

# CURRENT STATE OF COPPER STABILIZERS AND METHODOLOGY TOWARDS CALCULATING RISK

M. Koratzinos, CERN, Geneva, Switzerland

## Abstract

This paper attempts to review the factors that limit the maximum safe energy of the LHC. It concentrates on improvements on knowledge gained since the author gave a similar talk in Chamonix 2010 last year [2]. A way of defining risk at different beam energies is also put forward, as well as a proposal for a qualification tool that might allow to run at higher energies in the future.

Another important development is that the idea behind a complete qualification tool, the so-called ‘thermal amplifier’ was for the first time tested in the lab (“proof of principle studies”) giving good agreement between simulation and actual measurements. This idea is important as it might be the only qualification tool that might allow us to run safely at an energy higher than last year’s.

## CHANGES SINCE LAST YEAR

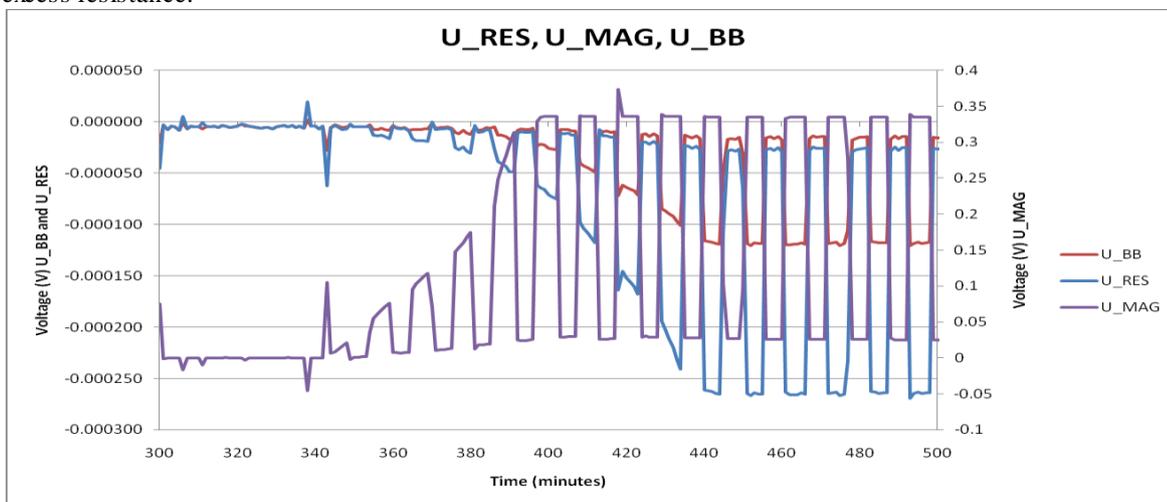
There have been two major improvements to our knowledge that affects the maximum safe energy of the LHC: the first comes from an improved simulation of burnout limits [1] that includes a detailed study of the case of interconnect quenches due to heat conduction through the busbar (“busbar propagation” quenches). This mode of quenching an interconnect had not been studied in detail before (instead a very crude model of quench propagation through the busbar was used) and leads to somewhat lower limits for the maximum allowed excess resistance of a joint.

The second improvement comes from an increase in the knowledge of the busbar segments themselves: an important quantity (related to the quality of the copper stabilizer), namely the residual resistivity ratio (RRR) has been measured for all sectors of the machine, and for both the RB and RQ circuits. The RRR measured is higher than the conservative number used last year and therefore leads to somewhat higher limits for the maximum allowed excess resistance.

## RRR MEASUREMENTS

The burnout limit of a joint of the main circuits of the LHC depends on the RRR of the copper stabilizer of the circuit. Measuring the RRR of the busbar stabilizer *in situ* in the tunnel is not an easy task, and previous attempts to measure it were not very successful. For this reason, a conservative value (of 100) was used up to now for the determination of burnout limits. If the RRR would be 200 instead of 100, the safe energy of the LHC would increase by a good fraction of a TeV.

It was therefore decided to measure the RRR in all sectors. A type test was performed exactly a year ago and a campaign for measuring the RRR was performed between 10/12/2010 and 27/1/2011. This is a combined effort between the EPC, CRG and MPE groups of TE – without the help of all these groups the measurement would not have been possible. The campaign took only two days per sector thanks to the efficiency of the cryogenics and power converter groups, but necessitated many hours of work in the LHC tunnel, designing and installing patches, etc.



