

Executive Summary:

The Large Hadron Collider (LHC) has raised the bar for particle physics, smashing together protons and lead ions at unprecedented energies. But even before the debris of the first particle collisions lit up the detectors in November 2009, researchers working on the project were confident that an upgrade that makes use of the latest knowledge and technology would help this accelerator go ten times as far. A higher density of protons in the colliding bunches - technically described as an increase in the 'luminosity' of the proton beams - produces more particle collisions. Several parts of the accelerator chain must be upgraded to provide denser proton bunches and raise the luminosity at the LHC and the detectors must also be upgraded to cope with the higher collision rate.

Improvements start at the very beginning of the chain of accelerators leading up to the LHC. The Linac4 accelerator will replace Linac2, injecting particles into the Proton Synchrotron Booster (PSB), after it has accelerated the beams from a standing start up to energies of 160 MeV. On its own, this faster acceleration brightens the beam by a factor of two. In a second step, the Superconducting Proton Linac (SPL) and an all-new accelerator ring, the Proton Synchrotron 2 (PS2), were to replace the old PSB and PS. The SPL would pack even more energy into the beam than the PSB, accelerating it up to 4 GeV. Work Package 7 tackled critical subjects for this accelerator, developing a new high-intensity particle source and the modelling tool to optimise the RF system. From the PS2, the beam would next enter the upgraded Super Proton Synchrotron (SPS) and, finally, the LHC. Here, a 'Phase-1' improvement to the 'Inner Triplets' - magnets that focus the beam - could deliver up to 40% more collisions inside the LHC's two largest detectors, ATLAS and CMS. Work Package 6 designed and has begun prototyping these magnets. Meanwhile, Work Package 2 developed accelerator coordination tools and standards in the context of a world-wide collaboration with multiple partners.

Once the high luminosity beams collide, the detectors need to be ready to capture ten times as many particles. Upgrade management teams in ATLAS and CMS, established in part through Work Packages 3 and 4, got the ball rolling on these improvements. They set up systems to decide what upgrades will be necessary and feasible and began coordinating research and development. From the outset, it was clear that the high-precision inner detectors of ATLAS and CMS would need more efficient powering systems. Work Package 8 contributed to the development and successful testing of two new powering strategies. To ensure that radiation associated with the LHC and its upgrades was minimised and safely managed, Work Package 5 simulated the radioactivity in several key locations and prescribed the best ways to achieve safe and environmentally clean operating conditions.

It is important to note that some of the goals changed during the three years of SLHC-PP. An accident in September 2008 meant the first collisions were not seen until November 2009. After months of smooth running, the CERN management adapted the upgrade schedule. In particular, the decision was taken to run the LHC for longer periods, punctuated by longer duration shutdowns. The Inner Triplet magnets, originally scheduled to be replaced in two phases during the shorter shutdowns, will be replaced in a single step instead. The upgrade will take place five years later than initially planned, providing more time for developing the Inner Triplet magnets and designing a suitable construction method. Finally, after a first analysis leading to a preliminary cost estimate, the CERN management also decided not to build the SPL and PS2, choosing instead to upgrade the existing PSB and PS. This changed the perspective of the SPL team's studies in WP7.

The SLHC-PP project, funded by the European Commission, took the first steps towards the LHC upgrade.

Project Context and Objectives:

Proton collisions are all about creating new particles. Albert Einstein showed that energy and mass were interchangeable currencies, and the Large Hadron Collider (LHC) is a bureau de change of sorts. The CERN chain of particle accelerators packs dense beams of protons into two beam-pipes of the LHC, beneath the Swiss-French border. Travelling in opposite directions through the 27-kilometre-long circular tunnel, the protons receive their final acceleration up to tera electronvolt (TeV) energies (3.5 TeV today, and 7 TeV after 2013).

The LHC then collides the protons together at the hearts of four particle detectors - ATLAS and CMS being the largest - that intersect the circular tunnel. When these collisions occur, the energy in the protons transforms into mass, generating a shower of new particles. The higher the energy of the collision, the heavier the particles that form in the debris. And they aren't just any particles - occasionally, the LHC will produce objects and interactions that have never before been seen by physicists, providing vital experimental evidence to back up the current theories of why the Universe is the way it is.

Though particle physics studies the very small, new observations could reveal the explanations for cosmic mysteries - what gives objects their mass, the nature of the 'invisible' matter that makes up most of the Universe, how the four forces that govern everyone and everything might unite into a super-force, and - maybe - why gravity is so much weaker than the other three forces.

The LHC is already the most powerful particle collider in the world, but for it to live up to its true potential, it must be upgraded to run at the highest possible luminosity. If the LHC is imagined as a torch probing the darkest recesses of the Universe, it's easy to see why a brighter beam would be more useful. A torch is made brighter by increasing concentration of photons in the beam, but the LHC beam will be brightened by making the proton bunches within it even denser.

Most of the proton collisions in the LHC will produce debris that is already familiar to physicists. The large or unusual particles they are most interested in are diamonds buried in heaps of common rubble. But while diamonds tend to stick around, the new particles are often extremely short lived: they disintegrate in fractions of a billionth of a second. To catch a glimpse of the fleeting particles, researchers take images of them through the four major detectors - recording particle 'signatures'.

Physicists rely on a complicated statistical process to separate the rare signatures of new physics from the huge number of similar signatures already expected to show up due to known physics. The larger the quantity of data physicists have to search through, the greater the chance of finding and confirming the signatures of rare particles.

The LHC upgrade will extend the lifetime of this scientific exploration into the unknown and produce ten times more data during the collider's second decade than during its first.

The main aim of the SLHC-PP was to prepare the first phase of this upgrade project for a formal proposal in 2011. Researchers needed to demonstrate new technologies, find solutions for radiation safety issues, form the collaborations that would carry out the work, and prepare the necessary management tools.

The SLHC-PP explored key upgrade pathways, demonstrating solutions to better focus the beams, design key components for the beam injectors, and power detectors more efficiently. A dedicated SLHC-PP team addressed radiation safety concerns by analysing the immediate and long-term effects of beam loss, in view of designing means to ensure safety and minimise the impact on the environment. On the collaboration side, the SLHC-PP contributed to the efforts in planning and initiating the upgrade projects of the ATLAS and CMS detectors. With the support of the SLHC-PP, existing tools for financial management, monitoring progress, and quality assurance have been improved, using the experience gained during the construction of the LHC.

The SLHC-PP activities already represent work from 18 institutions in 10 countries, but they also join with larger efforts coordinated through the detector collaborations and CERN's accelerator departments.

Work Package 1: Project management

The multifaceted SLHC-PP effort was managed by the WP1 team - comprising the Project Coordinator, the Deputy Project Coordinator and the Administrative Manager. Roland Garoby and Duccio Abbaneo succeeded Lyndon Evans and Lucie Linssen as Project Coordinator and Deputy, respectively, at the beginning of 2010. Mar Capeans has been the Administrative Manager throughout the project. The members of WP1 handled contracts and finances, monitored and reported progress in the other work packages, and took care of disseminating information both within the SLHC-PP and outside the collaboration.

In the first year, the management team set up a collaboration website (Deliverable report 1.2.1), much of which is available for anyone to view. They also commissioned a website to address the general public, explaining the purpose of the SLHC-PP, the upgrades, and how they are funded.

The WP1 team arranged annual collaboration meetings, at which representatives from each partner institution presented their progress to the whole of the SLHC-PP collaboration. In addition, each work package group had the opportunity to discuss different challenges and approaches, bringing together collaborators from many institutions (Milestone reports 1.1, 1.2, 1.3, and 1.4).

The project management team also served as a bridge between the SLHC-PP collaboration and the European Commission (EC). The meetings gave them a chance to share instructions from the EC about managing finances and preparing reports. Likewise, the collaboration prepared information to send back to the EC - the governing board of the SLHC-PP met to approve the deliverable and milestone reports, which demonstrated the progress of the work packages.

The Work Package leaders summarised each year's progress and financial information, and WP1 team edited these contributions into annual reports - always submitted to the EC on time (Deliverable reports 1.1.1, 1.1.2, and 1.1.3).

One of the SLHC-PP's most important roles was facilitating communication between accelerator and detector groups, and also between the two scientific collaborations working around the ATLAS and CMS detectors. Though these collaborations are rivals in the race for exciting discoveries, they face similar sets

of problems as they prepare to deal with a luminosity upgrade that will deliver ten times as many collisions to their detectors and produce ten times as much data as at present.

Public events, organised by the project management and held at CERN, aimed at informing as many people as possible about the latest ideas and plans for the LHC upgrade. "The meeting in 2010 was especially important for explaining the change of strategy resulting from the lessons learnt from the accident in 2008 and the repair in 2009," says Roland Garoby. These are detailed in the Impact section.

As Garoby mentions, cutting-edge scientific research never goes precisely to plan, and it fell to the WP1 team to explain changes to the schedule and to make alternative proposals when some work packages ran into difficulty.

"The Commission was very open to our proposals for adapting the SLHC-PP Project to the changing LHC context," says Garoby. He calls it a nice example of good and non-bureaucratic management, demonstrating an "understanding of what pushing the performance of a state-of-the-art scientific instrument like the LHC is all about".

Work Package 2: Accelerator upgrade coordination

Five accelerator sub-projects were underway as of April 2008 in preparation for the upgrade to the LHC. Many of the links in the chain of smaller accelerators that gradually speed up the protons before they are finally injected into the LHC's tunnel were scheduled for upgrades. New links would include Linac4, the Superconducting Proton Linac (SPL) and the Proton Synchrotron 2 (PS2). The penultimate accelerator, the Super Proton Synchrotron (SPS), also needed improvements to handle high luminosity beams. In the LHC itself, the Inner Triplet magnets that sharply focus the proton beam just before it enters each detector cavern were also due for replacement.

In building the LHC, CERN made individual agreements with each institution that provided components and personnel. However, because the LHC upgrade contains five different projects working towards a common goal, with funding provided by several bodies, keeping the efforts integrated requires attention. The WP2 team helped prepare this collaboration and improved communication and management tools.

"Work package 2 could draw upon CERN's decades of experience in shaping international collaborations for building and operating particle physics detectors," says Thomas Otto, coordinator of WP2. Such a collaboration has to assemble accelerator engineers and technicians "who look for high-tech solutions which are operable 24 hours a day, seven days a week - and which may last half a century", he says.

The WP2 group contributed to drawing up membership procedures for the future accelerator collaborations, and created a schedule and cost plan. They also joined the effort to set up systems for managing finances, monitoring work done by collaborators, and ensuring the quality of contributed parts and installations (Milestone report 2.1, Deliverable report 2.1.1).

Communication

Previously, the main decision-making on accelerator projects happened at CERN, and because most collaborators worked at CERN, they could frequently meet in person.

In contrast, the collaborators on the LHC upgrade are spread across Europe, Asia and North America. Although they cannot simply walk to a specific CERN conference room, video conferencing through the web allows them to continue meeting regularly. The upgrade groups also assemble for week-long workshops at CERN several times each year (Deliverable report 2.2.1). Between meetings, project websites connected with databases form common repositories of technical information, which are accessible worldwide at all times (Deliverable report 2.2.2).

Quality Assurance System

"The products that we need are quite complicated and depend on a lot of different suppliers and contributors," says Otto. While comparing the quality standards for different accelerator components is often like comparing apples to oranges, the procedures for acquiring these devices can at least be standardised through a Quality Assurance (QA) System.

Orders for accelerator components begin with market surveys: private companies are asked to propose how they could produce the parts, and the LHC team then chooses the best offer. Finally, the delivered products must meet certain criteria. These standards were first introduced in the building of the LHC, and they now cover components from collaborating universities and laboratories as well as those from private companies (Deliverable report 2.1.2).

Earned Value Management

The WP2 team tailored the Earned Value Management (EVM) software, which successfully managed the LHC construction after 2003, to the needs of the upgrade effort. EVM breaks down a large project into manageable chunks, achievable by one institution within a single year. The costs of material and manpower are fully incorporated into the project schedule as each step is defined in terms of money spent (Milestone report 2.2, Deliverable report 2.1.3). "In its revamped version, EVM statistics are available as a handy 'dashboard' within CERN's planning software, avoiding the need to transfer unwieldy data files from one application to another," says Otto.

With regular communication and close monitoring of progress through the EVM and QA systems, the many facets of the upgrade to the LHC combine in a single well-coordinated effort.

Work Packages 3 and 4: Coordinating the ATLAS and CMS upgrades

The present CMS and ATLAS collaborations each contain over 2000 people, scattered around the world. The upgrade effort, once it hits its stride, will be just as far-flung. In order to manage the upgrade research, design, construction, and installation, a management team needed to be in place from the beginning. These teams, defined as part of Work Packages 3 and 4 for ATLAS and CMS respectively, had two primary objectives: coordinating upgrade efforts among the groups in charge of various parts of the detector, and making sure any upgrade work is compatible with the existing detector.

Project management

The upgrade management teams assessed the scope of the upgrades, calculated their costs, and drafted schedules for how they might be completed (Milestone reports 3.1 and 4.1, Deliverable reports 3.2.2 and

4.2.3). In addition, they initiated the financial planning for the upgrade work. Teams to shoulder these responsibilities were appointed from among the ranks of the existing ATLAS and CMS collaborations (Deliverable report 3.1.1 and 4.1.1). An arm of the management known as technical coordination ensures that upgrades are feasible and compatible (Deliverable report 4.2.1).

ATLAS headed their upgrade project with an Upgrade Steering Committee while the CMS team is led by an Upgrade Management Board. These groups each drew up global agreements called the Initial Memoranda of Understanding (Deliverable Reports 3.1.2 and 4.1.2) which are documents to be signed by the collaboration management and the national funding agencies contributing to the upgrade. The funding agencies then provide resources to universities and laboratories that join the ATLAS and CMS upgrade collaborations, allowing them to contribute components or money to support the upgrade.

The management of each experiment defined the decision-making structure, or how upgrades go from an initial idea, to a design, to an accepted project - including detailed specifications, installation procedures, and safety considerations (see WP5 for more on radiation risks). They also decided on reporting procedures (Deliverable report 4.2.2, Milestone report 3.2).

