The Large Hadron Collider

Lyndon Evans

1. Introduction

It is a great pleasure to be here today to pay tribute to one of the founding fathers of CERN. I am sure that Edoardo Amaldi would be very happy if he could see what the Organization that he fought so hard for has become today. It is the prime example of international collaboration in science with strong European leadership. The LHC will be the frontier instrument in particle physics for the foreseeable future.

It is generally accepted that the birth of the LHC was at the Lausanne Workshop in March 1984 where particle physicists and machine builders under the leadership of Giorgio Brianti (who incidentally was a protégé of Amaldi) got together for the first time. In reality the first seeds were sown much earlier and, as usual, Edouardo Amaldi had an important role to play.

At the beginning of the 1960's a debate was raging about the next step for CERN. Opinions were sharply divided between a "large PS", a proton machine of 300 GeV energy or a much more ambitious colliding beam machine, the Intersecting Storage Rings (ISR). In order to try to guide the discussion, in February 1964, 50 physicists from among Europe's best met at CERN. They decided to transform themselves into a European Committee for Future Accelerators (ECFA) under the chairmanship of Amaldi. It took more than 2 years more before the consensus was formed. On 15th December 1965, with the strong support of Amaldi, the CERN Council approved the construction of the Intersecting Storage Rings (Figure 1).

Figure 1 - Intersecting Storage Rings.

The ISR was an important step towards the LHC for a number of reasons. First of all, many accelerator builders were skeptical that the machine could be made to work at all. Previously, a few storage rings had been built but they all stored leptons with strong synchrotron radiation damping to cancel out the machine imperfections which would lead to beam loss. The ISR was the first proton storage ring to operate without any damping. Stochastic cooling was first demonstrated in the ISR. This later opened the door to the next step towards the LHC, the SPS proton-antiproton collider. On the experimental side, it is generally acknowledged that the ISR missed some fundamental discoveries due to our lack of understanding of how to build detectors for such an innovative machine.

The final step towards the construction of the LHC was taken when Carlo Rubbia proposed that the SPS could be converted into a proton-antiproton collider using the technique of stochastic cooling demonstrated in the ISR to accumulate antiprotons. In the SPS, beam conditions were even more difficult than in the ISR. The beams were bunched, with stringent conditions on the level of RF noise required for a good beam lifetime. In addition, the beam-beam tune shift responsible for driving strong nonlinear effects was an order of magnitude stronger than in the ISR. In fact, our colleagues at SLAC tried to cheer us up by performing an experiment at SPEAR where they reduced the energy and therefore the amount of synchrotron radiation damping and measured the maximum achievable tune shift. Their conclusion was that when they extrapolated to zero damping, the ppbar collider could never work! But the collider did work. A new generation of detectors built by veterans of the ISR worked superbly and the rest is history. On the machine side, a number of effects were elucidated, including the maximum achievable beam-beam tune shift and the importance of intrabeam scattering that could be fed directly into the design of the LHC.

2. LHC approval

Through the 1980's until close to his retirement, Giorgio Brianti led the LHC study team. However, the approval of the SSC in 1987 put the whole project into doubt. It was only the resilience and conviction of Carlo Rubbia that kept the project alive. Carlo argued that in spite of its lower energy, the LHC could be competitive by having a luminosity an order of magnitude greater than could be achieved in the SSC. He also argued that the LHC would be more versatile. As well as protons, the LHC could accelerate heavy ions at little extra cost.

The SSC was eventually cancelled in 1993 (the year I took over from Giorgio as project leader), after a series of cost overruns. This made the case for building the LHC even more compelling, but the financial and political climate in Europe at the time was not conducive to the approval of a large project. Germany was struggling with the cost of reunification and a number of other countries were trying to get to grips with the problem of meeting the Maastricht criteria for the introduction of a single European currency.

During the course of 1993, an extensive review was made in order to reduce the cost as much as possible, although an accurate cost estimate was particularly difficult to make since much of the R&D on the most critical components was still to be done, In December 1993 a plan was presented to Council to build the machine over a ten-year period, reducing the other experimental programs of CERN to an absolute minimum with the exception of the full exploitation of LEP. Although the plan was generally well received, it soon became clear that two of the largest contributors, Germany and the UK, were very unlikely to agree on the budget increase required. They also managed to get Council voting procedures changed from a simple majority to a double majority where more weight was given to countries with the largest contributions. In this way a few of the countries with the largest contributions could dictate Council decisions.

On the positive side, after the demise of the SSC, a US panel on the future of particle physics recommended that "the government should declare its intention to join other nations in constructing the LHC". Positive signals were also being received from India, Japan and Russia.

In June 1994, the LHC project was presented once more to Council. Seventeen member states voted to approve the project. However, because of the newly adopted double voting procedure, approval was blocked by Germany and the UK, who demanded substantial extra contributions from the two host states claiming that they obtained disproportionate returns from the CERN budget. They also requested that financial planning should proceed under the assumption of 2% annual inflation with a budget compensation of only 1%, essentially an annual budget cut of 1% in real terms.

In order to deal with the new constraint, we were forced to propose a "missing magnet" machine, where only two thirds of the dipoles would be installed in a first stage, allowing the machine to start up with reduced energy, eventually being upgraded when additional funds became available. This would have been a very inefficient way to build the machine, costing much more in the long run but saving some 300 MCHF in the first phase. The proposal was put to Council in December 1994. The deadlock with the extra contributions from the host stated was broken with the proposal that they would award the full indexation of 2% compared with 1% from the other member states. This time, the proposal was approved with instructions to start negotiations with non-member states for extra contributions that would restore the single stage construction.

