ATLAS Silicon Cooling

Ben Wynne

Particle Physics Experiment Group

Performance Overview

The ATLAS detector contains 1744 silicon pixel detectors and 4088 silicon strip detectors, in total 62.8 m² of silicon producing ~60 kW of heat [1].

The detector modules are arranged along 204 cooling loops. The schematic below shows one such cooling loop.

Liquid C_3F_8 is used to cool the loops, with 160 kg circulating when all loops are operational. The liquid is split between loops by the distribution rack, and then cooled in a counter flow heat exchanger with the exhaust vapour. Once the liquid reaches the silicon modules it starts to boil, cooling them to an operational temperature of between 4 and -8 C, as required. The exhaust passes back through the heat exchanger, where it is heated further to ensure that only gas returns to the plant.

Time from cooling start to module power: 4 minutes. Stable module temperature 3 minutes later.

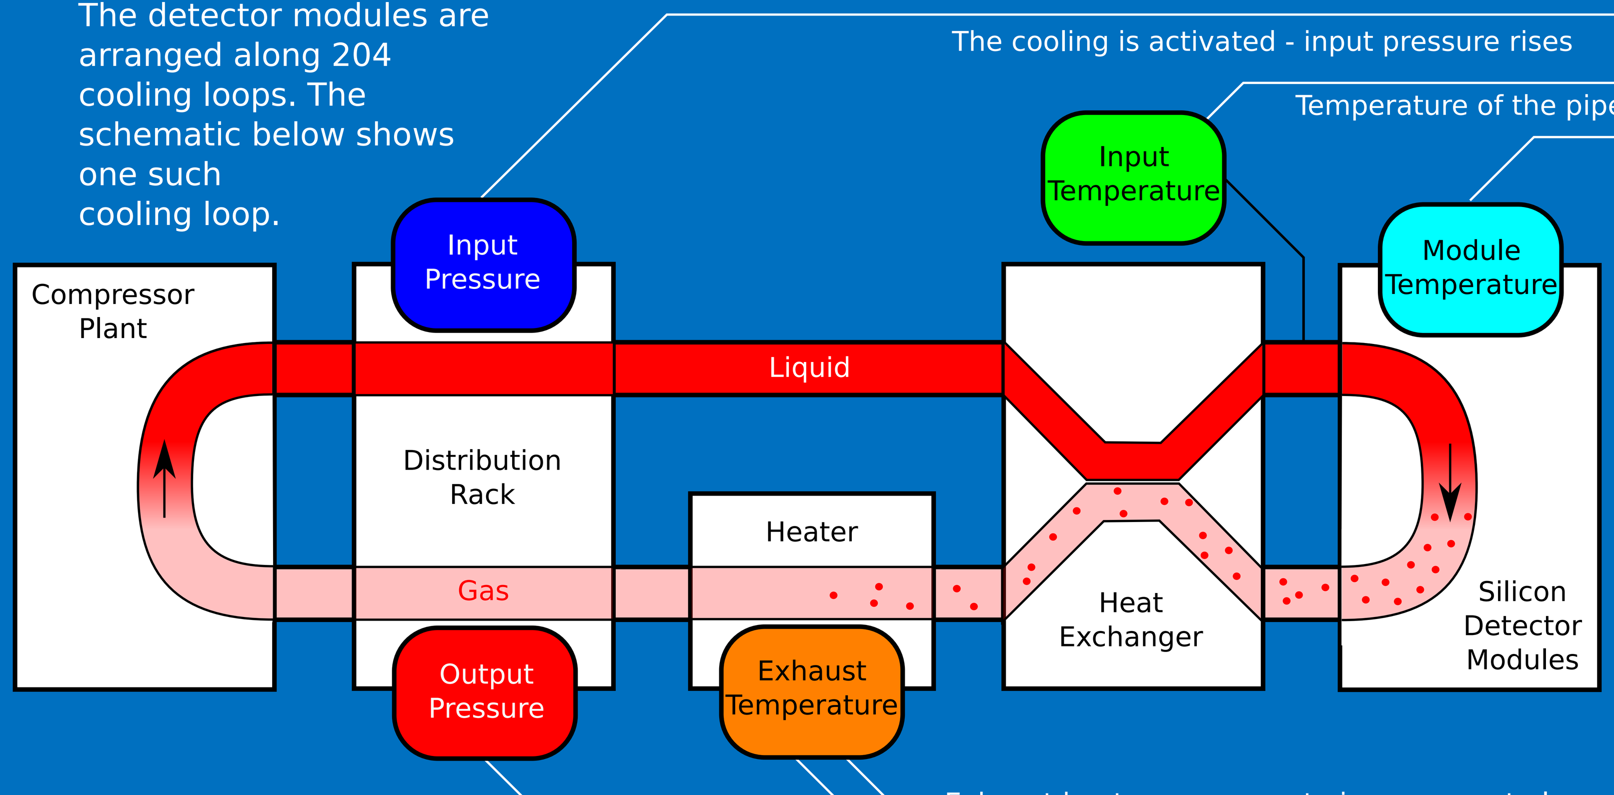