The CMS and ATLAS upgrade management teams are already making many important decisions about the scope and scheduling of the upgrades (see the Science and Technology section). They have assessed the cost of upgrading the detectors, estimated at about €160 million each for ATLAS and CMS. These figures include materials, engineering, and staff effort (Deliverable reports 3.1.3 and 4.1.2). The installations will be distributed among several LHC machine shutdowns over the next decade.

ATLAS and CMS each have teams of technical experts that, among other tasks, check for compatibility between the present experiment and proposed upgrades. These teams, called the Project Office in ATLAS and the Review Office in CMS, are arms of the technical coordination groups. As the upgrade work integrated into the existing experiments, the Project and Review Offices became increasingly central bridges between the upgrade teams and the existing technical coordination for the two collaborations.

Technical coordination faces a complex working environment. These teams are under significant pressure to remove old components and install new ones without damaging nearby parts of the detector - a difficult task when some subdetector components, such as the pixel detectors, allow only a few millimetres' clearance on each side.

"The technical coordination plays a very important role," says Steinar Stapnes, leader of WP3. "In some ways even more important than in the original experiment, as the constraints and complexity of the experimental environment are higher."

From idea to upgrade

For each proposed upgrade to an existing subdetector component within the ATLAS or CMS detectors, an important step in the process is a simulation to establish that the new device will bring the improvements to detector performance that it promises. At the same time, researchers begin developing the prototypes needed to demonstrate the feasibility of the proposed solutions. All this happens within the group in charge of upgrading that particular subdetector. Once these steps have been cleared, the final idea can be presented to the relevant Steering Committees or Management Boards.

At this point, the idea goes through the management of the existing experiment: the Project and Review Offices ensure compatibility with other subdetector components, and the resource management team ensures that funding exists to take the idea further.

If it all checks out, a formal proposal including schedule, delegation of work, and cost is submitted to the LHC Committee (LHCC) and Resource Review Board (RRB). These bodies provide oversight for CERN activities. The LHCC is composed of respected experts who review CERN science and technology proposals, selecting the most scientifically promising projects and ensuring that they are sound. Representatives of the funding agencies make up the RRB, and this group oversees the finances.

CMS upgrade project leader Jordan Nash calls the system "federal" as it gives the group in charge of each subdetector component a lot of latitude to choose its own best upgrades. At the same time, a second tier of management can coordinate the upgrades of each group to make sure they are compatible and serve the broader performance goals of the overall detector.

Mapping the present and future of the detectors

In upgrading ATLAS and CMS, the collaborations are taking state-of-the-art technology and pushing it further. Each detector, which is essentially a digital camera with hundreds of millions of pixels, fills a space half the size of the Notre Dame cathedral in Paris and contains as much as 13 thousand tonnes of material. In order to effectively upgrade the detectors, the collaborations need to keep close tabs on current and future equipment specifications.

To do this, the ATLAS Project Office and CMS's Technical Coordination team have set up central databases, accessible through the Web, containing the designs and locations of each detector component. Dozens of users are already adding the next generation designs for upgraded parts to these databases as they become available, including new designs for sensors, electronics, support, and cooling devices.

Many of the original drawings were made with previous generations of computer aided design programs. These old drawings have been converted to the new standard programs, ensuring that the descriptions and specifications for even the oldest parts of the detectors - developed in the 1990s - would not be lost.

Deputy project leader of the SLHC-PP Duccio Abbaneo says these drawings are crucial. "If you want to plan an upgrade of a detector that is inaccessible and in a radiation environment, as a starting point you must have a very detailed and accurate model of what you have built, because you can't go there and check."

The databases aren't just for detector components - they also include equipment for taking apart the present detectors and installing the new components as well as drawings to show how installations could proceed (Deliverable reports 3.2.3 and 4.2.2).

Work Package 5: Radiation Safety

Radiation strong enough to damage the LHC machinery deep below ground is an unavoidable side effect of accelerating and colliding particles. At CERN, assessing and managing this risk is a top priority. While the beam is running, physicists need regular access to auxiliary caverns, not far from the LHC tunnel and

detector caverns, which contain electronics and cooling pumps. And during downtime, when the beams are no longer circulating and colliding, they need access to the LHC tunnel and the detector caverns in order to carry out any repairs and upgrades. But the detectors and accelerator components continue to emit radiation for some time after the beam has stopped circulating, much like an oven takes time to cool down once it is turned off.

To gauge the increased radiation hazards associated with the luminosity upgrade, the WP5 team has simulated the expected radiation in areas where people are expected to work during beam. It has also assessed the degree to which magnets and detector components will be 'activated', continuing to radiate after the beam is stopped.

Key results:

- Safety upgrades were implemented to protect Linac4 construction workers without need of a radiation area
- Upon replacement, collimators and inner triplet magnets will be classified and disposed of as radioactive waste
- The ATLAS service cavern must be a controlled radiation area while the upgraded LHC is running
- The shafts delivering power from the surface to an SPL-like machine will need protection against escaping radiation
- The beam dump of a PS2-like machine can be made safe for access without cool-down periods
- The upgraded PS will not make significant amounts of air radioactive.

Work Package 6: Final focusing magnets

Just before the two beams enter the detector caverns from opposite sides, a set of focusing magnets known as the 'Inner Triplets' squeezes them down. By leaving less empty space between the protons in a bunch, these focusing magnets increase the number of protons that actually collide with those from an oncoming bunch.

The Inner Triplets will suffer significant radiation damage, the accelerator version of wear and tear, over the first decade that the LHC is running. By around 2022, they will need to be replaced. "If we replace them, we might as well make use of later technology and knowledge we have gained with the LHC project to increase the luminosity," says Stephan Russenschuck, leader of WP6.

With over a decade's head-start, the WP6 team is already developing new magnets that can squeeze the protons into more tightly packed bunches at the point where they collide inside the detectors.

Key results:

- The suite of final focusing magnets were designed
- New equipment was developed to make and test the new magnets
- Test coils for the main magnets have been made; these passed electrical tests
- Two corrector magnets are completed, with the last finished by the end of 2011

- Two 2-metre-long prototypes of the main magnets are under construction

Work Package 7: Towards a new accelerator

If the Superconducting Proton Linac was to be built, it would have followed on from Linac4, accelerating negative hydrogen ions from 160 mega electronvolts (MeV) to 4 giga electronvolts (GeV). As it was designed to accelerate longer, more frequent pulses of ions, the SPL project team needed to design a new negative hydrogen source to be placed at the beginning of Linac4.

These two linear accelerators would endow the beams with almost three times as much energy as the current Linac2 and Proton Synchrotron Booster do, leading to brighter beams entering the new synchrotron, PS2. The improved brightness relates to a peculiarity of Einstein's theory of special relativity. Particles effectively get more massive as they approach the speed of light. Since objects of greater mass are harder to move, the repulsion between the particles - brought on by their negative charges - seems weaker.

In 2007, it appeared affordable to design and build the SPL and PS2, but in 2010, after a detailed analysis and first cost estimate, the CERN management decided against it as part of the baseline upgrade to the LHC. Still, the WP7 team continued with the development of the negative hydrogen source and studies for developing control systems for the accelerator's radiofrequency cavities. This work is expected to benefit the nuclear and particle physics community worldwide in other ways, detailed in the Impact section.

"The construction of the SPL remains only as a back-up plan for the LHC," says SLHC-PP project coordinator Roland Garoby, "but even technologies developed through the SLHC-PP that are not of direct use in the new plans have important spin-offs".

Key results:

- A plasma chamber was designed to cope with high driving rates and constructed
- Two more prototypes study different magnet setups
- The first measurement of optical hydrogen emission lines in a pulsed plasma was made
- Two existing superconducting radiofrequency (RF) cavities were studied in depth, providing information for modelling
- Computer models were developed, showing that it is possible to power two RF cavities with a single amplifier
- Control boards for this system were developed and will be tested

Work Package 8: Power delivery at the trackers

The tracking detectors at the centres of the ATLAS and CMS experiments presently contain 86 million and 75 million individual sensor channels, respectively, to trace the paths taken by each particle created in the collisions that occur at the detector's core. The sensor channels are arranged in tens of thousands of modules, and each of these has its own 100-metre-long cable to supply it with power. With the upgrade

work in the ATLAS and CMS detectors (see WP3 and WP4), the number of sensor channels in the tracking detectors will rise by a factor of ten.

There are three good reasons to find a new strategy to power the three-quarters of a billion sensor channels in each upgraded tracker, detailed in the Science and Technology section. The most obvious is space: there simply isn't enough room for all those new cables.

The WP8 team has contributed to two new detector-powering schemes: DC-DC conversion and serial powering. The DC-DC conversion scheme allows a small current to run through the long cables at a high voltage, so the cables can be much thinner, making room for additional cables. A DC-DC converter then provides the small voltage and high current needed by the chips on the sensor modules that export data. Serial powering, meanwhile, delivers power to several sensor modules through a single set of cables - again helping to free up space and reduce the power dissipated in the cables.

Key results:

- A DC-DC converter design was selected and radiation resistant control systems were developed
- The converter board was optimised to avoid interference with the sensors
- Prototype converters were constructed for the future tracking detectors
- DC-DC converters were integrated with present detector modules
- Four strategies for serially powering detector modules were developed
- Timing signals coming into the detector module and data signals coming out were standardised
- A scheme was implemented to bypass any problematic module
- The method was successfully tested in ATLAS tracker assemblies

Plan for the future

The High Luminosity LHC (HL-LHC) and LHC Injector Upgrade (LIU) projects will implement the upgrade plans developed in part through the SLHC-PP. The present running plans for the LHC are shown in the figure to the left, superimposed on graphs which map the expected accumulation of data. Though the precise dates may change, engineers will make the 'splice repairs' which will allow the LHC to run safely at its full design energy of 7 TeV during the next long shutdown, starting early in 2013. ATLAS has also scheduled its first major upgrade for this time. Around 2017, accelerator engineers will turn their attention to installing new equipment for the upgrade of the accelerators. The CMS collaboration will also begin making major improvements to their detector - some making use of the new powering schemes.

After 2020, both detectors will finalise their upgrades, and new Inner Triplets will be installed in the LHC ring. The accelerator and detectors should run into the 2030s.

Project Results:

Science and technology results

A variety of technologies have been planned, designed, and prototyped through the SLHC-PP, and scientific studies have been undertaken. Results of this work include identifying key detector upgrades, the findings of the radiation studies, the designs and prototyping of the final focusing magnets, the research towards a new accelerator and source, and the development of two new schemes for powering the central subdetectors within the ATLAS and CMS detectors.

The futures of the detectors

Through the management bodies and procedures developed with contributions from the SLHC-PP, the CMS (WP4) and ATLAS (WP3) upgrade teams have identified key improvements needed to make the most of the extra particle collisions in the upgraded LHC.

WP4: Upgrades for CMS

The CMS collaboration will install the first major upgrades during the shutdown scheduled for around 2017. At this time, CMS will receive a new pixel detector - the subdetector component that sits at the core of CMS, closest to the path of the LHC particle beam. CMS collaborators will also replace the photodetectors in the hadronic calorimeter, a subdetector further from the particle beam that measures the energy of particles containing quarks. Then, in a later shutdown after 2020, the entire tracker subdetector will be replaced, and nearly every other system will receive some improvement. Of note, the system for recording the most interesting proton collisions will be upgraded.

Pixels

CMS's silicon pixel subdetector is the highest-resolution detector in CMS, tracking particles fresh from the collisions in three dimensions. When the LHC is operational, 10 million particles per square centimetre pass through the pixel subdetector each second. This rate is expected to increase by almost a factor of ten when the LHC is upgraded.

Added sensing layers in the silicon pixel subdetector will provide an additional handle for teasing apart the particle tracks streaming through the detector. An additional layer of pixels will be added, together with extra pixel-filled 'caps' on each end of the cylindrical pixel subdetector.

But the pixel upgrade involves more than adding extra sensors. The CMS pixel group is also cutting out detector dead space - things like cables, cooling, and structural elements. These components don't sense the particles from the collision, but may absorb particles or break up clean, high-energy tracks into a cascade of smaller particles that make interpreting those high-energy tracks more difficult.

Part of the material savings come from new power management schemes, detailed in WP8. Large cables and connectors will be replaced with streamlined versions that take up less space. The carbon fibre support structure for the pixel detector will also be lighter. By using carbon dioxide to cool the detector rather than the present fluorocarbon coolant, the CMS team can use thinner cooling pipes as well as smaller heat exchanger contacts.

This design is scheduled to be implemented in 2016. A third iteration of the pixel subdetector may be installed in the CMS detector in the early 2020s.

Photodetectors

The hadronic calorimeter in CMS is composed of brass or steel "absorber" tiles, layered with scintillating plastic tiles. The calorimeter tries to absorb all of a particle's energy - forcing it to create cascades of smaller particles as it passes through the brass or steel - and then measures the resulting particles with the scintillating plastic. Scintillators give off light when charged particles pass through, and this light is collected in optical fibres, which take it to a photodetector to be converted to an electrical signal for analysis.

Photodetectors are specially designed photodiodes that can operate in the exceptionally strong (4 Tesla) magnetic field that exists inside the CMS detector. However, the photodetectors currently installed in CMS sometimes spark when exposed to that strong magnetic field: charges build up inside them, and when discharged, the signals look confusingly like that of a high-energy particle. "It makes you think something happened and it didn't. You get a big spike of energy that isn't real," says Nash. The detector may waste time by exporting data from these 'phantom' collisions.

Silicon photomultipliers could avoid the sparking problem. This technology wasn't around when CMS was being designed - they were first developed in Russia in the late 1990s. Silicon photomultipliers operate at a low voltage and are resistant to the charge build-up seen in the photodetectors now deployed in CMS. Moreover, silicon photomultipliers have a signal-to-noise ratio ten times better than the present photodetectors, making the data they produce more easily interpreted.

