THE ARCHITECTURE, DESIGN AND REALISATION OF THE LHC BEAM INTERLOCK SYSTEM

B. Todd[1], A. Dinius[1], P. Nouchi[1], B. Puccio[1], R. Schmidt[1]
[1]*CERN, Geneva, Switzerland*

INTRODUCTION

The Large Hadron Collider at CERN (LHC) is designed to operate at 7 TeV/c with a luminosity of 10^{34} $cm^{-2}s^{-1}$. This requires two beams with approximately 3 10^{14} protons per beam, giving a stored energy of around 360 MJ, an unprecedented value for accelerators, enough to heat and melt around 500 kg of copper.

An uncontrolled release of the energy stored in LHC beams could lead to serious equipment damage. Major damage of superconducting magnets will result in long repair times as the equipment is delicate and difficult to access. It is therefore vital that the LHC is protected against damage due to uncontrolled beam losses [1, 2].

Protecting equipment against damage in the case of uncontrolled losses is challenging. A nominal injection into the LHC from the CERN Super Proton Synchrotron (SPS) at 450 GeV is already above damage thresholds [3]. The only equipment specifically designed to withstand the full energy of a nominal LHC beam is the LHC Beam Dumping System (LBDS). A Beam Interlock System has been designed for LHC, requesting rapid beam extraction of beam from the LHC by the LBDS in the case of failures or high beam losses being detected by “User Systems” connected to the Beam Interlock System. A dedicated Injection Interlock System ensures that beam from the SPS is only extracted, transferred and injected into the LHC when the conditions are correct. The LHC Beam Interlock System has to inform the Injection Interlock that beam cannot be extracted from SPS when LHC is not ready.

This paper discusses the overall architecture, design and realisation of the LHC Beam Interlock System, with particular emphasis on communications and control aspects of the system needed to realise the strict speed and safety requirements.

REQUIREMENTS OF THE BEAM INTERLOCK SYSTEM

There are two fundamental requirements in the specification of the Beam Interlock System which dictate architecture and design, namely *speed* and *dependability*.

Speed

The time after a failure until particle losses become unacceptable has been evaluated for many types of failures; this time depends primarily on LHC energy, optics and collimator settings. So called *Ultra Fast Losses* are a caused by single-turn failures (<100 µs) that can occur during injection and extraction of beam. Protection relies on collimators and beam dilutors being positioned correctly. The Beam Interlock System ensures no beam can be injected if the collimators and dilutors are not ready. *Very Fast Losses* (<5 ms) are failures causing the beam to become dangerous within some ten turns. A failure of a D1 normal conducting magnet has been identified as the fastest mechanism for these multi-turn failures. This kind of failure is detected by the Beam Loss Monitor System, and the Fast Magnet Current-Change Monitor [4].

To protect the LHC against such losses, the Beam Interlock System must act as a very fast communications system, transmitting information from User Systems connected around the full ~27 km circumference of the LHC to the LBDS within only ~100 µs [2].

Dependability

Coupled with the strenuous requirements of response-time is *safety*. To adequately protect the LHC from beam related failures, the Beam Interlock System and the Beam Dumping System are specified [5] as having Safety Integrity Level – 3 (SIL-3) meaning the Mean Time Between Unsafe Failure is between 1000 and 10000 years [6]. For the Beam Interlock System to meet this requirement, redundant signal paths have been used for mission critical signals, making the system both safe and allowing a full system test to be carried out on demand.

ARCHITECTURE

There are around 140 User Systems of the LHC Beam Interlock System. To accommodate these distributed User Systems, the LHC Beam Interlock System has 16 Beam Interlock Controllers (BIC) one situated to the right and to the left of each Insertion Region (IR). These 16 BICs communicate with the LBDS situated in IR6, and the permissions for beams to be present in the LHC, called Beam Permits, are carried around the machine between these BICs by dedicated Beam Permit Loops.

Beam Permit Loops

Figure 1 below shows the 16 BICs of the LHC connected to the LBDS at IR6. The communications from one point to another is carried out over four dedicated fibre optic channels. A clockwise and anti-clockwise link exists for each beam to be interlocked by the Beam Interlock System. This means that the request for a Beam Dump always takes the shortest path from one BIC to LBDS.