There followed an extensive round of negotiations. The first country to declare a contribution was Japan. This was quickly followed by India and Russia in March 1996 and by Canada in December.

A final sting in the tail came from Germany in June 1996, who unilaterally declared that in order to ease the burden of reunification, it intended to reduce its CERN contribution by between 8% and 9%. Confining the cut to Germany proved to be impossible. In December 1996 Council, Germany declared that "a greater degree of risk would inevitably have to accompany the LHC". At the same time, in view of the extra pledged contributions from the four non-member states and a declaration of interest from the US, the project was approved for single-stage construction with the deficit financed by loans. It was also agreed that the final cost of the project would be reviewed at the half-way stage with a view to adjusting the completion date. With such a budget reduction, it was inevitable that a financial crisis would occur at some time, and this was indeed the case when the cost estimate was adjusted upwards by 18% in 2001. Although this was an enviable achievement for a project of such technological complexity and with the initial cost estimate made in 1993 before a single prototype had been made, it created big waves in Council. CERN was obliged to increase the level of borrowing and extend the construction period (which was anyway necessary on technical grounds for both machine and detectors).

3. Design of the LHC

The fact that the LHC was to be constructed at CERN making the most of the existing infrastructure, not the least the LEP tunnel imposed considerable constraints. The first of these was that to obtain the maximum possible energy, the dipole field had to be pushed up to a level never achieved before in accelerator magnets. The nominal field of 8.3T is achieved by cooling the superconducting magnets to 1.9K, below the lambda point of helium. The second constraint was the small diameter (3.8 m) of the tunnel. This imposed the novel compact two-in-one design where the two apertures are integrated into a single yoke (Figure 2). This allowed the cryostat to be small enough so that magnets could be (just) transported through tunnel sections where the machine is already installed. Finally the constraint on using the CERN injector chain limits the injection energy into the LHC to 450 GeV.

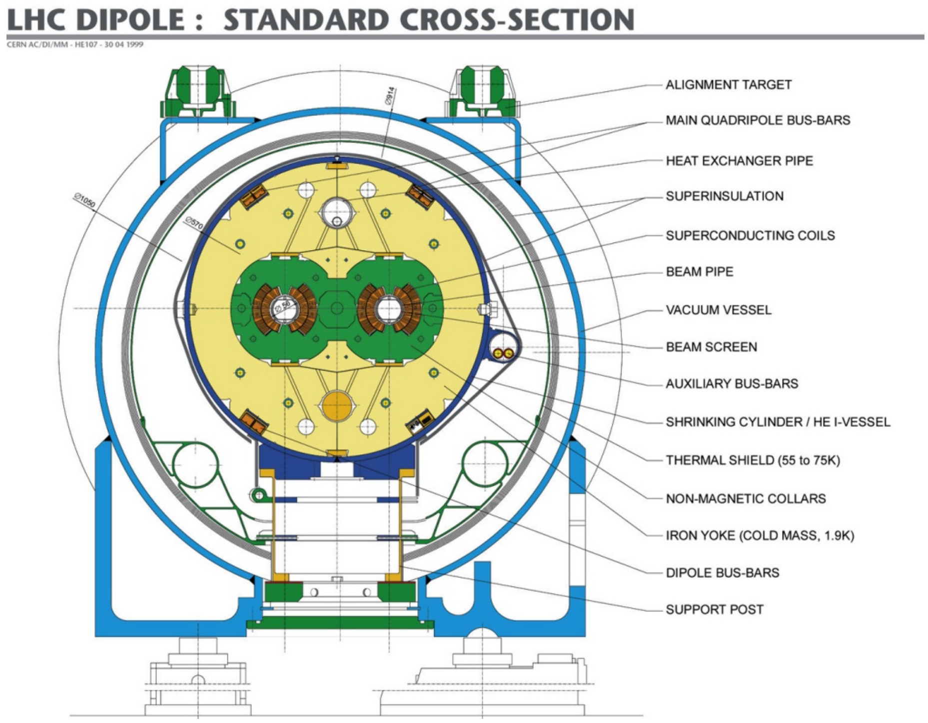

Figure 2 – LHC dipole cross-section

The LEP ring consisted of eight octants, each separated by 500 m long straight sections. Four of these (at Points 2, 4, 6 and 8) contained caverns for the LEP detectors, none of which was large enough to house the two general purpose detectors ATLAS and CMS, where new caverns would have to be excavated. Initially it was foreseen to build the CMS cavern at Point 6, but this would have resulted in the premature closure of the OPAL detector of LEP. Instead, in spite of known problems with ground water, CMS was allocated Point 5 and ATLAS Point 1. The other two detectors could be housed in existing caverns, ALICE naturally at Point 2 since it used the magnet of the L3 detector and LHC-b at Point 8. The four other straight sections were allocated to machine utilities. Point 4 contains the RF and Point 6 the beam abort systems. Finally, Points 3 and 7, which do not have underground caverns house the collimation systems.

4. Machine cooldown

2008 was a very eventful year for the LHC. During the first half of the year, the whole machine was cooled down (Figure 3). From room temperature to 80K the helium circulating in the magnets is cooled down by vapourising liquid nitrogen in a heat exchanger. In total, 1200 tons of LN2 is needed for a single sector, the whole process taking about 10 days with 60 trucks, each containing 20 tons of LN2 arriving every 4 hours. Between 80K and 4.5K, the helium refrigerators are used. Finally the cold compressors producing helium at 15 mbar pressure are switched on to reduce the temperature to the operating value of 1.9K.

