INTERLOCKING STRATEGY VERSUS MACHINE AVAILABILITY

L. Ponce, J. Wenninger, J. Wozniak, B. Todd and K. Fuchsberger, CERN, Geneva, Switzerland

Abstract

In addition of the hardware interlock system (BIS), the Software Interlock System (SIS) is providing a framework to program high level interlocks based on the surveillance of a large number of accelerator device parameters. Since its deployment in 2008, the LHC SIS has demonstrated that it is a reliable solution for complex interlocks involving multiple or distributed systems and when quick solutions for unexpected situations are needed. This paper presents the operational experience with software interlocking in the LHC machine and reports on the overall performance and flexibility of the SIS in the context of the overall interlocking strategy and availability of the machine

INTRODUCTION

The core of the LHC interlock system is the Beam Interlocks System (BIS) that is entirely implemented in hardware and designed to inhibit injection (Injection BIS) or dump the beams (LHC BIS) with extremely high reliability and availability requirements. As a complement of the BIS, the Software Interlock System (SIS) provides further protection by surveying and analyzing the state of various key equipment. Its open architecture allows for fast and easy configuration of more complex logic, which allows to anticipate failure rather than reacting to them. It is in particular possible to define complex interlocks that correlate the state of many different systems and that are difficult to implement as hardware interlocks. The system has been designed to be as highly reliable as possible for software, providing the flexibility for an easy reconfiguration of the logic to respond to the changing needs of the LHC operation. At the end of the LHC run period, it is worth to weight the interlocking strategy against the availability of the machine and to evaluate the need to move some interlocks from software into hardware or vice versa.

SIS FUNCTIONALITY AND PERFORMANCE

SIS structure

The central concept of SIS consists of boolean expressions represented as trees. The fundamental level is an Individual Software Interlock Channel (ISIC) associated to a reading of a state, value or property of a system based on the JAPC-Monitoring framework [1]. The acquired parameter is analyzed (tested) resulting in a logical state TRUE or FALSE. The logical states are then grouped into a tree-like structure and combined using logical operators (AND, OR, NOR,...) into intermediate nodes called Logical Software Interlock Channel (LSIC). The top of a tree corresponds to a so-called Software Permit, which itself can be TRUE or FALSE and is exported to Beam Interlock Controller devices:

- INJECTION (BEAM1, BEAM2 or BOTH BEAMS) PERMITS exported to inhibit the extraction(s) from SPS
- RING (RING1, RING2 or BOTH RING) PERMITS exported to the BIS to dump the beam(s)
- POWERING PERMITS (1 per octant) exported to the Powering Interlock Controller (PIC) to abort powering

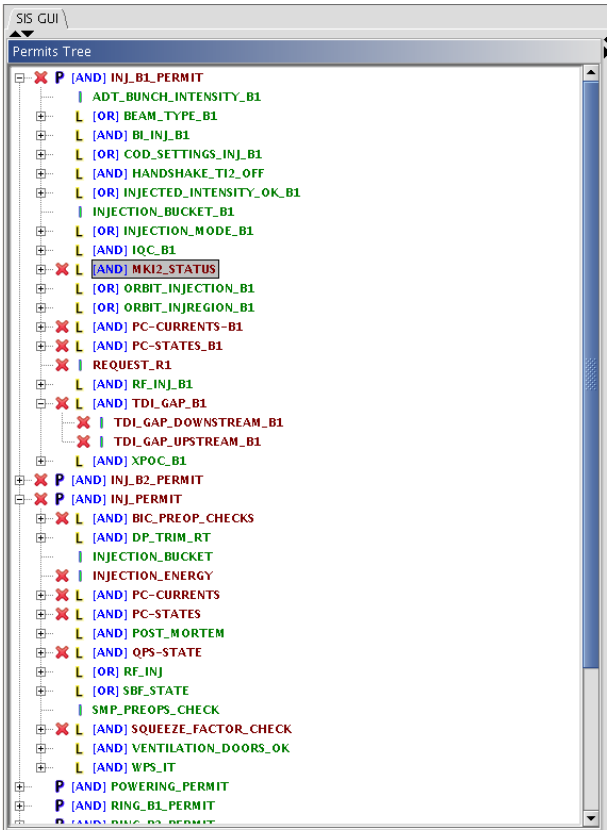

Figure 1: SIS GUI: the top permit like INJ_B1_PERMIT, INJ_B2_PERMIT are marked with a P; the tree can be expanded to see the channels.