**Figure 1:** Traces of four relevant signals close to magnet B12L6 during a superconducting to normal transition: *U\_MAG* is the voltage across the magnet (right scale). *U\_RES* is the raw signal of a bus bar segment and *U\_BB* is the actual voltage across a bus bar segment after corrections for the voltage drop of the QPS cable.

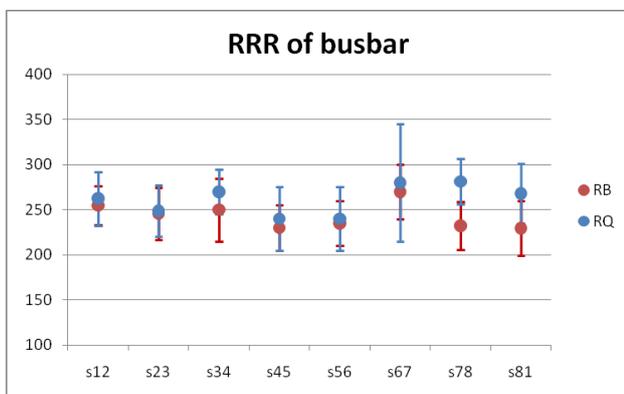
The voltage of all busbar segments in a sector (leading to the determination of resistance) is measured using the nQPS system. However, the nQPS system was not designed for this kind of measurements – corrections are necessary and are large. For this reason, apart from the measurement proper, an extra “calibration” step was taken.

Calibration is needed due the complex corrections necessary, amounting to more than 50% of the raw signal measured – see Figure 1. More than half the voltage recorded by the nQPS system is not due to a voltage drop in the bus bar segment, but due to a drop in the nQPS cables that take the signal to the voltage measuring device. This voltage drop is proportional to the cable lengths and their resistivity.

The calibration was performed by installing a series of ‘patches’ in about 10% of the bus bar segments of a sector that eliminated the voltage drops in the nQPS cables. This calibration gave a value for the resistivity per meter of the nQPS cables of  $87\text{m}\Omega/\text{m}$  with a preliminary uncertainty of 5%. The analysis presented here, still preliminary, uses one value for the resistivity per meter of the nQPS cables and thus does not take into account differences in the ambient temperature of the tunnel between different sectors. This (complex) analysis will be published in due time, but we can here present the results with systematic errors that will eventually be reduced. The results are shown in Figure 2. RRR values between sectors and between circuits are consistent and in the range 200-300.

The mean RRR of the copper busbar of the machine was found to be  $250 \pm 50$  (50 being the systematic error, taken as the maximum difference between RB and RQ circuits in one sector (40) and the maximum difference between different sectors (50) and is expected to be reduced in the final publication).

Therefore we can safely assume that the RRR for the whole machine is larger than **200**. The improvement on safe energy going from  $\text{RRR}=100$  to  $\text{RRR}=200$  is  $0.2\text{TeV}$  per beam for “busbar propagation” quenches and  $0.5\text{TeV}$  per beam for “gaseous He propagation” quenches.



**Figure 2:** The RRR values of the copper stabilizer measured for all sectors. The dot is the mean value and the spread the standard deviation of the measurements

## MAXIMUM SAFE ENERGY

The maximum safe energy of the LHC depends on:

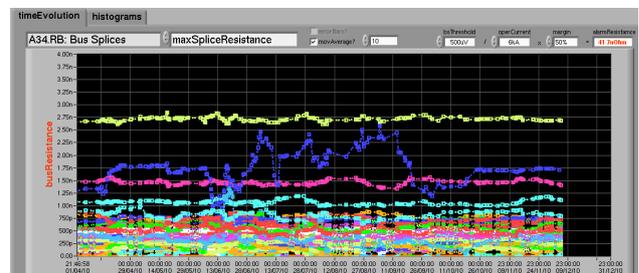
1. The condition of the soldering of the superconducting cables of the interconnects. As we shall see, this does not pose a problem for LHC operation.
2. The condition of the copper stabilizer joints ...
3. ...coupled to a quench of the joint

Points (2) or (3) alone are not sufficient to produce a serious incident. A bad copper stabilizer joint will not run away if there is no nearby quench. But the combination of (2) and (3) will produce a thermal runaway if the condition of the copper stabilizer joint is worse than the limits calculated in [1].