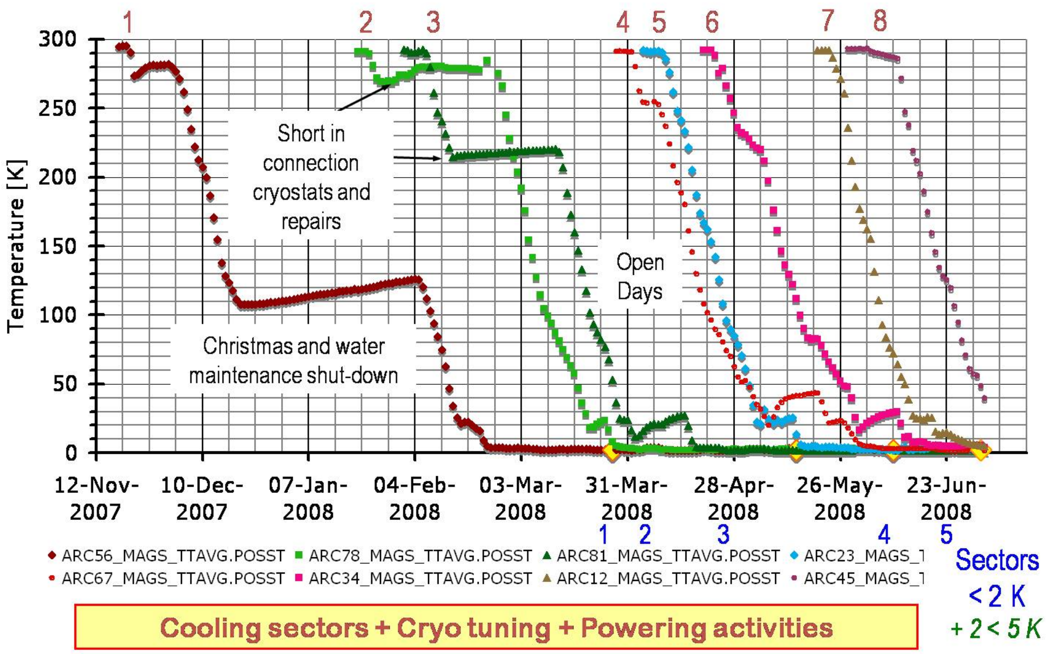

Figure 3 – Cooldown of LHC sectors.

5. First commissioning

By the 10th September 2008, seven of the eight sectors had been successfully commissioned to 5.5 TeV in preparation for a run at 5 TeV. Due to lack of time, the eighth sector had only been taken to 4 TeV. Beam commissioning started by threading beam 2, the counter-clockwise beam, (Figure 4) around the ring, stopping it at each long straight section sequentially in order to correct the trajectory. In less than an hour the beam had completed a full turn, witnessed by a second spot on a fluorescent screen intercepting both injected and circulating beams (Figure 5).

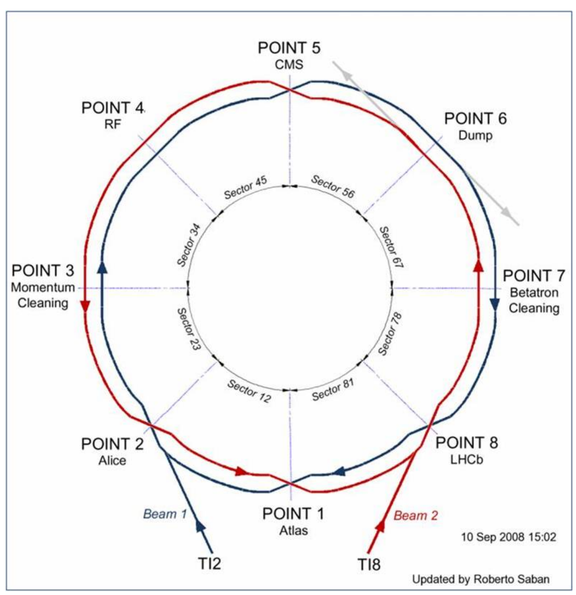

Figure 4 – LHC beam layout.

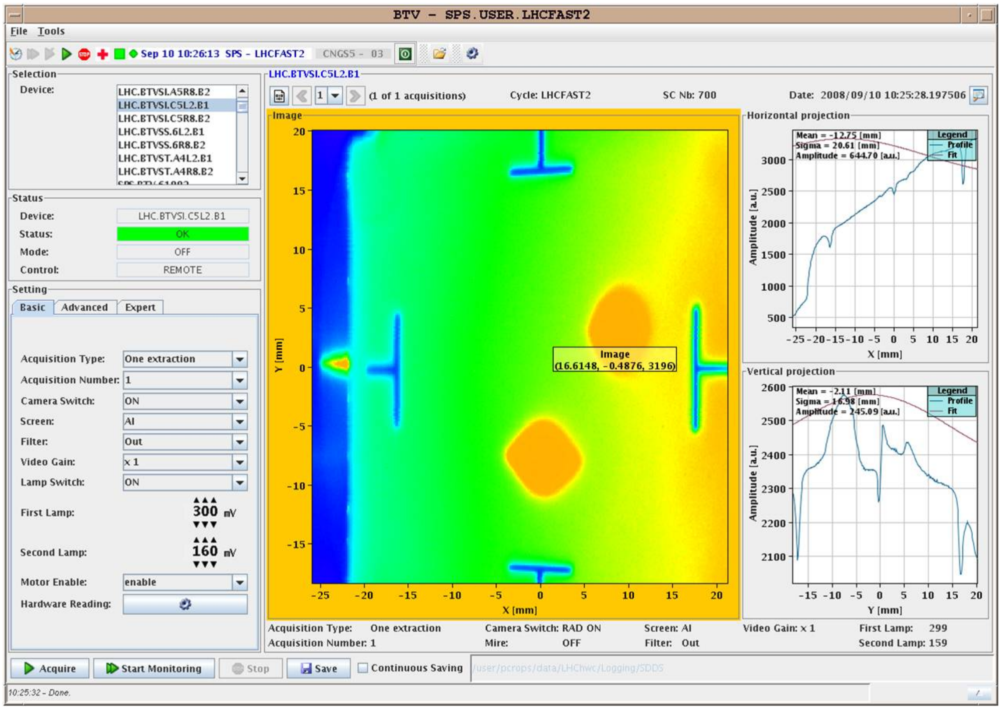

Figure 5 – Beam on turns 1 and 2.