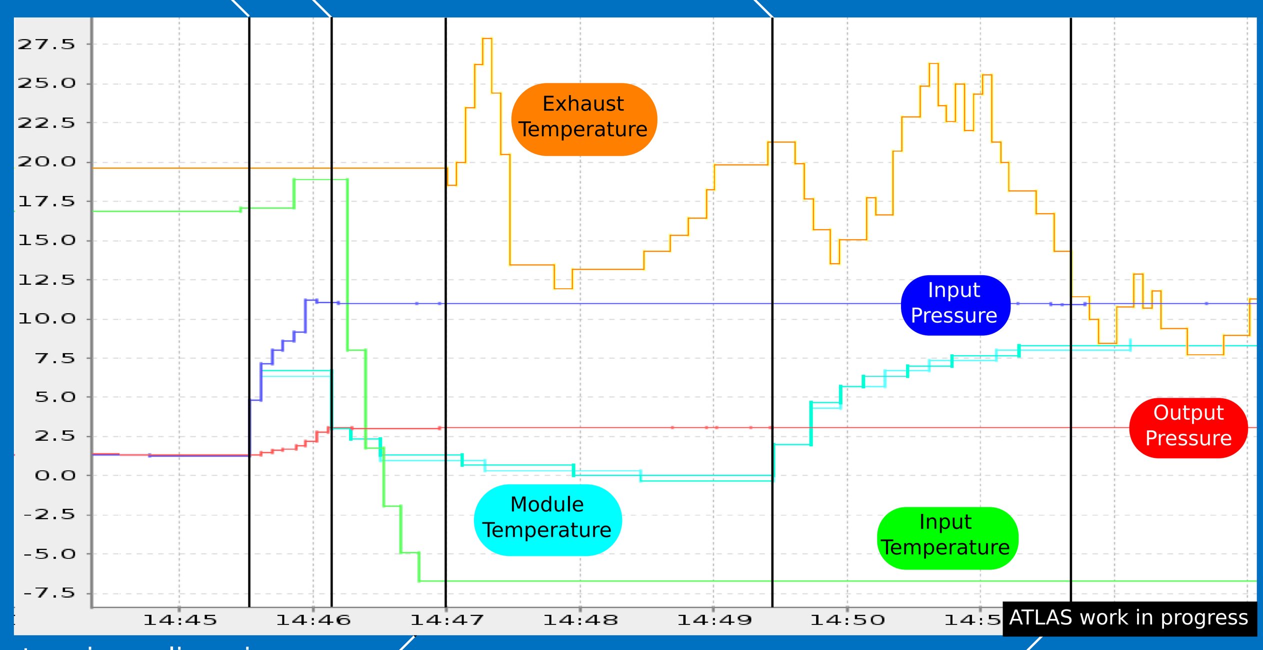

This plot [2] shows the sequence of events when switching on a cooling loop after it has had time to stabilise at ambient conditions.

The notable events are indicated by vertical lines, connected to the relevant object in the schematic. Corresponding sensors and plot lines are colour-coded

Reduced Exhaust Flow

On April 30th 2010, following the failure of a control sensor a compressor stopped, the cooling plant could not extract exhaust at the required rate and the pressure started to rise in the cooling loop outputs.

The desired silicon temperature is achieved by controlling the vapour pressure in the cooling loops (see the PT diagram for C_3F_8 [3] below). However, the pressure regulators can only work by limiting the flow of exhaust vapour.

Pressure rise: 0.45 bar
Temperature rise: 3.5 C

Consistent with C_3F_8 phase diagram, assuming temperature offset. Output pressure was crucial factor.

