PROTECTION CONTROLS FOR HIGH POWER ACCELERATORS

J. Wenninger, CERN, Geneva

Abstract

The next generation hadron accelerators will operate with MW beams or store beams with an energy of many 100 MJ. Such accelerators must be protected by fast and very reliable interlock systems to avoid damage due to uncontrolled beam loss. Machine protection will constrain operation, but some operational flexibility is still required for commissioning and performance optimization. This is a substantial challenge for control systems and application programs. New tools are developed to face those challenges: critical settings management, software interlocks, role based access to equipment, automatic accelerator mode recognition etc. This talk presents some of the challenges and tools. Experience with novel approaches are discussed.

INTRODUCTION

Present day high energy hadron machines like SPS, TEVATRON (and formerly also HERA-p) accelerate beams with stored energies of a few MJ. The Large Hadron Collider (LHC), presently in the hardware commissioning phase, will store beams of 360 MJ at its design luminosity of 10^{34} cm^{-2}s^{-1} [1].The energy stored in the LHC beams is more than a factor of 100 higher with respect to the existing machines, see Figure 1. At 7 TeV the damage level for accelerator components is four orders of magnitude smaller than the nominal beam current [2].

A number of hadron accelerators in operation or under construction reach beam powers in the range of 1 to 10 MW at PSI [6], at the Neutron Spallation Source (SNS) [3] located at ORNL as well as JPARC in Japan. Future neutrino factories are aiming for beam powers on target in the multi-MW range. The IFMIF (International Fusion Material Irradiation Facility) facility aims for deuteron beams for neutron production of around 10 MW.

High electron and photon beam powers are also achieved in third and fourth generation light sources, FELs and in high power electron accelerators at Jefferson Lab [8].

Machine protection systems (MPS) are an integral part of the design of such high power machines that are protected by large numbers of interlock channels. Uncontrolled release of even a small fraction of the stored beam energy may cause serious damage to equipment or excessice activation. SNS for example runs in a loss dominated regime with a maximum allowable loss rate a 1 W/m which corresponds to a relative loss rate of $\approx 10^{-4}$/m of the nominal power [3]. Fast interlock systems based on hardware links are used to achieve high reliability in transmitting interlocks signals to dump or extraction kickers, guns etc. The reaction times of such system range from microseconds to some ten's of milliseconds depending on the hardware platform and on the criticality of the equipment. Besides the hardware aspects of the MPS, there is a growing amount of control aspects that have to be managed and that become more critical as the stored energy and power increase. The is particularly important when the MPS incorporates some flexibility for setting up and commissioning of the accelerator.A modern MPS should be integrated into the accelerator control system to ease diagnostics, provide reliable testing, limit undesired access to MPS related equipment settings. High availability dictates automated analysis and recovery processes after MPS triggered faults.

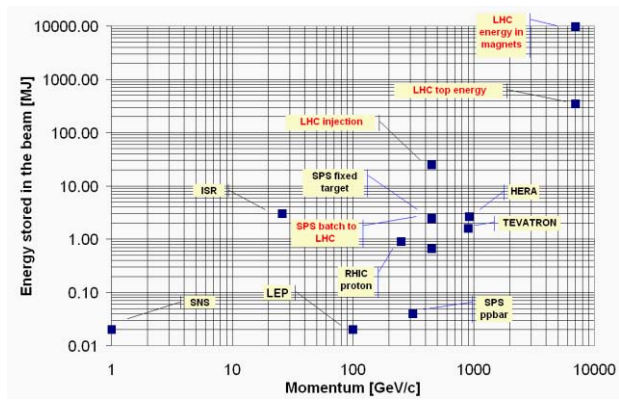

Figure 1: Stored beam energy as a function of the momentum for various accelerators.

ACCESS CONTROL

Securing the access to the accelerator control system is an essential part of the protection strategy. The access must be limited to avoid undesired changes to machine settings by outsiders, and to ensure that the accelerator operation crews remain in full control of the machine. The access control issue is most critical for controls system devices and parameters that are part of the MPS, for example interlock thresholds. While on one hand access to the accelerator must be strictly controlled, experts, but also software developers must be granted access to the accelerator devices for interventions and debugging.

A first level protection of the control system consists in separating the accelerator network from the general internet. Access to the accelerator controls network is then only

possible through a limited number of gateways with strict access control lists. Such a strategy is used for the CERN accelerators only since 2006, but it is also used at light sources, for example at the Swiss Light Source. At CERN access is presently also granted to a number of trusted hosts connected to the general CERN network. But in the near future access to the accelerator control system will only be possible through the special gateways.