The initial configuration of LHC SIS, using mainly AND and OR hard-coded logic, has been extended to allow more and more complicated Interlocks written in JAVA extension pulling together multiple signals and database references. All vital components are defined as Spring Beans [2] in an

XML file and managed in a bean container by this framework. Solutions from Spring have also been applied to provide remote access to the SIS core service via a Java Message Service (JMS) channel.

The SIS has a layered architecture which reflects the two major tasks of the system: Data Acquisition and Data Processing. The first layer deals with data subscriptions, providing values used later for the tree calculation. The second layer holds the definition of the trees and it is activated upon the tree calculation event, taking already prepared values from the internal buffer. One main architectural goal was to make the analysis as reliable as possible and thus as independent as possible of the data acquisition.

As the SIS is a server application, a (Swing) graphical user interface (GUI) was developed to show the system state to the operators in the control room. All permit trees are visible and dynamically updated; channel states are expressed with colors and markers (Fig. 1). Basic functionality for channel status analysis have been implemented in the GUI and will be extended for next run in order to improve the interaction with operators. The GUI also provides an interface for some user actions and a sophisticated analysis mechanism capable of identifying several typical fault scenarios, such as missing data or data with incorrect values.

To perform its job during last run, the LHC SIS was handling 2665 device/parameter subscriptions representing some 5500 checks grouped into 7 permits. All interlocks trees are evaluated every 2 second for the LHC (the typical update rate of the parameters being in the range of 1 to 10s) but can be faster if needed (in the order of 10-100 ms for the injectors SIS).

SIS AVAILABILITY

The LHC SIS core runs on dedicate HP server equipped with a timing receiver card (CTRI). Since the beginning of the operation in 2008 for the SPS and LHC SIS instances, only few crashes of the SPS server were observed during the 2009-2010 shutdown period. The problem was traced back to a concurrency problem in the timing library and was quickly fixed.

Table 1: Interlocks channels leading to dump.

SIS DUMP cause	**Ratio**
Communication problem	20%
Orbit feedback issues	20%
Power converter faults	15%
Beam position measurements	10%
Beam Loss monitors HV	10%
Others (wrong settings, masks)	25%

For the LHC machine, any time the beam is aborted a Post-Mortem file is produced tracing the root cause (first trigger) of the beam dump. Extracting the data of the Post-Mortem database for the 2012 operation period, 77 dumps are flagged with the LHC SIS as first input to the BIS. All events are real interlocking conditions (see Table 1), i.e. one of the ring permit changed from TRUE to FALSE status: none of the dumps are due to SIS failures, the programmed logic was always followed. In the "communication problem" case, data were not received by SIS and the logic of some of the interlocks (power converters, orbit, ...) is programmed to dump the beam in case there is no update for a time defined by programmer.

Ring permit

The Ring permit value is exported to the LHC BIS to trigger a beam dump (either for a given beam or for both beams) in case the evaluation result is FALSE. Taking into account the lengthy injection process and beam cycle in the LHC, dumping a circulating beam is really costly in terms of efficiency. Thus the reliability of ring permit should be as high as for the BIS system.

The initial configuration from 2010 (Table 2) has been

Table 2: Initial configuration of the ring interlocks.

TEST	**Coverage**
SMP energy	All RBs, SMP energy
SMP energy distribution	ALL BLM crates
BETS	Q4 and MSDs in IR6
TCDQ-Beam	Beam center in TCSG TCSG gap TCDQ-TCSG retraction
COD integral	All arc Hor CODs
Orbit	All ring BPMs
COD settings	All CODs in STABLE BEAMS
COD trips	60A CODS (not in PIC)

Table 3: Configuration of the ring interlocks at the end of 2012 run.

TEST	**Coverage**
RF voltage	Energy $>$ 3.4 TeV
BLM High Voltage	All BLM crates
Feedback Mask	during RAMP & SQUEEZE if $>$ 20 % BPM disabled
Ref Orbit	RAMP & SQUEEZE if zeroed/wrong orbit reference
COD integral	All arc Hor CODs
PC interlock	All 60A CODs

extended with several interlocks to fill the potential holes in machine protection that were discovered all along the operation period (Table 3).

Injection Permits

The LHC injection Permits are connected to the SPS Beam Interlock System in order to prevent injection into the LHC. In case one system is not ready or not in a nominal configuration for injection, an inhibit is sent to the the injector complex to prevent extraction or even production of the beams in the the injectors allowing a more efficient operation of the complex. As loosing one injection in the machine is not so critical for the operational efficiency, the interlock policy can be very strict and with a large number of checks.