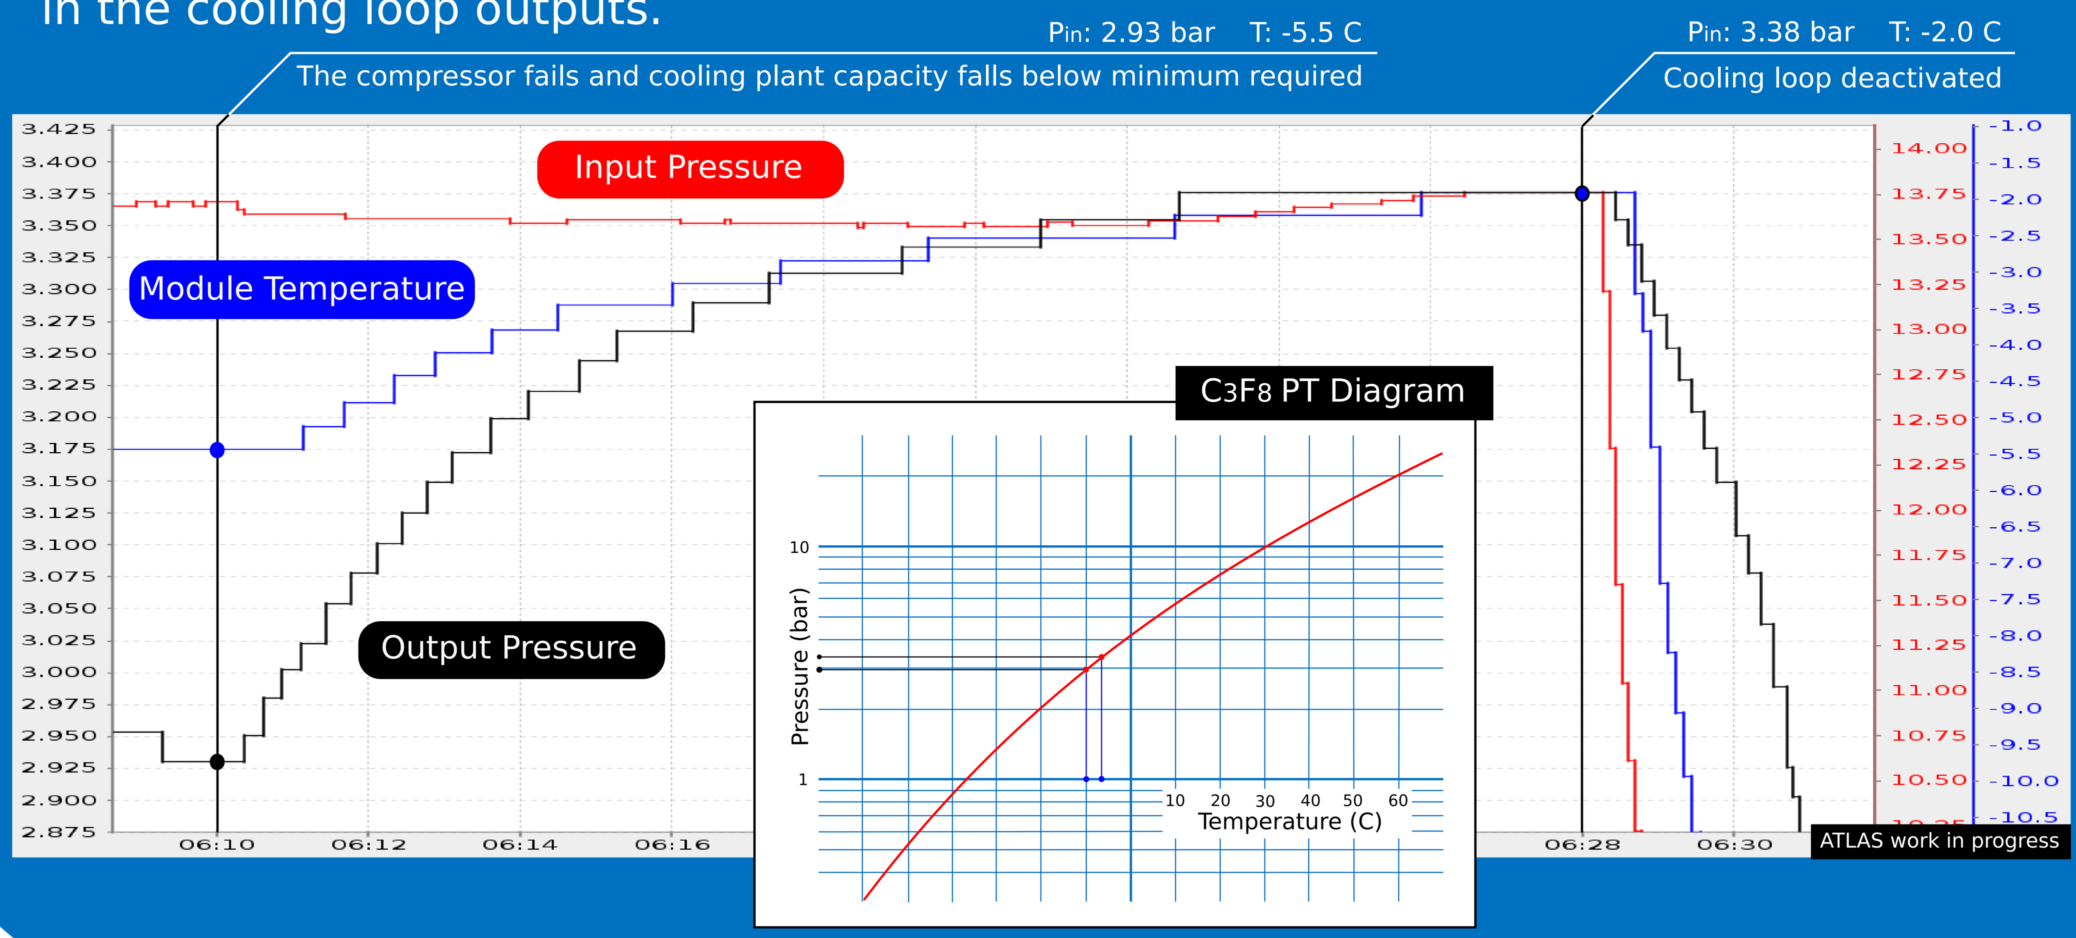

The increase in output pressure caused a corresponding rise in module temperature, although the detector continued to operate for 18 minutes. The fault conditions were detected by the monitoring software, and the cooling plant was manually shut down before either temperature or pressure reached levels that might be dangerous to the detector. Particularly useful were the new pressure sensors - installed in summer 2009 on the distribution racks - which meant the problem could be swiftly traced back to the cooling plant.

The problem was initially attributed to a faulty valve, before the reduced plant capacity was noticed. As the cooling plant had been designed with plenty of spare capacity (7 compressors, of which 4 are needed at any one time), no specific warning had been created to indicate when the number of available compressors was insufficient. This warning has now been implemented.

However, it was encouraging that even without it, this unanticipated problem was correctly handled by the cooling control software.

Operation Without Cooling

On March 9th 2010, the cooling plant shut down due to a short in a pressure sensor which also stopped communication with the detector control. Because the detector control system was not aware that the plant had stopped the silicon modules continued to operate.

The immediate result of the plant shutting down was that the coolant stopped flowing to the cooling loops. Additionally, there was no suction to draw the vapourised coolant exhaust out of the loops. As a consequence the input pressure fell, but not to the level of the output pressure.

Pressure rise: 0.40 bar
Temperature rise: 18.4 C

Greater temperature change than predicted by phase diagram - instead caused by lack of input coolant