[The burnout limit depends (amongst other things) on the RRR of the main busbars, as already discussed.]

### Condition of superconducting cable joints

Superconducting joints are measured by estimating the resistance of a busbar segment at cold. The largest resistance seen is  $\sim 3\text{n}\Omega$  which poses no problem for operation (compared with  $300\text{n}\Omega$  of the joint responsible for the accident of 19 September 2008). However, a question remains: do they deteriorate with time? Monitoring the resistances for a year has shown that no deterioration of cold resistances has been seen during 2010 (Figure 3) [3].

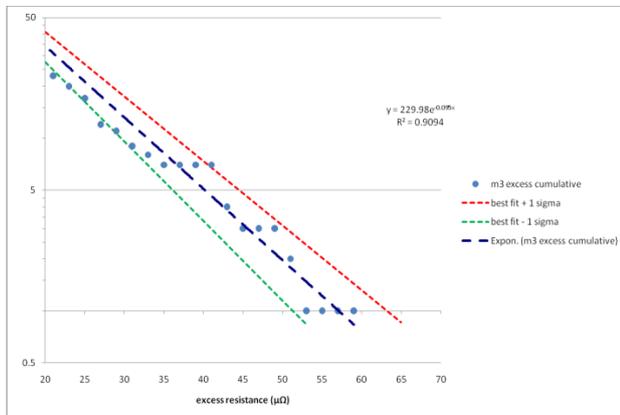


**Figure 3:** Evolution of bus bar segment resistances at cold during 2010. No deterioration is seen. The blue trace belongs to a noisy QPS measuring board.

### Condition of copper stabilizer joints

The maximum safe energy of the LHC critically depends on the condition of copper stabilizer joints [1]. Unfortunately, our knowledge of the condition of these joints comes from a small and most probably biased sample of 134 joints (out of 10000 total in the machine) measured accurately in 2009. Out of these, 23 are above  $20\mu\text{Ohms}$ . The analysis fits a functional form on the distribution of these 23 values and this functional form is used to estimate the number of joints exceeding a specific threshold in all subsequent discussions on maximum safe energy of the LHC. The distribution and the fit are shown in Figure 4. The sample of the aforementioned 23 joints is from the RB bus only, and from 5 out of 8 sectors. We

have no measurement of the stability of these joints with time (indeed there are scenarios where a joint resistance can deteriorate). This analysis is unchanged since last year [2].



**Figure 4:** The tail of the cumulative distribution of excess resistances measured accurately in 2009 (values larger than 20 $\mu$ Ohms). An exponential function fits the distribution well (blue line). The  $\pm 1$  sigma fit (red and green lines) is also shown.

### Quenches in 2010

One of the pleasant outcomes of this year's running, is that we have observed zero unintentional beam-induced quenches. This was mainly due to a well behaving, reproducible machine and a BLM system that intervened in time and prevented magnet quenches due to beam losses.

If we were sure that we will get zero quenches next year, running at ANY energy would be safe from the splices point of view. However, non-beam-induced quenches have been seen in 2010.

### Quench statistics

A quick attempt has been made to look at how many quenches happened in 2010 (thanks to H. Reymond): The PM system is interrogated for cases where a/ there was a QPS signal present in the post-mortem files denoting a quench and b/ there was high current in the sector at the time (>5000A).

This gave a total of 27 events. Manual inspection of those revealed that 9 were not real quenches (no heaters fired). Three more quenches were identified manually that were missed by the automatic analysis. The categories identified were:

- Trip during PGC1: 11 quenches
- Provoked quench: 3 quenches
- Power cut: 5 quenches
- Other: 2 quenches

In conclusion, in 2010 we had about 20 quenches of the RB circuits above 5000A, due to various reasons.

For next year, we simply do not know how many quenches we are going to get. Therefore in what follows, the number of quenches is left as a free parameter.