Limiting access to critical devices and their settings to the experts through some form of login is the second level of control that does not necessarily require highly specialized controls solutions. Such schemes are in use at almost all facilities, even though some operate without particular protection. Those isolated cases involve however small(er) teams and are not part of the accelerators at the top of the power hit list.

Access control by location (of the computer) is another simple scheme, used for example at FNAL. Depending on the location of a computer, the user will have full access to all settings (main control room), or only limited access and for a limited duration.

Role Based Access Control

A sophisticated access control relying on a role-based system (RBAC) is put in place for the CERN accelerator network in collaboration with FNAL [5]. In this scheme access to a specific parameter is limited to persons holding a given role. A role is typically a job function, such as LHC-Operator, BLM-Expert, MPS-Expert etc. A user can have multiple roles at any time, and any user can be a member of any role. To access a device within the control system, a user has to login with his normal CERN userid/password and the RBAC system authenticates the user. If authentication is successful, the user can select one of his roles to access a set of accelerator devices. Each device property may be protected with access rules. These rules specify what roles can access settings. In addition to the roles, three other parameters can be specified in an access rule: the location (i.p. host address), the application, and the accelerator mode. This access scheme will be used by physicists and engineers of FNAL to remotely access certain LHC devices from FNAL: magnets that are part of the US contribution to LHC, beam instrumentation data, etc.

RBAC is presently in a progressive deployment phase within the LHC control system.

MPS CONFIGURATION

The correct functioning of the MPS can only be fullfilled when the configuration of the devices that are part of the MPS is correct, for example correct assignment of settings, activation of the appropriate device etc. In many places the configurations are stored in (or derived from) a 'master' machine database that holds the description of the accelerator, see for example Ref. [9]. An essential operation during commissioning, but also after periodic shutdowns, is the validation of the MPS configuration or of the changes to the configuration. The validation frequently involves 'manual' verification by one or more experts. Changes to the configurations are frequently decided by a laboratory or machine safety authority.

During commissioning and controlled machine experiments phases part of the MPS constraints must sometimes be relaxed. This may in particular be the case for the commissioning of the MPS with beam, for example to calibrate BLMs or setup collimators and protection devices. Relaxing the constraints may be performed by either masking certain interlock channels or by adapting MPS parameters like thresholds etc.

Machine and Beam Modes

A commonly used scheme to adapt the MPS to the running conditions is to define a number of machine and beam modes [3, 4, 8, 11]. The modes are adapted according to the state of beam dumps, target positions, gun configurations etc. A pre-configured interlock channel map is defined in advance for each mode, and the MPS automatically applies the maps and masks out a certain category of interlocks. In some cases, like for example SNS, the threshold of certain BLMs are automatically increased before a flying wire is send through the beam for an emittance measurement. At the CERN SPS the operator defines the mode for each of the beams. A software interlock system [11] verifies that the declared modes match the actual state of the accelerator.

Such mode-based schemes are very powerful, but sometimes do not provide sufficient flexibility. They are usually well adapted to cope with standard machine operation modes, but are not always sufficiently flexible for machine commissioning or setup phases.

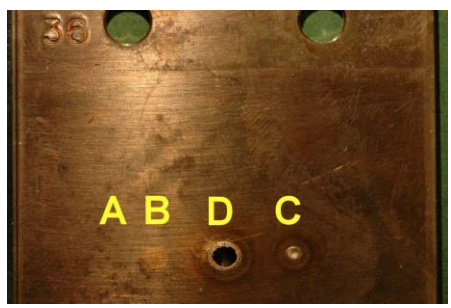

Figure 2: Impact of a 450 GeV/c beam from the CERN SPS onto a Copper plate inserted into a 20 cm long target [2]. The position of this plate within the target corresponds approximately to the peak energy density. Four beams of different intensities were send to the target: 1.2×10^{12} (A), 2.4×10^{12} (B), 4.8×10^{12} (C) and 7.2×10^{12} (D) protons.

Safe Beam Concept

The LHC MPS does not use a mode scheme but relies on the concept of a 'Safe' beam. Each hardware interlock

connected to the LHC MPS is preconfigured to be either maskable or unmaskable [1]. Unmaskable interlocks will always remain active, irrespective of the machine or beam state: they include vacuum valves, dump system state, critical power converter states. Maskable interlocks may be masked provided the beam is considered to be safe. As soon as the beam becomes unsafe, the interlock system will automatically re-active any masked input.

