CERN Council in June 2008 and published in November in the renowned Journal of Physics G.

This study shows that LHC collisions present no danger. Even though the LHC will unquestionably be capable of attaining a level of energy never before reached in a particle accelerator, this will still be lower than energies commonly produced by Nature in cosmic ray collisions, for instance. The phenomena generated by the LHC have already been produced by Nature countless times since the Earth and the other celestial bodies came into being.

The report was reviewed by the Scientific Policy Committee (SPC), a body comprising 20 independent external scientists that advises Council on scientific matters. The SPC approved the methodology adopted by the authors, who based their arguments on irrefutable observations and concluded that new particles that could be produced by the LHC will present no danger.

vii. People in the organisation and student programs

Altogether, CERN's research programme involves some 9000 researchers from over 550 institutes and 113 countries. All the experiments are international collaborations and the big ones have more than 2000 members each.

CERN's international connections and its large collaborations provide a valuable environment for educating and training young people, not only in particle physics, but also in engineering of many kinds, information technology, and administration. The Laboratory is a temporary home for students from around the world, whether pursuing work towards postgraduate qualifications or as participants in the renowned programs.

Out of the 9000 people working at CERN, only a third are actually employed by CERN. And surprisingly, CERN employ very few physicists! It is because the aim of CERN is to build the machines and software needed for the experiments, and not actually do the research work.

CERN has many different job programs. The first is the normal staff. These people are usually professional and experienced in their field. When CERN recruit a new staff member, the person has to go through a tough selection procedure. CERN want to employ only highly skilled and outstanding people. For this, it got multiple advantages to attract such "brains". The working conditions are really good. There is no pressure on results but everybody is doing his job. The environment is attractive, with Geneva and the Leman lake just aside, Jura and Alps easily reachable within hours. This offers a lot of sportive or entertainment possible combinations.

CERN salary is really rewarding too. The salary scale is divided in eight major grades, each subdivided in approximately 30 steps. Every year a personal review is done in which you can earn a few steps, increasing your salary, and sometimes your responsibilities.
The starting salary for an engineer at CERN is 7418 CHF per month (~4 900 euros). A senior engineer with some management duty can earn up to 16843 CHF per month (~11 100 euros).
As it is an intergovernmental organisation, it is quite difficult to speak in terms of gross or after-tax salary. A contribution to CERN's own health insurance is automatically made. Aside of that CERN has also its own pension fund. Otherwise, no tax are deducted from a staff salary. And the person got a diplomatic status according to its function in the organisation (in french grade: K for technical or administrative workers, CD (Corps Diplomatique) for important managers).
Most of the staff contracts have a limited duration of usually four to six years. At the end of the contract the person can either leave or continue at CERN, according to his wills and also to CERN needs. CERN is working by project, which mean after some years a position might not be needed anymore. That is why getting an indefinite contract at CERN is really difficult and tend to be exceptional. This also allows for internal mobility of people within the organisation's projects.

The lowest staff salary is 3678 CHF per month (~2500 euros) for mechanic/craftman with no experience.

But other programs are aiming for young people. The first relevant is the Summer Student program. A three month internship in which every morning, students have open lectures about physics, computing or electronics aspects used at CERN. There is about 150 official summer students every year at CERN, and about 1 000 candidates. I started working at CERN with this program two years ago, for my I2 internship. The selection procedure is through a committee, no interview. If someone is selected, the person is proposed to CERN's section leaders who pick up the ones they want. The student is then proposed to accept the work topic or reject it.
Paid 2000 CHF per month (~1500 euros). It seems a lot for a low-level internship but in the CERN world it is one of the lowest stipend...

Next is the Technical or Doctoral student program. This is the program I am currently in. It is intended for students between a bachelor and an Master in Science. The Technical Studentship is for engineer or computing students. Doctoral studentship is for physicists. It last between six months and one year. The selection procedure is similar to the Summer Student program, except for a phone interview. There is about 300 to 500 candidates, and around 90 selected per year.
Stipend is about 3000 CHF per month (~2000 euros), no tax. Still, this "low salary" is not even ranked in the staff salary scale (stipends are considered as subsistence allowances and not salary, it makes things easier).

Then come the Fellowship program. This is a one year program, extensible a second year (and exceptionally a third). It aims for freshly graduated engineer for the junior version, or for physicists that already made a PhD for the senior version. Selection procedure is a bit tougher and longer. Salary is according to the diploma and experience of the person. From word of mouth I heard it may vary between 5000 CHF per month (3300 euros) for the unexperienced junior up to ~9000 CHF per month (6000 euros) for someone holding a PhD and various exceptional references. No tax, and with a pension fund.

Similarly to the Fellowship program exists the Openlab program which has almost the same conditions at the difference that an external company is paying the person, not CERN. This happen within really predefined conditions, work topics and a few selected partner companies: HP, Intel, Oracle and Siemens. Openlab program is currently aiming at four specialised fields: Automotion and Control, Database, Networking and Platform competences.
CERN openlab provides a framework for collaboration with industry, welcoming major industrial partners. These partners bring equipment, solutions, and experts, as well as funding to hire young engineers and scientists. In turn, CERN provides a demanding computing environment to stress-test solutions and push cutting-edge technologies to their limits. Significant results were achieved in 2008 in several areas, in particular energy-efficient computing and thermal measurements. Disseminating the expertise and knowledge created in the openlab is an important ongoing objective. In 2008 these activities included regular training sessions, contributing to the activities of the CERN School of Computing, and organizing a dedicated programme of lectures for the openlab Summer Student Program.

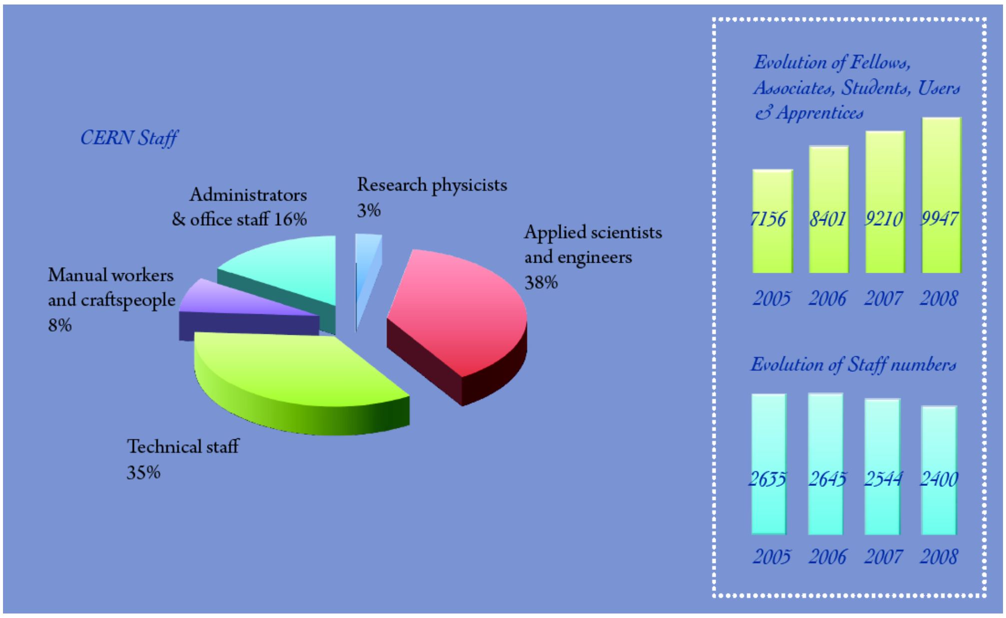

Staff statistics in 2008

viii. Global structure of CERN

CERN has the following organs.

The Council: Supreme authority, the council is responsible for the strategic goals of CERN. It has the last word on all major decisions. Scientific, technical and administrative decisions are within its very specific jurisdiction.
Each member state is represented by two persons at the council, one as the government and the other as the science spokesperson of the country. Each country count for one vote at the Council. Decisions are taken by strict majority, even though usually consensuses are found.

The Scientific Policy Committee (SPC): This committee is composed of scientists of any nationality (even non-member states). The committee is mandated to verify and propose long-term scientific strategies for CERN.

The Financial Committee: Composed of spokespersons from government of every member state, the financial committee is managing CERN's budget. It also decide the financial contributions of member states as well as checking expenses in the organisation.

The Director General: Nominated for five years, the DG is administrating CERN thanks to hierarchical structure detailed afterwards. The Directorate is assisting the DG in his duty. Currently the DG is Rolf-Dieter Heuer.

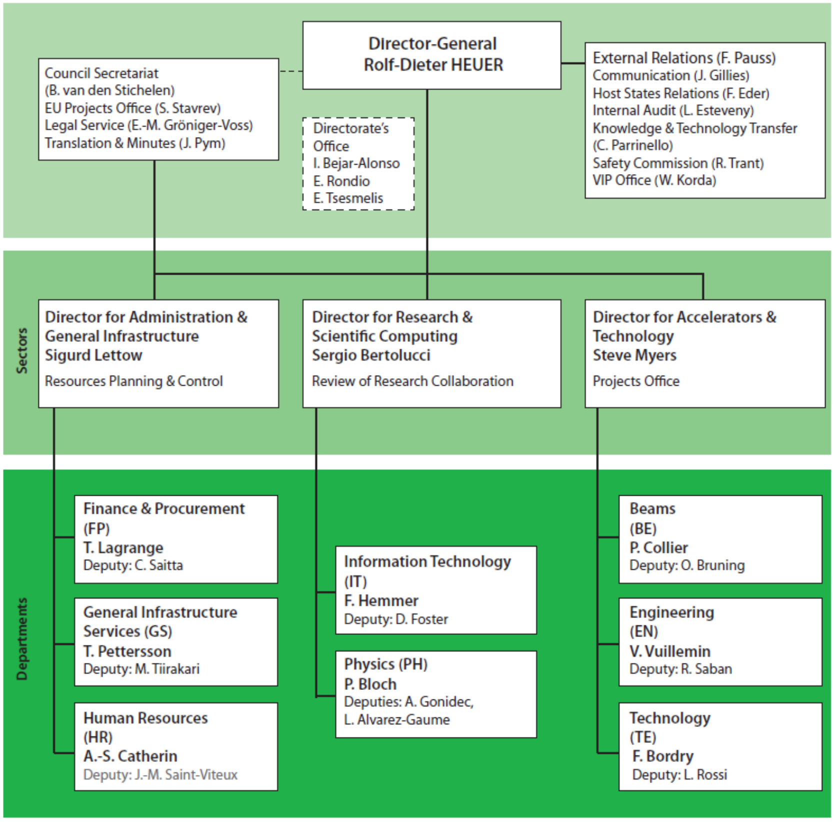

Each department is then divided in groups, themselves divided in sections. I am in the Accelerators and Technology sector, within the Engineering department, Industrial Control and Electronics (ICE) group, SCADA systems (SCD) section.