Very quickly, a beam circulating for a few hundred turns could be established (Figure 4). The decay in intensity is due to the debunching of the beam around the ring since the Radiofrequency system was not yet switched on. Figures 6 to 10 show the RF capture process. Each horizontal line on the mountain range display records the bunch intensity every 10 turns. Without the RF the beam debunches as it should in about 250 turns, or 25 msec. In the next figure, the first attempt was made to capture the beam, but as can be seen, the injection phase was completely wrong. Adjusting the phase allowed a partial capture, but at a slightly wrong frequency. Finally, adjusting the frequency resulted in a perfect capture.

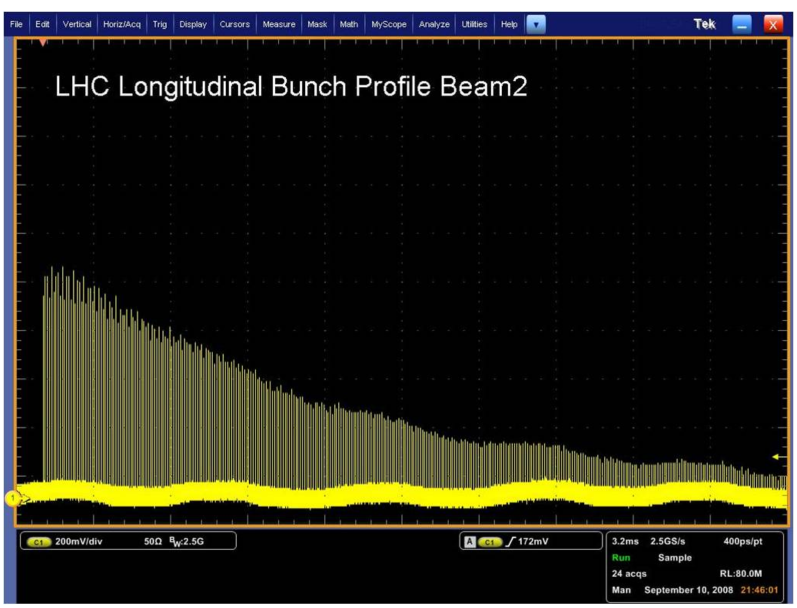

Figure 6 – A few 100 turns.

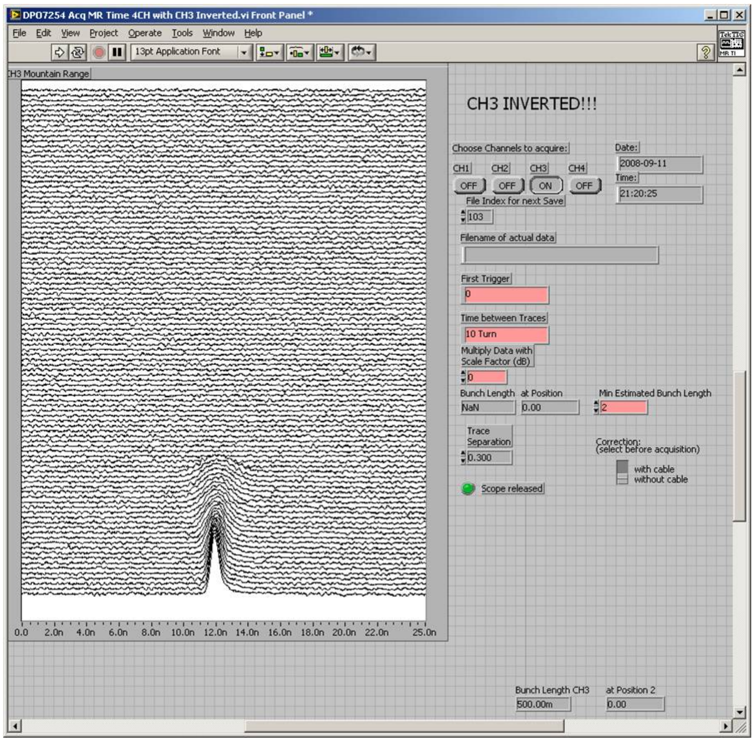

Figure 7 – No RF, debunching in ~ 25*10 turns, i.e. roughly 25 ms.