## SUGGESTED STRATEGY

As the probability of an accident, and hence the risk, is a function of both the condition of the copper stabilizer joints and the number of quenches (which is largely an unknown quantity for the year(s) to come), we found that the best way to present how safe it is to run at a specific energy is to calculate how many quenches we need per year before reaching a pre-defined incident probability level.

I have used 0.1% as the acceptable incident probability level (one incident every 1000 years), but the reader can easily translate the conclusions presented here to his favoured level.

In this framework, if the number of quenches we can "afford" is large, then operation is safe, and we would only need to re-assess the situation when this number of quenches is reached. If, however, the number of quenches we can "afford" is less than one, getting a single quench would put us above our pre-defined 'risk' level.

### Types of events

Quench events can be broadly divided in two categories:

Prompt quenches, where the interconnect quenches first or at the same time as an adjacent magnet, are the most serious as they happen at the highest current. However, they are also the most unlikely (we had zero prompt quenches in 2010).

The other category is quench events where the magnet quenches first and the quench is propagated to the interconnect either through the busbar connecting the (rapidly heated) diode to the main busbar, or through the propagation of warm gaseous Helium. The gaseous Helium propagation speed has not been measured accurately, so a conservative value of 20 seconds is taken.

Quenching a quadrupole magnet has less severe consequences than quenching a dipole magnet due to the longer distances between the diode and the main busbar and due to the fact that less energy stored. A quench of a dipole or quadrupole magnet affects both the RB and the RQ circuits but it is important to note that the time constant of the energy extraction of the two circuits is very different: currently around 50 seconds for the RB and 10 seconds for the RQ.

In summary, the types of quench events and their corresponding consequences are:

- **Prompt quench** of the interconnect region (beam induced). This is very unlikely since calculations show that the adjacent magnet has a quench sensitivity which is a factor  $10^3$  higher. The number of interconnects that will quench in

such an event are 2 RB interconnects and 4 RQ interconnects.

- **Magnet quench** of a dipole magnet:
  - 1 RB interconnect will quench first from heat propagation through the busbar from the diode. This calculation is new this year and leads to more strict limits than the simpler calculation of last year's [1].
  - 3 RB and 8RQ interconnects will quench ~20 seconds later from heat transferred through gaseous Helium [this is only relevant for the RB circuit, as the RQ circuit will have most of its energy extracted in 20 seconds]
- **Magnet quench** of a quadrupole magnet:
  - 2 RQ interconnects will quench first from heat propagation through the busbar from the diodes. This calculation has not been performed yet in detail, but it is believed that it is less important than the RB case.
  - 4 RB and 6 RQ interconnects will quench in much more than 20 seconds later from heat transferred through gaseous Helium i.e. this failure mode is safe at the energies we are considering.

The table below shows in summary the maximum allowed defect at different energies depending on the failure mode, the number of joints affected per quench event, the number of joints above the limit using the best knowledge of the state of interconnects taken from the fitted function of Figure 4, and finally the number of quenches to reach a predefined accident probability level of 0.1%. Some numbers presented in the second and third columns are taken from last year and have not been recomputed with the new values of RRR measured in the machine. Therefore they are conservative (denoted by a plus sign next to the number). For 4TeV operation, the numbers corresponding to two energy extraction constants are shown, 50 and 67 seconds [in square brackets]. The colour code of the last two columns is as follows: if one quench already takes us above the predefined accident probability, the number is green. If not (and therefore operation at this energy is risky after one quench) the number is red. In the case of prompt quenches or gaseous helium propagation quenches, it makes no sense to talk separately about the RB and RQ buses, so the last two columns are combined. Busbar propagation quench calculations in the case of the RQ circuit have not been performed, but the limits are believed to be much larger than the equivalent prompt quench cases (denoted by a ++).