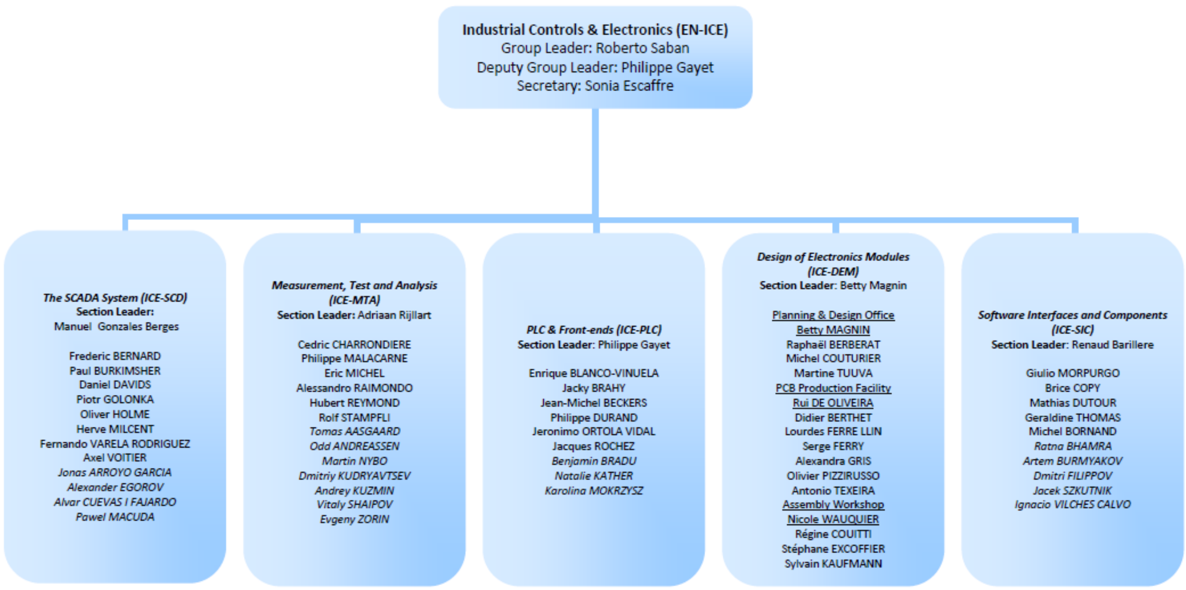

ix. All this thanks to your taxes: financing

CERN budget is shared between members states, following their Gross Domestic Product (GDP). Here is the repartition of contributions from member states in CERN's 2009 budget (in percentage, million of Swiss Francs and the equivalent in million of Euro):

Country	Contribution	MCHF	MEUR
Germany	19.88%	218.6	114
France	15.34%	168.7	111.2
United Kingdom	14.70%	161.6	106.5
Italy	11.51%	126.5	83.4
Spain	8.52%	93.7	61.8
Netherlands	4.79%	52.7	34.7
Switzerland	3.01%	33.1	21.8
Poland	2.85%	31.4	20.7
Belgium	2.77%	30.4	20.1
Sweden	2.76%	30.4	20
Norway	2.53%	27.8	18.3
Austria	2.24%	24.7	16.3
Greece	1.96%	20.5	13.5
Denmark	1.76%	19.4	12.8
Finland	1.55%	17	11.2
Czech Republic	1.15%	12.7	8.4
Portugal	1.14%	12.5	8.2
Hungary	0.78%	8.6	5.6
Slovakia	0.54%	5.9	3.9
Bulgaria	0.22%	2.4	1.6

From http://dg-rpc.web.cern.ch/dg-rpc/Scale/Scale.html

The total is 1 098.49 million CHF (723.91 millions Euro).
The LHC cost 4.6 billion CHF (about 3 billion Euro) just for building the machine itself (excluding detectors and running costs).

The organisation's expenditures in 2008 were as follow:

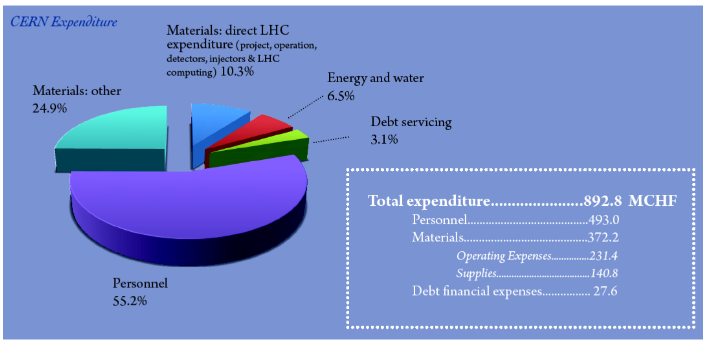

CERN expenditures in 2008

II. Test suite for the archiver of a SCADA system

To help the reader in understanding the key aspects and requirements of my work topic, important points have been enlightened in the text **like this**.

The introduction about CERN needed to be extensive to ensure the reader understand how big things are in the laboratory. CERN does not have the usual scale of scientific experiments, and that is why engineering is an aspect as important as research in CERN experiments.
Bear in mind that machines at CERN need a precise and complex control. The amount of controlled devices and their geographical repartition, as well as the amount of different subsystems needed to operate the machines lead to very large scattered and distributed control system. In addition, these systems have to be resilient. Even more, when there is a problem, it should be possible to **inspect the conditions of incidents** and then correct the design of the machines if needed.

We are now going to see a part of control systems at CERN, and follow these basic requirements about **scalability** and **reliability**.

i. Control systems at CERN

A very simple control system is a process running in a device and aiming at monitoring, acting on or switching state of an external piece of hardware. The system that is intended to be controlled can be of any kind, from a voltage generator to an electro-valve for instance.
Industrial control systems usually gather several of these subsystems that need to be controlled altogether in a smart way to achieve the industrial mission that a machine should do.

At CERN the different machines, from accelerators to detectors, are composed of electric systems needing to be powered, cryogenic system flowing any kind of cold liquid, gas system for creating a vacuum inside beam pipe, electro-mechanic system like ventilation or valves, etc.
These systems have to work altogether in order to put the machines in abstract status such as "pre-cycling a beam", "colliding", "taking physics data", "cooling down/warming up", "standby", "safety access", "calibration", etc.
To do so, the usual method is to engineer the subsystems in a way they are independent and are taking care only of their own local status. Then, each subsystem report a set of command and relevant monitoring signals to an upper layer that will be in charge of managing a more global state, without bothering about the low level details of subsystem control processes.

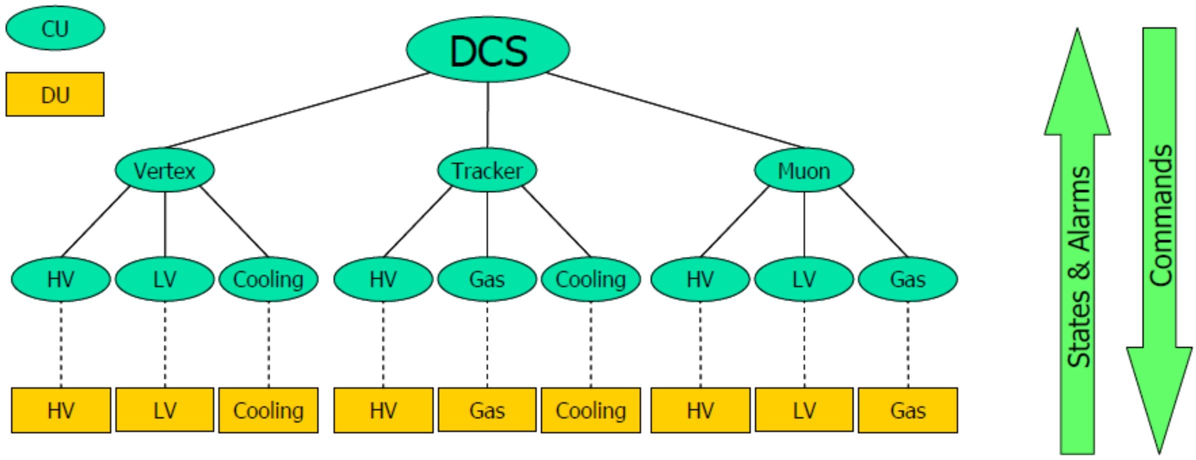

A simple hierarchical system example. CU: Control Unit. DU: Device Unit.

From the operations point of view, a control system must ensure **continuous monitoring** and control of the process: the operator deals with the dynamic aspects of the process and must be able to acquire and interpret information at every level in order to take the correct decisions in a timely manner. To satisfy this requirement, the control system architect must design a system that is **reliable, available, flexible** and safe. These qualities must be designed into all the levels of the Control System.

This usually lead to have an architecture composed of three layers:

- The field layer: This layer is the closest to the field, to the hardware, to the equipment making up the installation. It is populated by devices, which allow the interaction with the process either via measurements (e.g. temperature, pressure) or actions (e.g. to throttle a valve, to turn a heater on, to start a pump) as well as by the equipment electrically interfacing to them. These devices are either directly cabled to the interfacing electronics via digital or analog signals or, via field buses for the so-called smart sensors and actuators. The field buses are digital communication busses implementing robust protocols. The main role of this layer is to enable communication between the equipment in the field and the next layer up.

- The process layer: This layer implements the control of the process by interacting with the equipment in the field layer. It runs real-time algorithms to handle regulation, automatic procedures, emergency situations, etc. These algorithms are coded in embedded application programs running in VME modules, PXI controllers (PCI eXtensions for Instrumentation), PLCs (Programmable Logic Controller) or industrial Front-End Computers (FEC). A very high level of reliability and fast reaction times are expected at this level. In addition, acting as an intermediary between the field and the supervision layer, the process layer handles large amounts of data and generates status information and alarms. This is passed to the next level up via high level communication mechanisms generically called middleware.

- The supervision layer: This layer is often equally called the decision layer. At this level, the Operator receives combined and summarized information through synoptics, trend curves, alarm and status screens, etc. This information supports the decision making process. High-level services like **data archiving and retrieval**, email notification, interoperability with other supervision systems and web-based remote monitoring are provided in this layer. This level can either use measurement and instrumentation tools for developing and commissioning a control system, or use a monitoring and operating tool for normal operations.

In the latter case this kind of tools are usually called **SCADA systems**, for Supervisory Control and Data Acquisition.

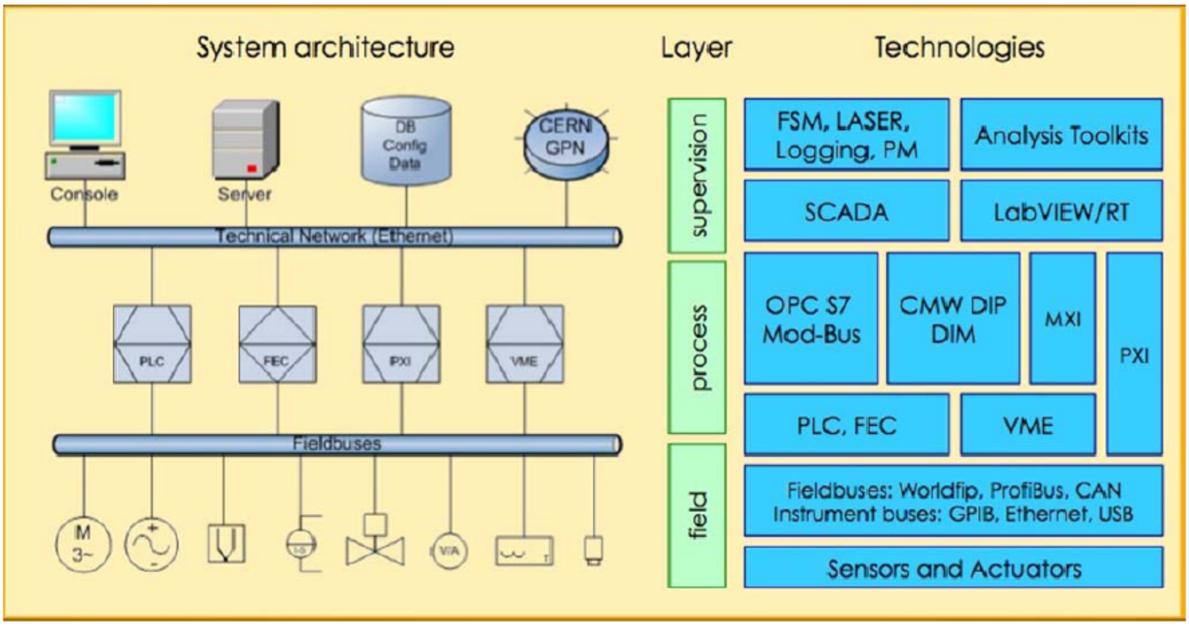

Layout and component of a control system at CERN.

ii. The EN-ICE group

EN-ICE is the name of the group I am working in at CERN. It stands for Engineering department, Industrial Control and Electronics.
The ICE Group develops solutions and provides support in the domain of large and medium scale industrial control systems as well as laboratory measurement, test and analysis systems and promotes their use in the organisation's projects.

