WHAT ARE THE WEAK POINTS OF LHC OPERATION?

B.Goddard CERN, Geneva, Switzerland

Abstract

The few weeks of LHC operation in 2009 was a resounding success with extremely fast progress in the beam commissioning. Nevertheless, the period also revealed a number of possible weaknesses in various aspects of the LHC operation, including procedures, tools, discipline, equipment and organisation. The weaknesses are discussed with a view to feeding the experience back to improve the machine operation for 2010 and beyond. The distinction is made impacting efficiency, and those points which are more serious and might impact machine protection.

INTRODUCTION

You cannot run away from a weakness; you must sometimes fight it out or perish; and if that be so, why not now, and where you stand?†.

The short period of LHC operation in 2009 without and with beam revealed a number of weaknesses which could impact on machine efficiency, or potentially on machine protection. This paper describes some of the weaknesses, grouped rather arbitrarily into Preparedness, Injection, Experiment-machine interface, Sequencer, Ergonomics, Discipline, System specifics, Procedural and General. No coherent attempt is made to propose or explore remedial issues for each point, but those weaknesses which are judged to have a potential machine protection impact are highlighted with the potential implications. Some of the more general issues highlighted are certainly in the 'motherhood' category of statement; however, it is hoped that specific examples will actually be of practical use.

PREPAREDNESS

In general the degree of preparation for the operation of the LHC with beam in 2009 was extremely good, with thorough checks of systems, functionalities, software, applications and operational scenarios. This was reflected in the very rapid progress with the beam commissioning, which allowed all the target commissioning steps to be reached, and some non-planned steps to be added. The weaknesses were:

No clear definition of detailed conditions/results needed to move to next commissioning step

Although the coarse-grained target conditions were defined, the lack of fine detail and the rapid progress led sometimes to heated ad-hoc arguments about whether it was safe to proceed to a new set of beam conditions.

† Robert Louis Stevenson

Informal approach to planning and test tracking

The loose approach to the beam commissioning was well adapted to the need for fast progress with very low intensity beam, and indeed proved very efficient and gave fast progress. However, this approach will not be suitable when beam intensities and energies increase in 2010, since much more discipline and formality will be needed to validate each change in conditions. In addition this fast-moving style relies on an excellent and wide ranging communication and awareness which must be provided by the machine coordinators. This level of intensity may not be sustainable by these few people for a full 9 months.

Some things not fully tested in Dry Runs or Machine Checkout

Despite the very complete Dry Runs, some systems were not adequately tested, or for different reasons were not included in the Dry Runs. Some examples where full tests were not performed were:

- Inadequate stress-tests (concentrators, proxies, CMW subscription);
- Settings (knobs for tune, chromaticity not available for squeeze);
- Partial tests only for some systems (alarms, collimators, feedback…).

A general effort should be made to include all systems and functionalities into the general Dry Run and Checkout framework; at the present this is not enforced.

INJECTION

Injection into the LHC was improved from 2008 but was too often a source of error or inefficiency for operation. A rather large number of specific issues were identified – making the injection more solid will need a large number of relatively small issues to be solved, with sequencer, sequences, IQC, LBDS arming, SIS, etc.

Frequent accidental over-injection

Over-injection onto an already circulating beam happened many times. This was due to mistakes by EIC/experts, having wrong/multiple sequences running, injecting into the wrong ring, injecting with circulating beam present, IQC faults, debunched beam present etc. This was inefficient since generally an emergency dump resulted from the beam losses, but could also have a machine protection impact, since accidental over-injection with high intensity would stress the protection in the sense that the correct positioning of the TDI absorber is the last line of defence.

Injection Quality Check application (IQC)

The IQC checks were commissioned in 2009 and proved work well under most conditions. However, it needs to be faster (CMW subscription data latency), should not be 'optional' in the Injection Sequencer, and sometimes gave the wrong result ('no beam injected') if the injected intensity was too low, which led to over-injection. Also the MKI acquisition (IPOC) was sometimes missing.

Injection with ring with wrong settings

This is potentially dangerous and occurred several times, for example with several sectors still set to pre-cycle current; with collimators and protection devices retracted. This could easily lead to a quench with pilot bunch. An overall surveillance or sanity check in SIS or with the sequencer may be possible.