**Table 1:** Maximum excess resistance allowed, number of bad joints and number of quenches to reach an accident probability of 0.1%

Energy (TeV), T <sub>dump</sub> (s)	Allowed R <sub>excess</sub> RB (μΩ)	Allowed R <sub>excess</sub> RQ (μΩ)	No. of bad joints above R <sub>excess</sub> RB	No. of bad joints above R <sub>excess</sub> RQ	No. of quenches RB to reach P <sub>accident</sub> =0.1%	No. of quenches RQ to reach P <sub>accident</sub> =0.1%
<b>Prompt quenches – 2 RB and 4 RQ joints affected</b>						
3.5, 50s	85	87+	0.1	0.2		5
4.0, 50s [4.0, 68s]	72 [56]	71+	0.4 [2]	0.9		1.3 [0.6]
4.5, 68s	47	56+	4	4		0.2
5.0, 68s	44+	48+	6	8		0.1
<b>Busbar propagation quenches 1 RB or 2 RQ joints affected</b>						
3.5, 50s	90	any	0.07	0	46	∞
4.0, 50s [4.0, 68s]	74 [55]	71++	0.3 [2]	0.9	10 [2]	6
4.5, 68s	46	56++	5	4	0.7	1.3
5.0, 68s	44?	48++	6	8	0.6	0.6
<b>Gaseous He propagation quenches (RB) – 3 RB and 8 RQ joints affected</b>						
3.5, 50s	115	any	0.01	0		164
4.0, 50s [4.0, 68s]	98 [75]	any	0.03 [0.3]	0		33 [4]
4.5, 68s	63	any	1	0		1.2
5.0, 68s	48+	95+	4	0.1		0.3

In conclusion, regarding the maximum safe energy of the LHC for 2011:

- The most stringent limits come from prompt, beam induced, quenches, considered unlikely. If we are confident we will get no prompt quenches, we can ignore the ‘prompt’ quench category.
- The most relevant limits come from the ‘bus propagation’ quenches, whose calculations are new this year.
- Some calculations (for 5TeV or for the RQ) have not been updated with the latest information (higher RRR) and will become less stringent.
- 4TeV operation with an extraction time of 50 seconds gives us some margin for all types of quenches, therefore 4TeV operation cannot be ruled out from the information we have about the state of the machine.
- However, 4TeV limits are more stringent than for 3.5TeV operation. Therefore, running at a higher energy needs to be balanced against pushing for higher luminosity this year, something that might result in a larger number of quenches than in 2010.

## THE “THERMAL AMPLIFIER”

### *The need for a qualification tool*

Up to now, with the splice consolidation campaign to take place in 2012, it was clearly not high priority to investigate ways of increasing the safe energy of the LHC. However, if 2012 will be a year of LHC operation, it makes sense to try and see if a test that fits in a period of (an expended perhaps) shutdown can be performed which might allow us to run at a substantially higher energy in 2012, say 5TeV.

The main reason for not being able to go higher in energy this year is that our current knowledge of the state of the copper stabilizer joints in the main circuits is very poor since:

- Not all sectors have been measured
- Different sectors that have been measured seem to have very different joint quality.
- Time degradation has not been studied and is a worry.

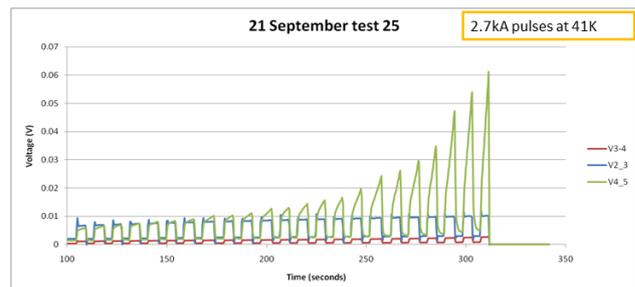
For this reason, the current safe energy analysis is based on (mostly) pessimistic assumptions, to counterbalance the above lack of knowledge.

A promising possibility is the ‘Thermal Amplifier’ [provisional name]. It is a qualification tool that can qualify a sector to the maximum current it can safely withstand. Since the idea was first conceived (by Howie Pfeffer and others), a lot of conceptual work has taken place (with a lot of input from A. Verweij), that has simplified the original idea considerably making it easier to implement. There are still a series of engineering challenges to be solved before we can put such a method

into production, necessitating close collaboration of many different groups (EPC, CRG, MPE, etc.). Also, the very first ‘proof of principle’ test of the idea was performed – so this now is not simply an idea on paper.