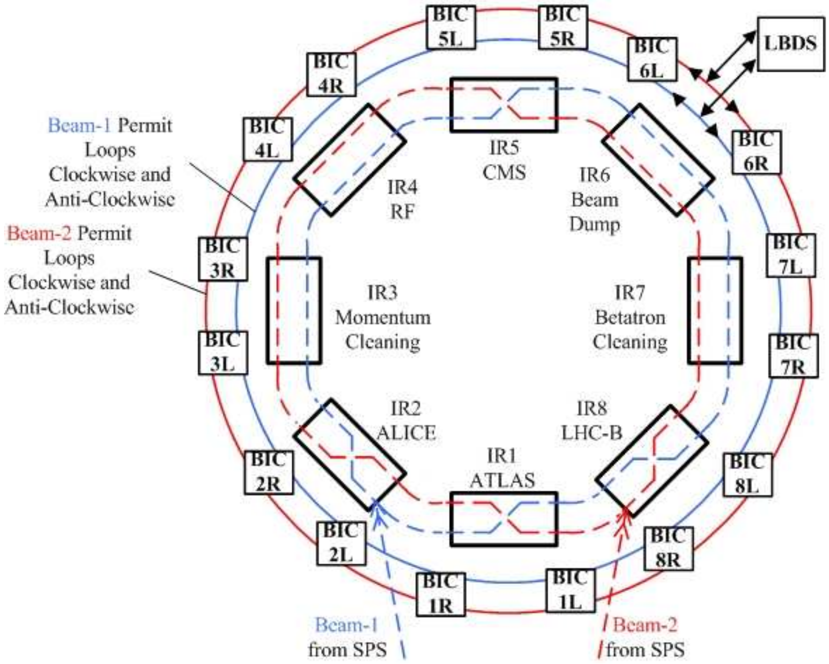

Figure 1: Architecture of LHC Beam Interlock System, showing Beam Interlock Controllers and Permit Loop Positions for all LHC Insertions

On each of the Beam Permit Loops, a 10 MHz square wave signal is generated at IR6 beside the LBDS, this signal travels around the Beam Permit Loop and through each BIC. Each BIC can monitor the Beam Permit Loops and open the Loops to request a Beam Dump. A correct frequency being detected at the end by the LBDS represents a *true* Beam Permit given for all User Systems connected to that loop. If the LBDS detects a change of Beam Permit from *true* to *false* during the course of an LHC mission, then the corresponding beam is removed from the LHC as soon as possible. The Beam Permit Loops representing the two beams operate completely independently.

Beam Interlock Controllers

The BICs simply collect and route User Permit signals from User Systems situated in their vicinity to the Beam Permit Loops, as shown in Figure 2. Connected User Systems can interlock the two LHC beams either independently, or simultaneously, and the BIC combines these to act on the relevant Beam Permit Loops. To reach the high level of safety the links are duplicated (labelled A and B), with the A link serving the Anti-Clockwise Beam Permit Loop, and the B link serving the Clockwise. The relevant User Permit signals are ANDed together to allow or deny the propagation of the 10 MHz square wave. The BIC also returns Beam Permit Info back to the User Systems (not shown), this is a non-critical signal representing the status of the Beam Permits.