Started mainly with checks of statuses and values in range for different equipment (Power Converters, Quench Protection, RF...) in 2010, as shown in Table 4, the injection interlocks have been extended to include more operational settings checks using the possibility to combine parameters published by different systems (see Table 5).

Table 4: Initial 2010 configuration of the injection interlocks.

TEST	**Coverage**
PC state	All PCs
PC current	ALL BLM crates RB, RQ, RD, MCBX
QPS_OK	All circuits with QPS
RF	Synchronization Cryo maintain
BTV position	Ring and dump lines BTVs
Orbit	All ring BPMs
Injection bucket	Abort gap and over-injection protection
Injection mode	
Energy	
(Pre)-ops Checks	XPOC, PM, IQC, BIC, SMP
triplet alignment	WPS in all IPs

Most of these additional interlocks were implemented to improve the machine protection level following some initially un-foreseen operational conditions. One example is related to the very large range of bunch intensities that has been used in the LHC during the 2012 operation: the configuration of some components, like the transverse feedback system, depends on the peak bunch intensity and the equipment may be damaged if the injected beam intensity is higher than the pre-configured one. No hardware interlock is available on the extracted bunch intensity but this information is available in time before extraction from the injector, so a software interlock comparing the actual feedback settings and the intensity just before extraction in the SPS has been added as part of the injection permit.

Table 5: Configuration of the injection interlocks at the end of the run. The lines marked with a $*$ covers holes in the machine protection of the initial configuration.

TEST	**Coverage**
ADT bunch intensity *	SPS intensity vs ADT settings
Beam type *	Check ions/proton beams
TL handshakes	IP 8 and 2 allow extraction on the TED
Injected Intensity *	SPS intensity versus circulating beam
Injection orbit *	All BPMs
Orbit in inj. region *	BPM around inj. IPs
TDI gaps *	
RF RT trims *	Radial modulation OFF
MKI vacuum	Magnets and interconnect
MKI temperature	MKI magnets
Ventilation doors	Non LASS interlocked doors

Interlock masking

Masking is a mechanism that allows operators to ignore an individual ISIC or LSIC. Masking a channel means overriding its real state and evaluating it always to TRUE. The ability to mask a given ISIC/LSIC is defined for each channel individually and the Permit signals are not allowed to be masked. The masking itself is done from the SIS GUI by operators. The role based access control framework is used to define the right to mask, however, the masking rights apply only to channels defined as maskable. Two roles are used: LHC-EIC (used by LHC Engineer in Charge) and MCS-SIS (SIS developers involved in the machine protection). After LS1, we should consider the possibility to make the masking more role-dependent and create different roles for different group of interlocks.

When applied, a mask is always active, independent of beam conditions or Set-up Beam flag. Therefore another way of "masking" interlocks automatically for a given period within the beam cycle or a given energy/intensity range is largely used in the SIS via the OR logic. A beam intensity, beam mode (describing the time in the beam cycle) or beam energy test is added as a ISIC with a OR logic to the channel that should be masked, see for example Fig. 2. As soon as the intensity, energy or mode condition is TRUE, the interlock is *de facto* masked.

After the initial commissioning period, a long list of interlocks have been made UNMASKABLE: Orbits in physics, XPOC, machine protection Post-Mortem permit,

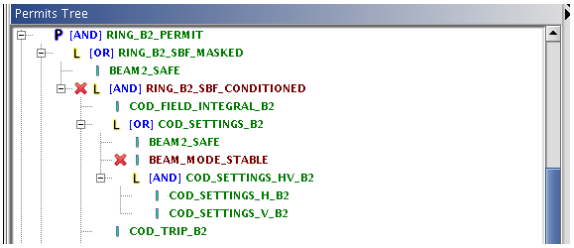

Figure 2: Example of masking using the beam intensity: the 60A power converter settings LSIC is combined in an OR logic with the BEAM_SAFE ISIC (intensity) and the BEAM_MODE_STABLE ISIC (collision period in the beam cycle).

IQC injection oscillation permit, etc. Even more will come for the 7 TeV operation.

SIS IMPROVEMENTS

CMW communication

During 2012 operation, several dumps were caused by a stop of the data streams. The most affected data source was the Power Converters Function Generator Controller publication, which was not received by the SIS data acquisition task for several minutes on some occasions. In such a case, the tree is evaluated to false if the data are not updated before a pre-defined time-out to avoid being blind with beams in the machine. SIS time-out was increased from 20 s (in 2010) to 120 s (end of 2012). The programmed logic was correctly followed. The problem was traced back to a problem within CMW, which was not protected against "slow clients".