Finally, the silicon multipliers can be installed in a way that gives information about the rate at which a particle deposits its energy in the calorimeter. In the current design, the particles essentially travel 'up' a tower of alternating scintillator and absorber tiles. The photodetectors have to sum up the total signal from the scintillators in the tower to give the original particle's energy. But the silicon photomultipliers can differentiate between individual scintillator tiles, providing information about how much energy is deposited in each 'floor' of the tower.

Trigger and data acquisition - recording the important particles

Most digital camera users don't print every photo they take, and nor do particle detectors record every collision they see. In fact, the vast majority of collisions - over 99% - are discarded the moment they are measured. This is because they contain only well established particles and interactions, which are uninteresting in the search for new physics.

The first port of call for data is the trigger, a system which decides which particle collisions to keep and which ones to discard. Presently, the electronics that export data from the detector are connected to the

trigger through electrical wires. In order to cope with higher data rates, the CMS team will replace these with an optical system. "You can send a higher volume of data through an optical fibre," says Nash - one reason why internet broadband networks around the world increasingly rely on optical fibres rather than traditional copper wires.

By increasing the processing power of the data acquisition system, CMS will be able to look at the finer details of the data coming in. The detector systems currently assess whether a particle has passed through a sensor or not, but they can squeeze more information out of the signal by looking at nuances of the particle's energy and location. To harness the power of the upgraded processors, the CMS team is developing new software.

For more information on the scope of the upgrade see this draft of the plans:
http://www.hep.ph.ic.ac.uk/~nashja/plutpla_auto.pdf (over 45 MB).

WP3: Upgrades for ATLAS

The silicon sensors, CMOS chips and controllers within the ATLAS detector, and the optical fibres for exporting data, are all designed to be "radiation hard" - they can withstand the flux of high-energy charged particles flying out of the proton and lead ion collisions. However, even these durable devices eventually break down after a few years to a decade of exposure, so replacing them will be a necessary part of any upgrade work. At the same time, the electronic systems for sending data out of the detector and choosing interesting events may be improved, streamlining data analysis. The 'end cap' regions of the detector, which close off the barrel-shaped central parts, will also need upgrades to cope with higher radiation rates.

As in CMS, most ATLAS subdetectors are expected to eventually receive some upgrades. While the Insertable B-Layer is the current main project, planned for the shutdown period 2013-14, others are under various stages of planning and consideration.

Insertable B-Layer

The pixel detector contains silicon sensors, each not much larger than a hair's breadth to a side. Eighty million of these, distributed across three cylindrical layers and four end-cap discs, precisely track the paths of particles as they leave the collision point.

The innermost layer of the pixel detector, known as the B-layer, is the first port of call for particle debris leaving the beam pipe and entering the detector, and the concentration of particles endured by this first layer of sensors is higher than anywhere else in the detector. Originally, this layer was to be removed and replaced, but removing the old layer was too risky, and the operation would take too long.

It is easier to instead remove the beam pipe and replace it with a slightly smaller version. This leaves enough space to add a fourth cylindrical layer, nesting inside the current B-layer. ATLAS upgrade project leader Steinar Stapnes says this new pixel layer will "add a very important early tracking point of particles leaving the interaction region and increase the robustness of the entire pixel system".

Further upgrades

The ATLAS collaboration is planning several other upgrades, in various stages of study and approval:

- The trigger team, which develops the automated systems for choosing which particle collisions to record, is developing an improved trigger which draws on data from the particle tracking systems. "It will help select interesting events both faster and more efficiently," says Stapnes.
- In general, the trigger team is trying to combine information from many parts of ATLAS as early as possible. The trigger runs through a chain of selection criteria to decide whether to keep or discard a proton collision, and combining multiple detectors earlier can allow ATLAS to adjust the selection system to more effectively seek out or ignore specific kinds of collisions.
- The subdetector devoted to muons, electron-like particles that often herald interesting collisions, will need upgrades for its end cap 'wheels'. The revamped LHC will produce muons in such abundance that more than one will regularly come through a sensor at the same time. An extensive research and development effort is looking into the best type of replacement.
- The upgraded LHC will expose the end caps of the calorimeter system, which stops particles and measures their energies, to the maximum barrage of particles that they were designed to withstand. "Two solutions are being considered, replacing it with version more robust with respect to high rates, or putting another calorimeter in front of it where a significant part of the energy of the incoming particles is deposited," says Stapnes.
- To keep the proton beams from colliding with air molecules, they are kept inside pipes which have been emptied of air. To make this 'beam pipe' more transparent to particles passing through, ATLAS is looking to replace it with something made of lighter materials. And, so that sensors can be placed as close to the collisions as possible, the upgrade team is also trying to reduce the beam pipe's radius near the centre of the detector.

The largest and most costly upgrade needed is the replacement of the entire ATLAS Inner Detector, expected to happen in a decade's time. "A very comprehensive research and development program is underway," says Stapnes. This program of R&D will produce new designs for silicon sensors, data export systems that can withstand high levels of radiation, cooling systems, and support structures. The upgrade will comprise both the silicon strip tracker and pixel detector, for which new powering schemes were developed in WP8.

"The planning and detailed research and development for all these upgrades are well underway," says Stapnes. "The formal steps needed to move them one by one into realisation will follow the procedures established for the Insertable B-Layer project, and these have been developed with the support of the SLHC-PP WP3 funding and activities."

WP5: Radiation Safety

Higher luminosity beams and collisions have one unfortunate side-effect: they create more radiation. Some of it is along the accelerator chain as protons occasionally stray from the beam, and some of it is in and around the detector caverns, where ten times as many protons are colliding in tiny explosions of new particles. The increase in radiation needs to be predicted ahead of time so that additional safeguards can be implemented to ensure that neither personnel nor the environment will be harmed.

Radiation studies also looked into the radiation risks associated with implementing upgrades, such as the construction of Linac4 and the replacement of the accelerator components nearest the collision point. These studies are ongoing at CERN, and will eventually cover all radiation concerns connected to the upgrades.

Originally, WP5 intended to study six areas: the point where the Linac4 construction site approaches the running Linac2 accelerator, the radiation that might escape the hypothetical Superconducting Proton Linac (SPL) and Proton Synchrotron 2 (PS2) accelerators, the radiation that enters the service caverns of ATLAS and CMS, and the activation of the Inner Triplet magnets, which are slated for replacement after 2020.

However, because the PS2 and SPL are no longer part of the luminosity upgrade at the LHC, in 2010 WP5 instead began to consider how much radiation is likely to be emitted by the older accelerators running at higher luminosity. Specifically, focus turned to the Proton Synchrotron (PS), which WP5 Project Leader Thomas Otto calls, "the more than 50 year old centrepiece of CERN's accelerator chain."

A good starting point for radiation safety is the law: anyone expected to come into contact with ionising radiation through their work must not be exposed to more than a 20 milliSievert (mSv) radiation dose per year. This applies to personnel at hospitals or nuclear power plants, and to those at CERN working near the running accelerators and particle detectors. For comparison, the French Institute of Radioprotection and Nuclear Safety estimates that the annual dose for flight personnel is 5 mSv on the routes with the highest exposure to radiation from space, for example New York to Tokyo or Paris to Tokyo.

CERN then takes radiation safety further, following the ALARA principle, which ensures that radiation doses are As Low As Reasonably Achievable. The upshot: it's preferable to avoid exposing personnel to radiation by shielding even those areas that pose a very low risk of contributing significantly to an annual radiation dose.

One of the first steps for the WP5 team was to identify the most important factors that contribute to radiation risk. In accessible areas while beam is running, these are: the energy of the beam, the number of particles accelerated with each pulse, and the average amount of beam passing through an area at any given time. After beam has stopped, activated materials can also emit radiation (see box: Activation).

The team also identified the locations in the accelerator chain where particles from the beam were likely to escape and pose a radiation risk. One such point is at the collimators, blocks of heavy metal (often tungsten) that strip straying particles off the beam by passing it through a relatively narrow hole, which are spaced at regular intervals around the accelerator.

The beam is less stable at the points where it enters and leaves accelerators, so these areas are also more likely to be exposed to escaping particles. Finally, beam dumps - massive blocks of tungsten, iron, concrete, and other materials designed to absorb the beam when it is no longer of use - are liable to become activated (Milestone report 5.1).

Once the WP5 team had identified the risk factors, they assessed the radiation risks on a project-by-project basis.

Activation

When charged particles run into atomic nuclei, they can change the identity of a nucleus, making it into a radioactive isotope. These atomic nuclei have unusual neutron numbers for the element, and the unstable forms eventually decay by emitting other particles. Those with short half-lives will disappear quickly, settling back to a non-radioactive isotope after emitting an energetic photon, known as a "gamma ray", and a particle - most often an electron or its antimatter partner, the positron. However, the longer-lived isotopes are more problematic, and can require equipment to be stored for tens to hundreds of years underground or behind thick walls of concrete.

Linac4

The new Linac4, currently under construction at CERN, will need to feed into the Proton Synchrotron Booster (PSB) through the same injection line used at present by Linac2. The construction workers digging the new tunnel were present on the site from December 2008 until October 2010, overlapping with the LHC's first run at the end of 2009 and its second from March to December 2010.

The question for WP5 was: how can CERN ensure that the construction workers excavating the earth and building the housing for Linac4 will not be exposed to radiation from the running Linac2?

Initially, CERN planned to mitigate the risk by building the link between Linac4 and Linac 2 in winter, when the LHC was not running. But the civil engineers quickly pointed out that digging in frozen ground is impractical. Building work would have to occur during the warmer months, with the LHC operational. At the onset of excavations, the workers were definitely safe - 20 metres (m) of earth separated them from Linac2.

However, the construction workers needed to dig right up to the side of the Linac2 building, which meant that as the excavation progressed, eventually just two 2.5-metre-thick concrete walls separated them from Linac2. Only the last metre of excavation was likely to pose any significant risk, so the WP5 team simulated the radiation levels likely to come through the two walls.

The simulation showed that without additional shielding, radiation levels would be higher than the 0.5 microSievert per hour ($\mu\text{Sv/h}$) limit for sites accessible by the general public. To bring the radiation down to acceptable levels, CERN added a third concrete wall, 40 centimetres (cm) thick.

With this addition in place, Otto says: "A number of calculations with the same Monte Carlo code were done in order to demonstrate that the radiation levels would be so low that there is never the need of defining a limited access radiation area" (see box: Simulations). The excavation workers were at such low risk that they were not required to wear dosimeters to measure radiation levels.

The WP5 team also added a radiation monitor inside the Linac2 tunnel, which can shut down the accelerator if the radiation level there gets too high. The sensor was set for a limit 25% lower than that prescribed by the simulations in order to provide an ample safety margin. These precautions ensured that the radiation dose received by the construction workers was well under 1 mSv per year, the legal dose limit for non-radiation workers - a success for the WP5 team.

The CERN technicians and engineers who will install the new accelerator in the finished Linac4 tunnel will spend more time near the wall approaching Linac2, and so will still be required to wear dosimeters -

but for these workers, such precautions are a standard aspect of their daily work (Deliverable report 5.2.1).

Simulations

To forecast the amount of radiation in various parts of the accelerator complex, the WP5 team modelled the effects that the beam and collisions would have on equipment and work spaces. The code used to simulate the radiation and activation of materials, FLUKA 2008, is a type of Monte Carlo code made for the purpose of estimating the effects of ionising radiation and how it moves through materials. The name Monte Carlo comes from gambling algorithms, which attach probabilities to a number of outcomes.

In the FLUKA code, for example, researchers can evaluate the probability that a negative hydrogen ion collides with a nucleus after travelling a certain distance through a magnet. Provided a collision occurs, probability also rules the fate of the outcome: if the starting nucleus is copper it may become radioactive cobalt after the collision, through the removal of a few protons and neutrons. The code runs through these scenarios for millions of particles, randomly selecting the outcome of each step along the way in accordance with the probabilities set through experiments and theoretical models. It simulates the effect of years of radiation exposure.

Inner Triplets

The sets of magnets which provide the final focusing of the LHC beams before they enter one of the detectors (see WP6, p. 13-14) are on the front line for absorbing the shrapnel of the particle collisions that occur inside the detectors. When high energy particles strike the magnets, they can turn ordinary atoms into radioactive isotopes. For a while, these isotopes act like tiny time-bombs, ready to release radiation when they decay into ordinary atoms again.

Provided that the level of radioactivity in materials becomes negligible within 30 years, CERN will be able to store them and then dispose of them as ordinary waste. However, WP5 showed that the inner triplets will stay radioactive for a longer period.

Ideally, these magnets could be reused, but workers can't afford to spend long hours carefully dismantling them inside the tunnels, where the radiation levels are likely to be high. With the right strategy for disconnecting the magnets, though, Paolo Fessia of CERN's technology department reckons that half of them could be saved and used as spares for the ALICE and LHCb experiments. The rest of the material will be scrapped at a special site for radioactive materials.

The WP5 team recommends that considering how to dismantle the radioactive accelerator parts should become a consideration in the designs of future triplet magnets. Components should fit together "like bricks in a Lego game" rather than being bolted to one another, says Otto. Alternatively, robots could be used to take the magnets apart remotely, ensuring worker safety no matter how long the dismantling procedure (Deliverable report 5.2.1).

Service caverns and surface buildings

A quirk of the local geology required CMS to leave a 6-metre-wide pillar of rock between the experimental cavern and the service cavern. The corridor between the two winds around it. This natural

shielding is much better than that installed around the ATLAS experiment. The ATLAS service cavern, which contains electronics and cooling equipment, is protected by just 2 m of concrete.

Presently, while the LHC is running, the dose rate inside the ATLAS service cavern is 2-5 $\mu\text{Sv/h}$. With ambient dose equivalents ranging between 0.5 and 15 $\mu\text{Sv/h}$ at the highest, the area can be classified as a supervised radiation area. This means that only those certified as radiation workers have access, but they may take their time accomplishing their tasks as it is impossible to exceed annual dose limits with the quoted radiation levels.

When the LHC is upgraded for high luminosity, the dose rates should increase by ten times, reaching 20-50 $\mu\text{Sv/h}$. Because there is no space to add extra shielding, the cavern will have to become a controlled radiation area, which means that the personnel who enter will need special training and closer monitoring. At that point, their work in the service caverns will be subject to time restrictions because in the absence of additional shielding, working more rapidly is the only option to minimise the radiation dose. The workers are unlikely ever to approach the legal limit for exposure of 20 mSv per year, but it is worthwhile to train more people to repair or replace components in the electronics, cryogenics, and other services.