### *Thermal amplifier principle*

The thermal amplifier applies a pulse of high current (of order 3000A at 40K or 6000A at 20K) for about 10 to 20 seconds in a sector which is kept at non-superconducting temperatures (tests can be done between 20K and 40K). Any bad splices selectively warm up, whereas good splices remain cold. A joint that warms up from 20K to around 200-300K, has its resistance increased by a large amount (order 100), hence it is easy to detect. What makes the operation safe and an important ingredient of the principle is that the current flows through the diodes, so very little energy is stored in the circuit. The other big advantage of the method is that it is a direct measurement of a thermal runaway at the exact conditions of a joint so no further assumptions are needed.



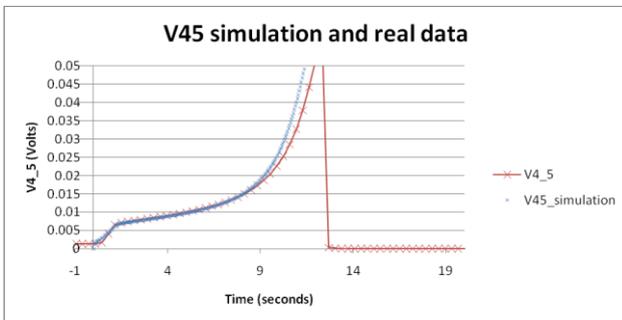
**Figure 5:** Voltages across busbar segments during a number of 7s long pulses (every 10s) of 2.7kA at 41K. The green voltage contains a 50uOhm defect and runs away after about 30 pulses. The blue and red voltages are across perfect joints.

### *The block 4 tests*

A series of tests were performed in September 2010 in the Block 4 test facility (special thanks to M. Bajko, C. Giloux, J. Fevrier). A defective joint (about 50μOhms) and a series of perfect joints were monitored as high current (2000A to 6000A) passed through the copper at temperatures varying from 20K to 40K. What was measured agreed with simulations (Figure 7) and valuable experience was gained.



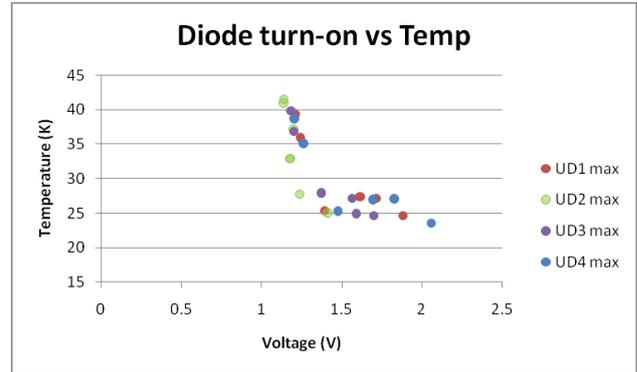
**Figure 6:** A picture of the defect machined out of a solid copper bar for the Thermal Amplifier tests in Block 4. This represents a defect of about 50uOhms at room temperature.



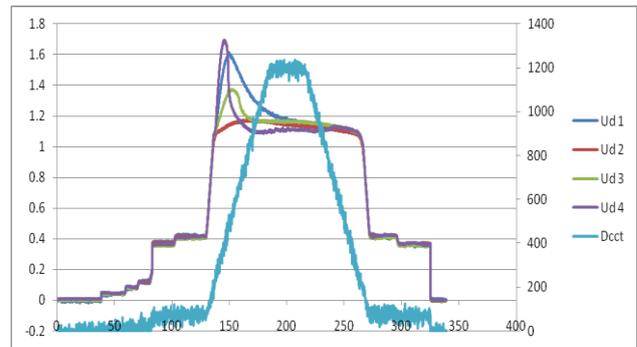
**Figure 7** Comparison of simulation (blue) and real data (red). The single current pulse in this test was 3200A. Initial temperature was 42K.

A second test at Block 4 had to do with another important ingredient of the Thermal Amplifier idea: the diode turn-on voltage. For the RB, the total voltage budget of the power converter is 190V for 154 diodes (this corresponds to 1.23V per diode). The diode turn-on voltage has been measured extensively in the past, but not at the temperatures we are interested in. For this reason four diodes were tested in Block 4 at temperatures between 20K and 40K. The four diodes exhibited slightly different characteristics, with the maximum diode turn-on voltage at a specific temperature varying by as much as 0.5V depending on the diode (see for instance Figure 9 showing the voltage across the four diodes during a test performed at 35K).