A beam is considered safe at LHC injection energy of 450 GeV/c if its intensity is below 10^{12} protons for the nominal emittance of 3.5 μm. The intensity of the safe beam is scaled with energy according to a $E^{-1.7}$ scaling law obtained from simulations. The safe beam limit depends on the material, and for the case of the LHC it is defined for Copper, which is common in many accelerator components. The simulation results have been benchmarked in an experiment with a 450 GeV/c beam in an SPS extraction line [2] as shown in Figure 2. The safe beam information is transmitted to the components of the SPS and LHC MPS in the form of a logical flag.

The LHC safe beam concept has been designed to provide a high level of safety. A possible drawback is the missing granularity of the concept with only two states, safe or not. In certain conditions a higher granularity might be desirable.

Management of Critical Settings

Handling the large number of MPS settings is very critical for the LHC. To provide a safe, but relatively simple and homogenous solution for all MPS parameters that are subject to be changed (trimmed), the concept of 'Critical Setting' was introduced [5]. A critical setting is any machine setting that is relevant for the safety of machine operation or for the MPS system: for example BLM thresholds, calibration tables of the beam dump kickers, fine delays of the injection kickers etc. There are two issues for such parameters: first, only a limited number of persons must be allowed to change the parameters and secondly it is important that the values are not corrupted, either when sending the value to the device, or due to data corruption or data loss at the level of the device, for example after a reboot.

The solution adopted for the CERN control system in view of the LHC is based on the private-public key encryption mechanism for which algorithms and implementations are readily available. Each critical parameter set is assigned a digital signature that is obtained by encoding a hash of the data with a private key. Both the parameter and its digital signature are stored in the controls database, and the parameter values are visible for anyone able to access the application that can visualize the settings. Both the parameter and its digital signature are send to the front-end computer (FEC). On reception of the critical parameter, the FEC verifies the integrity of the data by comparing the data with its digital signature using the public key. The parameter setting is only accepted if data and signature agree, else the setting is refused and an exception is thrown. This concept is illustrated in Figure 3. Storing the signature together with the data ensures that no changes may be performed deliberately or accidentally in the database. The integrity of the setting in the database, of the signature stored in the database and of the setting value stored in the FEC are verified periodically, typically once or twice per day before beam is injected into the LHC. In case of inconsistency injection of beam into the LHC is inhibited.

The critical settings concept will be fully integrated into the control system, and software developers do not have to make any additional effort to define a critical setting. Generation and verification of the digital signatures are transparent for the developper. Each critical setting is associated to a role of the CERN RBAC system (see previous chapters) and only users that are assigned to the appropriate role may obtain the keys to change and properly encrypt the data. The system is presently in a first deployment phase, but must be made operational for the startup of the LHC.

MPS VALIDATION AND VERIFICATION

An essential task during the commissioning and after accelerator maintenance periods is the validation of the MPS functionality. For large MPS systems such tasks cannot be performed 'by hand' by operators or experts, and automated test sequences must be employed. The advantage of automated sequences is that they can be run periodically. Careful specification, development and validation of the sequencing itself must be performed by the system experts before such a sequencing task may be used to validate MPS functions.

Automated schemes were foreseen at SNS [3] and are already employed at the SPS for the high intensity transfer lines due to the large number of interlock channels (over 500). Automated test sequences, based on a 'sequencer' tool used for LHC commissioning and operation, will be used to commission some large LHC MPS sub-systems, in particular for magnet powering. The sequencer has been employed very successfully in the commissioning of the protection systems for the LHC superconducting magnets [10]. The sequencer not only triggers the tests, it also collects data (including post-mortem data) and automatically enters test results into the controls database.

Periodic and continuous checks of the MPS are performed at the SPS and in the future at the LHC by a Software Interlock System (SIS) [11]. SIS surveys the accelerator and its transfer lines by software monitoring of a large number of parameters and states, including the configuration of the MPS itself (for example thresholds etc). In case a problem is detected SIS may inhibit injection, extraction or even dump the beam. SIS is also sending information across the machines to the injectors to ensure that the beam is stopped at the source. One of the roles of SIS is the integrity check of parts of the MPS system. It is also used to implement complex interlocks that require input from a number of systems, in particular when such systems are distributed over the machine. As an example