At the surface, the two shafts leading down to the ATLAS experimental hall are each covered with a 1-metre-thick 'plug' made of concrete. The narrower of the shafts emits 1 $\mu\text{Sv/h}$ while the wider emits 3 $\mu\text{Sv/h}$. An additional metre of concrete - a little more for the wider shaft - will be added to the plugs after the upgrade. The extra shielding will be enough to keep the radiation rates below 2.5 $\mu\text{Sv/h}$, which is the limit for a public area that is not permanently occupied. However, safely setting that much concrete over a hole up to 18 m in diameter is not an easy task, and engineering work will be needed to design the structures that will hold and move the concrete (Deliverable report 5.1.2).

Superconducting Proton Linac

The SPL would be built 20-30 m underground, allowing the soil to provide effective shielding against any lost negative hydrogen ions from the beam, which will have energies of up to 5 giga electronvolts (GeV). The accelerator needs about 16 shafts, each 2.8 m wide, to run bundles of cable-like radiofrequency (RF) waveguides from the klystron RF amplifiers at the surface to the RF cavities that accelerate the beam. "These vertical shafts are also a way out for neutrons which are produced in the accelerator," says Otto. Neutrons may be generated if stray negative hydrogen ions hit an RF cavity or other equipment.

Simulations of beam lost in the tunnel revealed that more work is necessary to minimise the radiation exposure inside the klystron buildings. Since technicians need access to the klystron building while the accelerator runs, how to protect them from radiation leaking up from below will be a concern for the future developers of SPL-like machines, such as the team working on the European Spallation Source (ESS) in Lund, Sweden (Deliverable report 5.2.1, also see Impact section, p. 27).

Proton Synchrotron 2

The WP5 team also simulated the radiation risk expected for the beam dump near the PS2, before plans to build the PS2 were shelved. When negative hydrogen from the long, straight SPL entered the oblong ring of the PS2, it would have run through a thin piece of carbon which strips away the two electrons, leaving the bare proton. The PS2 would then alter the trajectory of these protons to travel around its accelerator ring. However, in some cases only one or neither of the electrons would be removed, and the hydrogen

would not make it into the PS2 ring. The negative hydrogen would bend the wrong way in the magnetic field, and the neutral hydrogen would go straight.

To catch these wayward particles, accelerator physicists would need to add a beam absorber. Ideally, the absorber would be able to catch both the negative and neutral hydrogen. To capture the diverging streams of particles in a single, relatively small absorber, the beam dump would have to lie close to the accelerator.

The dump near the PS2 would need to catch about 6.4×10^{19} particles per year, with each particle at an energy of 4 GeV. It would have a design reminiscent of the "layers of an onion", says Otto. A carbon core would be enveloped by aluminium and then tungsten. This would then be surrounded by a block of iron, 1 metre high and 3.2 metres long. Finally, the entire structure would be encased in a 20 cm layer of concrete.

If this was to go into the PS2 accelerator vault, it would need to give off less than 50 $\mu\text{Sv/h}$ per hour so as not to pose a radiation risk to the workers making repairs to the accelerator during shutdowns. The beam dump would breach this limit only at the opening where the beam enters the block of absorbers. A simple mechanism to block that beam entrance would be enough to keep the radiation level sufficiently low (Deliverable report 5.2.1).

Although the PS2 is no longer planned for construction, this study of the PS2 beam dump can serve the ESS in Lund, as the accelerator there will employ an identical injection scheme, requiring similar equipment (see Impact section, p.27).

Validating the simulations

Simulations are a valuable tool for forecasting the activation of materials at the LHC and, eventually, the radiation emitted by the LHC after upgrade work has been completed. However, to rely on this information for safety purposes, WP5 needs to give evidence that the simulations are accurate. "The best way of proving that you are not making a big error somewhere is to conduct some experiments, some validation," says Otto.

One way to do this is to see whether the simulations give accurate predictions about the LHC. Before the accelerator started up in 2008, WP5 placed hundreds of radiation detectors in the LHC tunnel, the ATLAS and CMS experimental halls, and the service caverns. Most of these are passive detectors, which record the total radiation to which they are exposed, to be assessed after collection. A few send out a live feed of information. One kind, the MPX detectors placed in the hottest positions of ATLAS experiment, have been already able to register a short-term activation induced during some relatively high luminosity proton collisions in the LHC.

The plan was to collect the passive detectors after the LHC's first year of running to make sure that the simulations were accurate. Unfortunately, the first year of LHC beam came later than expected. At the end of 2010 the Paul Scherrer Institut (PSI) team removed and evaluated their neutron detectors, indicating radiation levels in the same order of magnitude as expected for the luminosity accumulated so far. Another set of detectors, placed by different collaborators coordinated by the CERN radiation protection group, is still underground, waiting until the end of 2011 to be removed.

These detectors will eventually give vital feedback about the amount of radiation that will reach accessible areas during beam as well as the amount of activation that materials will pick up (Deliverable report 5.1.1).

Impact study

The WP5 team also needed to explore ways that radiation may leave CERN and possibly impact the environment. There are three ways this could happen: through the release of radioactive air or water, and through the inappropriate disposal of radioactive waste.

The inner triplet magnets are set for replacement, as are the collimators. Because simulations show that both types of equipment will be significant sources of radiation for over 30 years, they will need to be sent to a special site for radioactive waste storage and disposal.

During the operation of an accelerator, it is also possible for the oxygen, nitrogen, and other components of ordinary air to become radioactive, and so the WP5 team studied this activation process. Of special concern is the chance that the 52-year-old Proton Synchrotron accelerator may activate the air inside the tunnel and then release it into the environment.

Even though the PS will handle beam intensities 1000 times higher than it was originally designed for, the simulations of WP5 show that the release of radioactive air should be 100 times smaller than the legal limit. Nevertheless, the ventilation system will be renovated to control airflow from the PS and measure the actual amount of radioactivity released through the air.

On the other hand, the simulations did reveal weak points in the PS shielding that will need to be reinforced with concrete or iron to reduce the likelihood of stray radiation escaping into the CERN campus. This is hardly surprising - although scarcely any part of the actual accelerator is original, its circular tunnel and earthen hill of shielding date back to the late 1950s.

CERN is generally fortunate in its location when it comes to water risks - the tunnels are far from large stores of ground water. However, rainwater may percolate through the ground near the accelerators and so pick up radiation. Water which reaches the present accelerator tunnels at CERN is not activated beyond the levels of natural radioactivity. Even so, Otto says: "In the vicinity of future beam dumps, built for higher beam intensities, a sensible precaution would be to separate the evacuation paths of water ingress from the ordinary gutters. In the case of a slight elevation of radiation levels in the water, this would allow CERN to treat it separately."

As a result of the radiation studies in WP5, CERN has taken steps to reduce the radiation exposure of personnel, the upgrade planners know how much the radiation from the LHC and its accelerator chain will increase, and safety officers have recommendations for how to mitigate this future risk.

WP6: Final Focusing Magnets

Although the protons that circle in opposite directions through the LHC's 27-kilometre-long tunnel are focused into two beams, the protons inside those two beams are still widely dispersed. Colliding the

beams would be rather like trying to smash two rainclouds together: they would simply pass straight through one another.

To increase the chances of collision, focusing magnets are used to sharpen each of the two beams down to the width of a human hair just before they enter one of the LHC's detectors. This essentially squeezes the two 'clouds' so that there is a smaller space between the protons within them, and so more chance that protons travelling in opposite directions will hit one another in the collisions that give the LHC its name.

These focusing magnets, also known as Inner Triplets, are to be upgraded around 2022. In preparation, CERN and its partners (CEA Saclay, IN2P3, STFC/RAL, and CIEMAT) have produced engineering designs (Deliverable reports 6.1.1 and 6.1.2, Milestone reports 6.2, 6.3 and 6.4) for the new, longer focusing magnets and their correctors. The 10-metre-long prototype hasn't yet been built, however; such long magnets are too big to build with existing techniques, and the need to focus on a new way to construct the magnets set the project behind the initial schedule.

The magnet team designed and constructed a new set of tools and equipment for building the longer magnets, and they have also started construction work on two 2-metre-long prototypes. The prototype magnets should be built and tested by the end of 2011 (Deliverable report 6.3.2).

To get the most from the new magnets, they will have to be used in conjunction with a set of 'corrector' magnets. Even a flawless magnet would suffer from some resonances that can deflect the beam, so the corrector magnets help to counteract the resonances and keep the beam on course. Two corrector magnets have already been constructed, with another to be completed later this year.

Presently, the LHC schedule plans for the installation of the upgraded Inner Triplets around the year 2022.

Focusing the beam

The series of focusing magnets that squeeze the particle beams just before they enter one of the detectors are called quadrupoles. They are cylindrical and hollow: the protons or heavy ions that form the particle beam run through the middle - an open space called the aperture.

Each quadrupole squeezes the beam in one direction - either up and down or left and right. Squeezing the beam from the top and bottom makes it expand to the left and right, and vice versa. To counterbalance this expansion effect, a total of four quadrupoles are arranged in a line, focusing the beam in two spatial dimensions. Because of the powering scheme, the beam only feels three distinct magnets, so this set of four magnets is called a triplet.

The upgrade team will deliver four quadrupole magnets for each side of ATLAS, to focus the beams entering from both sides of the detector. Another four magnets on each side of the CMS detector perform the same role. A final set of four quadrupoles serve as spares, for a total of 20 magnets.

The quadrupoles have one function: to focus the proton beams to a finer point where they collide with one another at the centre of the ATLAS and CMS detectors. The tighter and more densely spaced the proton bunches are at the moment of collision, the greater the chance that some of the protons in one beam will collide with those in the other beam.

But there is an unwanted side effect associated with prepping the proton beams so that they focus to such a fine point: the beams swerve, deviating from a straight line, by the time they exit the opposite side of the detector (see box: Taming the wobble). To make sure that the beam doesn't miss the exit hole, the magnet team has to widen the aperture on all of the quadrupoles from 70 to 120 millimetres (mm). This adjustment gives the swerving beams the leeway they need to exit the detectors smoothly.

This adjustment itself has a knock-on effect: a larger aperture weakens the quadrupole's magnetic field, which compromises its ability to focus the beam tightly. To compensate, the magnet team need to make sure that the magnetic field generated by the quadrupoles is as strong as possible. "The design of the new quadrupole magnets brings the superconducting cable to its limits," says Russenschuck.

Within each quadrupole are four coils of niobium-titanium cable that generate magnetic fields when electric currents pass through them. These four fields create one field with four magnetic poles.

At very low temperatures, the niobium-titanium coils enter a special state called superconductivity, which means that the coils have no electrical resistance. This allows very high currents to run through the cable, creating a very strong magnetic field.

However, if the current density or the magnetic field becomes too strong, then the cable will stop superconducting. The magnetic fields that are produced in the new quadrupoles will be about as high as is possible to achieve without placing the niobium-titanium at regular risk of losing its ability to superconduct.

The fragile superconducting state is also at risk from heat. The reason for focusing the beams more tightly within the ATLAS and CMS detectors is to generate even more particle collisions. But again there is a knock-on effect from this heightened activity. There is a greater chance that some of the debris from these high-energy collisions will fly out of the detector and heat up the quadrupole magnets. With the increase in particle debris, simulations show that the magnets will need to dissipate about 500 Watts of heat due to particles from the collisions in order to avoid losing their ability to superconduct.

There are two potential ways to tackle this additional heat: add more shielding around the magnets to keep the particles from hitting the superconducting coils, or make it easier for the liquid helium that helps keep the coils cool to penetrate the niobium-titanium cable.

Adding more shielding is difficult because there is so little space to spare in the already cramped tunnel. But increasing the flow of the liquid helium to the cable is possible. To do this, the team has developed more porous electrical insulation for the cables.

Taming the wobble

"If you want to squeeze the beam to a very small size at the interaction point, it will deviate from the orbit a bit more in the Inner Triplet magnets," says Russenschuck. He compares it to a ball rolling down a gutter. Ideally, the ball would run straight down the centre, but it may instead swing slightly from side to side as it rolls along.

This swerve is conserved. The proton beams are prepped to be tight and centred in the beam pipe when they meet in the middle of the detector, producing more collisions. But by the time they reach the other

side of the cavern, the beams tend to swing towards the edges of the beam pipe. As this is the point where the beam re-enters the tunnel, coming back under magnetic control, it needs a wide berth to keep from smashing into the side of the beam pipe and damaging the magnet.

Originally, it seemed that the amount of wobble on the tightly focused beams could be cut by half at the collision point, but this proved impossible to accomplish by upgrading the focusing magnets alone. However, the new magnets should at least be able to reduce the wobble by 20-30%. This reduction is achieved in part by lengthening the focusing magnets, designed at 10 m long rather than the 7 m length of the present focusing magnets.

The cable

The magnets contain superconducting cable, roughly half of which is niobium-titanium; the other half is made up of pure copper. The composition is a safeguard against damage caused by a sudden destruction of the cable's superconducting state - known as a 'quench' - either by heat, surges in the strength of the magnetic field, or electric current flowing through the cable. For all of the precautions put in place to avoid problems such as overheating that would compromise the superconducting state, occasional quenching is unavoidable.

Copper is an excellent conductor and so can take the burden of the current when the superconducting state collapses and suddenly makes the niobium-titanium component of the cable far more resistive. Without the copper to ease the flow of current though the cable during quenches, there would be a risk of overheating that could irreparably damage the cable.

To help prevent quenches, the magnet team employed a 'barber pole' wrapping style. Three layers of polyamide insulation spiral in opposite directions across the cable, leaving gaps tiny enough for the liquid helium coolant to worm through without leaving direct routes which admit electrical sparks (Milestone report 6.1).

Superconducting coils

With the design of the cable in place, the next step was to wind them into the coils to be used in the magnet prototypes. A CERN contingent made short trial coils before a team at CEA Saclay made larger coils for the prototypes.