The domain of action of the ICE Group covers the accelerator, the experiments and the technical infrastructure systems where it develops solutions (the frameworks and the interface components) and applications which encompass the different layers (supervision, process control, communications, etc.) of the systems used in a variety of different types, sizes and complexity of controls applications.

Unlike many groups building equipment for the machines and other groups specifying and procuring installations of varying sizes, the ICE Group is in a partnership relation for the definition, implementation, procurement and support of the controls and **data acquisition** parts of these installations. Its activities are defined below:

- In collaboration with the controls community, identifying and selecting commercial solutions (platforms, communications and software systems) which are both widely used in industry for process control and match the CERN needs.
- Negotiating supply and licensing contracts, including a system-lifetime support, with the industrial partners.
- Providing both product support and support for the implementation of systems and applications based on these solutions.
- Designing, implementing and supporting a set of frameworks, components and applications to ease the development of complete control system solutions with

the objective of re-using components and cutting down the cost of development, maintenance and support.

- In collaboration with the Controls Groups of the accelerator and the experiments, designing, implementing and supporting interfaces which allow the exchange of data and providing the tools necessary to run the installations of the accelerators and the experiments.
- Designing, implementing and supporting complete control systems or parts of such systems for the experiments and the accelerators based on selected industrial solutions.

All these activities imply the responsibility of the ICE Group in providing **reliable, efficient and fully functional** control systems.

Inside the ICE Group, I worked for the section EN-ICE-SCD, dedicated to the SCADA part of CERN's control systems.
EN-ICE supports PVSS as the SCADA tool of choice at CERN. SCADA tools address "slow" controls, which includes a human-being in the loop (as opposed to time-critical control, which is addressed by the other layers).
PVSS is a commercial product provided by ETM, which is a company in the Siemens Group. It was first selected by the LHC experiments in 2000, following an extensive comparison of commercial and open-source products. It has since been adopted across CERN as its strengths have become widely recognized. PVSS is:

- Scalable: it can be used to implement extremely large control systems.
- Robust: the design supports the containment of problems preventing failure of the entire control system. Redundant systems are also supported.
- Multiplatform: it works on multiple platforms, Windows and Linux being most popular at CERN. In particular, PVSS supports mixed-platform implementations, allowing Windows and Linux to inter-operate as components of a single control system.
- Open and extensible: application programmers may make use of the programming API as used by ETM.
- Modular: only the functionality required by a particular application need to be loaded and run.

PVSS provides the complete range of standard SCADA tool functionality, including:

- User Interface construction.
- **Data archiving**.
- **Alarm generation and handling**.
- Drivers for a wide range of hardware (the open specification also provides the possibility to write custom drivers).
- **Action and Event logging**.
- Graphical Trending and Textual Reporting.
- Access Control.

The JCOP Framework has been developed at CERN to extend PVSS in those areas important for CERN but hitherto not addressed by any other SCADA tool (e.g. Finite State Machine functionality).

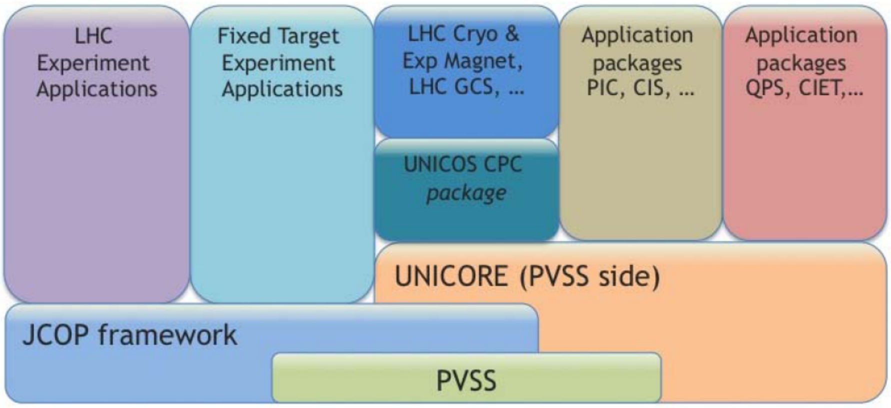

JCOP and PVSS, along with other frameworks and applications of the ICE Group.

PVSS runs on commodity off-the-shelf PC computers, running one of the CERN supported operating systems (Windows or Linux). For standard development, a modern desktop-class computer is sufficient to run PVSS. For production, mission-critical systems, where reliability is of high importance, CERN uses industrial PCs with redundancy assured at the hardware level (e.g. hot-swappable power supplies and hard-disks).

PVSS is provided by ETM as a set of feature packages. Some are essential as they form part of the core functionality of PVSS. Others can be installed to extend this core functionality.
In addition, EN-ICE provides a series of extensions to PVSS that extend the core functionality. Two of them were important in my work:

- **CtrlRDBAccess**: this is a PVSS Control Extension that allows to access **relational databases (RDB)**, namely Oracle. It is meant to be a functional replacement for the CtrlADO extension shipped with PVSS. It allows exchange of data with Oracle databases even if the ODBC driver is not present. It is optimised for high performance and assures perfect compatibility between Linux and Windows.
- **PL/SQL API for PVSS RDB Archiver**: a high level interface allowing extraction of the data stored by the **RDB Archiver** into any user application. It hides the complexity of the database structure from the application developers.

PVSS is architectured with a set of manager. Each manager is handling a specific function, and all are centered around the event manager (EV). Driver manager provide connection with external devices (PLC, FEC, VME, etc.) through a communication protocol like OPC (OLE for Process Control, a Microsoft protocol) or DIM/DIP (CERN's own protocol) to access devices through Ethernet for instance.

PVSS has also its own programming language named CTRL. It looks like C code, but without pointers and with an extended function library specific for PVSS use. It is an interpreted language, thus no compilation is required. A CTRL script can lie into its own file to be executed in standalone, or within the user interface panel tools provided by PVSS.
Devising a user interface with PVSS is really simple as everything can be done graphically with drag and drop actions.
It took me a few weeks to get in touch with all PVSS world. In addition, as the EN-ICE-SCD is providing courses on PVSS, I followed one.

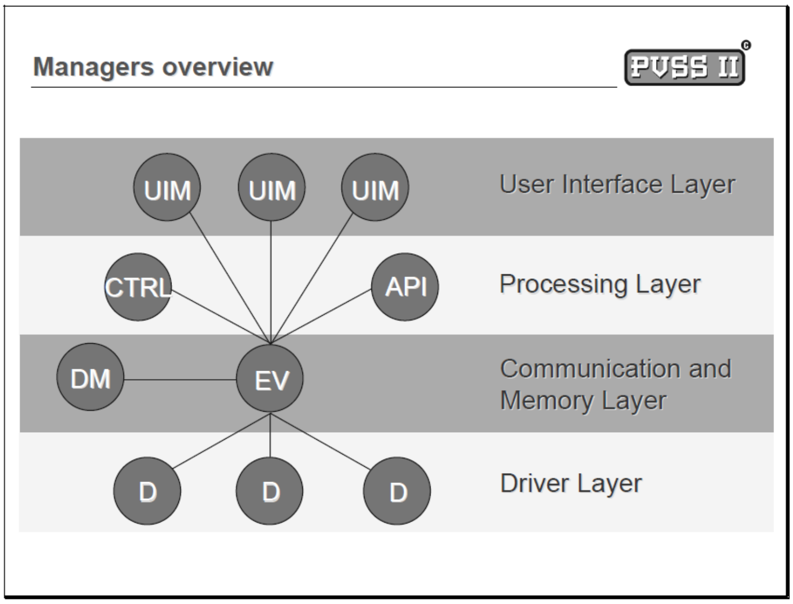

DM: Database Manager

Data coming from a device enter a PVSS system through a driver manager. It then go to the event manager who decide to send the information to concerned managers. It will for instance send it to **the database manager for archiving**, and to a CTRL or UI (User Interface) manager because they subscribed to this information coming from this device.

Underneath reside two main concepts of PVSS. The first one is that every data flowing from one manager to another is an event. Then it follows the event programming paradigm with all its usual attributes, like asynchronous communications and locking functionalities on data for instance. Also, a manager can register at the event manager that it want to be kept up to date when new data is coming from a specific input. This allows for instance to have a reactive user interface without having to constantly loop over an input to track and react to its changes.

The second important concept of PVSS is about how to access specific input and output, as well as virtual data actually. The device data in the PVSS database is structured as, so called, Data Points (DP) of a pre-defined Data Point Type (DPT). PVSS allows devices to be modelled using these DPTs/DPs. As such it allows all data associated with a particular device to be grouped together rather than being held in separate variables.

A DPT describes the data structure of the device (DPTs are similar to Classes in Object Oriented terminology) and a DP contains the information related to a particular instance of such a device (DPs are similar to Objects instantiated from a Class in Object Oriented terminology). The DPT structure is user definable and can be as complex as one requires and may also be hierarchical as shown in the example on the right.

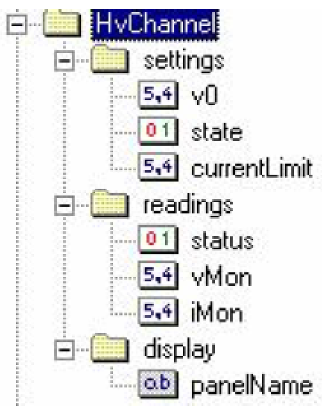

In this example we see a DPT representing a simple high voltage channel. This has a set of read (readings) and write (settings) parameters as well as associated display information - in this case the name of the panel (user interface) to be used associated with it. Each of the folders (settings, readings and display) as well as the individual parameters (e.g. v0, vMon, panelName) are called Data Point Elements (DPEs) and are user-definable, i.e. the structure can be fully defined by the user. After defining the data point type, the user can then create data points of that type which will hold the data of each particular device.

Once instantiated each DPE got some "config" elements hierarchically associated to it. These configs are used to enable specific behaviours of a DPE such as **alarm handling** or **archiving**. They also can hold attributes of that DPE sush as its current value or its hardware address.

In the previous example the data point elements mapping to the device data are only of types **float**, **boolean** and **string** but several more data types are available. In particular **dynamic arrays** of the simple data types, like **dyn_int, dyn_float, dyn_string**, etc.

PVSS allows an extremely flexible definition of an alert. The JCOP Framework implements the consensus of the four major LHC experiments on these matters. Alerts are configurable on any numeric or boolean DPE. The user basically has to say the different ranges in which a value can be. An alert can cross several range and of course went back. An alert can be configured as acknowledgeable.