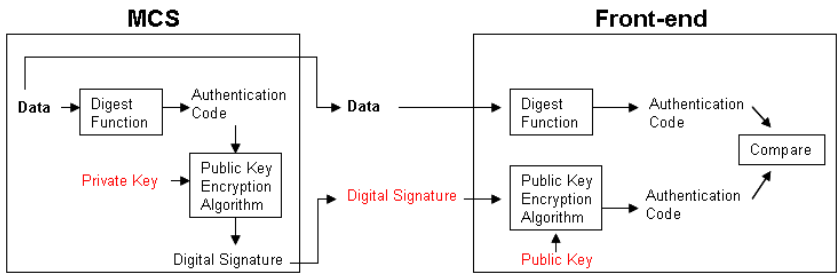

Figure 3: Schema of the CERN critical settings management system.

the SIS will ensure that the LHC horizontal orbit correctors do not change the beam energy by more than $\pm 0.2\%$, since this contribution to the beam energy is not taken into account by the LHC Beam Dumping System.

MPS DIAGNOSTICS

Whenever the MPS system is triggered it is important to trace the source of the problem as precisely as possible. This is required to ensure that the MPS reacted correctly (detection of near-misses or incorrect handling of faults) and to optimize the machine uptime by understanding frequently re-occurring faults.

Depending on the type of machine, linear or circular, the data collection is different. For a linac or a transfer line, data is normally collected for every beam passage. This provides both the data necessary to understand the 'steady-state' operation of the line and any abnormal event, including the events that lead to MPS triggers. Having the possibility to compare events with MPS triggers and with other beam passages is usefull to diagnose possible trends that led to the MPS events. Such continuous logging is for example performed at the SPS high energy transfer lines [12] and at FLASH [4].

Post-mortem Diagnostics

For circular accelerators where the beam may be stored for many hours, a continuous logging system is important to understand the performance, but not sufficient to diagnose the transients around a MPS event, for example a quench or a BLM trigger. The recording of the transient data around the MPS event is of utmost importance, and this data recording is usually designated as Post-mortem (PM) data. The time resolution that is required ranges from nanoseconds for feedbacks or RF systems to machine turns for beam instrumentation and to few milliseconds for slow power converter status information. A precise and fast trigger mechanisms must be provided, and circular buffers must be filled at the appropriate rate by the monitoring devices. The trigger distribution for the PM event must also be foreseen, since in many cases it is necessary to collect data from devices that did not directly trigger the MPS event. Post-mortem system for all or part of the diagnostics is nowadays available in all accelerators [3, 6, 7, 13].

At the LHC a post-mortem system has been foreseen from the start [13] and circular PM buffer have been integrated into all essential accelerator systems. The trigger for the PM buffers is distributed by the LHC timing system, except for the devices that are self-triggering, like for example quench protection systems. Very precise timestamping must be performed for all the data to ensure that the proper time sequence may be reconstructed. The timestamping information is distributed by the machine timing system. Under certain conditions it is possible to suppress the collection of post-mortem data for injection studies and when only one out of the two beams is dumped for a machine experiment.

Table 1 lists data volumes for post-mortem events at a number of accelerators. For the LHC the data volumes are in the multi-GB range, which poses problems of data collection and requires automated data analysis. Such automated analysis is already provided for the commissioning of the LHC powering and quench protection systems [10]. For the LHC startup only a limited number of simple analysis are foreseen for beam data. It is expected that the analysis and display requirements will become clearer once the LHC will be running with beam.

Table 1: Post-mortem data volumes for diagnostics at different facilities.

Facility	Data volume
HERA	5 MB
RHIC	130 MB
LHC	2-4 GB
FLASH (5 Hz)	1 MB / shot
CNGS (0.1 Hz)	0.1 MB / shot

Post Operational Checks

At the LHC the beam dumping and interlock systems play a vital role for the protection of the machine since they are unique and have no operational spare that can replace their function in case of problems. To ensure that the systems are 'as good as new' after each beam dump (MPS event), both systems will perform automated Post Operational Checks (POC) based on their post-mortem data.

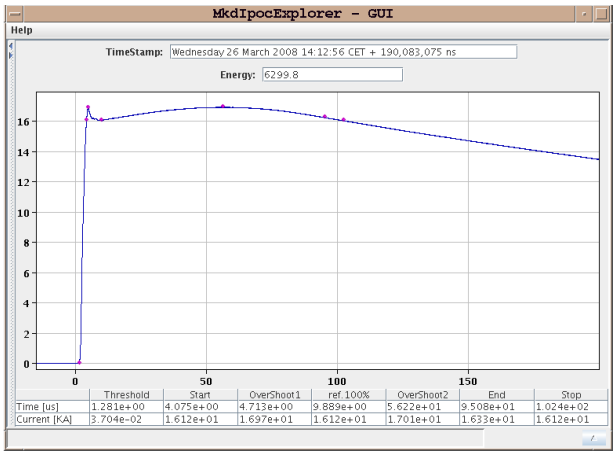

Figure 4: Recording of the field in an LHC beam dump kicker. This trace is analyzed as part of the LHC Beam Dumping System IPOC to validate the correct functioning of each magnet. The markers (magenta) are the reference points used for the analysis. The horizontal scale in in μs.