The coils are 'racetrack' shape: two straight sections are joined by arcs at each end (see picture above). They don't lie entirely flat but instead form long gutter-shaped structures. Four such gutters can be arranged to form a cylindrical cavity, through which the particle beam travels.

To make these racetrack coils, a coil drum - like a giant bobbin from a sewing machine - spins on an axle, supplying the cable to a rotating platform. A long, specially machined piece of steel, shaped a bit like an overturned racing rowboat, helps wind the cable into its racetrack shape.

In order to keep the magnetic field the same strength all the way along the magnet, it is important that special 'end-spacers' are added between the cables where they loop around at each end of the magnet. These end-spacers are made of steel or epoxy-glass. They are difficult to design because at the ends of the

'gutters', the cables are not simply looping around, but also dipping down and up again to leave space for the beam.

But with new design techniques and a three-dimensional printer that can rapidly make shapes from a powder, the team could quickly develop end-spacers that closely match the complicated contours of the cables.

"We are really happy with this outcome, how they fit," says Russenschuck. "In the past, if gaps developed in between the superconducting cables, they were closed with epoxy-glass resin." But epoxy-glass resin also fills in the tiny crevices that allow the cooling liquid helium to squeeze between layers of cable. The new end-spacers are such a good fit that epoxy-glass isn't needed, allowing the liquid helium to work more efficiently.

Once the coil is wound, it is cured at 190°C and a pressure of about 1000 atmospheres. "The shape is actually determined not by the winding process but by this curing process," says Russenschuck. The exact shape influences the magnetic field of the final product, so the entire process - from winding the cables to curing the finished coil - is high precision work. Just the winding tool and the arch that sits atop it cost €82,000.

The first production coils from CEA Saclay have undergone five electrical tests: coil resistance, insulation resistance, inductance, leakage tests, and pulse tests. These demonstrated that the coils could withstand voltages more than twice those to be expected during operation of the magnets, testament to the quality of the construction (Milestone report 6.6).

Collaring the coils

The electromagnetic energy stored in a racetrack coil leaves it inclined to push itself apart. This could create enough friction to cause part of the magnet to heat up and locally trigger a quench. The magnet becomes resistive and no longer produces magnetic fields strong or precise enough to control the particle beam. To prevent the beam veering off course and damaging the magnets, systems kick in to automatically dump the beam, bringing collisions at the LHC to a temporary halt.

To minimise the chances of the cables pushing apart so abruptly that they cause a quench, the coils are held under about 50 million Pascals (MPa) of pressure - about 490 times the pressure of Earth's atmosphere. As the coil shrinks when it reaches the deep chill at which it operates - -271°C, or 1.9 Kelvin (K) - that pressure must be 90 MPa at room temperature.

The collar provides this pressure with the help of shims which slide into the collar alongside the coil. The collars and shims also hold the four coils together to form one electromagnet with four-fold symmetry.

The collars themselves are in four pieces, and they must be assembled around the coils, encircling them in wide rings. Each is three millimetres thick, and they cover the entire length of the magnet.

The easiest way to assemble these 3000 collars is by standing the coils on end and stacking the collars vertically around them. Unfortunately, the length of the upgraded quadrupole magnets means that CERN would need to dig a pit in one of the buildings to accommodate them. The cost of this excavation is

excessive given that only 20 magnets in total are needed. Instead, CERN has designed a collaring press that can function with the coils in a horizontal position.

A coil is placed on each of the four sides of the beam aperture (see picture). They are set into position on an assembly mandrel while the collars are locked into place around the outside. To ensure that all the collars are perfectly aligned, each collar carries a square chink. When several collars are tightened, these chinks align and a long box-shaped key can be inserted to lock several of the collars into place.

To make sure this operation produced a highly symmetrical magnet, the WP6 team developed a collaring press. This device evenly tightens the collars with twelve anvils connected to a hydraulic system. An additional four anvils insert the keys to lock the collars. Once the collars are secured, the spring-loaded assembly mandrel can be taken apart and withdrawn from the coils.

Under pressure

In order to give the coils 50 MPa of pressure when they are cold, the magnet team needs to thoroughly understand how the coils, collars and shims bend under pressure.

Teams at CERN and CEA Saclay calculated the stretching of the collar pack, the expansion of the coils as they come off the collaring press, and the way the coils contract when cooled to 1.9 Kelvin. The CERN team also designed and made a device that compresses the coils from the sides, revealing how much their sizes change after the application of pressure. "This needs to be done for each and every coil because each coil may be different," says Russenschuck.

This pressure testing showed that the porous insulation on the cables, which allows the liquid helium to pass through and keep the cables cool during operation, is less stiff than the insulation of the typical LHC magnets. Putting so much pressure on the new coils produces a permanent change in their sizes after testing. To establish their final size when they are installed in the LHC, the magnet team will need to keep track of each coil's initial size after manufacture and its history of testing under pressure (Deliverable report 6.2.2).

The CERN team developed collar packs fitted with pressure sensors to study the stress on the coils after the magnet is assembled and cooled. These sensors will reveal how much shimming is needed to give 50 MPa of pressure.

The CERN team also developed a collar pack to measure heat transfer, to ensure that the porous insulation around the cable and between the coils still allows the helium to carry away the heat when the coils are under extreme pressure (Deliverable report 6.2.1).

Corrector magnets

Although the well-formed coils of the upgraded quadrupoles will make the purest field yet for superconducting magnets of their size, they aren't quite perfect enough to control the beam without a little additional help. A series of corrector magnets already fix errors in the current quadrupoles of the LHC, and they will also be replaced as part of the new Inner Triplet system.

Corrector magnets fall into three main categories. Orbit correctors guarantee that the beam is indeed steered through the centre of the focusing magnets. Skew quadrupole correctors cancel out any bias in the

quadrupole field that might steer the beam along a curve rather than in a straight line. And multipole correctors cancel out the additional magnetic fields that arise due to tiny manufacturing imperfections in the quadrupole. The corrector package sits further from the particle detector - ATLAS or CMS - than the quadrupoles, but an extra set of orbit correctors placed between the two middle quadrupoles will help bring the veering beam under control more quickly.

Simulations have shown that these corrector magnets will absorb even more radiation than the quadrupoles, owing to their location. However, this dose can be cut by 50% if the aperture through which the beam runs is widened from 120 mm to 140 mm. A block of steel between the beam pipe and coils could also shield them from particles, further reducing their exposure to radiation by a factor of three.

Researchers at CIEMAT in Madrid, Spain, have built superconducting sextupole and octupole corrector magnets, while CERN has been developing superconducting corrector magnets made with the same copper and niobium-titanium cable that is found in the quadrupoles. These are the first corrector magnets ever to be made with this type of cable. The magnets are already assembled, but they cannot be tested at superconducting temperatures until the measurement equipment has been built. CERN, CIEMAT, and STFC/RAL in the UK are committed to completing the skew quadrupole magnet and cold test by the end of 2011 (Deliverable report 6.3.1).

WP6 made major contributions to the Inner Triplet upgrade effort, which continues through the High Luminosity LHC project. The new magnet designs bring niobium-titanium cable to its limit, so research and development in this area is also looking into cables of niobium and tin, which has the potential to produce even higher magnetic fields.

Work Package 7: Towards a new accelerator

The Superconducting Proton Linac (SPL) would have helped to produce brighter beams by accelerating them to 4 GeV - approximately three times the energy of the present Proton Synchrotron Booster - by the time they reached the next accelerator, the Proton Synchrotron 2 (PS2). In addition, accelerator physicists could also have taken advantage of the fact that the SPL accelerates negative hydrogen ions rather than protons.

Negative hydrogen ions are simply protons chaperoned by two electrons. Within the PS2, packets of negative hydrogen would have their electrons stripped away to leave a proton beam.

As those first protons in the beam started their second circuit around the PS2 ring, the SPL would add a second pulse of negative hydrogen ions directly on top of the first protons. By stripping the electrons from the second pulse only after it joins the original protons, the pulse would be brighter and more intense but no larger.

This process relies on the fact that the negative hydrogen wears its electrons like a disguise to mask its positively charged proton. Whereas it is not possible to inject a pulse of positively charged protons directly on top of the protons already circulating in the PS2, negatively charged hydrogen ions can infiltrate the existing proton bunches, says WP7 project leader Richard Scrivens. Finally the electrons

from the additional negative hydrogen would be stripped away to leave a proton beam that is much brighter.

In 2010, after a detailed analysis and first cost estimate, the CERN management decided to focus on upgrading the older accelerators rather than constructing the SPL and PS2.

Nevertheless, the WP7 team continued with its work and found a synergy with those developing the European Spallation Source (ESS) in Lund, Sweden, a facility for science with neutron beams that will need a high-intensity proton accelerator much like the SPL was supposed to be. More details may be found in the Impact section.

The SPL would require a high-intensity negative hydrogen source, providing pulses three times the length of those generated by even the Linac4 source, and 25 times as often. Half of the WP7 team set out to develop this high-intensity negative hydrogen source.

The SPL itself would be composed of radiofrequency (RF) cavities which accelerate these negative ions. Klystrons - RF amplifiers which power the RF cavities - are expensive pieces of equipment, each costing hundreds of thousands of Euros. For this reason, a second team on WP7 explored the possibility of powering multiple RF cavities at once from a single klystron, which would help reduce the cost of building the SPL.

Negative hydrogen source

"The first part of any particle accelerator is its particle source," says Scrivens. "For an accelerator like the SPL, this would be a negative hydrogen ion source." Sources have two major components: a plasma generator, which provides the particles, and a system for extracting these ions out of the plasma and beginning the acceleration process. A simple, static, high voltage applied to the extraction system usually provides the initial acceleration.

Plasma generator

The first step in building a negative hydrogen ion source is making a plasma generator. For the SPL this consists of a ceramic chamber, just 5 cm across, emptied of air to create a vacuum. A copper coil winds around the outside, carrying pulses of an alternating electrical current that switches direction about 4 million times a second, cycling at 2 megaHertz (MHz) - a frequency in the radio range. The alternating current in the coil creates an alternating electromagnetic field within the chamber.

Around the outside, eight or more bar magnets form a ring fence around the barrel of the cavity, creating a magnetic field that herds electrons into the centre of the chamber.

Hydrogen gas, made up of pairs of hydrogen atoms bound together, is injected into this chamber. Then, an electron gun, like a spark plug on a petrol engine, injects electrons into the chamber. The electromagnetic field drives them in spirals through the hydrogen gas. Some of the hydrogen molecules are ionised, falling apart into a plasma of free electrons and protons.

But not all of the hydrogen molecules break up into protons and electrons. Sometimes, the electrons in the molecular bond get into an excited state, making it easier for the molecule to capture an extra electron. When this happens, the molecule can fall apart into a negative hydrogen ion and a neutral atom.

Handling the heat

The Linac4 source needs only send out a pulse of negative hydrogen ions 0.4 milliseconds (ms) long every half-second. Even with a peak power of 100 kiloWatts (kW), Scrivens notes that the average power required to send out these short bursts is just 80 Watts (W) - enough to illuminate a light bulb. It is easy for the plasma generator to radiate the leftover heat away.

However, if the SPL - the next link in the accelerator chain - was being built, those pulses would need to be lengthened to 1.2 ms and occur fifty times every second to bring these accelerators to their full potential. The pulses would still require a peak power of 100 kW, meaning the average power of the plasma generator would now be 6 kW.

As Scrivens points out, this is enough to heat a house, and probably more than the plasma generator could radiate away on its own. Using the current design specifications for the Linac4 plasma generator, the WP7 team confirmed that this was true by simulating the effect of a 6 kW heater set in the middle of the chamber - without improvements, the cavity would store heat until components began to melt (Deliverable report 7.1.1).

To help the plasma generator rid itself of heat more efficiently, the source team decided to make the chamber out of aluminium nitride, a material which conducts heat better than the aluminium oxide ceramic currently used for the Linac4 source. They soldered some of the critical components together as well, allowing heat to flow more freely. Carefully optimised water cooling channels running through the chamber are designed to carry the heat away. The source team also added more cooling channels to the front and rear of the chamber, where the hydrogen and electrons are injected and where the negative hydrogen ions are extracted.

The copper coil, which winds around the chamber, would also generate more heat in the SPL scenario as it would have to carry a higher average current, which would therefore experience more electrical resistance. For this reason, the negative hydrogen source team used a hollow copper tube to generate the electromagnetic field so that they could run cooling water through it.

Finally, the RF pulses of electromagnetic field produce "eddy" currents in the magnets that encircle the chamber. Because the SPL scenario would require more frequent pulses, these currents would heat the magnets until they lost their magnetism. To prevent this, the WP7 group decided to encase the magnets in copper. The copper is easier to cool, and as the eddy currents run through it, the magnets are shielded from the effects of the alternating electromagnetic field (Milestone report 7.1, Deliverable report 7.1.2).

Prototypes and testing

The first prototype, completed in September 2010, used 12 magnets to corral the electrons. These magnets have a strong field near the edge of the ring, but in the centre of the chamber, they cancel one another out. The more numerous the magnets in the ring, the quicker the field drops off towards the centre. Once the pipes that cool the aluminium nitride chamber had been added, the WP7 team found that they would need to move the magnets further from the plasma (Deliverable report 7.1.3). In order to keep the field stronger near middle of the chamber, the team tried reducing the number of magnets to eight.

The WP7 team's tests on the plasma source continue. They are now using a different type of magnet configuration, known as a Hallbach type, in which each bar magnet is composed of three pieces. The different orientations of the magnetic fields from individual pieces may also help to produce a stronger magnetic field nearer to the centre of the plasma chamber. The source team has simulated these fields for 8 and 12 magnets, and they are also measuring the fields with an experimental setup.

The WP7 source team also tested the 12-magnet prototype of the plasma generator with fifty 1.2 ms pulses per second at 50 kW peak power. These tests were successful, but upping the ante to a maximum power of 100 kW created sparking through the epoxy insulator around the copper coil, even when the frequency of pulses was reduced to just 1 or 2 pulses per second. Adding reinforced insulation helped reduce this sparking.