The diode test was inconclusive due to the large spread of values of turn-on voltages seen by the four diodes tested. If they are a representative sample of the diodes installed in the machine, then at 40K the current RB power converter might just be sufficient to ignite the diodes, simplifying the setup considerably. For operation at 20K, more measurements will be needed to determine if the power converter needs to be modified or boosted by an ‘igniter’ power supply.



**Figure 8:** diode turn-on voltage (max.) versus the temperature during the time of the highest voltage. At 40K, diode turn on is around 1.1V whereas at 25K 1.5 to 2V.



**Figure 9:** A typical diode test. Diodes start conducting at 1.1V, their voltage initially rising before it finally drops to 1.1V after they heat up (left scale). The average maximum turn-on voltage is 1.4V. Right scale: current flowing through the diodes (in blue). Horizontal scale is time in seconds. This test was performed at 32K.

### Engineering challenges

As we have mentioned, there are a series of engineering challenges that need to be met in order to be able to implement the thermal amplifier concept, some of them briefly discussed in the previous section. For completeness, the challenges we need to solve before this idea goes to production are the following:

- The DFBS might need to be cooled to a different temperature than the rest of the arc (CRG group)
- The voltage budget of the power converters to be used needs to be sufficient (EPC group).
- Safety should be ensured either with a very robust procedure or with an interlock (MPE group)

It is worth mentioning here the worst case scenario, if for any reason the safety systems should fail and thermal runaway ensues during a Thermal Amplifier test:

- The joint that run away will be destroyed
- The energy stored in the circuit is minimal, so no other damage will take place

- The vicinity of the joint (a cryogenic subsector) would have to be warmed up, opened up and repaired

Although such a scenario is clearly undesirable, the fact that the test is performed with such small energy stored in the magnets means that the consequences are not catastrophic.

### *Resources and time estimate of a possible project*

If such a project is approved with the intention to perform tests on all sectors during the 2011-2012 Christmas shutdown, a series of milestones would first need to be achieved. Such milestones include:

- Proof of principle. This is already done.
- Verification of the voltage budget; what is available now with no modification might just be sufficient, but some more tests are needed.
- An interlocking scheme needs to be defined.
- The best temperature to perform the test would need to be defined. This will be between 20 and 40K.

Here we should also mention that the RRR measurements performed during the 2010-2011 Christmas technical stop have good synergies with an eventual thermal amplifier test, as many components are the same (nQPS, etc.). For this reason, the experience gained during the RRR tests will be very useful to the project.

Regarding the time estimate for performing such a test, we estimate we need ~4 weeks for a type test (the first sector) and then ~3 weeks for all other sectors, with the possibility of sector overlap. From the point that cryogenic conditions have been established and after a type test has taught us how to use the system, we believe we can perform the measurements of two sectors in parallel in one week (the figure of two sectors comes from the assumption that special QPS cards for the interlock will be manufactured for two sectors).

The project could possibly be divided into three phases: an in-depth study (about 3-4 months); preparation (about 3-4 months); and production (~ 3 months). After phase 1, an external review would be desirable. Manpower resources estimate for phase 1: a minimum of 1.5FTE of an engineer. For phase 2: 1.5FTE of an engineer plus 1.5FTE of a technician. The production phase will take more resources (but for a shorter period).

## CONCLUSIONS

Recent RRR measurements indicate that we can safely assume an RRR of 200 for the copper stabilizer of the main circuits of the LHC.

Given the (lack of) knowledge of the excess resistance of the copper stabilizer joints, running at any energy carries a certain risk.

The maximum number of quenches before we reach a pre-determined level of accident probability has been presented as a function of energy.

A qualification method (the thermal amplifier) has been presented that will allow us to run at the highest possible energy with very little risk.

A thermal amplifier campaign can fit in the 2011-2012 shutdown.

## REFERENCES

- [1] Arjan Verweij, "Update on calculations of max. excess resistance allowed as a function of energy for the case of prompt/semi-prompt/adjacent quenches", these proceedings.
- [2] M. Koratzinos, "Do the splices limit us to 5TeV – plans for the 2010 run", Chamonix 2010.
- [3] Z. Charifouline, private communication.