There was a clear degradation during 2012 operation but the problem should be fixed with the upgrade of CMW planned during the Long Shut-down.

GUI and Post-Mortem

In order to ease the diagnostic after the SIS has triggered a beam dump, some improvements are also planned/needed for the GUI and the link with the Post-Mortem server. Indeed, the tree structure display could be more user friendly to ease the understanding of the more and more complex structure, especially with the use of the OR logic to mask some interlocks. It must also be made possible to monitor any kind of parameter, including parameters that are complex combinations of various other parameters that are hidden in the construction of the logic (JAVA class). Furthermore, some extra protection is needed on the subscription management panel to avoid accidental stopping of critical data subscription.

Even though a local Post-Mortem file is produced on the SIS server for every SIS trigger of the beam dump, the data mining is quite painful and *de facto* reserved to trained people. Only the channel which triggered is present and there is no details in case of complex JAVA coded interlocks, like orbit interlocking. A straightforward improvement is to export the Post-Mortem file to the PM server, where a dedicated SIS analysis module could extract details on the triggered test.

*Beta**

The SIS also provides a Beta* publication to the Safe Beam Parameters system,which is used by collimators for their interlocking logic. The SIS uses the quadrupoles current in the IPs to derive the actual Beta* at each IP, publishing and reading back the value from the timing to cross-check with a reference table. For the time being, the calibration curves are hard-coded in SIS configuration files, one file per IP. It worked very well with the 2012 optics (nominal physics optics, HighBeta) because for this optics all quadrupoles have monotonic current functions during the squeeze, but it does not work for the Achromatic Telescopic Squeezing (ATS) optics, which is using quadrupoles of adjacent IP.

The proposal to improve the handling of the data is to migrate the calibration curves and the list of used IP quadrupoles in LSA to allow flexibility for different squeeze by using the hypercycle and beam processes concepts.

Orbit interlocking and PC interlock

Another important ring interlock is the Orbit and Correctors Orbit Dipole (CODs) interlocking. The principle is to limit the global orbit excursions of the beams to prevent beam losses and catch undetected orbit bumps. It uses distributed systems to compare the settings of each COD and the reading of each Beam Position Monitor (BPM) included enable flag, with a reference and a tolerance stored in LSA. A beam dump is triggered when 10 BPM or 2 kicks per beam or plane are out of tolerance. The tolerances are defined as a trade-off between machine protection and availability and have been set so far quite strict in STABLE BEAM ($\pm$ 2.5 mm in IR 1, 2, 5 and 8, $\pm$ 2 mm elsewhere), but are more relaxed during the ramp and squeeze process ($\pm$ 6 mm in IR 1, 2, 5 and 8, $\pm$ 3-3.5 mm elsewhere). The SIS configuration allows to condition the reference with a beam mode or an energy through the AND/OR logic and also to read the reference from the database settings with a predefined periodicity.

The interlocking strategy worked very well for standard operation but several problems occurred during special fills like Van Der Meer scans or injection optics collision. Indeed, it was needed to open the tolerances via a trim in LSA (RBAC protected to MCS-SIS role), which is very flexible but possibly also too flexible for critical parameters.

The proposed improvement for post-LS1 operation is to remove the CODs settings check which is now redundant with the PC interlock, presented at this workshop [3]. The PC interlock is designed to check that the CODs current is within tolerance for each beam process and triggers a beam dump when 2 kicks per beam or plane are out of tolerance. The settings are stored in a reference beam process in LSA which is cloned from the Power Converter beam process.

To allow following the complex change of current occurring during the ramp and squeeze, the change of tolerances function is triggered by timing events.

CONCLUSION

LHC SIS is successfully used in operation since 2008. It is a reliable solution for different class of interlocks: injection interlocks (when high reliability is less critical), complex interlocks involving multiple systems or distributed systems (like orbit) or as a fast answer to un-expected situation like feedback problems. It is all software, soft real-time, but the reliability, even if it will never be SIL3, is remarkably high. During the long shut-down period the following software interlocks will be moved to hardware interlock: TCDQ interlocking, TDI gap interlock and CODs setting. However, many more new software tests will be introduced in the SIS.

REFERENCES

[1] J. Wozniak et al., "Software Interlock System", ICALEPCS'07, Knoxville, October 2007, WPPB03, p. 403.

[2] http://www.springsource.org

[3] K. Fuchsberger, "Software tools for MPS", these proceedings