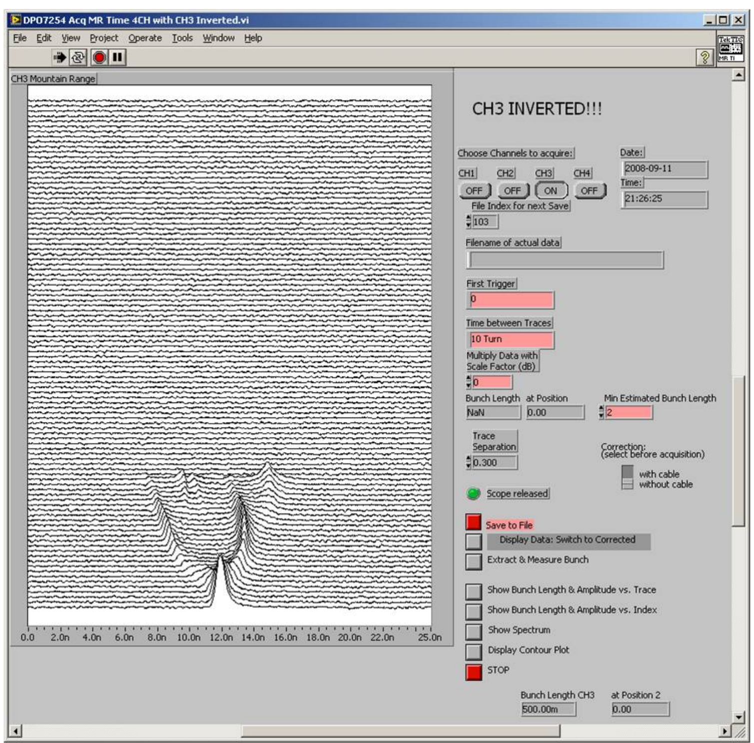

Figure 8 – First attempt at capture, at exactly the wrong injection phase.

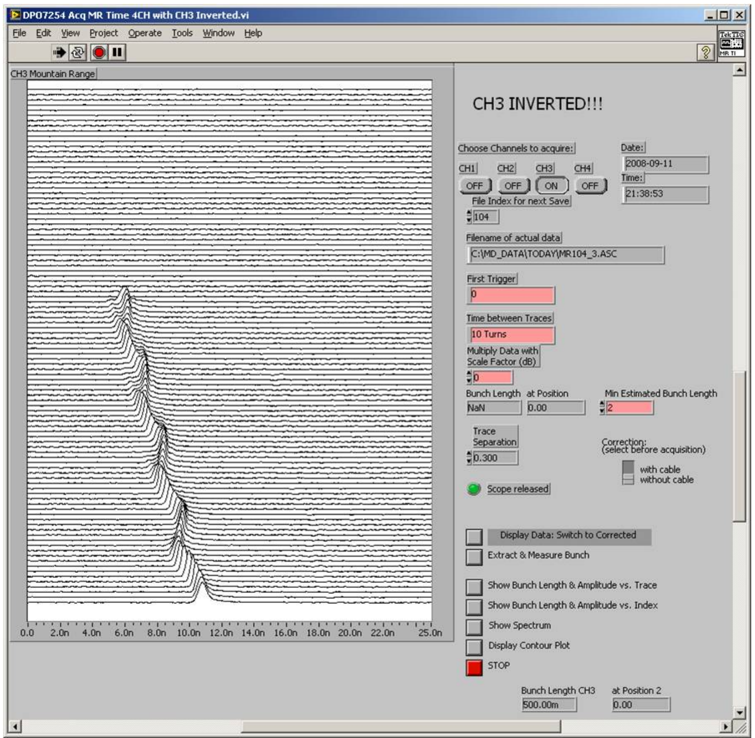

Figure 9 – Capture with corrected injection phase but wrong frequency.

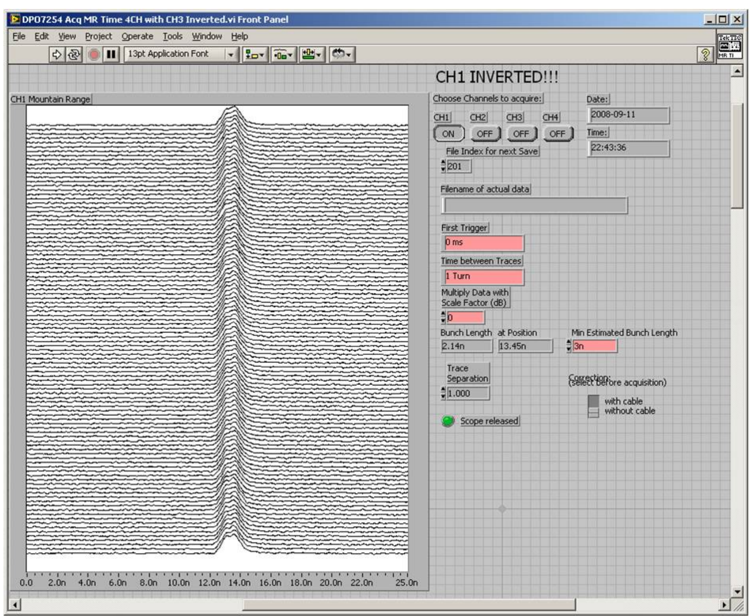

Figure 10 – Capture with optimum injection phasing, correct reference.

The closed orbit could then be corrected. Figure 11 shows the first orbit correction where remarkably at this early stage, the rms orbit is less than 2 mm. It can be seen that in the horizontal plane the mean orbit is displaced radially by about a millimeter, indicative of an energy mismatch of about 0.9 permil.

Figure 11 – Corrected closed orbit on B2. Energy offset of ~ -0.9 permil due to the capture frequency.

6. The accident

Commissioning proceeded rapidly with circulating beam in the other ring until, on 18th September a transformer failed at Point 8, taking down the cryogenics in that sector. Since it was impossible to circulate beam, attention turned to bringing the last remaining sector up to to.5 TeV like the others. On 19th September, the last remaining circuit was being ramped to full field when, at 5.2 TeV a catastrophic rupture of a busbar occurred causing extensive damage in Sector 34. These busbars are connected by induction brazing with three layers of Tin/Silver solder in a copper box. Initially it was foreseen to clamp these busbars mechanically, but this was discarded on the recommendation of an external review committee on the grounds that it would increase the hydraulic impedance in the interconnect region. A fact finding commission was established where it was concluded that the most probable cause of the accident was too high a resistivity in a superconducting joint due to the omission of the solder.