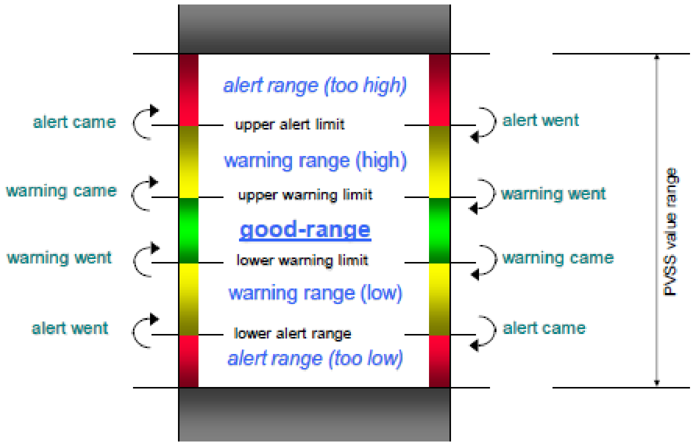

Typical alert range configuration

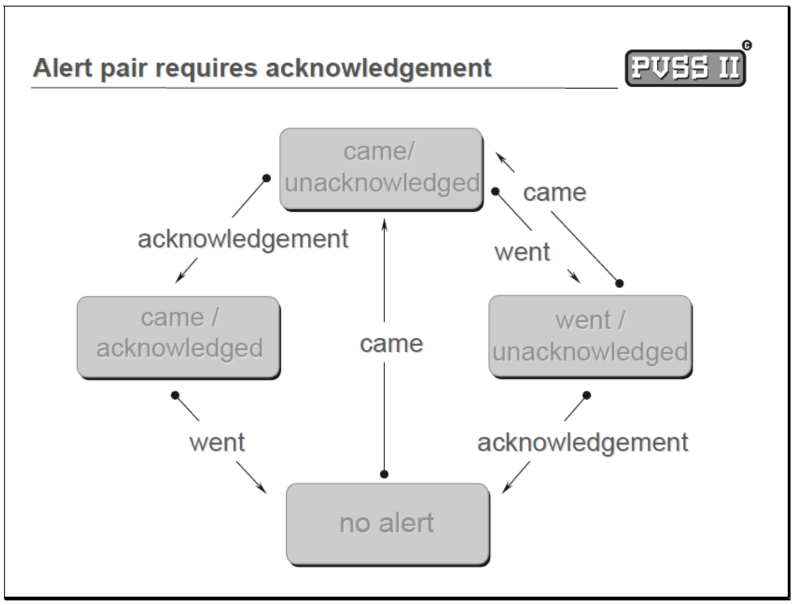

Alert state map

Finally, it is worth mentioning that a PVSS system can connect to another PVSS system thanks to a distributed manager. This allows to interconnect different control subsystems to a more global PVSS project that will be able to control an entire machine just from a few user interface panel used in the control centre. At CERN, machines operators are not experts in control system nor in the machine itself. It can be a student on shift for instance. For operators, controls and possible actions must be really simple interactions, but still able to control a very complex machine with them.

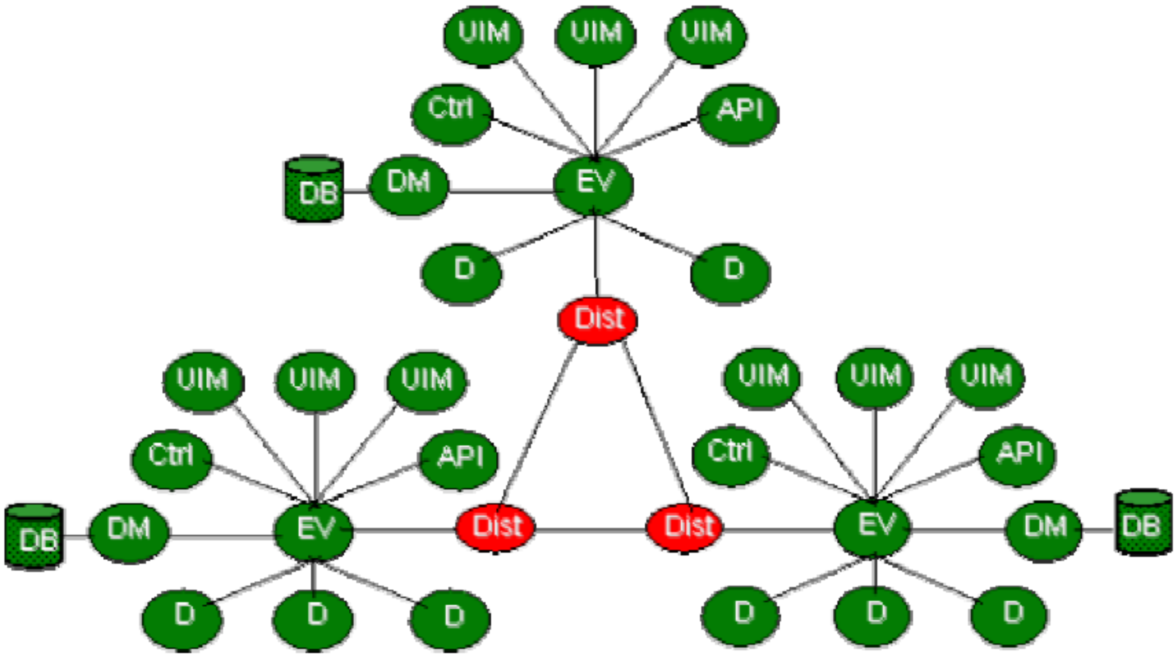

A PVSS distributed system with three PVSS projects.

iii. RDB, the Oracle Archiver

PVSS, as a SCADA system, can acquire data coming from various sensors and archive them. It can also archive the different commands sent to the hardware devices.
For CERN's applications using PVSS this archiving functionality allows to keep trace of the contexts of every controlled system. These data can afterwards be used for a post-incident analysis. They could also be used to statistically study how some systems performs and help to optimise their efficiency.

To do so, at the beginning PVSS was able to archive data in files on the local hard disk. But the control system working group at CERN decided it would be better to store these archives in the central database of CERN. The main reasons to do that were because central database is capable of ensuring an infinite archiving space, distributed over a cluster of PC, redundant and automatically backup.
Moreover, CERN is in partnership with Oracle. Oracle products are extensively used at CERN and are perfectly adapted to its scale.

Data collected in this archiving have nothing to do with physics data acquired at every particle collision. Here it is just about voltage values, temperatures, pressure, commands, etc. of all systems used to operate accelerators and detectors.
Are also archived all alerts that could be automatically raised in the different layers of control. An important thing here is the status update of these alerts made by an operator in the control room. By archiving also these changes, one is able to track down what have been the human decisions taken through time in case of an incident.

Even though the data acquisition rate is really smaller than the physics data taking rate, the Oracle archiver of every PVSS system has to be able to deal with up to 3 000 value changes per seconds.

The manager linking PVSS and Oracle is named RDB (Relational DataBase). This manager is directly communicating with the database manager of PVSS.
This manager, more than just communicating with Oracle, has also to buffer data in order to achieve the required performances. It has also to be capable of handling the situation when the database is not accessible or saturated. To do so, it can temporally keep data locally on the hard drive and send them as soon as possible to the database.

All data have to stay coherent with the original values. And RDB must never lost any data.

One could think that as PVSS is a professional and commercial control system, ETM should normally furnish a final product perfectly tested. But this is not the case...
Since the beginning of the cooperation between CERN and ETM, CERN never stopped to ask for more features. It also contributed a lot to the development of PVSS and tracked a lot of bugs. But ETM is a small company (even though Siemens hold it) and sometimes they struggle with CERN needs.
But PVSS has been chosen because it was the SCADA solution that fulfilled the most CERN requirements. There is no other solution than PVSS for CERN.

The Oracle archiver is a good example of this. It has originally been made only for CERN. And CERN contributed tremendously to this manager, especially by offering ETM resources that they would not been able to get by themselves. This is one of the reasons why ETM/ Siemens became an Openlab partner.
CERN continue to assist ETM for their common projects. Through the Openlab project first. But also with a tool dedicated to perform automatic functional testing of RDB. And this has been my work topic for this internship.

iv. Requirements

This functional testing suite has to collect as much as possible of RDB behaviours in various context and situation of usage, and do a report when an inconsistency appears.
A typical testing situation is by archiving a specific type of data (int, float, status bits, etc.) and check that when you retrieve it, it got the expected original value, and come with the right timestamps and other config fields.
Another test scenario is to trigger an alert and check that all steps of its life-cycle are properly archived (Went, Came, Acknowledge, Comment, etc.).

It should also be possible to test different archiving contexts. For instance, PVSS has some advanced functionalities allowing to filter, modify numerical values (translation) or reduce their occurrences (smoothing). Some of these features can have a direct influence on the archiving process.

Eventually, the archiving context could also be linked to the operating system (Linux or Windows), as much as with the connectivity to the Oracle environment.

This test suite will be used each time ETM will publish a new version or a new patch for PVSS. It should ensures to CERN PVSS users (experiments and accelerators) that RDB is performing and behaving as it should. RDB is considered as a critical component in control system topologies adopted by CERN. It has to be tested and certified that it will be a reliable component under CERN requirements.
A misbehaving of RDB during a run is not acceptable at all.

III. RDB Tester

i. Method and design

At the beginning of the project there were not really any objectives nor methodology proposed by my supervisors to tackle the problem. There was just a bunch of ideas about what such tester should be testing ideally.

I chose to divide the work in several prototypes. I defined them like this:

- The first prototype would be used to get in touch with RDB and understand superfically how it works.
- The second prototype would be used as a first architecture attempt able to simply automate tests.
- The third prototype would be used to rework the software architecture so that the tester become more abstract and generic (adaptive for all kind of tests). It should also implement some functionalities regarded as interesting in the two previous prototypes.
- The fourth prototype would be used to implement advanced functionalities to the tester.
- At last, the final implementation, ready to be used on a regular basis, and able to report bugs to ETM.

I chose here an iterative development process for two reasons. First, I did not known at all PVSS system and its programing language. So I had no idea on how to do things with it and its programming potential.
Secondly, the objectives could not be more defined as nobody known which tests would be useful to do. Requirements as given just before where not totally known at the beginning! The only thing I was certain was that I would have to identify and target the sensible points of RDB during the devising of the tester. So it was really about writing some code, testing if it works correctly and simultaneously testing how RDB works and if it is doing the right thing.

It is a methodology I am used to deal with. Because I usually do like this for most of my works, especially when I don't know where I am heading to. This methodology certainly imply to rewrite most of the code, if not everything, at least two times during the devising. These rewriting are usually due to a complete reviewing of the software architecture. In the end it ensures a clean code, reworked many times, without any inessential part, and correctly adapted to solve the given problem.

It is also a way to conceive software engineering where the programmer know from the beginning that it is not such a huge problem to rewrite completely some functions, and it is certainly not a waste of time too. Because it actually already been planned from the beginning. And thus the programmer is not bothered with very old functions or modules that become some old demons as time pass by.

ii. RDB workshop: performance testing

After two months we worked on something a bit off the RDB test suite. EN-ICE-SCD held a one week workshop about testing and improving performance of RDB. One developer from ETM came at CERN, as well as three Oracle experts.

CERN, ETM and Oracle are doing this kind of workshop regularly. The main advantage is to have someone from ETM able to modify RDB and PVSS code directly and following the advice of the Oracle experts.
CERN provided a cluster of 62 machines, all on a dual quad-core processor (so eight cores per machine) and 16 Gb of RAM. In addition we benefited from two Oracle cluster: one of three nodes, and the other of six nodes.