BLM sensitivity and saturation

This could be a major issue for the operation, since the fact that the BLMs are very sensitive and saturate at low intensity means that even the loss of a single pilot saturates the BLMs on the fastest time integration scales. This problem already prevents over-injection onto the IP8 TDI [1] and affects the procedure for the setup of TDI/TCDI. The local issue of the TDI has to be solved without compromising machine protection functionality. In addition higher intensity will produce losses on the TCDIs which are near LHC [2], and which are also expected to generate interlocks at intensities well below nominal. To address these issues, in addition to the BLM thresholds, the beam tails need to be investigated, together with the systematic use of the SPS scrapers.

Injection and circulating beam for many turns with screens in

This happened at least once, with the SIS masked, and could be dangerous for screens even for low intensity. The limits for the screens are clearly defined and the interlocking is available in the SIS; however, the interlock must remain enabled under all conditions.

EXPERIMENT-MACHINE INTERFACE

The experiments themselves had a long and useful list of problems which occurred [3], including poor communication with the CCC. From the operational side there were several issues which require improvement:

Tacit acceptance of unofficial 'quiet beam' mode

The experiments took physics data in poorly defined conditions, when the minimum required machine protection setup and checks had not been made. This was with the complicity of operation, and was driven by an understandable eagerness to get 'first' data, from the experiments, but also from the management and the machine. The practice was sloppy and should not be tolerated, since it encouraged bad habits and posed a direct danger for the experiments. Similarly, there were instances of experiments turning detectors on without informing operation.

Communication with experiments

Several communication problems were encountered, including failure by operation to signal End-of-fill, machine mode changes forgotten, slow handshake (e.g. injection, before dumping beam, …). The manual mode change from 'stable' after a dump could be replaced by an automatic mode change; this was sometimes forgotten and is needed by at least ATLAS to switch off

Experiments slow giving injection permit after Post Mortem

Some experiments were systematically slower than the others to give the injection permit after a post-mortem event - ATLAS in particular was much slower than all others and hours of beam time were lost in some tests. A method of quickly giving the injection permit when the beam dump was not triggered by the experiment should be investigated.

SEQUENCER

The LHC sequencer is a key system for being able to efficiently operate the machine. It also has a bearing on the machine safety, since although not a safety system itself, incorrect sequences or actions can stress the machine protection systems and multiply the risk of a dangerous incident. Some of the problems identified were with the operation of the sequencer, which were sometime exacerbated by ergonomic aspects, some problems were with the sequences themselves, and some problems were with the functionality provided.

Sequences left running unintentionally

On occasion another sequence was left running which interfered with the desired sequence, or a sequence was running twice on different machines. This was particularly the case during injection, for special modes like Inject and dump or Circulate and dump. An effort to check multiple instances of running sequences could help.

Playing wrong sequence at a bad time

Some sequences were launched by mistake, despite the "are you sure?" panel which seems to be acknowledged as a reflex on occasion. The Softstart of the MKI kickers was launched by accident which dumped the beam in the machine, the sequence to set all collimators to parking launched by mistake was difficult to recover from, etc. Errors of this nature are not strictly the sequencer's fault – but it might be possible to add some 'protection' when trying to launch certain sequences in an abnormal beam mode or unusual time in the overall operational cycle.

Wrong timing tables loaded for given sequence or mode

This was a frequent source of confusion during injection, and sequences should include tasks to unload and load relevant default timing tables at the end of sequence, to return the machine to a well-defined state. There could be an issue with looping sequences which might need special entrance and recovery sequences for this purpose.

Sequence structures not optimal

The flat representation in the Sequencer application means that entry points into sequences are not clear, and that sequence and subsequence hierarchies are not well defined (or at least not apparent in the application). An effort to be able to load and view sequences in a proper hierarchical way, with collapsible layers, is needed.

(Too) easy to skip any line, and also to edit sequences

The protection of some sequences and tasks is an issue for the most safety critical requirements, like the loading of the collimator thresholds. At present it is trivial to skip or alter such tasks, and the sequencer needs an effective way to prevent this, with RBAC.

ERGONOMICS

The ergonomics of the applications and displays in the control room make a big difference to the efficiency of operation and can, as already mentioned for other issues, result in circumstances which produce an unnecessary stress on the machine protection systems when not optimised. Some of the particular issues identified are:

Fixed fixed displays

The setup of the fixed displays which give essential overview information to the operators varied greatly, and depended on the personal preferences of EIC. This meant that important information was sometimes not to hand, and also that time was spent between shift changeovers in reorganising the displays. It must be possible to fix given displays to given consoles, and also to prevent other applications or windows from being overlaid.