While the plasma chamber itself is small, it requires considerable additional equipment to function. The 100 kW RF generator resides in a cabinet about 2 metres high, converting the mains electricity to high-power pulses that run through the copper coil.

A vacuum pumping system removes all the air from the chamber before the hydrogen can enter, and pumping systems drive cooling water in and out of the plasma generator. The WP7 team purchased and installed these devices through the SLHC-PP project (Deliverable report 7.1.3).

An optical spectrometer, which measures the intensity of different colours of visible light emitted by the plasma, was crucial for monitoring the plasma's composition and density. Excited electrons in hydrogen atoms jump back to lower levels, and the energy difference between the upper and lower level prescribes the colours - or more precisely, the wavelengths - of the photons that the electron will emit on the way. By measuring the amount of light emitted at each wavelength, the researchers can deduce the temperature and density of hydrogen atoms and ions in the plasma (Deliverable report 7.1.4).

Aside from tuning the plasma, these measurements can also reveal more about how hydrogen behaves when excited in pulses. Pulsed plasmas, like those studied in WP7, have received little detailed attention by physicists before now, and the new measurements showed that the proportions of light coming in different wavelengths changes considerably over the course of a pulse. This new information, presented at the 2nd International Symposium on Negative Ions, Beams and Sources and soon to be submitted to a journal, has yet to be fully understood.

Accelerator system

"In order to accelerate charged particles along the beam path, an electric field must be applied," says Wolfgang Hofle of the accelerator team. Most accelerators create alternating electric fields within resonators called RF cavities.

The resonators studied for use in the SPL accelerator hold stationary electromagnetic waves that swap polarity, cycling at a frequency of 704.4 megahertz (MHz). The waves manipulate the metal of the RF cavity, making one end of the cavity positively charged and the other end negatively charged. The switching must be carefully timed so that as bunches of negative hydrogen ions travel from one RF cavity to another they always see an attractive, positive potential ahead of them and a repulsive, negative potential behind them.

Each cavity is about 21 centimetres long, a length that is defined by the distance that particles moving at near light speed can travel within a single polarity cycle. At a frequency of 704.4 MHz each cycle is just 1.4 nanoseconds (ns) long.

Like drums, these RF cavities resonate with a certain frequency. The electromagnetic waves that feed energy into the cavity match this frequency as closely as possible. Scrivens compares it to timing the push of a child's swing. Done right, each push adds on to the last with only a small amount of energy loss - one millionth of the stored energy per oscillation, in the case of these RF cavities. The resonance allows accelerator physicists to maximise the energy stored in the standing waves, allowing them to push and pull proton bunches more forcefully and accelerate them to greater energies faster and more efficiently.

Superconducting cavities, which allow the electrons to travel without resistance, accumulate extremely high electric fields. This makes them particularly efficient accelerators over short distances. But the benefits of superconducting cavities come at a cost: they are trickier to control.

RF cavities can bend under pressure, and these slight changes in shape change the resonant frequency. If the deformation isn't accounted for, the cavity's frequency falls out of sync with the RF power fed into it. In extreme cases, the cavity's electromagnetic field fades as the energy flows backwards out of it. The window of frequencies that a superconducting cavity resonates with is much narrower than that of an ordinary RF cavity, so the resonance is easier to lose.

"When we want to accelerate a beam pulse, the electric field in the cavity has to be pushed to a high level - so high that forces associated with the electric and magnetic fields will act on the cavity, mechanically deforming it," says Hofle. The field compresses the bellows-shaped cavity along the beam line and expands it at the outer edges, but tuners simultaneously counteract these forces, keeping the RF cavity resonating with the power fed into it.

Because superconducting cavities are especially susceptible to losing sync, their tuning must be extremely fast and accurate. Piezo crystal tuners, which contract or relax depending on the presence or absence of an electric control field, fit this bill. They are used to counter the deformations caused by the high magnetic field in the two superconducting RF cavities studied through WP7.

In exploring how to best control the tuners and the electric field in the superconducting cavities, the SLHC-PP has picked up the baton from previous efforts.

Characterising existing cavities

Teams at the INFN Laboratory of Accelerators and Applied Superconductivity in Milan, Italy, and CEA Saclay in France, made superconducting cavities for accelerating pulsed beams as part of the "HIPPI" program in the European Commission's Sixth Framework Programme. Through WP7 of the SLHC-PP, they went on to test these cavities at CEA Saclay, running them at their design powers and frequencies while the cavities were cooled to superconducting temperatures.

They measured several factors which impact the development of a system for controlling multiple cavities with a single RF amplifier. Among the most important are the range of frequencies over which the cavities can be tuned, how quickly the tuners respond, the stiffness of the cavities, and how much energy is available for accelerating particles.

Storing and maintaining energy

The cavity must be prepared for each 1.2-millisecond-long pulse of negative hydrogen ions. Half a millisecond before the beam arrives, energy is supplied through special radiofrequency amplifiers called klystrons. This energy brings the electromagnetic field in the cavity up to the desired level. "At this point the stored energy in the cavity is as small as the nutritional energy stored in a grain of rice," says Hofle. "However, this process needs to happen very quickly so although the amount of energy is low, an impressive amount of power is needed - 1 million Watts, or the equivalent of the power of 20 mid-size cars, per cavity."

When the pulses of hydrogen ions arrive, they take some energy and continue on their way. "The trick is now to continue to feed the right amount of energy into the cavity," says Hofle. The field strength and timing must stay extremely stable, in perfect sync with the arrival of the next particle bunch. This requires a feedback loop which compares the existing field with an ideal field. With this information, the RF wave is adjusted before it is sent into the klystron for amplification to ensure that it hits the cavity's resonant frequency correctly.

Designing a cost-effective accelerator scheme

Because klystrons are so expensive, the accelerator community has long thought that it might be more efficient to run multiple RF cavities on a single klystron. The SPL would require 200 to 300 RF cavities, and one klystron for each cavity would amount to tens of millions of Euros. The potential cost savings by running two or more cavities from a single, larger klystron is attractive, but two important questions need to be answered. First, is it possible to control multiple RF cavities with a single klystron? And secondly, do the costs of this control system outweigh the cost of the extra klystrons? The WP7 team focused on answering the first question.

As outlined above, the RF cavities are controlled by manipulating the RF power coming from the klystron, so the present one-to-one relationship is valuable. If the klystron feeds two or more cavities which are behaving differently, this control system suddenly becomes far more complex.

"Through the klystron, you can control the average voltage of two or several cavities," says Matias Hernandez Flano of the simulations team. The voltage is the electric potential across the cavity, which accelerates the particles. The trouble, Hernandez explains, is that controlling the average voltage may see the voltage on one cavity overcompensated while the voltage on the other is undercompensated, potentially compromising each cavity's ability to control the hydrogen ion pulses passing through. "At the moment, through simulations at least, it seems like it's possible to have one power source and one control loop for two cavities, provided that you have some special equipment to tune each cavity," he says.

The piezo crystal tuners can help. These rapidly deform the cavity by small amounts, helping to maintain the right resonant frequency needed to control the hydrogen ions. In addition, a motorised tuner takes care of larger, slower drifts in resonant frequency. These changes are caused by pressure fluctuations in the liquid helium coolant which surrounds the cavity, keeping it at a temperature of 1.9 Kelvin (K) or -271°C. These are also controlled through the feedback loop.

With such modifications, the simulation team has shown that two cavities can be fed from a single klystron. Feeding more than two cavities from one klystron might not be feasible, however (Deliverable report 7.2.2).

The SPL poses an additional challenge, namely the high repetition rate of the pulses, which means that a cavity has little recovery time. The deformation caused by the pulsing RF field can fuel mechanical resonances in the cavity. These resonances come in many different frequencies, and not all of them have time to die out before the next pulse arrives just 0.02 seconds later. Hernandez compares the phenomenon to striking a metal bell: as the sound fades, it becomes less complex. This is because some resonances disappear faster than others. In the cavity, this 'ringing' with mechanical resonances also alters the electromagnetic resonance.

If the mechanical resonances aren't damped out by the time the next beam comes, they can build up, disrupting the acceleration. The simulation team found that, if the resonances are left unchecked, the negative hydrogen ions bunches begin to lose synchrony with the field in the RF cavities.

The piezo tuners are again useful. They can counteract the mechanical resonances by deforming the cavity, stifling the ringing at the moment the new pulse arrives. Because the pulses come too fast to correct the ringing within a single pulse, the piezo tuners must instead be controlled with an 'adaptive learning algorithm'. This feeds information from one pulse forward to the next, cancelling the vibrations quickly enough to prevent a catastrophic build-up.

Building control electronics

Like a program on a computer, the adaptive learning algorithm and feedback control run on electronics boards. These are also under development at CERN. In addition to being capable of controlling two cavities attached to a single klystron, they are also more streamlined and modular than previous RF electronics designs, with each RF cavity controlled by only two boards (Deliverable report 7.2.3).

"It's basically the same as a radio," says Daniel Valuch, who is working on the boards. The high frequency signals from the cavities are converted to a lower frequency, just as a radio takes high-frequency signals from the air and converts them to the lower-frequency sound we hear coming through its speakers. In the boards, the signals from the cavities are then made digital.

In this form, the control algorithms can compare the electromagnetic fields in the cavities to the ideal values, and then send out instructions to the piezo tuners. The signal that goes to the klystron is made analogue again and then converted back to radiofrequency so that it can be amplified and returned to the cavities. "The algorithms which run on these boards are based on the simulations, which tell us how to control the field," says Valuch.

These boards are not yet finished and fully tested, but the work continues at CERN. A CERN team is also designing and constructing a test stand for the superconducting RF cavities, housing them at 1.8 K in liquid helium tanks built specifically for the cavities developed by CEA Saclay.

Because the steps leading up to a test of the electronics with the klystron and RF cavity at CEA Saclay have taken longer than expected, the last goal of the WP7 accelerator study has not yet been achieved (Deliverable report 7.2.4). However, through collaboration with CEA Saclay and the ESS, the test will be

completed before the end of 2011. Moreover, the work is intended to continue at least through 2014, building on the developments made through the SLHC-PP by integrating the electronics with four RF cavities and testing the system at full power.

WP8: Powering the trackers

When higher luminosity beams circulate around the LHC, there will be more collisions inside the ATLAS and CMS detectors, and hence more particle tracks to tease apart in the detectors. "To cope with the higher luminosity of the beam, the number of sensor channels will be increased by a factor of five to ten in the silicon strip trackers," says Georges Blanchot of the DC-DC conversion team.

The 86 million and 75 million individual sensor channels in ATLAS and CMS, respectively, are spread across tens of thousands of detector modules, each module with its own power cable. With a significant increase in the number of sensor channels, there are three good reasons to find a new powering strategy. The first is space. These cables, and the cooling pipes that export heat, already crowd the central subdetectors. Giulio Villani of the serial powering team says that adding so many more is "logistically impossible."

The present system is also highly inefficient. "Something like 60% of the power is actually being dissipated along the cables that bring the power to the modules," says Villani.

The last reason is data quality: particles from the collisions can be deflected or absorbed by atoms inside the cables - or they can collide with a nucleus within one of these atoms and create a cascade of new particles, muddying the picture of the original collision. Add more thick cables, and the powering teams would face the wrath of their collaborators for adding more disruptive, dead material to the detector. WP8 project leader Wladyslaw Dabrowski hyperbolises: "We are beaten for every gram of additional material that we put into the detection volume!"

For all these reasons, the particle physics community has sought to develop new powering schemes which could take up less space and produce less heat. The contribution made by WP8 focused on two methods: DC-DC conversion and serial powering.

DC-DC conversion

In theory, each detector module can continue to receive power through its own private cable, but for all the cables to fit, they must be much thinner. DC-DC conversion, which converts a direct current of one voltage to a direct current of another voltage, can make this possible.

This trick has been used by scientists for well over a century. The power running through a cable is related to both the voltage and the current. This means the same amount of power may be delivered through a cable by increasing the voltage and simultaneously lowering the current.

The advantage is that cable diameter is proportional to current, and so the high voltage cables can be thinner. What's more, the amount of resistance falls with the reduced current, so the thinner cables lose less energy.

When the cable reaches the detector module, the DC-DC converter can step down the voltage and increase the current to the levels needed to power the sensors and electronics. "In this way, the current carried by the long cables is significantly reduced, as well as the power dissipated by them in the form of heat," says Blanchot.

Although these converters are widely commercially available, they need a special design to cope with the high radiation and high magnetic fields within the ATLAS and CMS detectors.

Designing the circuit

The first question was what type of DC-DC converter to use. After looking through the available designs in science and industry, the DC-DC team settled on a buck converter. This is the simplest type of converter, offering the best combination of small size and relatively high efficiency: 80-90%.

An inductor - a circuit element that stores energy in a magnetic field - is a key part of a buck converter. While a voltage is applied to an inductor, the current flowing through the inductor increases. When the voltage is removed, the inductor is shorted out and the current flowing through decreases. The voltage reduction comes from the fact that the voltage is applied only part of the time.

A bank of capacitors - circuit elements that store charge - turns the variable current from the inductor into a constant, average current, which is delivered to the sensors. Stefano Michelis of the DC-DC conversion team compares the role of the converter to a tank of water with a hole near the bottom. It doesn't matter if water is poured into the tank sporadically - the stream of water coming out the bottom will stay roughly constant so long as the water stays above a certain level. "You are providing almost a fixed current, and therefore a fixed voltage," he says.

Typically, inductors are constructed by wrapping copper wire around a ferromagnetic core. The spiral pattern generates a magnetic field when a current runs through the wire, and the presence of the magnetic iron strengthens the effect as its atoms align with the field.

But this design cannot work inside the ATLAS and CMS detectors. The superconducting solenoid magnets within these detectors produce powerful magnetic fields of up to 4 Tesla, which would take the ferromagnet in each converter to its maximum state of magnetism, leaving it unable to respond to the much weaker magnetic field formed by the wire wrapped around it. To solve the problem, the inductor needed a customised 'air core' design, which required a larger coil - about a centimetre across - rather than one a few millimetres wide.