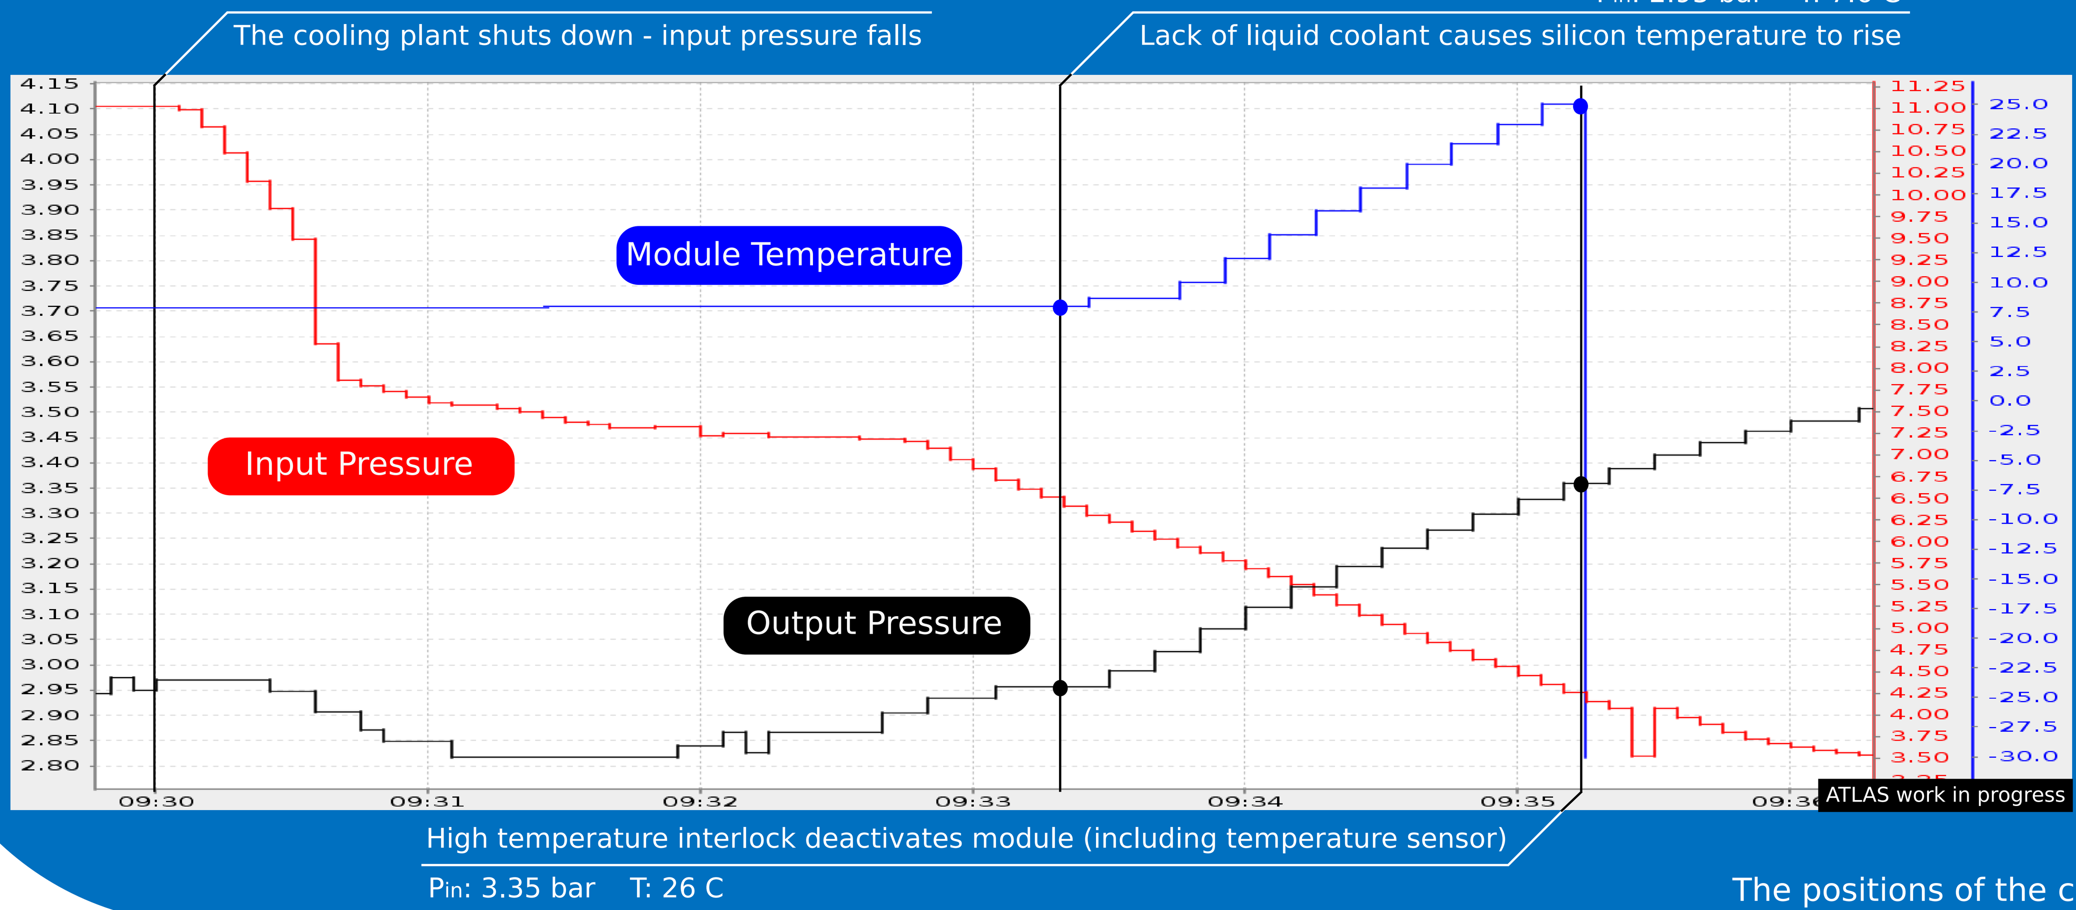

For 200 seconds the silicon temperature stayed constant as the remaining liquid coolant boiled off and the vapour pressure equalised throughout the exhaust system. Once this had finished the silicon temperature rose rapidly, going up 18 C in 2 minutes. Hardware interlocks were triggered by the high temperature and the modules were powered down automatically, protecting the detector from overheating: the detector was never at risk.

In response to this incident an additional, separate set of alerts from the cooling plant - via a separate line of communication - are now monitored. While the incident was unfortunate, it does suggest a new mode of operation. Since the silicon temperature was stable for so long after the cooling stopped, why not keep the modules active for a