7. Diagnostics

An urgent priority after the accident was to sift through most mortem data to see if any precursors of the accident could be detected, in particular any anomalous temperature increase in the affected area.

Detecting temperature rise in the superfluid helium is made difficult by two factors. The first is the enormous thermal conductivity of superfluid helium (Figure 12). This provides good cooling of joints initially, but the thermal conductivity is a function of flux density (Figure 13), so as the heating increases, the cooling capacity quickly collapses, especially in the region of the splices with high hydraulic impedance.

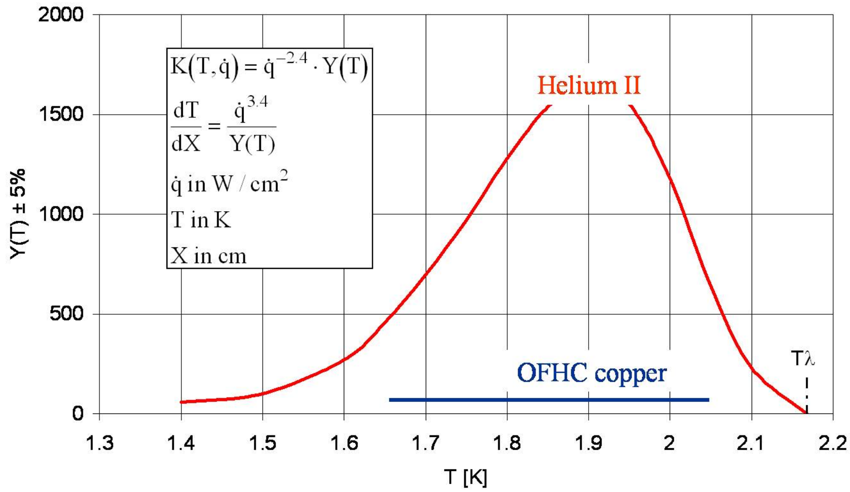

Figure 12 – Equivalent thermal conductivity of He II.

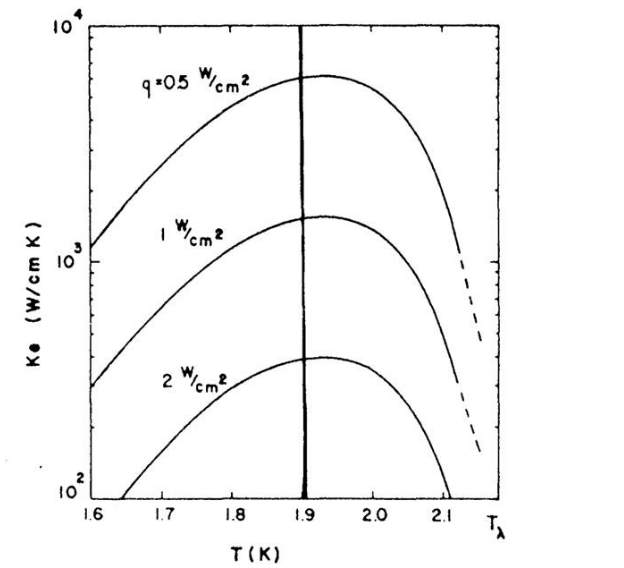

Effective thermal conductivity of He II.

Figure 13

The other reason why it was impossible to observe a temperature rise was the configuration of the superfluid cooling circuits themselves. Figure 14 shows one cryogenic cell containing two 106 m long periods of the machine. The primary superfluid flows through bayonet heat exchangers, the flow rate being controlled through Joule-Thomson valves (CV910 in the diagram). These valves are in a servo loop which keeps the temperature constant.

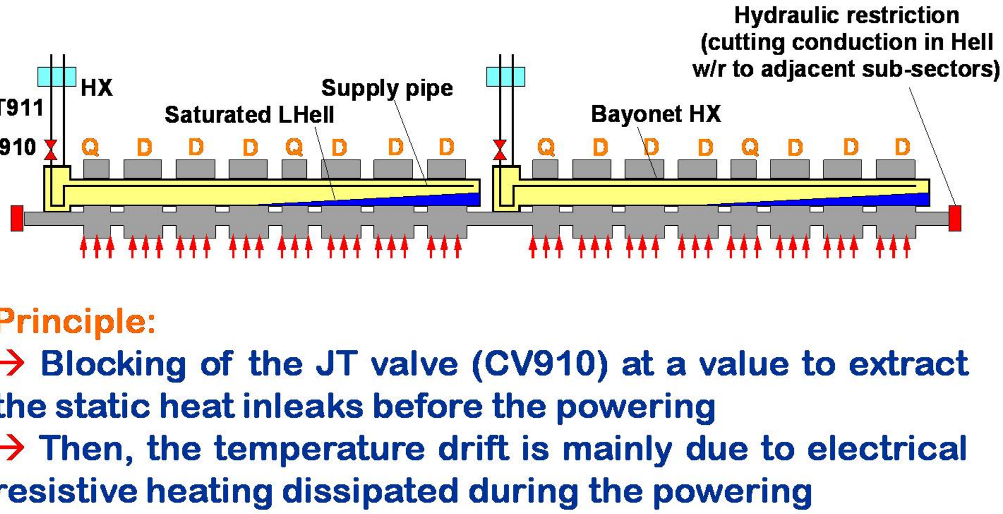

Figure 14 – Sub-sector magnet cooling scheme.