The other consideration for the inductor is the control of stray magnetic fields. These fields could interfere with the electronics, ultimately creating errors in the data recorded by the experiments. The traditional solenoid inductor looks a bit like a spring, forming a straight cylinder, but its magnetic field stretches out from the coil. The DC-DC team chose instead to wind the copper wire around a doughnut-shaped piece of plastic. This creates a more introverted magnetic field, contained within the closed loop of the doughnut.

However, the current moving around a ring creates a secondary magnetic field. "A plastic box covered with a thin copper layer, a few tens of micrometres thick, is able to completely shield this residual magnetic field," says Michelis.

Managing noise

Electromagnetic 'noise' was one of the most important concerns for the DC-DC conversion strategy, says Dabrowski. Stray fields from the converter boards can interfere with the nearby sensors and readout electronics, potentially drowning out the data signals that physicists want to analyse. The inductor isn't the only culprit - the switches in the ASIC also emit electromagnetic noise.

To cope with the noise, the DC-DC conversion team designed the layout of the components on the converter board so that the noise-emitters were oriented away from the sensor and readout electronics. Combined with the metal casings atop the worst noise emitters - namely the inductor and switches - Dabrowski says: "We minimised the noise to a level that is lower than the intrinsic noise of our sensors and our electronics." Noise this subtle cannot affect the particle measurements.

A radiation-resistant chip

The buck converter uses two switches: one to turn the voltage on and off for the inductor, and the other to drain the current out of the inductor. These buck controller circuits are readily available commercial goods, but they wouldn't survive the high levels of radiation present in the central regions of the ATLAS and CMS detectors.

Planning for a 10-year lifetime, the converter must be tough enough to stand 250 megaRad of absorbed radiation, together with a similar onslaught of neutrons, protons and other heavy particles. For this reason, the DC-DC team developed custom-designed controller chip converters, implemented as application specific integrated circuits (ASICs).

Radiation can, for instance, disrupt the silicon crystal structure within the chips, deposit electric charges, change information encoded in the chip, or cause a short circuit. "The response of the power transistor to radiation exposure is very specific for the given technology and it not really predictable," says Dabrowski. Therefore, the CERN team tested different commercial technologies using X-rays and proton beams at CERN as well as heavy ions at the Cyclotron Facility at Louvain-la-Neuve, Belgium.

The WP8 team identified two radiation-tolerant technologies, one from On Semiconductor, a company headquartered in Phoenix, Arizona, and one from the Institute for High Performance microelectronics (IHP) in Frankfurt, Germany.

The power converter ASIC known as AMIS2, which was designed with the On Semiconductor technology, passed all these tests successfully. The IHP ASIC passed the X-ray test and was partially qualified with protons and ions, with some operating voltage restrictions.

Using the AMIS2 controller ASIC, the WP8 team developed a DC-DC power converter that could plug into newly developed front-end electronic modules for any experiment. Fed from an input voltage of 10 Volts (V), it delivers a 3 ampere (A) current with a voltage of 2.5 V. This converter aims to power tracker systems that are now under development for the ATLAS and CMS upgrades but not yet available for tests.

"In the meanwhile we would like to test the DC-DC converters," says Michelis, "so we use commercial components that can stand 5 amps." This makes the converters compatible with back-up modules from

the present trackers of CMS and ATLAS, which have higher power demands. These test converters, replacing the AMIS2 ASIC with a chip from Linear Technology, a company headquartered in Milpitas, California, are known as the SM01C design (see picture).

CERN engineers produced both the small converters with the AMIS2 ASIC and larger ones using the Linear Technology chip. These successfully powered ATLAS tracker sensor modules at the University of Liverpool in the UK and University of Geneva in Switzerland.

At RWTH Aachen University in Germany, another team produced a second set of buck converters with the AMIS2 chip, specially designed to power modules of the CMS silicon strip tracker, and one known as PIX_V7, which has already been chosen for use in the upgrade of the CMS pixel detector.

With the concept proved, the next step is the design of detector modules with the converters already built in. For the ATLAS silicon strip tracker, sets of four modules powered through DC-DC converters are under production at Rutherford Appleton Laboratory (RAL) in Didcot, UK. Collaborators at the University of Geneva are building a different version of those tracker modules, powering eight double-sided sensor modules. These efforts are bigger than WP8 alone, but the SLHC-PP has contributed significantly.

Although the DC-DC converters were designed for the trackers, they may also find homes in calorimeters, detector elements which measure the energies of particles. The CMS, ATLAS, and LHCb collaborations are all considering DC-DC conversion for powering parts of their detectors.

Serial powering

If the trackers within the Inner Detectors of ATLAS and CMS were Christmas trees, with the detector modules on each like a string of decorative electric lights, then both trackers would have wires for garland: each light has a direct line to a power source. Christmas tree lights don't work like that, of course; instead, a single pair of wires feeds all of the bulbs, saving considerable space as a result. The WP8 team believes that a scheme based on this 'serial powering' approach could save space and reduce power losses in the detectors, too.

The silicon pixel and silicon strip detectors within the trackers are built of modules - arrays of sensors connected to electronics which export the data. As an example, the current ATLAS silicon strip tracker contains 6.2 million channels for sensing particles, crossing 4088 detector modules - each module with its own 100-metre cable. When upgraded, the detector is expected to boast some 60 million channels. But with serial powering, it's possible that the tracker won't need any extra cables.

One major down side of a serial powering approach used to be apparent whenever one of those Christmas tree lights burnt out: the whole strand of lights would suddenly stop working. Modern strings of lights get around the problem by providing alternate pathways so that electricity can bypass burned-out bulbs - an approach that the WP8 team also employed.

Passing the power on

In order for serial powering to work in the tracking detectors, a pair of cables needs to run a constant current through a series of detector modules, and each module must be capable of turning that constant current into a constant voltage.

This is accomplished with a shunt regulator, which consumes enough current to create the desired voltage and sends the rest down the cable towards the next module.

For WP8, the serial powering effort is headquartered at RAL, but it also includes groups at the University of Bonn in Germany, AGH University of Science and Technology in Krakow, Poland, and other institutes around Europe and the US. "Each of these groups covers different aspects of the serial powering approach," says Giulio Villani of the WP8 serial powering team.

Each silicon strip detector module is powered by twenty or so application-specific integrated circuits (ASICs), of the type ABCN-25. These ASICs were designed by the members of the ATLAS collaboration - including members from including AGH Krakow, CERN, the University of Geneva and the University of Pennsylvania in Philadelphia - and manufactured by IBM.

One way to allow each ASIC to perform this powering role is to provide it with its own shunt regulator, so 20 conversions from constant current to constant voltage - one for each ASIC - take place on a single module. But there are other possible approaches. For instance, a Serial Powering Interface (SPi) custom chip, developed in a close collaboration between Fermilab in Batavia, Illinois and RAL, could regulate the voltage for the whole module. And a third method splits up the conversion process so that part of it happens at the module level, finished off by each of the 20 ASICs.

The serial powering team also developed a version of the technology for powering the pixel modules within the tracker, as an FE-I4 ASIC rather than ABCN-25. The sensors can run on a lower voltage than the one needed to process and transmit data out, which reduces the noise in the particle signals as well as the total power needed by the detector. To provide two voltages, one for the sensors and one for handling the data, each pixel chip was equipped with two voltage regulators, developed at the University of Bonn.

Standardising the signals

The next problem is standardising the signals coming out of the modules. Although the voltage drop across each module is the same, the voltage reference point of each module in the chain changes - in the same way that water flowing down a flight of stone steps drops the same distance with each step, but is each time brought a little closer to the bottom of the flight: a fixed reference point.

Because each module is at a different voltage compared to this fixed reference point, communication with the data acquisition and timing systems needs to be separated from the DC current. These signals come through AC currents fed in and out of each module, delivering clock information and carrying out data. Extra circuits separate the AC signals from the voltages so the information stays standardised.

Protection from blackouts

To keep the whole chain from failing when trouble arises in a single module, the WP8 team developed a way to bypass individual modules that may begin to create too much noise in other sensors or stop working entirely. A separate protection circuit powers transistors, gateways that open or shut in response

to a voltage, that sit in front of each module. When a transistor is powered, it sends the current uninterrupted past the faulty module, towards the remaining chain of modules.

The power protection boards which contain these electronics, supplied by Brookhaven National Laboratory in Upton, New York, are compatible with all three serial powering schemes considered for the strip detector modules. The University of Bonn developed different protection boards for the pixel modules.

These boards can sense problems in modules, such as unusually high voltages, and disable them automatically. Alternatively, the modules may also be manually switched off through the remote control system; for instance, a module might produce a lot of electromagnetic noise which corrupts its own data and affects nearby modules, so the detector is better off without it.

Putting serial power to work

To test out the new serial powering electronics, the WP8 group teamed up with the ATLAS Inner Detector upgrade effort, which seeks to improve both the pixel and silicon strip detectors at the heart of the ATLAS experiment. At RAL, the collaboration developed a stavelet comprised of eight silicon strip modules. With 20 ABCN-25 ASICs on each module, this meant that 160 ASICs were powered on a single set of cables, reading out over 20,000 individual sensor elements.

This feat represents a more than tenfold reduction in the number of cables needed to power the ASICs. Moreover, the close collaboration with those designing the next ATLAS silicon strip tracker has meant that the serial powering is already integrated with the improved sensor design and data readout scheme.

These test stavelets proved that serial powering does not introduce extra noise that would interfere with the smooth running of the detectors, and that individual modules can be switched off while the others continue to run. The next step is a set of 24 modules powered in series, 12 on each side of the stave. This 'supermodule' would be a full-size component of the upgraded ATLAS tracker.

"The supermodule projects are more extensive than just the SLHC-PP," says Dabrowski. "Many other groups are contributing."

The ATLAS pixel detector upgrade is also testing the serial powering scheme, using the ASIC FE-I4. The University of Bonn and Lawrence Berkeley National Laboratory in California have collaborated on developing pixel staves for the outer cylindrical layers of the detector, in which they integrate a single, low-mass power and signal cable directly into the light carbon foam support structure.

The modules consist of two FE-I4 chips bonded to a sensor containing 53,760 individual pixels, and eight of these sensors are connected in series. This method is under investigation for the replacement of the ATLAS pixel detector, which could occur in 2017. With only four power lines needed to feed the 32-module outer layer of the upgraded ATLAS pixel detector, serial powering promises many fewer cables than the present pixel detector.

"Within this project, we worked on common problems between CMS and ATLAS," says Dabrowski. Steinar Stapnes, WP3 project leader for ATLAS, sees the development of new powering strategies

through cooperation between the two detector collaborations as one of the most valuable aspects of the SLHC-PP from the detector point of view.

Deciding between serial powering and DC-DC conversion depends on the needs of the detector, but Dabrowski is confident that the findings of WP8 will be put to good use. "One thing is obvious - one of these powering schemes will be used for sure," he says.

Potential Impact:

SLHC-PP Impact

The plan to upgrade the LHC is an investment in future scientific discoveries. The management and coordination work packages were not designed to generate results in themselves but to facilitate the smooth upgrade of a major scientific facility. The technical work packages as well were more concerned with solving problems than probing the unknown. Nevertheless, work packages 6-8, as well as the detector upgrade plans, have advanced the technology of high energy physics even at this early stage.

Moreover, the various teams in the SLHC-PP have been fastidious about communicating their goals and progress to both the particle physics and the wider scientific communities. They have also reached out to interested members of the general public. Developments made in part through SLHC-PP projects are already contributing to other experiments within the field of particle physics and could eventually prove important in other areas of scientific endeavour, including space exploration and medicine.

Communication and outreach

The meetings organised through WP1 gave the accelerator and detector collaborations a window into the preparation for the upgrade to the LHC. The three annual meetings, where collaborators discussed the project status at a technical level, were open to the whole of the particle physics community. They were held at CERN, CEA Saclay in France, and CIEMAT in Madrid, Spain. The host institutes also included talks on other particle physics and accelerator efforts within these conferences. Agendas for each of the two-day meetings can be found below:

- 2009: <http://indico.cern.ch/conferenceDisplay.py?confId=49777>
- 2010: <http://indico.cern.ch/conferenceDisplay.py?confId=70620>
- 2011: <http://indico.cern.ch/conferenceDisplay.py?confId=116438>

The SLHC-PP public events brought together key experts from CERN's highest levels to discuss long-term strategies and prospects for the LHC. Notable speakers from beyond the SLHC-PP team include CERN research directors Jos Engelen and Sergio Bertolucci; Steve Myers, Director of the Accelerator Division; and Phillippe Bloch, head of the CERN Physics Department. These events were always held at CERN to maximise attendance, Mar Capeans explains. Webcasting allowed interested individuals worldwide to view the proceedings.

Capeans describes the 'kick-off' public event marking start of the SLHC-PP project, held on 9 April 2008: "In a packed auditorium, the event began with a speech by CERN's chief scientific officer, Jos Engelen. He emphasised the importance of developments towards the SLHC within the European particle physics strategy and commended the position taken up by SLHC activities within CERN's overall program."

Three overview speakers summed up the scientific potential that an upgraded LHC promised, the accelerator upgrade plans and rough schedule, and the upgrade plans for ATLAS and CMS. "The event concluded with lively discussions about the impact of the announced gradual luminosity increases on the present physics, operation, and upgrade plans of these experiments," says Capeans.

Members of the public who are curious about the upgrade to the LHC but missed the webcasts can still view the slides on the event web pages:

- 2008: <http://indico.cern.ch/conferenceDisplay.py?confId=30583>
- 2009: <http://indico.cern.ch/conferenceDisplay.py?confId=49774>
- 2010: <http://indico.cern.ch/conferenceDisplay.py?confId=95580>
- 2011: <http://indico.cern.ch/conferenceDisplay.py?confId=117411>

As slides are not as effective without a speaker, the SLHC-PP also provided a stand-alone website which describes the upgrade effort: <http://slhcpp.web.cern.ch/SLHCPP/>. It explains the various upgrades and the motivations behind each of them at a basic level, in the lively style of a popular science article.

