Discovery machine

CERN's amazing achievement in September, with its large hadron collider launching the first beam, was a triumph of engineering. Brian Tinham talks to engineers behind the plant

CERN's Large Hadron Collider (LHC) – the planet's most powerful subatomic particle accelerator – made it into the record books when, at 10.28am on Wednesday 10 September, in front of the world's hyped-up media, the first beam was successfully steered around its 27km ring. Then, just nine short days later, there was an incident that resulted in a large leak of helium into the beam tunnel, forcing shutdown of this amazing project until spring of next year.

So much we all know. And we probably also understand that one of the key reasons for the wait was the sheer scale of plant warming required before inspection, repair and subsequent re-cooling could take place – as well as CERN's facility-wide winter maintenance closure. But what's really behind this outstanding development, dubbed the word's most ambitious 'discovery machine'? What exactly does it take, in terms of plant engineering, to get as far as CERN (the European Organisation for Nuclear Research) has with this project, which started more than two decades ago – while the earlier LEP (large electron and positron) accelerator was still running?

Dr Paul Collier, who heads up operations, first makes the point that this is one of the few achievements worthy of the term 'unique'. "It is the only plant ever to have launched a beam of subatomic particles circulating 11,245 times per second. And that beam and our control room are the only elements in this incredible structure that see everything working together."

It's a profound observation. Such is the scale of this plant – or, rather, several thousand machine elements – that, prior to September, even full-scale testing had only been done on individual segments.

Ultra-cool plant

So let's get some detail. Collier starts with the LHC's huge cryogenic plants, numbering eight in all (one serving each 3km arc of the giant sub-surface ring and its associated 600m straight section) and together taking its 6,000 beam-bending cryomagnets down to –271°C, just 1.9°C above 0K, the point at which molecular motion ceases.

"At each of those plants, cooling is performed in two phases. First is the compressor and turbo-expander plant, which cools helium down to 4.5K (-269°C, the point at which it becomes an incompressible liquid). That uses heat exchanger techniques similar to those of conventional refrigeration. Then, second, we have separate distributed plant that takes the temperature down to 1.9K, at which point helium becomes a superfluid, with huge thermal transfer capabilities – making it ideal for cooling our superconducting [niobium-titanium alloy] cryomagnets, because it extracts so much heat for such small temperature differences."

Serge Claudet, who leads the cryogenics

Automatic control

Much of the plant described is self-regulating – only exception data being transmitted to the control room, using conventional SCADA (supervisory control and data acquisition) systems. "We certainly don't observe every single parameter from the control room desks," laughs head of operations Dr Paul Collier. "There are over 30,000 temperature sensors alone. But we do manage all plant from here, with two operators looking after the beam and two more monitoring all eight cryogenic plants."

It's a classic control pyramid, but on a grand scale. "As a general rule, we start from the equipment itself, with local controllers connected to specific plant front-end computers, designed to manage the local loops automatically. Those are connected to the CERN plant network and so up to the control room," he explains. And that includes all data logging and control, as well as automatic safety systems.

We're talking about highly distributed controls and, interestingly, Collier makes the point that CERN has taken advantage of modern control technology throughout, including the use of digital fieldbus networks – and several of them. Profibus, for example, has been used for digital signalling with the range of intelligent valves and transmitters, for example on the cryogenic plants. Then, the WorldFIP standard (foundation Fieldbus) has been installed, mainly for instrumentation associated with the power circuits. And there's also a huge amount of Ethernet.