The point of this workshop was to stress-test RDB as much as possible and to judge its scalability performance. We set up three PVSS projects per machine in the cluster, each having their own RDB manager. All projects were configured to generate random values and alerts at the maximum pace. We were then monitoring the throughput of RDB as well as its behaviour in time. We especially targeted its buffers. Because if all projects cannot send their data simultaneously to the Oracle nodes without buffering a lot, it would then means that PVSS lack of scalability. Oracle nodes were ensured to be sufficiently powerful to not be the bottleneck.

We found out that the throughput of one PVSS system is 13 000 value changes per second without archiving, 3 000 value changes per second with archiving, and 12 000 alerts per second with archiving.

My duty in this workshop was to assist one of my supervisor in the administration of the Linux cluster. It concerned managing the automatised deployment of all 186 PVSS projects. But also managing the binding of processors' cores to PVSS managers.
This has been possible by using SSH (Secure Shell) and Bash scripts. Thanks to these scripts we were able to dynamically and quickly reconfigure the cluster to try different combinations of parameters.

I also managed the deployment of a remote monitoring program to visualize the cores usage, which helped to decide how to reconfigure the cluster and also to identify which manager (binded to a core) was the bottleneck.

Resources usage monitoring of a small part of the 62-machine cluster.

iii. Prototypes

The first prototype has been made in a few days. It served at understanding how the different internal buffers of RDB are working. It also served at picking up which set and get functions I should use in the tests to ensure that PVSS is really accessing the Oracle database and not just a local memory cache.
Functions are: dpQuery, dpGetPeriod and alertGetperiod.

The dpQuery function is actually using a kind of SQL language adapted for the PVSS data model. *GetPeriod functions are inherently using dpQuery, but they have to be validated too because PVSS users are more likely to use them rather than the dpQuery function. And it is actually the case of the trending tool of PVSS. The trending tool is displaying graphs of values over time. It is constantly used in the control rooms by operators. Hence, we really need to be sure that these graphs are displaying correct information!

The second prototype aimed at being able to do the first automated tests for all the basic data types (char, uint, int, float, bool, bit32, string, time). It also served at trying a first approach to accomplish the test procedure: Set - Wait flush - Get - Test.
Prototype 2 introduced also a first attempt for implementing a generic test description and reporting system using only PVSS data-points to store data.

This prototype spotted the first eight bugs in RDB found by my tester. And this only by doing a few restricted tests... (not extended at all possible test scenario).
Eventually I extended the test case to the dynamic types which are arrays of basic types with a variable length. This spotted an important bug in the handling of every dynamic data-point in RDB.

Only one of the bugs has been reported to ETM to avoid reporting possible false bug. As my program was still very young, I could not ensure bugs could not come from my code. The bug we reported concerned a behaviour of RDB's buffers that were randomly freezing during a long period (longest measured: ~10 minutes), whereas RDB should have flush them in any case after a few seconds.

ETM claimed this bug has also been reported by other clients. But we were the first ones to discover it. It helped that other of their clients confirmed it because ETM was not able to reproduce this (random) bug. When multiple similar bug reports arrived, they accepted this bug exists, and started to work on it. In the end this bug has been fixed by one of my colleague involved in the Openlab program about five months later. This colleague was having the rights to modify the code of RDB to do so.

These first bug experiences with RDB and ETM started to frighten me about the quality of RDB and how difficult it is going to have them corrected... It changed my inner view on the project, and I started to be less optimistic and candid about it. It influenced the new software architecture adopted in the following prototype.

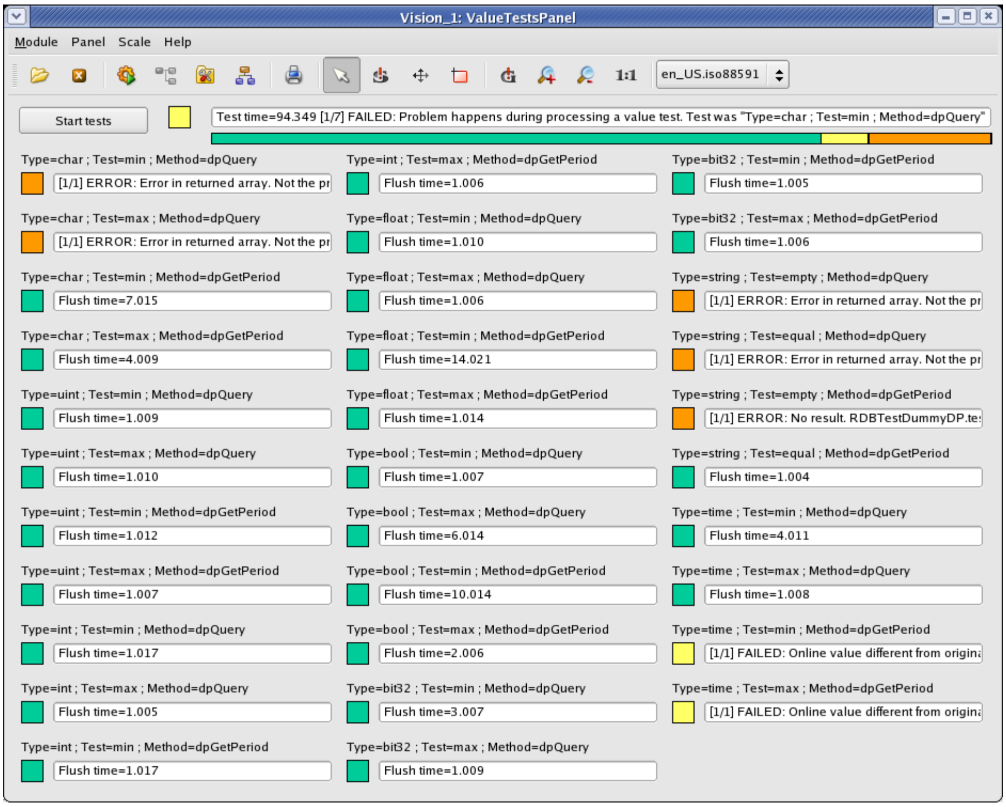

The first version of the control and report panel of my tester, prototype 2.

<u>The third prototype</u> goal was to rework completely the architecture of the test suite. The purpose was to make it more versatile and efficient. Versatile to be able to realise any kind of new test without having to edit the tester code in itself, but just by adding a new test description and model. Efficient to parallelise as much as possible the execution of all tests.

I adopted the following architecture:

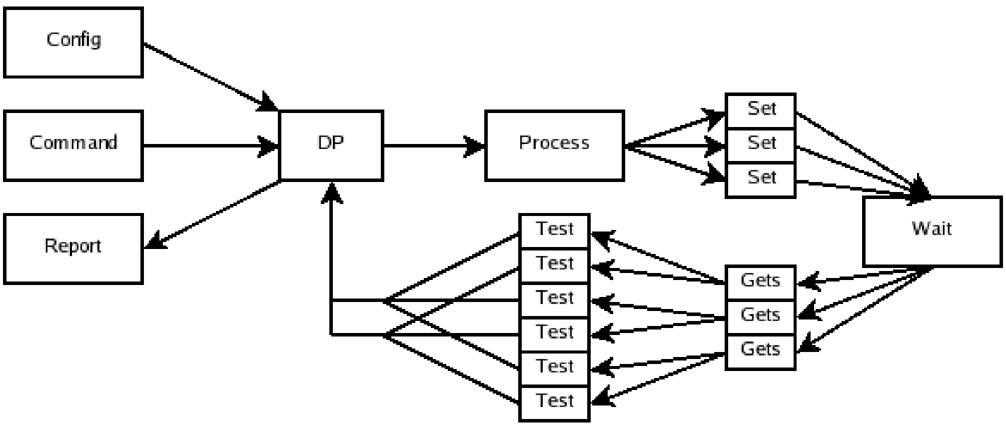

Prototype 3 schematic. DP: Data-point.

For each test, the tester retrieve the configuration of the test case from a data-point. It then generate a CTRL script by mixing this configuration and a model file. It then execute this script which implement the test procedure (Set - Wait - Get - Test). The script then report the result in a data-point and an overview panel can display it.

For some technical reason the wait step has to be synchronised between all running test scripts. This wait mechanism that allows to wait a data-point to be flushed in the Oracle database does not exists by default in PVSS. There is no function for that. So I had to create my own mechanism which query directly the SQL schema without going through any PVSS function.
Later, this functionality might also be used in a system overview tool project made by some of my colleagues. The point will be to monitor the status of the control system in itself. And one problem with RDB is that there is very few ways to have informations about the current state of the manager. So my wait-flushed module will help this tool to check that the Oracle Archiver is still alive.

To be able to realise the parallelisation of test scripts I had to device a sub-component. PVSS has some threading functions for CTRL scripts. Unfortunately ETM implemented its own scheduler rather than using the operating system one. And this scheduler is having a problem with RDB and the dpQuery function. When several threads are calling the dpQuery function, it blocks completely. I don't know the reasons behind, but this bug was already known.

I used the fact that PVSS managers are real and native processes. It is possible to add tenth and tenth managers to a PVSS projects, especially CTRL managers. A CTRL manager is merely a manager executing one CTRL script. I just needed to control which script these managers were using. I devised a "script pooler" following this architecture:

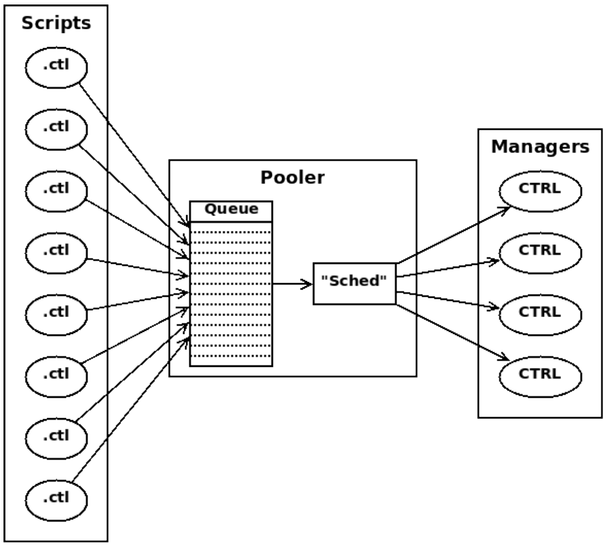

ScriptPooler schematic. ".ctl": CTRL script files

This component receives paths to scripts waiting to be executed and put them in a queue. As soon as a CTRL manager from the "manager pool" is available, the scheduler of the component pick up a path in the queue and send it to the CTRL manager.
It seems quite simple. But this mechanism was not existing neither in PVSS than in JCOP. And I hit some problems with mutual exclusions (mutex) made on some data-point used by different managers at the same time.

Debugging these mutex problems became pretty interesting by using PVSS normal tools. As a SCADA system it can plot over time the value of a variable (a data-point). With this tool I was able to track rush conditions on my mutex... up to a millisecond scale. It is quite rare in programming to be able to plot the content of internal variable with such precision. I was then able to tune precisely my interlock mechanism.