CERN publications have discussed progress on the LHC upgrade, including interviews with Steve Myers and CERN Director General Rolf-Dieter Heuer in the CERN Bulletin. The CERN Courier marked the start of the SLHC-PP with an article entitled "The Super-LHC is on the starting blocks". It also highlighted efforts made through WP8 with another piece, "Electronics experts connect in Aachen", published in January 2011. The ATLAS e-News has been following the upgrade effort from initial funding to the design of new tracking detectors. Links to these articles are listed below:

Articles in the CERN Courier

- "The Super-LHC is on the starting blocks" (<http://cerncourier.com/cws/article/cern/34932>)
- "Electronics experts connect in Aachen" (<http://cerncourier.com/cws/article/cern/44854>)

Articles in the CERN Bulletin

- "CERN in FP7: very successful participation so far" (<http://cdsweb.cern.ch/record/1159929?ln=en>)
- "A global view" (Rolf Heuer) (<http://cdsweb.cern.ch/record/1154954?ln=en>)
- "Stephen Myers - More collaboration for accelerators" (<http://cdsweb.cern.ch/record/1156859?ln=en>)

Articles in the ATLAS e-News

- "Super-LHC preparation gains momentum - and some funding" (http://atlas-service-enews.web.cern.ch/atlas-service-enews/2007-8/news_07-8/news_slhc.php)
- "ATLAS upgrade gains intensity" (tracker workshop) (http://atlas-service-enews.web.cern.ch/atlas-service-enews/2009/features_09/features_upgrade.php)
- "Thumbs up for Insertable B-Layer" (http://atlas-service-enews.web.cern.ch/atlas-service-enews/2010/news_10/news_Insertable%20B-layer.php)
- "The ATLAS tracker upgrade: Silicon detectors for the sLHC" (http://atlas-service-enews.web.cern.ch/atlas-service-enews/2011/features_11/feature_ATLAS_tracker_update.php)

Although scientific research was not the main purpose of the SLHC-PP, the explorations of upgrade possibilities are chronicled in 43 peer-reviewed papers published in over ten different scientific journals

and conference proceedings. The upgrade work has been discussed at 49 workshops and 60 international scientific conferences. These workshops and conferences ranged from the highly focused - such as a meeting concerning the air core inductors needed for the DC-DC converters - to general assemblies of national physics societies.

Researchers in academia and in industry, attended these conferences and workshops in Austria, Canada, Croatia, France, Germany, Greece, Italy, Japan, the Netherlands, Portugal, Spain, Switzerland, the UK, and the US. A full list of publications, workshops, and conferences may be found on the SLHC-PP website.

WP6: Magnetic education

Stephan Russenschuck, leader of WP6, says that the one of the most important forms of dissemination is the training of young physicists and engineers. "After having designed and built the LHC for the last 20 years, we must now ensure the expertise for the operation, upgrade, and repair of components for at least the same time span," he says. "With the flagship accelerator in operation, CERN has been able to attract a record number of technical and doctoral students, not only in physics but also in the engineering sciences."

The magnet system - including the liquid helium cooling system and vacuum systems - is an ideal training ground for physicists and engineers. The carefully tuned magnetic fields are an aspect of good beam physics, and the magnet protection system, which ensures that the magnet will be safe if part of the superconducting coil becomes resistive, falls into another discipline: computer and numerical simulations of real systems. The 'barber pole' insulation, the interlocking collars, and the development of new manufacturing techniques and tools combine mechanical and electrical engineering.

Four PhD theses described developments made in part through WP6. "The expertise gained and the developed methods serve not only for the building of accelerator magnets but can also be applied to a number of technical devices such as motors and generators, and energy storage," says Russenschuck.

Other applications include the potential to apply superconductivity more broadly. The technology could ultimately have a huge impact on Maglev trains and could revolutionise energy storage and delivery with resistance-free circuits and a lossless power grid. These potential applications became popular in the early 1990s, with the discovery of high temperature superconductivity. But 'high temperature' is a relative term - these materials need to be well below -100°C to begin superconducting. Until superconductivity reaches more ordinary temperatures, exotic applications remain a dream.

Even so, low-temperature superconductivity remains hugely important in areas like high-energy particle physics, nuclear fusion reactors, high field nuclear magnetic resonance (NMR) and full-body magnetic resonance imaging (MRI) machines. The magnetic fields needed for these applications are so strong that they can be achieved only through superconducting electromagnets, typically cooled with liquid helium.

"The communities are all facing the same problems: saturating performance, long development lifecycles, cost, reliability of operation, supply chains, and knowledge transfer," says Russenschuck. Projects like the

upgrade to the LHC, which push magnet technologies further and train new experts, are valuable to the whole field.

In Russenschuck's view, CERN and other national and international laboratories, as well as their funding bodies, have certain responsibilities: "to keep the expertise for the design and construction of such devices, to push the development of superconducting material such as niobium-tin, to ensure the technology transfer to the industrial partners and to develop reliable, cost efficient, and accurate manufacturing techniques."

Because the market is so limited and does not invite mass production, it is hard to find industrial partners for superconducting technologies. Even so, Russenschuck notes that along with the WP6 efforts, CERN has qualified new industrial partners in the UK, Bulgaria, and Portugal to produce components of the superconducting magnets.

WP7: Synergy with other European projects

Like WP6, WP7 also contributed to four PhD theses. But since its developments towards the Superconducting Proton Linac no longer serve the LHC directly, the collaborators found other projects for which their expertise is valuable.

The negative hydrogen source developed with the high-intensity SPL in mind can help serve Linac4, as the engineers are finding it difficult to coax the current source up to the beam intensity and high reliability that is needed at the start of the LHC injector chain. "Aspects of the SPL source design will be a basis for the Linac4 source, and the SLHC-PP allowed us to build the source test area that can be used for its further development," says WP7 project leader Richard Scrivens.

Ion sources like the one designed in WP7 are used in many particle accelerators, and different version of ion sources have a variety of purposes. Some are used in electronics manufacture to pepper semiconductors with 'doping' elements. Done in a controlled way, this changes the semiconductor's properties, essential for their use in electronic circuits. Ions beams can also work at the surfaces of materials, etching patterns into chips or modifying the properties of powders.

Protons and ion beams can also take aim at tumours in cancer patients. This treatment, called hadron therapy, offers much more concentrated delivery of destructive energy to the tumours. Unlike neutrons or x-rays, charged particle beams can be tweaked so that they stop inside the tumour. This is better for the patient because most of the energy gets deposited at that stopping point rather than evenly distributed along the beam's path, where it might damage the healthy tissue surrounding a tumour.

Ion sources are also used to initiate fusion reactions - still seen by some as offering an alternative to the nuclear fission reactions occurring within today's nuclear power plants. Unlike those fission reactions, fusion promises to leave no long-lived radioactive waste in its wake [reference: <http://www.iter.org/safety>]. "The closest synergy for the ion source development of the SLHC-PP project is with the neutral beam injectors for Tokamak Fusion Reactors," says Scrivens. ITER, the project formerly known as the International Thermonuclear Experimental Reactor, seeks to make the fusion reactor a commercially viable, sustainable energy source.

The Tokamak reactor would fuse deuterium and tritium, forms of hydrogen containing one and two neutrons, respectively. This reactor's ion source is similar to that designed through WP7, using a spiral antenna to carry a current which alternates at a radiofrequency. The beam from this plasma would then be shot into the reactor, where it would serve to heat the plasma, helping the hydrogen atoms to fuse into helium.

The SPL studies also resonate with another project. "The European Spallation Source (ESS) project in Lund, Sweden, is extensively interested and is co-supporting this research and development in view of using it in its own accelerator," says SLHC-PP project leader Roland Garoby.

This facility will provide large numbers of neutrons using a similar accelerator, converting the energetic ions to neutrons by ramming them into a 'target', typically a tile of heavy metal. The resulting high-intensity neutron beam will then be used to probe molecular bonds, exploring chemical reactions and the structures of new materials and medicines. Neutron beams can also help archaeologists analyse ancient ceramics without damaging them, revealing the materials under the glaze and even the temperature at which the pottery was fired.

For this reason, the development of systems which can run two RF cavities from a single klystron continues. However, the project is not solely for the benefit of ESS - CERN is still considering building the SPL sometime in the future, though not as part of the LHC. This version of the accelerator is designed to send out many more pulses of ions each second than Linac4, so it could be used instead to generate neutrinos or produce rare and radioactive nuclei for the exploration of physics at the heart of the atom.

Neutrinos are ghostly particles that travel at near the speed of light and hardly interact with the atoms that make up most of the universe. There are three different kinds, and although their masses seem to be zero - within experimental error - indirect evidence has shown that they have some small mass. Measuring their masses, or even just figuring out which is heaviest and which is lightest, is a major goal in neutrino physics.

Furthermore, experiments at Fermilab in 2010 hinted that there might be one or two undiscovered neutrinos. The number of neutrinos, and their masses, affected how galaxies formed and so could give more insight into the early universe. With a different kind of target, an SPL-like machine could be turned to generating neutrinos, which could feed experiments that are designed to solve these mysteries.

Although 2011 marks 100 years since the discovery of the atomic nucleus, they still confound researchers. In December last year, CERN's own Isotope Separation On Line Detector (ISOLDE) collaboration announced a surprise fission of an isotope of mercury - rather than dividing into two identical zirconium nuclei as the best theories predicted, it instead divided into krypton and ruthenium.

In addition to providing experimental ammunition for improving nuclear theories, rare isotopes generated with an ion accelerator such as that developed by the WP7 collaboration can be used to study the production and radioactive decay of atomic nuclei that may form in the hearts of stars or in cosmic explosions. Earthly beams can tell nuclear astrophysicists more about space.

To push the boundaries of nuclear physics, experimenters are increasingly turning to the creation of radioactive beams. These are formed by smashing an initial ion beam into a target to create a shower of radioactive nuclei, collecting the desired radioactive isotope, and then accelerating it using another

machine. Slamming this beam into another target produces even rarer forms of atomic nuclei. A high-intensity, SPL-like source could produce the ions for the initial beam in greater quantities, allowing new exotic isotopes to be produced in measureable quantities.

WP8: Proving new powering systems

The low-mass power distribution system developed in part through WP8 will serve the upgraded LHC as well as the future experiments on the linear collider that will study LHC discoveries in greater detail. One of the major successes of the effort was simply to demonstrate to the community that these powering schemes were possible in particle detectors.

"When we started the WP8 program there was a lot of scepticism concerning the feasibility of either the serial powering or using DC-DC converters on the detectors," says Wladyslaw Dabrowski, WP8 project leader. "The impact of our development was clearly visible at the yearly Topical Workshop on Electronics for Particle Physics (TWEPP) conferences, where presentations of WP8 results drew the attention of people from other detectors in LHC experiments and from other experiments."

Although DC-DC converters were already widely available in industry and consumer electronics, Dabrowski says that the low mass, radiation-resistant converters that WP8 helped develop may find homes outside earth's atmosphere. Research and communication satellites to be launched into space must be as light as possible but also radiation resistant since they no longer have the atmosphere to protect them from cosmic rays.

The long view

The SLHC-PP's real impact will be the upgraded LHC itself. Over 4000 physicists, engineers, and technicians, from CERN's 20 European member states and beyond, are expected to contribute to the upgrades and analyse the resulting data. The management and coordination work of the SLHC-PP will help cement the relationships between CERN and the contributing institutions in the context of the upgrade projects. Likewise, the safety and radioprotection issues addressed through WP5 ensure that the more powerful accelerators meet high safety standards of the contributing institutions as well as legal requirements.

The improvements to the LHC - including the LHC Injector Upgrade, High Luminosity LHC, and detector upgrade projects - will allow Europe to maintain its leadership in the field of high energy physics into the 2030s, and the SLHC-PP has laid some of the crucial ground work. Once the machine is running at high luminosity, the abundant data reaped by the detectors will allow physicists worldwide to describe known particles and forces more accurately, refine new discoveries made by the LHC, weed out versions of exotic theories from the realm of the possible, and extend searches for new particles and forces.

But the SLHC-PP also bore early fruit, contributing to the education of new scientists and engineers, forging new links between science and industry, and connecting with other scientific experiments within the field of particle physics. Through improvements on common devices such as superconducting magnets, particle sources, and sophisticated particle detectors, the gains made in part through the SLHC-

PP could impact experiments outside particle physics, and may even prove useful tools in the diagnosis and treatment of cancer.

List of Websites:

- <http://slhcpp.web.cern.ch/SLHCPP/>
- <http://info-slhcpp.web.cern.ch/info-SLHC-PP/>
- <http://project-slhcpp.web.cern.ch/project-slhcpp/>

List of beneficiaries (Institution, Institute representative for SLHC-PP):

- CERN, Roland Garoby <Roland.Garoby@cern.ch>
- AGH-UST, Wladyslaw Dabrowski <W.Dabrowski@ftj.agh.edu.pl>
- CEA-Saclay, Jean-Michel Rifflet <jean-michel.rifflet@cea.fr>
- CIEMAT, Fernando Toral <fernando.toral@ciemat.es>
- CNRS-IN2P3, Daniel Gardes <gardes@ipno.in2p3.fr>
- CTU, Stanislav Pospíšil <Stanislav.Pospisil@utef.cvut.cz>
- DESY, Guenter Eckerlin <guenter.eckerlin@desy.de>
- ETH Zurich, Guenther Dissertori <Guenther.Dissertori@cern.ch>
- FOM-NIKHEF, Nigel Hessey <nigel.hessey@nikhef.nl>
- GSI, Georg Fehrenbacher <g.fehrenbacher@gsi.de>
- Imperial, Jordan Nash <Jordan.Nash@cern.ch>
- INFN, Paolo Pierini <Paolo.Pierini@mi.infn.it>
- PSI, Wernli Christian <Christian.Wernli@psi.ch>
- STFC, Stephen Haywood <s.haywood@rl.ac.uk>
- UBONN, Markus Cristinziani <cristinz@uni-bonn.de>
- UNIGE, Allan Geoffrey CLARK <Allan.Clark@cern.ch>
- USFD, Ian Dawson <Ian.Dawson@cern.ch>
- RWTH Aachen, Lutz Feld <Lutz.Feld@physik.rwth-aachen.de>