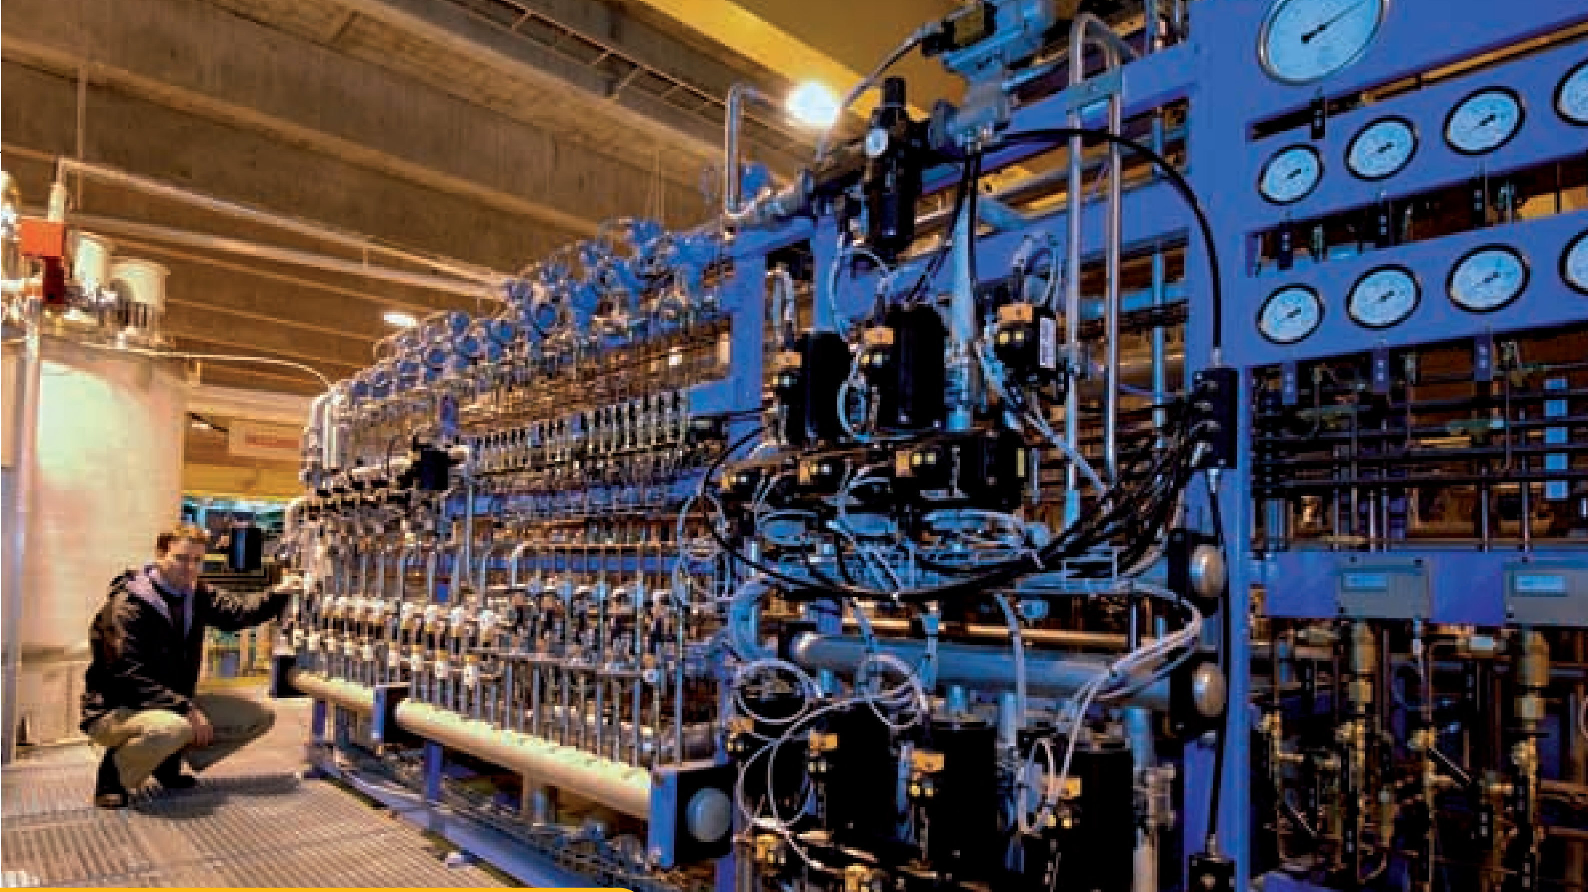

One of the eight cryogenic plants at CERN's LHC: conventional plant, extended and cascaded to achieve the most extreme ends

operations team, explains that most plant is effectively highly developed versions of commercially available equipment. "Consortia involving Air Liquide in France and Linde in Switzerland built the helium cold boxes – cascades of heat exchangers and expansion turbines – for the four new 4.5K refrigerators and the other four re-engineered units from our earlier LEP plants. They also partnered with companies such as Howden Compressors, York, Kaiser in Germany and Mycon in Japan for the screw compressors that drive them.

"Twenty years ago, we could only work with refrigeration compressors, because we needed their high tolerance standards, but the compressed air industry has caught up and there are economies. However, all of them have had to be adapted for helium, with special attention to preventing leaks. So the teams have engineered improvements in three directions: the lube oil, which is a synthetic compound; the gaskets and joints, which now use materials compatible with this oil; and the shaft seals, which maintain a leak-tight film of oil."

Getting some idea of scale, each compressor unit weighs in at 100 tonnes of equipment sitting on 120 tonnes of concrete, and consumes 4.2MW of power. That drives electric motors ranging in size from 500kW to 1.8MW and delivering 40,000m^3 of water-cooled helium per hour at 20bar to the turbo expanders. Meanwhile, each of the cold boxes takes 5MW, weighs some 100 tonnes, occupies 100m^3 and runs under local automatic control, with the computers connected to 850 instrumentation and control loops, and driving pneumatic control valves designed for cryogenic plants, with long stems to minimise thermal losses.

And that's just phase one cooling. Phase two works by rapidly reducing the liquefied helium saturation pressure down to 15mbar in an isothermal process, using what Claudet describes as "sophisticated hydrodynamic centrifugal cold compressors with special magnetic bearings, able to function at 1.9K, as well as screw compressors operating at ambient temperatures". Incidentally, in terms of control, we're talking about circulating a total of 15 tonnes of liquid helium around each 3km

Incident explained

Investigations at CERN have confirmed that the incident on 19 September was caused by a faulty electrical bus connection between two of the accelerator's magnets. That resulted in mechanical damage and release of helium from the magnet cold mass into the energy tunnel.