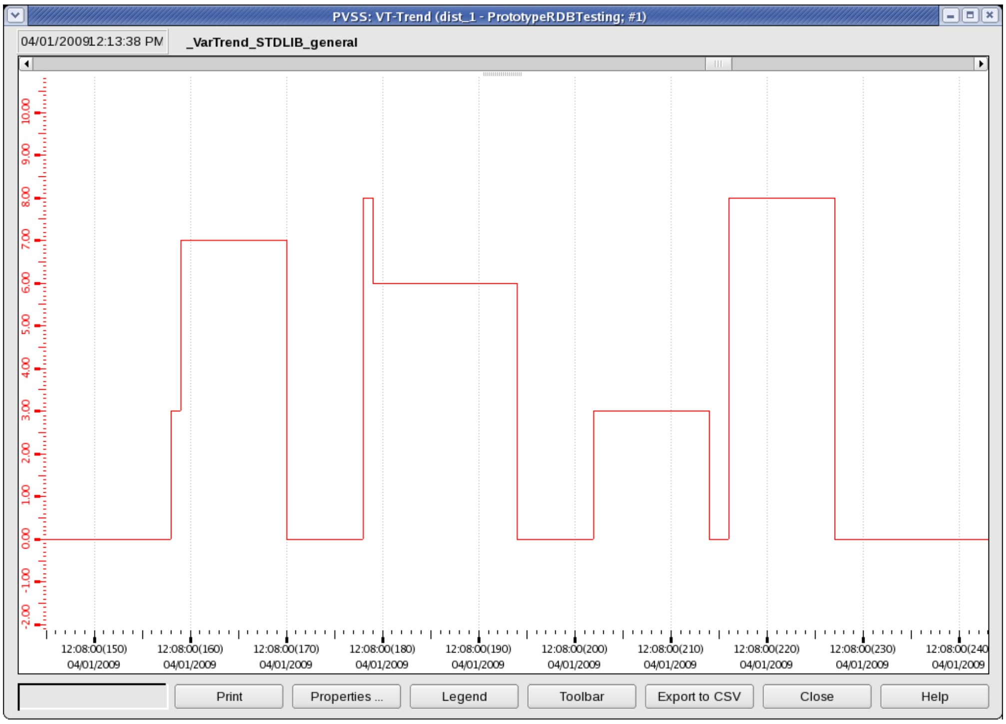

X axis: time in ms. Y axis: the process number trying to get a semaphore. Date is wrong.

On this graph you can see that glitches of a few milliseconds happen sometimes. It is when a process is trying to get the semaphore but another process "stole" it right after. This is due to the delay between the availability test of the semaphore and its actual modification.
Modern processors implement atomic operation to do this kind of thing. But I could not access from a CTRL script. Thus, I had to find a workaround by adding verification delays.

It was really interesting to debug a PC program with a method usually used in microelectronics for instance.
More than offering real parallelisation to CTRL scripts, this component allows to run each test in a dedicated script. A .ctl file is equal to one test case. Thus, when reporting a RDB bug to ETM, one can just attach the generated script to the bug report. It simplify a lot the communication with the support team of ETM. It also isolates the faulty code and reduce the possibility that the bug come actually from something else than RDB.

The last goal of this prototype was to implement every kind of possible test known until now:

- Basic types
- Dynamic arrays of basic types
- Alert scenarios
- Bit status field (control bits associated to each data-point to give information about its viability, filter it has been through or any other boolean information coming from the hardware device with a value)

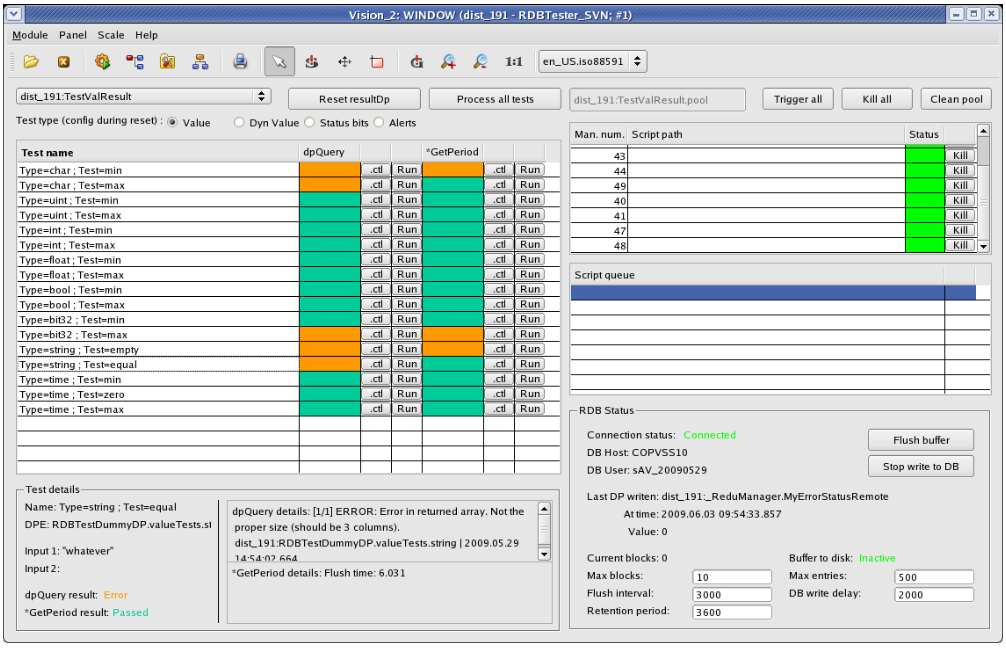

The new overview panel to control the tester. Left side: test results. Right side: script pooler.

This third prototype started to spread within the ICE Group. One of my colleague, another Openlab student, started using it to write his own test cases. His work topic is about... completely rewriting the RDB manager!
At the end of this prototype I became much more enthusiastic about it. It is pretty rewarding to know that my test suite will be used to ensure the quality of the future RDB manager. Not only by reporting error with it, but also by using it during its devising. I used inputs and remarks from my colleague to start thinking about what I will need to do in the fourth prototype.

To be able to ship my test suite out of my development environment I needed to package it as a JCOP component. These components are installable on any PVSS projects thanks to a tool made previously in my section.

We also started to report RDB bugs at ETM. Out of 166 test scripts, seven "fundamental" bugs has been reported, concerning only basic types. Which means that we did not reported every bug stupidly. Instead we tried to gather the bugs that could possibly had the same source. A patch released a few months after contained corrections for four of them.
Also, all alerts test scripts were in the end not working properly. Which is not surprising considering they are much more complicated than the other ones.

<u>The fourth prototype</u>, which leaded to the final version, was about adding advanced functionalities to the test suite.
In the end the only functionalities added were:

- a panel to ease the creation of a RDB schema. This means that the structure of the database used by RDB is now created from this panel, as well as all necessary user and rights on these schema. The panel also handle the automatic configuration and synchronisation of the PVSS project to use the newly created schema.
- a panel to automatically deploy and configure a list of CTRL manager for the script pooler. This is done thanks to a "recipe" system. Thus, when deploying the managers, the user can choose to configure how many managers he wants and for which test case it allocates them, according to the kind of automatic test procedure he wants to do.

These two panels are intended to be used at the installation of the RDB tester component.

Alerts script have been enhanced. But I faced a particular problem that would need another complete rewriting of the tester to address it. In brief, the script pooler has been made to parallelise execution of test scripts. But with alerts test scenario, they actually need to be started sequentially...

This is because an alert is a complex state with some hysteresis effects that change when the value associated to it is changed. For the other tests, type testing, it is just about setting one value at tracking it with its timestamps. With alerts, if scripts are running concurrently, they modify each other state and in the end no script can found the expected results.

Working test scripts for alerts, when executed sequentially, are mostly all positive.

Otherwise, ETM released some new patches since the last time, which make turn green all test cases related to basic types. Hurrah!

iv. Results

Now that all basic types test cases turn green, it is a first relief. Even though the other cases are still not solved (and not reported yet actually, but they should be before my contract finish). This is a definitive proof that working on this test suite improved RDB's quality in a few months.

Out of 34 test cases for basic types, we started with at least ten having problems (I cannot precisely say how much because their tracking system is too messy to find the relevant information and tickets). Plus the flush buffer bug that was really annoying for me at the beginning.

The most critical need was to have the four LHC experiments using a version of PVSS which does not show any bug with their current usage of the Oracle Archiver. This goal had to be achieved before the LHC restart this year. Otherwise the experiments would have run during more than a year with a bugged version of PVSS.

And this goal had been achieved thanks to a lot of efforts from ETM and EN-ICE-SCD. I think I did a fairly small part of it, by spotting bugs before experiments experience them. And also with my modest contribution during the RDB workshop which validated the performance and efficiency of RDB, at least for the current usage we have of it.

In the end no work has been done to facilitate the use of the program by my Openlab colleague. Firstly because ETM postponed the rewriting of RDB. Secondly because I have been allocated to other work topic that I will not present in this thesis, apart in the Gantt diagram.

v. Discussion

ETM might have postponed the rewriting of RDB for several reasons. But it is possible that the progress made thanks to my tester leaded ETM to think in the end it is not so urgent, now that CERN is complaining less and less about RDB.

Although RDB is getting better and better since its creation, I think it still need a complete rewriting. In addition CERN is complaining less about RDB because now the LHC restarted, most of the control system developments, testings and commissioning are finish or frozen for one year.

I have been able to know more about the architecture of RDB. I had to sign a Non Disclosure Agreement for that. I would say there is several things that need to be cleaned up. They never scraped it and restarted it from scratch. It's all made of patch and various piece of code written by many different people with no homogeneity at all.

If I get a fellowship at CERN and continue in my current section, I already know I would still have to work on the RDB tester, of course, but also on some other part related to RDB. On the other side of RDB there is some SQL procedures for the retrieval of data asked with the dpQuery and *GetPeriod functions. Some of these procedures would actually need... to be written!

We, my supervisors and me, had some discussions about the future of the test suite. It got the potential do more than functional testing of RDB. It could actually be used to test other things than just RDB. Starting with alerts, but by the normal process this time, not the archived ones.

To address the problem with parallelisation and alerts, there is two possibilities. Either do a mode where scripts can run sequentially. Or create the data-point used by the script at each run, and with a random name, so that each test case will affect only his data-point. The latter solution need a deep reworking of the configuration data-point part of the tester (which seriously need a reworking any way).

Another possible future path for the tester is to become a development tools for test cases, as my Openlab colleagues would like to use it. This would be most likely needed if the test suite become a more generic tester.

Other suggestions has been made, but most of them would need a dedicated tools to do them actually:

- Delete and recreate data-point while value are changing in order to test the flush buffer interval side effects as well as testing the ability of RDB to associate "ghost" in the database.
- Test the change of alias or comment for data-point.
- Change the structure of a data-point type and see if older data-points can be recreated.
- Test a possibility of RDB to use different archive group within the database.
- Trigger the table switch for values and events (alerts).
- Test trending.
- Automatic test of the alert screen.
- Use filter (wildcards) with the dpQuery function.
- Long-term stability testing.

And many other test cases and scenarios are imaginable...

Conclusion

This internship at CERN has really been fruitful and pleasant. First of all, I have been able to live one year inside this enormous laboratory. Working conditions really pleased me here, and especially in the EN-ICE-SCD section. Moreover, I had time and opportunities to learn more about high energy physics. I was free to go to some of the nice conference held at CERN. And I have been able to visit impressive place in the laboratory.

One of the greatest thing at CERN is that you can meet a lot of people from different cultures, or professional origins. But everybody speak English, all with a different accent. Most people are scientist and are really easy to talk to. You can meet young physicists, well known physicists, retired physicists (they are still at CERN...) and even Nobel Prizes. The student cohesion is also predominant.

The environment around and at CERN is great. Close to the Jura and to Geneva with the Alps not so far way. It makes CERN a pleasant place to be, to work and to live. Services available inside CERN are really appreciated (restaurants, cafeteria, bike renting, hostel, etc.).