Both systems are highly redundant, but to ensure the required safety levels it is essential that the redundancy is not lost. For that reason the post-mortem data is automatically analyzed to detect any malfunctioning or loss of redundancy of the systems. For the case of the LHC Beam Dumping System, the POCs included internal data from the kicker and triggering system (IPOC, Internal POC) and external data from beam instrumentation (XPOC, eXternal POC). The beam dumping system will go to a fault state in case the IPOC or XPOC fail, requiring an intervention by an expert to analyse the situation. Figure 4 shows the data for a kicker magnet deflection.

CONCLUSION

Protection of the present day and future high intensity accelerators is a challenging task that not only involves fast hardware interlock components but also many control system aspects. Access control is becoming a critical issue and configuration changes of the MPS system must be well controlled and tracked to ensure that no information is lost on configuration changes.

Flexibility is provided in many places with the concept of machine and beam mode that allows automatic configuration changes of the MPS system configuration for different running modes. Settings control based on digital signatures using the public-private key principle are now implemented at CERN for the SPS and the LHC. The digital signature of the settings is used to prevent and uncontrolled change or data corruption.

The large number of interlocks of modern accelerators require automated testing of the main functions of the MPS system. At the LHC sequencer program is executing carefully designed task sequences to validate the protection functionality for large systems.

Diagnostics of MPS triggered event is provided everywhere by Post-mortem systems or shot-by-shot logging and analysis. For the case of the LHC the most critical components of the MPS will perform self-checks based on their post-mortem data to ensure that no malfunctioning may be undetected for some time.

ACKNOWLEDGMENTS

The author would like to thank A. Drees (BNL), M. Boege (PSI), K. White (SNS), E. McCrory (FNAL), M. Staack and M. Lomperski (DESY) for help and information.

REFERENCES

[1] R. Schmidt et al., *Protection of the CERN Large Hadron Collider*, New Journal of Physics 8 (2006) 290.
R. Schmidt et al., *LHC MACHINE PROTECTION*, Proc. of PAC07, Albuquerque, NM.

[2] V. Kain et al., *Material Damage Test with 450 GeV LHC-Type Beam*, Proc. of PAC05, Knoxville, Tn.

[3] C. Sibley et al., *The SNS Run Permit System*, Proc. of ICALEPCS01, San Jose, Ca.
C. Sibley et al., *Machine Protection Strategies for High Power Accelerators*, Proc. of PAC03.
C. Sibley et al., *The SNS Machine Protection System: Early Commissioning Results and Future Plans*, Proc. of PAC05, Knoxville, Tn.

[4] L. Froehlich et al., *First Experience with the Machine Protection Syste, of FLASH*, Proc. of FEL2006, Berlin, D.

[5] S. Gysin et al., *Role-Based Access Control for the Accelerator Control System at CERN*, Proc. of ICALEPCS07, Knoxville, Tn.

[6] M. Heiniger et al., *Post-mortem Analysis of Failures in Accelerator Systems*, Proc. of ICALEPCS03, Gyeongju, Korea.
G. Dzieglewski et al., *Protection Mechanisms for High Power Accelerators*, Proc. of ICALEPCS05, Geneva, CH.

[7] J. Morris et al., *Status Report of the RHIC Control System*, Proc. of ICALEPCS01, San Jose, Ca.
J.S. Laster et al., *Post-mortem System - Playback of the RHIC Collider*, Proc. of ICALEPCS01, San Jose, Ca.

[8] K. Jordan et al., *Machine Protection for High Average Current Linacs*, Proc. of PAC03.

[9] M. Zerlauth et al., *From the LHC Reference Database to the Powering Interlock System*, Proc. of ICALEPCS03, Gyeongju, Korea.

[10] R. Saban, *LHC Hardware Commissioning*, these proceedings.

[11] J. Wozniak et al., *Software Interlock System*, Proc. of ICALEPCS07, Knoxville, Tn.

[12] J. Wenninger, *CNGS Extraction and Transfer Stability in 2007*, CERN-AB-2008-002.

[13] J. Wenninger et al, *LHC Post-mortem System*, LHC Project Note 303 (2001).