The good news is that the safety systems performed as expected. Also, CERN management believes it can prevent a similar incident in the future and says the organisation holds enough spare components to get the LHC restarted next spring.

The less good news is the scale of damage. CERN says that the electrical fault led to arcing, which, within one second, punctured the associated magnet's helium enclosure, leading to its release into the insulation vacuum of the cryostat. Three or four seconds later, the beam vacuum also degraded and then the insulation vacuum started to collapse in the two neighbouring magnet subsectors.

Spring-loaded relief discs on the vacuum enclosure opened when the pressure exceeded atmospheric, thus releasing helium into the energy tunnel, but were unable to contain the pressure rise below 0.15 MPa in the vacuum enclosure – resulting in large pressure forces on the vacuum barriers separating the central subsector from its neighbouring subsectors.

Apart from contamination over a length of the beam pipes, the investigating team found damage to the multilayer insulation blankets of the magnet cryostats. Also, the cryostats housing the quadrupole magnets involved broke their anchors in the concrete tunnel floor, while the electric and fluid connections also pulled the neighbouring dipole magnet cold masses from the cold internal supports inside their still intact cryostats – resulting in damage to the jumper connections, but not the insulation vacuum.

arc of the ring, cooling 5,000 tonnes of magnets – and with temperature maintained throughout the cryogenic pipelines to a stability of better than 0.05°C.

Magnetic attraction

That's the 'conventional' plant: most of the rest is distributed around the machine segments and, again, the engineering is impressive. For example, the cryomagnets themselves: "The main dipoles that bend the accelerator beam each weigh 35 tonnes and are 16 metres long, producing a field up to 9 Tesla – and there are 1,322 of them," explains Collier. And there are more, including 500 quadrupoles that focus the beam (each about 10 tonnes and around six metres long), as well as other multi-pole magnets for chromatic, coupling and energy correction – every one equipped with cryogenic cooling and dc power.

Quite apart from the numbers, think about the issues of installing all that massive, but fragile, equipment into the confined space of a 27km long circular tunnel. Collier says key developments were a tunnel vehicle, built by a German consortium of BNN (Babcock Noell Nuclear) and MAFI, as well as a transfer equipment set (TES), developed by Slovak firm ZTS VVU Kosice. The former was built on a platform just one metre wide and 15 metres long, while the TES (comprising two modules) measures just 80cm by two metres. Both are remarkable engineering innovations, with the TES, in particular, capable of transferring each huge magnet laterally, with millimetre precision – if necessary between previously installed cryomagnets.

And then there's the tunnel electrical installation: around 2,000 circuits to install, commission and test. "The current has to be very precise to meet the beam-bending requirements, and the quantity of energy we're talking about is substantial," says Collier. "Superconducting magnets act as very large inductors, so, with 154 main dipoles connected in series per arc, we run at 12,000 amps and 10V. When we're flat out, each sector has a capacity of 1.1 gigajoules, so we also have sophisticated protection systems able to extract all that energy fast and safely in a fault condition."

It's a similar story with the vacuum systems: three in all – one for the two intersecting pipes containing the counter rotating beams around the ring (effectively 54km maintained at 10^{-10}bar); another to provide thermal shields for each of the superconducting magnets; and the third offering insulation for the cryogenic feed lines to the magnets (the latter two at about 10^{-4}bar). "There are various vacuum systems – ion pumps and roughing pumps, and all the warm sections of the beam vacuum are coated with a material developed here at CERN that acts like fly paper for gas molecules," comments Collier.

So much for the technology; what about commissioning and testing? Claudet reiterates Collier's point that verification was necessarily conducted on a unit-by-unit basis – explaining that it was the only way, given the phased build programme, wide range of plant types and spread of suppliers. "We had to establish boundaries between plants and assess the performance of each in isolation, before it was connected. It would have been impossible to unravel any errors once the plant units were fully connected," he says.

And hence that substantial timeframe. For each of the main cryogenic plants, for example, equipment installation alone took around six months, followed by between 24 and six months of commissioning and testing. "There are so many quality tests to be completed before starting up a plant like this," observes Claudet. "It's not just about the services plant. There are the cryomagnet installation and commissioning in the tunnel itself; installation and testing of the cryogenic lines and valve systems; helium leak detection; and electrical QA, particularly for the superconducting circuits.

"Finally, we got the green light for cooling down, which took about six weeks, and was followed by another battery of tests – for example, to re-verify the superconducting circuits and check for energy extraction. Only then were we ready to qualify the system as part of the accelerator and pass the energy beam into the tubes."

Pointers

- CERN's LHC energy beam rotates at 11,245 revs per second around 27km
- Eight plants cool 6,000 superconducting magnets to 1.9K in two phases: adiabatic and isothermal
- Compressor technologies have been significantly extended for the challenge
- Each compressor unit weighs 100t and stands on 120t of concrete, and consumes 4.2MW
- There are 1,322 main magnets, each weighing 35t and 16 metres long