No overview of collimator positions

On several occasions it was necessary to open many individual control applications to try to determine what the status of the overall system was. A need was indentified for an easy way to determine whether the global set of collimator positions is correct.

Fill number clocking erratic

The fill number clocking was not consistent, which made retrieving data difficult. This will be essential in the future for being able to recover and reanalyse operational, commissioning and MD data.

Page 1 complicated and sometimes frozen

Much of the complication on Page 1 is certainly needed, but overall it could be streamlined, for instance the inclusion of four separate BTV screen shots for much of the time seems gratuitous.

DISCIPLINE

My weakness is wearing too much leopard print[‡].

Discipline by operations and experts alike is an area which will probably always leave room for improvement. Many protection and operation systems are well-defined but rely on discipline in their deployment or operation to be effective. Some systems are less well defined, and these need even more rigour. Some specific examples or areas where discipline is particularly important are:

Parallelism

The intermingling of commissioning and luminosity operation should stop, since it gives rise to confused situations where priorities and pressures are different, and leads to compromises which could be dangerous for the machine.

Parallel MD or commissioning studies were also sometime problematic, for instance Beam 1/Beam 2 interference problems, and also from alleged ‘parasitic’ studies. Several times the other beam was dumped after a manipulation by the ‘parasitic’ partner, due to cross-talk between BLMs.

Machine state after MD or commissioning tests

Returning to operation from MD, or the machine state at the handover from one MD to another, provided a rich source of problems arising from indiscipline. For example, on occasion corrector settings were wrong, protection devices were left at wrong settings and the RF system was left off. For tests with safe beam this might only prove an inconvenience, as long as a reproducible way to reset the LHC to the run configuration can be applied. However, for tests and MD with unsafe beam (for example proposals to setup collimators with higher intensity), indiscipline could be a direct danger to the LHC.

Updates of Page 1 and machine mode not made systematically

This could be improved by making automatic updates where possible, although this might further reduce the frequency of update for the non-automatic changes.

Reset of TCDQ energy interlock after precycle

This interlock latches by design, but was sometimes overlooked in the precycle. A task in the pre-cycle sequence should reset this automatically

(Too) easy to disable ‘required’ functionalities

The disabling of SIS channels, IQC checks and tasks or steps in sequences was understandably frequent; however, this sometimes became routine and must be avoided in future, since re-enabling was not always done.

[‡] Jackie Collins

TCDI settings not updated to agree with deployed thresholds

This led to some inefficiency through unfamiliarity with the update mechanism and details, and a simpler (but still safe) method of a coherent update of settings and thresholds after setting up could be investigated.

PROCEDURAL

Perceived procedural weaknesses concerned the execution of the agreed commissioning plan in 2009. Several of these concerned machine protection, and how to defined, maintain and enforce the conditions required for the present operational phase:

No tracking or enforcement of allowed 'operational envelope'

For 2009 where the beam intensity remained low this was not a major issue, although any competent person was free to the fill machine with 'high' intensity. In 2010 this will become crucial, and interlocking is needed to prevent injection of higher intensity than approved.

No enforced Post Mortem analysis or systematic offline analysis of emergency dumps

The EIC should check and acknowledge the PM data from an emergency dump, to make sure that any potential problems are not missed. In addition, the Machine Protection team should setup an offline mechanism for checking emergency dumps in more detail, to look for long term trends, correlations or patterns of faults.

Patchy tracking of MP commissioning tests, with no central overview

Formal tracking of the Machine Protection tests was left to the individual system experts and was patchy or non-existent. There were many culprits, including the LBDS and the BIS; and even when records were taken they were frequently kept privately in varying formats – this meant that it was not possible to obtain a snapshot of the status of the machine protection system. Effort should be made to use the available tracking framework, which will allow reviews of, and decisions on, progress to be made on the basis of more complete information.

Systematically checking tunes, orbit, lifetime, chromaticity, filling pattern, etc. after injection

Before ramping and putting the beams into physics a set of checks are needed on the machine state and the beam parameters. A procedure or sequencer checklist is needed to ensure that the correct checks are made systematically.

Information exchange between EICs, operators, experts and coordinators

The daily 17:00 meeting was the *de-facto* forum for exchange of some of the relevant operational information, with other details transferred in the shift hand-over and some by other channels. A simple, SPS-style wiki site for a basic knowledge base and day-to-day issue collection would centralise the information and be easy to maintain.