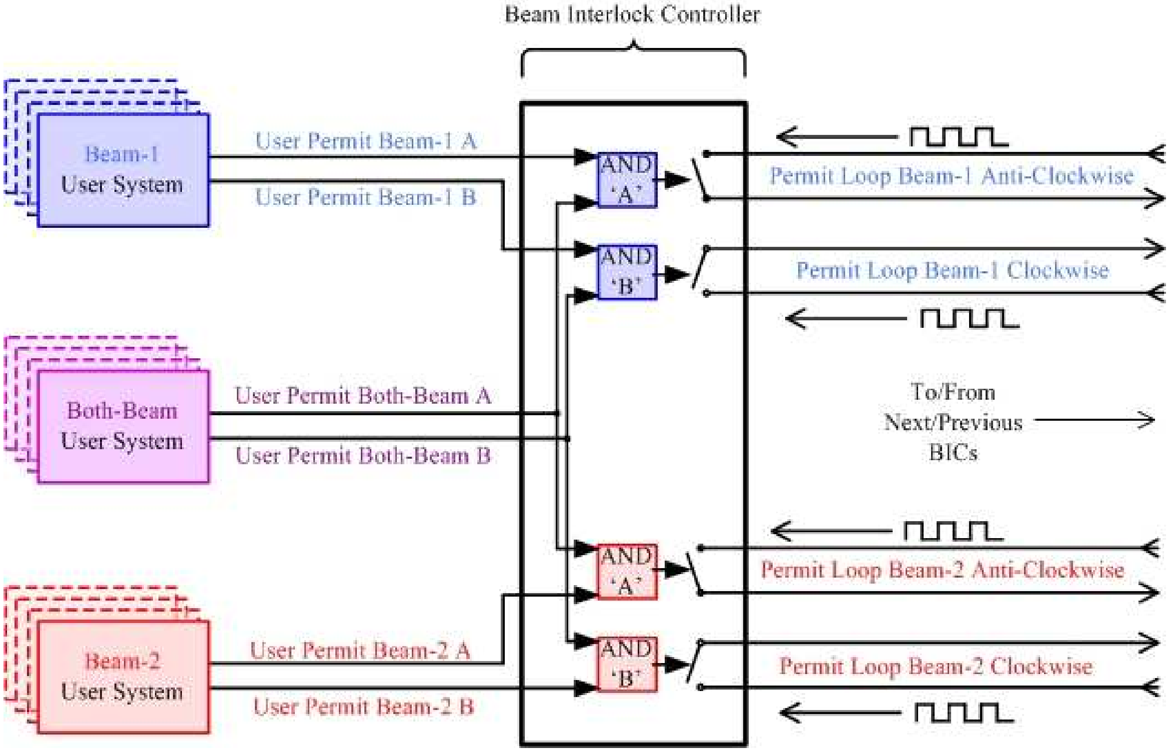

Figure 2: Beam Interlock Controller Critical Functionality

One Beam Interlock Controller consists of a Versa Module Europa (VME) chassis with numerous custom-made cards. VME was selected as the hardware platform as it is one if the standards supported at CERN, allowing a simple integration into the control network, and operating with sufficient speed to meet the design requirements. The key elements of the VME Chassis are:

A *Redundant Power Supply Unit* (PSU) vastly increases the reliability, as Power Supplies are amongst the most unreliable components of electronic equipment [7, 8].

A *Manager* module performs the ANDing of User Permits to make the Beam Permits; it also contains a robust monitoring device along with fibre optic receivers and transmitters for the Beam Permit Loops. The *Manager* records a history of changes, with micro-second precision, this can be read out remotely at any time. After a Beam Dump the Beam Interlock System provides the initial Post-Mortem information. In a single controller there are two *Manager* modules, each controlling the Permit Loops for one of the LHC beams. The *Manager* can also have a mask setting written to it allowing some User Permit signals to be ignored when LHC beam is less than predetermined damage levels. Half of the inputs to the BIC are not maskable.

A *Test* module: this provides the additional functionality needed to perform online testing and remote monitoring of the Beam Interlock System connections to User Systems. There are two *Test* cards in each BIC, one coupled to each *Manager*.

User Interface

User Systems operate with a variety of voltages and hardware platforms, so a generic User Interface has been designed, capable of interfacing any signal from TTL to PLC voltage levels. This signal is then transformed for long distance communication to the nearest Beam Interlock Controller. The longest distance that a User System is situated from a BIC is around 1 km. Each User Interface has dedicated full-duplex communications to enable online testing and monitoring of links on request. The User Interface is a small rack mounted module, having redundant Power Supplies that are remotely monitored. The User Interface is also equipped with a Complex Programmable Logic Device (CPLD) which controls the testing and monitoring channels, this is programmed to ensure a safe transition into test-mode, and allows in-field upgrades to be rapidly prototyped and implemented.

COMMUNICATIONS

The Beam Interlock System has been designed as a very robust and deterministic distributed control system, with full remote monitoring and 100% testing of critical signals. The fundamental technology of the BIC interconnections is the standard TIA/EA-485-A, originally known as RS-485. This allows bi-directional links to be created of up to ~1200 m, with a maximum length link having a maximum data rate of around 100 kbaud [9]. The signals exchanged between the User Interface and the Beam Interlock Controller fall into two categories, *critical* and *non-critical*, as shown in Figure 3.