The main topic on which I worked on always appeared to me interesting and motivating. Even though it is not exactly what I asked when I applied for a Technical Studentship, that is embedded electronics. I also appreciate the irony of the situation of me working in a control systems group for one of the most demanding industrial application in the world. Whereas I failed several time my control module at ESEO. It is completely different of what we had to learn, but still you feel the same feeling of connection knowing "these stuffs" are linked to hardware. More over when the hardware is these gigantic machines at CERN.

By coming at CERN I known what I could find here because this is already the second time I find an internship here. Now I am currently waiting for the response to a fellowship application at CERN. Most people are telling me that I maximised my chances to get the position. But I know that sometimes CERN can be really though with positions. Just for one reason I cannot control, it can be refused. And this is the "law of CERN".

This does not afraid me, and after four internships all in various place (a city-hall, a school lab, a company and a research centre) I am know quite sure that I want to start my professional life in an environment like CERN. CERN is among the biggest ones, but there is plenty of other exciting projects out there!
I certainly don't want to go back in a company. I would consider it as a personal depreciation. And I have not been convince by public fonctionnaire or school system.

```
// #uses "fwRDBTester/fwRDBTester.ctl"

void main () {
  int err;
  dyn_string exceptionInfo = makeDynString();

  string dpe = dpSubStr("RDBTestDummyDP.valueTests.int", DPSUB_ALL);
  int testedValue = maxINT();

//  fwRDBTester_initTest(true); // Standalone mode
  fwRDBTester_initTest(false, "dist_191:TestValResult", "DpQuery", 6);

  // Set
  time timed = getCurrentTime();
  err = dpSetTimedWait(timed, dpe, testedValue);
  if(fwRDBTester_checkError(RDBTESTER_SET_ERROR, err, timed, dpe, testedValue))
    return;

  // Wait
  time flushTime;
  err = fwRDBTester_wait_onEvent(dpe, timed, flushTime, exceptionInfo);
  if(fwRDBTester_checkError(RDBTESTER_WAIT_ERROR, err, exceptionInfo))
    return;

  // Get
  dyn_dyn_mixed tab;
  time from, to;
  from = timed;
  to = getCurrentTime();
  string timerange;
  sprintf(timerange, "TIMERANGE(\"%04d.%02d.%02d %02d:%02d:%02d\",\"%04d.%02d.%02d %02d:%02d:%02d\",1,0)",
          year(from), month(from), day(from), hour(from), minute(from), second(from),
          year(to), month(to), day(to), hour(to), minute(to), second(to));

  string sql = "SELECT '_online.._value', '_online.._stime' FROM '" + dpe + "' " + timerange;
  err = dpQuery(sql, tab);
  if(fwRDBTester_checkError(RDBTESTER_DPQUERY_ERROR, err, from, to, dpe))
    return;

  mixed onlineValue;
  if(lookupFromDpQuery(tab, timed, onlineValue))
    return;

  // Test type
  if(getType(testedValue) != getType(onlineValue)) {
    fwRDBTester_raiseError(RDBTESTER_TESTTYPE_ERROR, getType(testedValue), getType(onlineValue), onlineValue);
    return;
  }

  // Test equality
  if(testedValue != onlineValue) {
    fwRDBTester_raiseError(RDBTESTER_TESTVALUE_ERROR, testedValue, onlineValue);
    return;
  }

  // Test ok
  fwRDBTester_testPassed(flushTime);
}

bool lookupFromDpQuery (dyn_dyn_mixed tab, time timed, mixed &onlineValue) {
  if(dynlen(tab) >= 2) {
    for(int i=2 ; i <= dynlen(tab) ; i++) {
      if(dynlen(tab[i]) != 3)
        return fwRDBTester_raiseError(RDBTESTER_GET_BADARRAY_ERROR, 3, tab[i]);

      if(tab[i][3] == timed) {
        onlineValue = tab[i][2];
        return false;
      }
    }
    return fwRDBTester_raiseError(RDBTESTER_GET_NOCORRESPONDINGRESULT_ERROR, tab);
  }else
    return fwRDBTester_raiseError(RDBTESTER_GET_NORESULT_ERROR, tab);
}
```

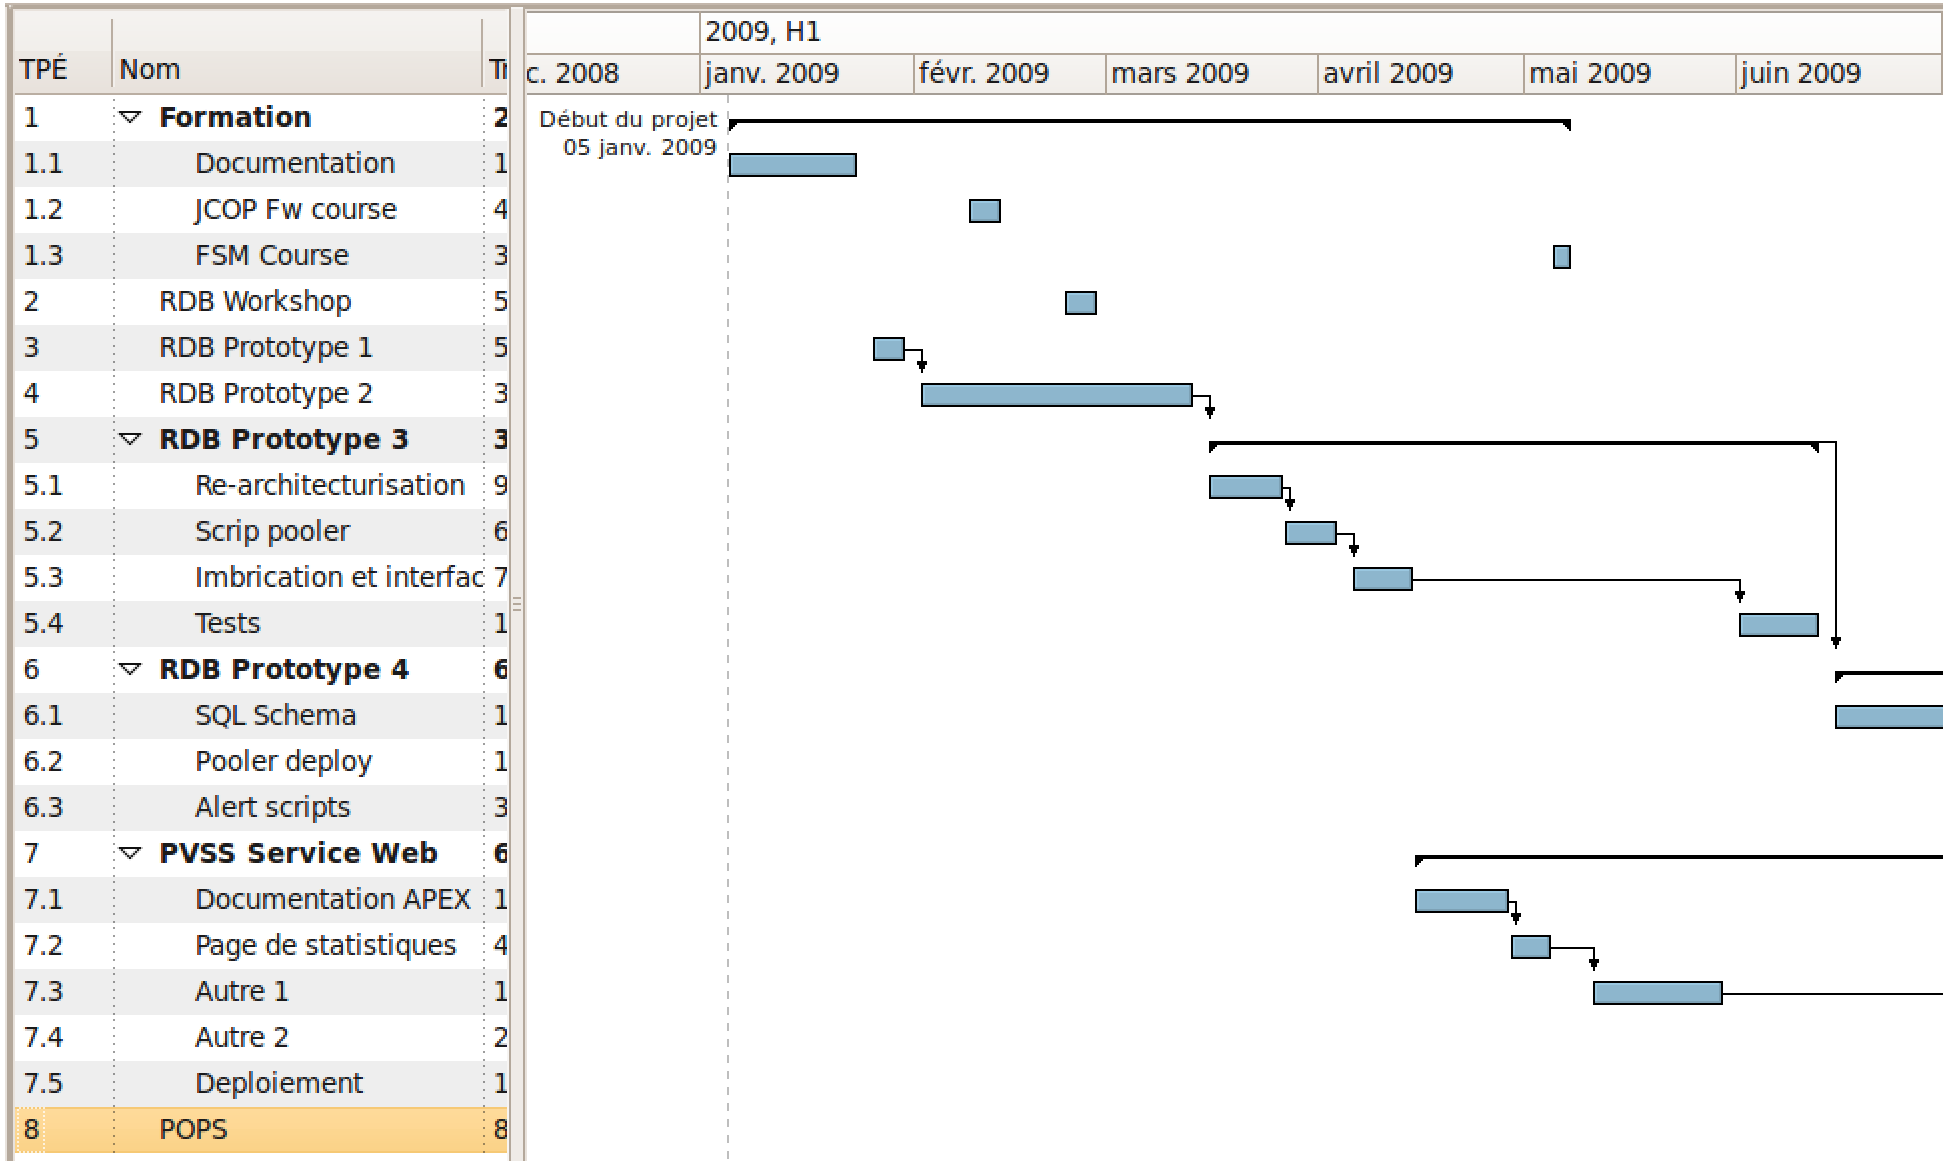
TPÉ
Nom
Ti
2009, H1
c. 2008
janv. 2009
févr. 2009
mars 2009
avril 2009
mai 2009
juin 2009
Début du projet
05 janv. 2009
1 Formation 2
1.1 Documentation 1
1.2 JCOP Fw course 4
1.3 FSM Course 3
2 RDB Workshop 5
3 RDB Prototype 1 5
4 RDB Prototype 2 3
5 RDB Prototype 3 3
5.1 Re-architecturisation 9
5.2 Scrip pooler 6
5.3 Imbrication et interfac 7
5.4 Tests 1
6 RDB Prototype 4 6
6.1 SQL Schema 1
6.2 Pooler deploy 1
6.3 Alert scripts 3
7 PVSS Service Web 6
7.1 Documentation APEX 1
7.2 Page de statistiques 4
7.3 Autre 1 1
7.4 Autre 2 2
7.5 Deploiement 1
8 POPS 8