short time after the loops are shut down? This would allow uninterrupted data taking in the event of a minor cooling interruption, and would reduce thermal distortion of the detector.

The positions of the carbon cylinders supporting the SCT barrel are monitored by frequency scanning interferometers. On the right [4] we see ~5 micron distortions of the detector after the cooling is shut down, measured by the FSI.

Similar distortions were observed as the module temperature rose in the March 9th incident, and whenever the cooling is interrupted.

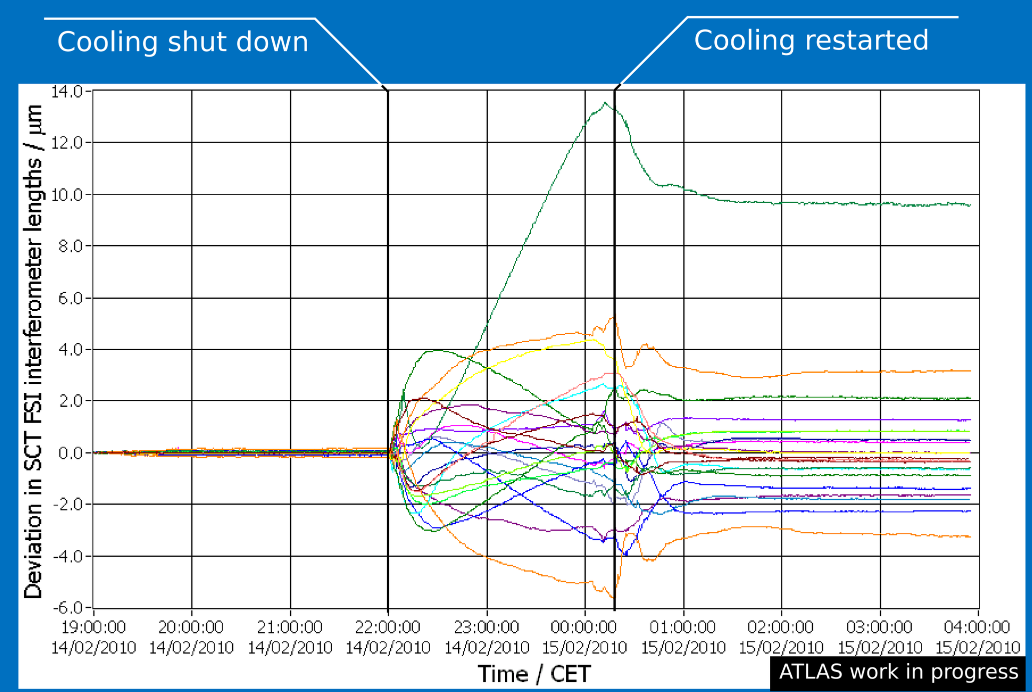

References:
[1] ATLAS Inner Detector TDR
[2] DCS plots using DCS Web Viewer tool by Luca Luisa
[3] Vic Vacek, CERN, July 1999
[4] Gibson et al. ATL-COM-INDET-2010-055

Ben Wynne <bwynne@cern.ch>
Supervisor:
Phil Clark <pclark@cern.ch>