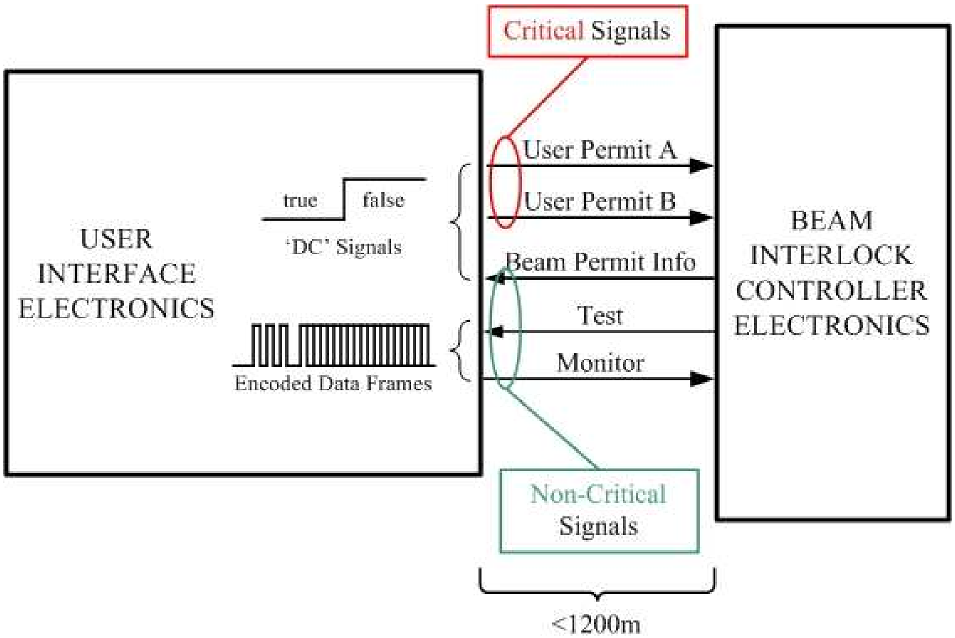

Figure 3: Critical and Non-Critical Signals between BIC and User Interface

Critical Signals

User Permit A and User Permit B are critical to the operation of the LHC, so transmission of this information from the User Systems to the BIC has to be extremely safe. These redundant signals are transmitted through fail-safe circuits ensuring that a single component failure cannot put the interlock system into an unsafe state. RS-485 is an excellent basis for robust links in terms of Electromagnetic Compatibility (EMC), as it can withstand large common mode electrical noise before data corruption is observed.

The User Permit link is made twice, for the redundant A and B links, through completely separate channels. The information communicated consists of only one bit for each link, representing *true* or *false*. A DC level is used to represent the value of User Permit, using Fail-Safe electronics means that *false* can be guaranteed to be read for almost all failures on these individual links. Coupling this with the redundancy means that an already unlikely event has to occur twice before the system fails unsafe, hence creating a very safe system.

Non Critical Signals

The non-critical signals are Beam Permit Info, Test and Monitor. Beam Permit Info operates in a similar way to the User Permit signals: a DC value is transmitted, exploiting fail safe links to make a dedicated communications channel. As this link is not critical, no redundancy is employed, meaning any hardware failure causing the link to fail unsafe will result in incorrect information being read on the output of the User Interface Electronics. This hardware failure will only be detected during a full system test, including the User Systems, which is expected to be performed at the beginning of each machine run.

Test and Monitor form a full-duplex RS-485 communications link from User Interface to BIC. When test commands are received, the User Interface switches the physical input circuits of the User

Interface to its own internal circuits, where it can set a *true* value on either User Permit A or B whilst simultaneously asserting *false* on the other redundant link. This has three key functions:

1. It tests the full link between User Interface and Beam Interlock Controller, including all active and passive components in the User Interface. Only the interconnecting cable between User System and User Interface remains to be tested, which is subject to a higher level test.
2. It forces both User Permit A and User Permit B to be read as *false* for around one millisecond when the links are being switched to the User Interface internal circuits, this means that beam cannot possibly be in LHC when the User Interface is tested.
3. The Beam Interlock System cannot in any way assert User Permit A and B simultaneously to *true*, so it can be tested safely.

The assertion of a Test on the User Interface, while not being dangerous, can lead to unnecessary beam dumps if it is not controlled. For this reason protection is built into the test mechanism to ensure it cannot be activated by accident.