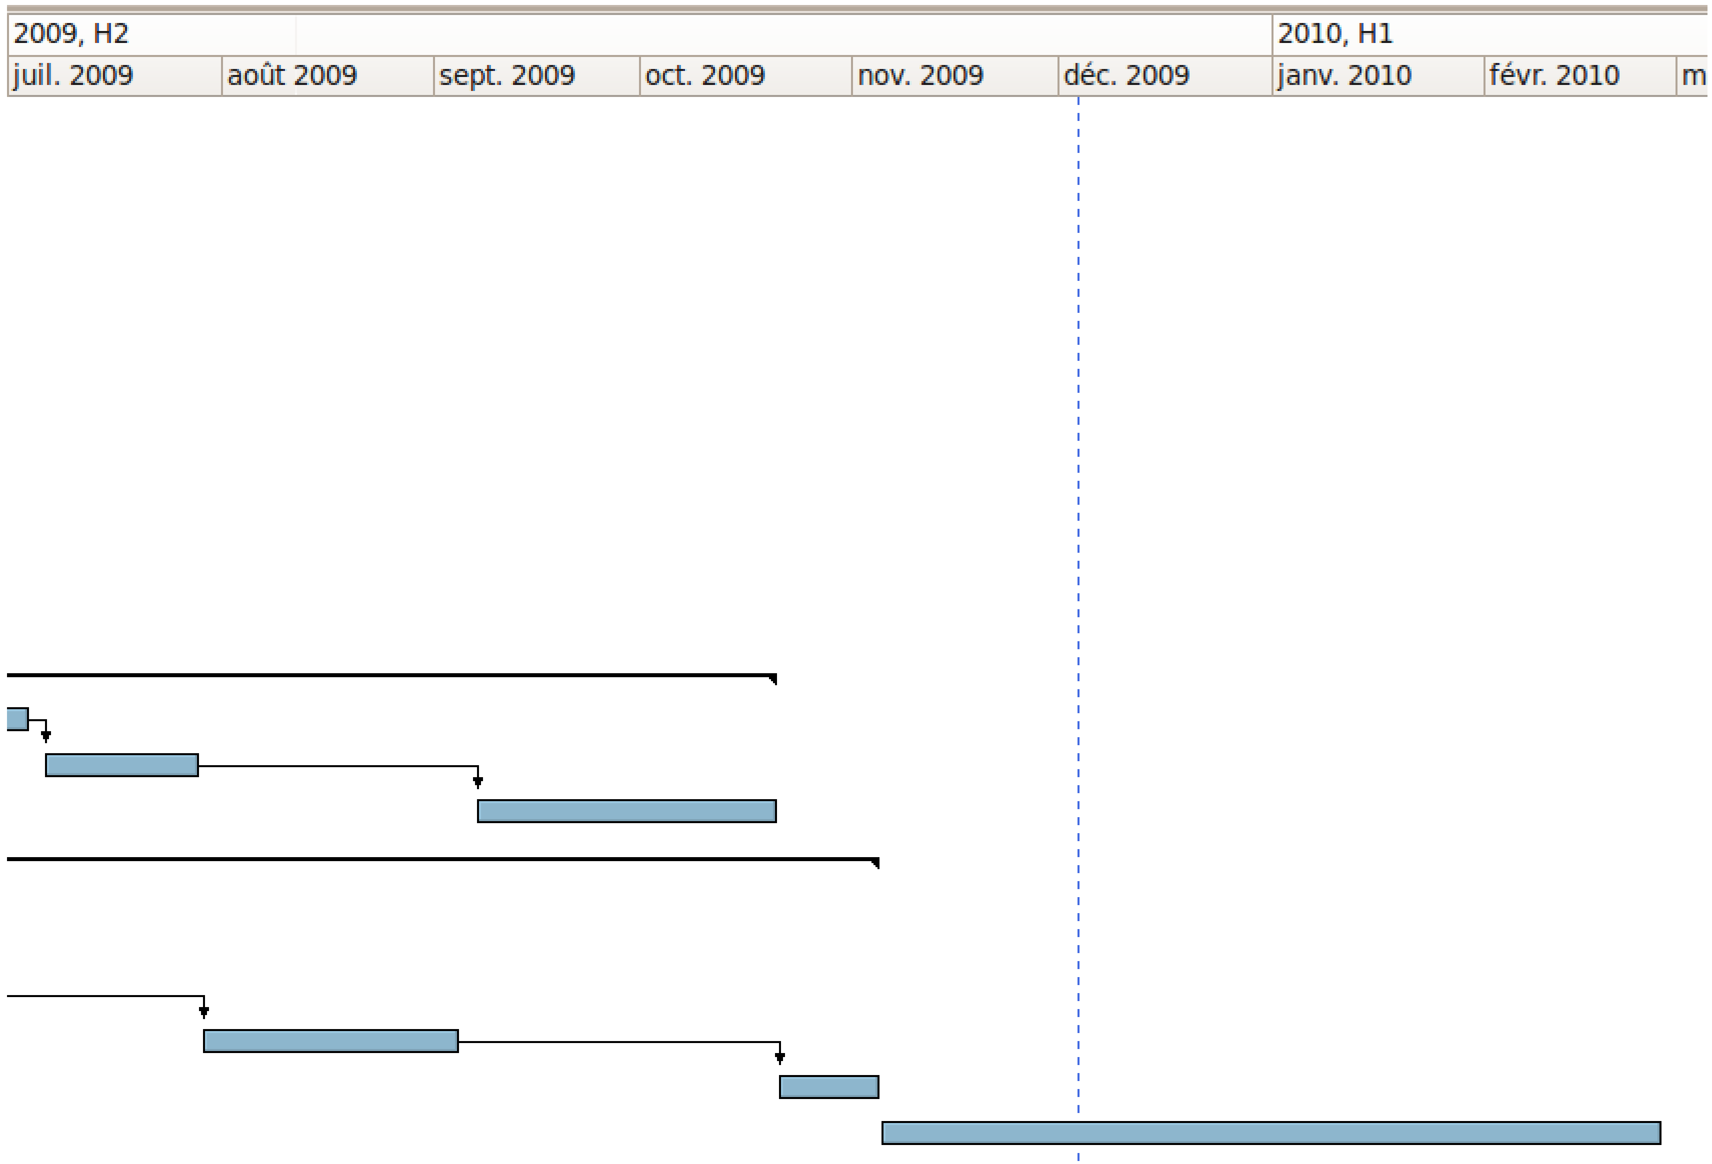
2009, H2
2010, H1
juil. 2009
août 2009
sept. 2009
oct. 2009
nov. 2009
déc. 2009
janv. 2010
févr. 2010
m

Axel VOITIER

4 rue Marcel Pinotie
51510 Fagnières, France
Phone (Fr) 0033 608 05 76 03
axel.voitier@gmail.com

23 years old
French
Single and non smoker
Mobile to work anywhere

To be an engineer in embedded systems and robotics

Education

- ⇒ 2007 / 2008 – Master of Science in Robotics at the University of Plymouth – *www.plymouth.ac.uk*
  Publication of accomplished works: *plymouth.tinigrifi.eu*
- ⇒ 2003 / 2008 – Master Degree to be an electronics and computer sciences engineer – ESEO, Angers, Fr
  Specialisation in embedded systems. – *www.eseo.fr*
- ⇒ 2003 – BAC scientific in France, equivalent to « A » levels in scientific subjects.

Knowledge

Languages: **French:** Mother tongue.
English: First level C and TOEIC 765 before one year of study in England.
Spanish: Able to read and write.

Robotics and Artificial Intelligence:

- ⇒ Neural networks, evolutionary computation (GA, GP), autonomous robots, fuzzy logic, speech and visual recognition, semantic web, information visualization.
- ⇒ **Software:** WorkSpace, MelfaBasic, Nuance, NetLogo, Cortex, WinCE.
- ⇒ **Hardware:** Bioloid Kit, anthropomorphic arm Mitsubishi RV-2AJ.

Computer skills:

- ⇒ **Environments:**
  - o **OS** – Linux (Ubuntu, Gentoo, Slackware, Debian and Mandriva), Windows: Administration, security, deployment, current use (console and graphic), maintaining, competitive intelligence.
  - o **Software:** Matlab, Netbeans, Gcc, WinDev, MySql, Oracle, SolidEdge, Subversion, Gimp, Blender.
- ⇒ **Computer languages:** C/C++, J2EE, UML, Design Patterns, VHDL/Verilog, SQL, XML, XHTML, PHP...

Electronic skills:

- ⇒ Analogue, digital and HF electronics, programmable devices (FPGA, ASIC, µC)
- ⇒ **Software:** Cadence (Analog Artist, Virtuoso, OrCAD), ModelSim, Agilent ADS, ISE, Quartus II
- ⇒ **Architectures:** dsPic, Colibri XScale, 68k, AVR, PIC18. **Controllers:** LCD (HD44780 et T6963C)

Work Experiences

- ⇒ 2007 / 2008 – Master project in robotics at the University of Plymouth, United Kingdom
  Bioloid – C on ATMega and XScale targets – Java – PhysX – Scientific researches.
  Devising of a walking controller for a bipedal robot, based on primitive movements.
- ⇒ 2003 / 2007 – Eurobot robotic association in ESEO – *www.robot-eseo.fr*
  C on dsPic target – CAN – Subversion – Project management – UML – Java.
  Devising of a standalone and intelligent robot that could compete for the Eurobot Cup
  - ⇒ President of the club in 2007, and also main board manager and quality chief.
  - ⇒ 2nd and 11th of France in 2005 and 2007, 1st and 2nd of Belgium in 2005 and 2006.
- ⇒ 2007 (3 months) – Internship at the European Organization for Nuclear Research (CERN), Geneva (CH)
  Python – SQL – Web development.
  Summer student. Integration of a complex workflow for approving scientific papers in the CERN digital library: *cds.cern.ch*
- ⇒ 2006 (2 months) – Internship at the Research Laboratory of ESEO, Angers (France)
  UML – Java – XML – Processors theory.
  Realisation of pedagogical software. It is an emulator for a simplified microprocessor, seen in a first engineering cycle lecture of the ESEO group.
- ⇒ 2004 (2 months) – Internship at the IS department of Châlons en Champagne's city hall (France)
  WinDev – WebDev – SOAP – SQL – HTML.
  Development of a set of tools integrated into the IS, allowing e-communications between city's schools and the city hall's educational department.
- ⇒ 2003 / 2007 – Anjougame and Salon Angevin du Multimédia associations in ESEO
- ⇒ 2001 / 2007 – Free software *www.ThacidLCD.ath.cx*, founder: C – Java – Io – LCD controllers

Interests and Hobbies

- ⇒ Music, Guitar, Art of cooking
- ⇒ Architecture
- ⇒ Mechanic Sports
- ⇒ Lego Technic and Mindstorms