It was obviously very important to find a way to be sure that there are no more such bad joints in the machine. Two methods were developed. The first of these relied on calorimetry. With the servo loops open, the valves could be adjusted to just balance the static heat inleak. Under these conditions it was shown that a calibrated heat inleak of 10 W through a resistor could be detected and by measuring the rate of temperature rise (Figure 15), the original 10 W could be reconstituted purely calorimetrically. Note the temperature axis with 5 mK ticks! Once this calibration was made, a sector was powered to 5 kA with the J-T valves in open loop. The normal signal to be expected during a current cycle is a slight heating during ramp and deramp due to eddy currents and slow cooling on flat top. Figure 16 shows a cell in which this was not the case. The slow monatonic heating on flat top was consistent with a 100 nano-Ohm resistance somewhere in the cell.

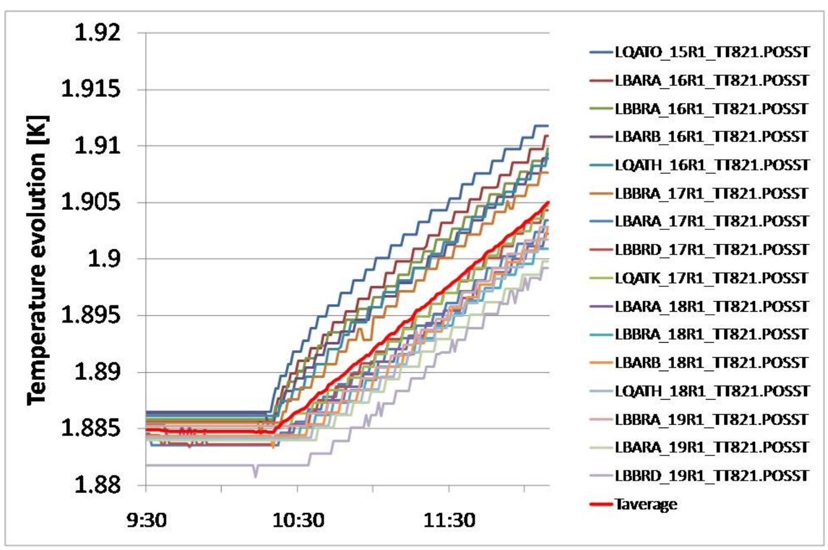

Figure 15 – Experimental validation: temperature evolution.

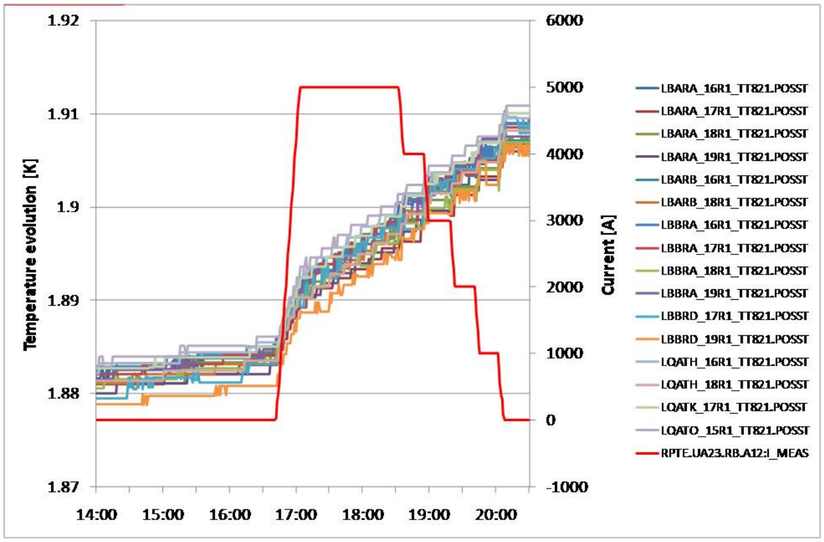

Figure 16 – Powering example: 15R1 powering @ 5000A.

Every magnet is equipped with a “post mortem” card containing an ADC and a buffer memory in order to measure voltages, usually on a trigger due to a quench. It was realized that these cards could also be used to improve the signal-to-noise ratio in measuring voltages in DC conditions by averaging, thereby opening up the possibility of making Ohmic measurements across each splice. Figure 17 shows such a measurements of all the joints in the dipole chains of Sectors 67 and 78 during a stepwise current ramp to 5 kA. In Sector 67, there is one anomaly visible with a resistance of 47 nano-Ohms. It was possible to locate exactly which splice was responsible. Both the 100 nano-Ohm splice previously mentioned and the 47 nano-Ohm splice were inside magnets which had already been tested to full current. They have both been removed and the bad splices confirmed. No other such splices have been detected anywhere else in the machine.

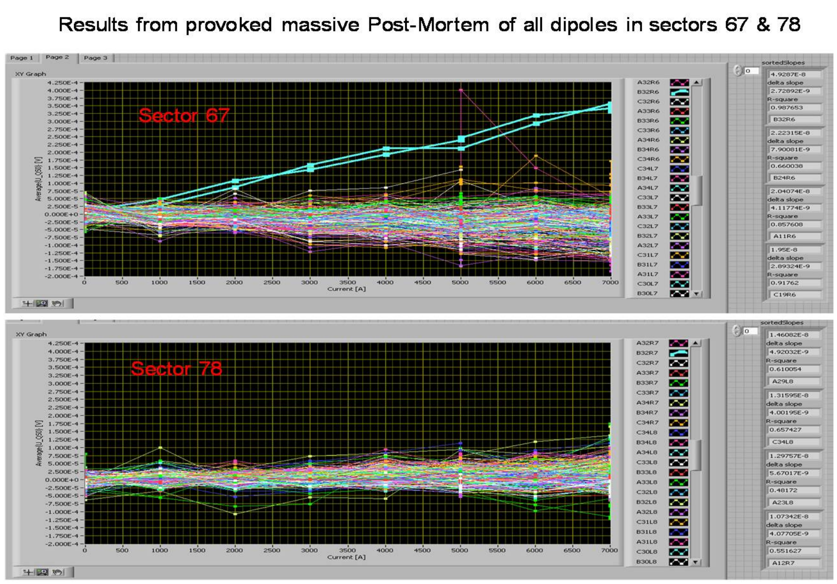

Figure 17 – Snapshots in S67 and S78 on all 154 dipoles - B32.R6 with a high (47 nΩ) joint resistance between the poles of one aperture.