Coordination and decision making

The difficult coordination of the very fast commissioning in 2009 could maybe have been improved by keeping the planning, status, results and relevant information more up to date, and in a central place. Full-time coordinators would possibly have been able to manage this.

In addition, sometimes decisions on the progress or priorities were subject to pressure from experiments anxious for collisions, machine experts anxious for beam time with their system and also from management anxious for progress. On occasion this resulted in inefficient last-minute changes of plan, discrepancies between information given to different people, operational configurations which had not been fully evaluated and aberrations like the 'quiet beam' mode.

SYSTEM SPECIFICS

A number of weaknesses with specific systems were noted. Many of these have already been addressed or are in the process of being fixed for the 2010 startup, although some are more difficult general or conceptual issues and will require longer to solve:

Feedbacks

The tune feedback gave problems after an emergency beam dump, since the system carried on correcting on noise and put the machine into an unknown state. A 'recover' button is needed, and some way of stopping the feedback as soon as the beam has been dumped.

For the orbit feedback erratic sending of real-time data to correctors led to problems, for example the MCBX.

RF systems

The recovery of the 18kV cells after access needs a clear procedure and possibly a SW check. Injecting with RF off is possible; this should be checked and prevented if not specifically requested. Resynchronisation of the RF system gave problems and it was not always clear whether the task in the sequencer worked or not. The BQM could be used to check this explicitly.

Interlock BPMs

The IR6 interlock BPMs gave wrong results and tripped when the beam intensity is too low. This affected some measurements and ramps. In addition a systematic reset of the latched interlock should be added in the sequencer.

Synchrotron radiation telescope BRSA

The undulator state is not surveyed yet in any sequence or application (or if it is, it is not widely known about).

Collimator/protection device threshold management

With the present system for managing and changing the collimators there is no protection on who opens the thresholds and when, since this is achieved via a beam process which can be driven by EquipState. The same is true for the jaw movement. Although the energy interlock remains in place and is less easily modified, there is still no redundant protection for the squeeze (no β^* factor).

LBDS XPOC

The Beam Dump eXternal Post Operational Check software is integrated into the PM framework but still gave a number of problems. The system was not stable enough (timestamp errors, MKD energy discrepancies etc.), and the majority of the BI checks remain to be deployed. Threshold management is still too rudimentary, and needs to be integrated into a dedicated application.

LBDS internal trigger latency to BIS

A conceptual weakness in the LBDS triggering system could lead to injection into machine with the LBDS unarmed, in the event of an internal emergency dump from the TSU a few ms before an injection.

Inputs to SMP need to be rock-solid

The Safe Machine Parameters depend on reliable acquisition; the CMW and JMS broker instabilities affected the BCT data. There is also the more general issue of the SPS/LHC BCT reliability and redundancy, since a false positive signal on the Beam Presence flag could allow injection of a full batch into an unready LHC.

Underground access systematically requires PCs OFF and power permit removed

The safety requirement to remove the power permit for *any* underground access adds a large (2 hour) overhead to all accesses, even when limited to UAs. This condition should be reviewed, since it adds a huge overhead to all accesses, especially since the subsequent precycle is often problematic, leading to a cascade of delays.

QPS resets in tunnel frequently necessary

The inability to perform remote resets of the QPS system was a source of much inefficiency.

Glitches in communication for interlock on powering/access status

Several times the whole LHC machine was flat-lined due to a transient fault in the communication of the state of the LHC. A way to improve the SW logic or the communication is required.

SUMMARY

Despite the above litany of potential issues, there are many, many **strong** points of LHC operation. These include the expertise, motivation, dedication, preparation, communication, coordination, teamwork and experience of the teams, together with the support groups, controls infrastructure and instrumentation. Many of the weaknesses identified are minor, which can only affect machine efficiency, and which should be followed-up in the 'normal' way by dry-run teams

A few points are more important or difficult to solve, and warrant a concerted effort before the 2010 3.5 TeV luminosity run:

- Make injection solid, including transient losses;
- A priori agreement of target parameters for commissioning phases;
- Method and discipline for maintaining 'operational (MPS) envelope';
- Machine Protection System progress tracking discipline/enforcement;
- Is the present "informal yet intense" planning/tracking sustainable?
- A simple knowledge base for operational aspects (SPS model?);
- Systematic beam PM acknowledge and Emergency dump analysis.