8. Conclusions

Initial commissioning of the LHC went extremely smoothly. Circulating and captured beam were achieved in record time. The 2-in-1 structure of the magnets works exactly as predicted. The machine optics already looks extremely good with the closed orbit corrected to less than 2 mm rms.

The unfortunate splice incident has created a lot of damage which has to be repaired. Two powerful diagnostic tools have been developed to detect bad splices and to allow a permanent monitoring during operation. The repair is estimated to take about 9 months. The LHC should be back in operation in Autumn 2009.

The construction of the LHC is the culmination of more than 40 years of accumulated knowledge obtained in building and operating hadron storage rings at CERN. The seed was set by Edoardo Amaldi and his colleagues in the 1960's. I am sure that they would be very proud to see CERN today as the world's greatest Particle Physics Laboratory.

9. Postscriptum 30th March 2010

Since this manuscript was first written, the necessary repairs have been made. The LHC was once more ready to inject beams in November 2009. Before the end of the year, first collisions were obtained with a center-of-mass energy of 2.36 TeV (corrisponding to 2 kA in the main dipoles), taking the high-energy frontier away from the Tevatron at Fermilab. The first two months of 2010 were taken with further consolidation of the quench protection systems and commissioning the hardware to allow 7 TeV operation. Today, first collisions were obtained at this new energy. It was a moving experience for some of us who have worked so long on the project to see the luminosity counters climb as the beams were brought into collision.

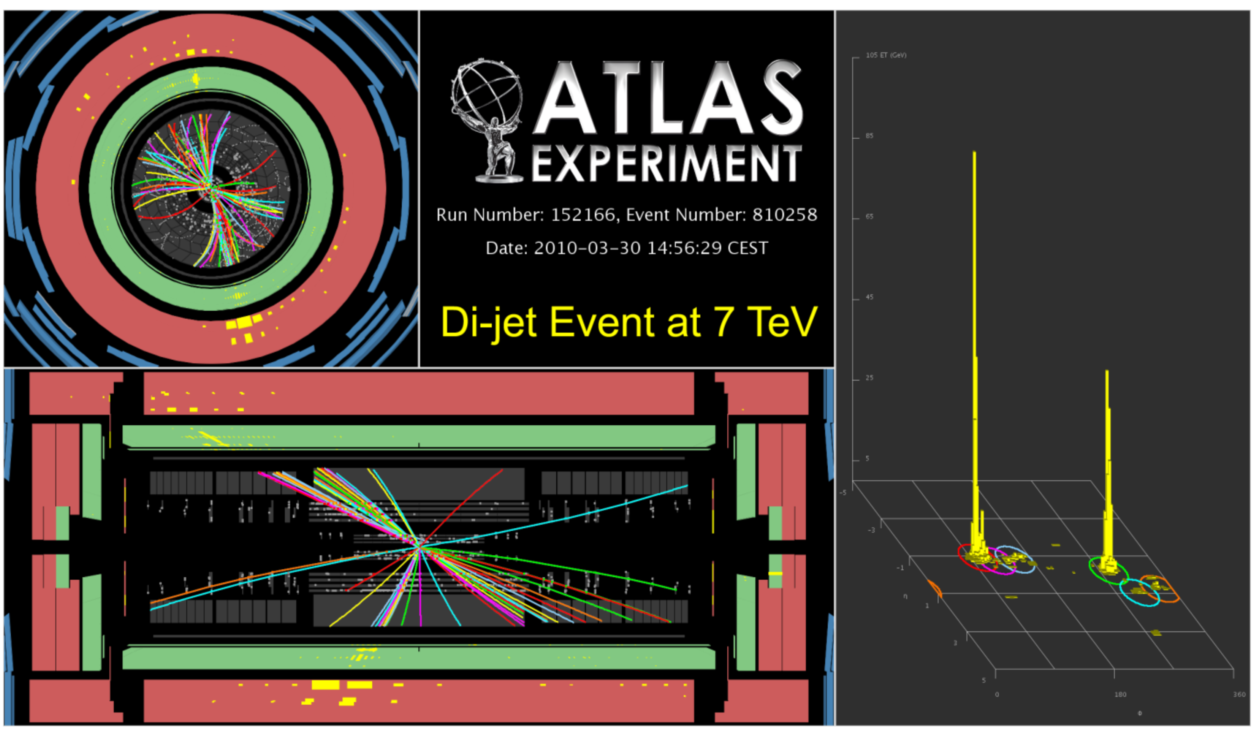

Figure 18 – A spectacular event collected during the very first run at 7 TeV. The two jets each have an energy of more than 300 GeV. (Courtesy ATLAS)

There will be no doubt further consolidation needed before the machine can be pushed to its full potential but already it is exploring a new energy regime never before seen by mankind. The adventure of LHC construction is finished. Now let the adventure of discovery begin.