2010: RUNNING WITH SCISSORS

The very fast progress in 2009 meant that the machine already has the capacity to 'easily' allow an increase the beam intensity and energy, and to reduce β^*. There will be a high and natural pressure to get useful luminosity, and to reach important and public milestones. However, in developmental terms the LHC machine and its operation crews and experts are still at the very start of a long learning curve - having just learnt to walk, the baby will soon be running with scissors.

It is very important that Machine Protection should not passively follow the progress; it should dictate the progress, or at the very least limit progress at strategic points in the commissioning program. This is not yet the operational paradigm.

We must also remember that the Machine Protection system is almost certainly not perfect - bugs and problems are still coming, and in an installation of such complexity, can be expected to keep occurring.

Operation of the LHC in 2010 above the safe beam limit will require much more discipline than in 2009, and Machine Protection should be central to the commissioning strategy.

ACKNOWLEDGEMENT

The detailed input and comments from V.Kain, A.Mcpherson, R.Alemany, M.Lamont, J.Uythoven and M.Meddahi are gratefully acknowledged. The operation crews and commissioning teams provided much information through the logbook – any inaccuracies or misinterpretations are entirely the fault of the author.

REFERENCES

[1] M.Meddahi, *LHC Injection and Transfer Lines*, proc. LHC beam commissioning workshop - EVIAN 2010.
[2] W.Bartmann, *Injection and Dump Protection*, idem.
[3] M.Ferro-Luzzi, *Lessons from Physics Runs*, idem.

Perceived weakness	MP?
Preparedness	
Clear definition of detailed conditions to move to next step	yes
More formal approach to planning and test tracking	
Tests of all systems in dry runs and machine checkout	yes
Injection	
Make injection more solid	yes
Reduce frequency of accidental overinjection	yes
Improvements to IQC	
Sanity checks on settings before injecting	
Allow overinjection of pilot in P8	
BLM sensitivy and saturation for injection losses	
Scraping, tails and injection losses on TCDI	
Preventing injection and circulating beam with screens in	yes
Experiment-machine interface	
Stop using unofficial beam modes like "quiet beam"	yes
Systematic and timely updates of machine mode	
Speeding up handshake	
Speeding up reset of injection permit from ATLAS after PM	
Sequencer	
Prevent incompatible or same sequences running together	yes
Method to prevent wrong sequence being launched	
Improve handling of timing tables in sequencer	
Optimise sequence structures with heirarchies and entry points	
Improve sequencer GUI	
Improve protection of tasks and subsequences	
Ergonomics	
Devise a way of keeping some fixed displays really fixed	
Add easy overview of collimator positions	
Improve fill number clocking	
Simplify and improve reliability of Page 1	
Discipline	
No finely interleaved commissioning and luminosity running	yes
Avoid parallel activities which can affect each other	yes
Returning machine to default state after MD or tests	yes
Systematic update of Page 1 and machine modes	
Automatic reset of latched TCDQ energy interlock after precycle	
Prevent easy disabling of some important MP functionalities	
Improve method for updating collimator settings and thresholds	
Procedural	
Define, track and enforce operational parameter envelope	yes
Enforce PM check and acknowledge after emergency dump	yes
Set up a procedure to review all emergency dumps in detail	yes
Improve and enforce tracking of MP commissioning tests	yes
Provide central overview of commissioning status including MP	
Define proecure for systematic checks of beam parameters in fill	
Improvement of operationally relevant information exchange	
Strict adherance of machine and experimental coordination roles	
System specifics	
Stop tune (and orbit) feedback immediately that beam is dumped	
Solve problem of erratic real time data sent to correctors	
Define procedure for switching back on the 18kV RF supply cells	
Prevent injection with RF off, and interlock if RF trips	yes
Dedicated check that RF synchronisation works correctly	
Easy access to undulator state in CCC	
Improve collimator and protection device threshold management	yes
"Safe' squeeze factor for collimator threshold interlocking	yes
Improvements to XPOC stability	
Deployment of BI checks in XPOC	yes
Improvement of XPOC threshold management	
Solve LBDS internal trigger latency issue	yes
Improve reliability of data for SMP	yes
Investigate reliability and redundancy of BCTs for SMP	yes
Investigate access to UAs without removing powering permit	
Make remote resets of QPS possible	
Improve logic of interlock on powering and access status	