1. If Beam Permit Info is true, the test command will be ignored by the User Interface.
2. The test command must be sent twice to the User Interface meaning a single badly received instruction with a badly formed parity bit is not sufficient to force a test mode.
3. These independent test instructions must arrive within 8 ms of each other, vastly reducing the chances of two spurious frames being intercepted correctly, and in sequence.

The integrity of the Test and Monitor communications channels is further enhanced by using Manchester encoded signals for the basis of communication; this reduces the bandwidth of the channel to a maximum of half the raw bandwidth, but increases the overall reliability of the link.

The Test commands and Monitoring Data are robustly encoded into frames, with a typical frame transmission as shown in figure 4:

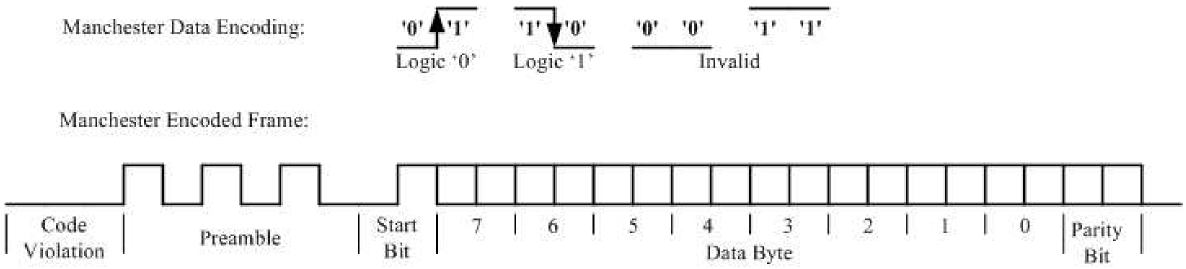

Figure 4: Manchester Encoded Data Frame for BIC to User Interface Communications

A small Complex Programmable Logic Device (CPLD) receives the encoded frames from the RS-485 link, and controls the Test assertion, it also replies with Manchester encoded information regarding the internal status of the User Interface. CPLDs are very limited in function, and hence much effort has been spent deriving compact and optimum circuits for the Beam Interlock System.

SAFETY

To determine the safety of the Beam Interlock System, US Military handbooks and manufacturer data have been used to derive failure rates and failure modes of components, through this it has been possible to develop an architecture meeting the safety requirements.

One of the key principles of designing a safe system is the ability to test and monitor it. The Beam Interlock System has enormous potential for debugging and self testing. All critical signals internal to the system are redundant and can be checked safely, this gives a clear indication of the system being "As Good As New" when it is being prepared for use. It also gives a very accurate diagnosis of problems, and can monitor developing issues that if left untreated could potentially leave LHC in an unsafe state.

User Interface Safety Example

Over 140 connections are made to the Beam Interlock System by User Systems, each one requiring a User Interface. This means that the majority of electronics in the system is in User Interfaces; making these safe and reliable has a massive impact on the overall safety and reliability of the Beam Interlock System. The User Interface safety has been thoroughly investigated, and results are shown in Table 1. The probabilities of failure are shown over a ten hour period, as this is considered to be the typical length of an LHC Physics run, and a good basis for reliability and safety calculations [7]. This gives an example of how redundant links serve to boost the safety of a system, requiring a double unsafe failure to occur in the same ten hour period to give an overall unsafe state for the User Interface.

Failure Mode	P(Failure Mode) in 10 hours
Any Failure Occurs	3.82E-05
User Permit A Fails Unsafe	4.91E-07
User Permit B Fails Unsafe	4.91E-07
Both User Permit A AND B Fail Unsafe	2.41E-13
Accidental Beam Dump	7.80E-06
Maintenance for next Mission	2.00E-05

Table 1: Failure Modes and Probabilities of a Single User Interface in a Single 10 Hour Mission

CONCLUSIONS AND THE NEAR FUTURE

In order to verify the design and the expected safety and reliability of the Beam Interlock System, the CERN SPS and Transfer Lines from SPS to LHC are to be retro-fitted with a Beam Interlock System for the next SPS start-up [10]. This coupled with the start of the CERN Neutrinos to Gran-Sasso (CNGS) project and the impending installation of the LHC Beam Interlock System for Beam operation in 2007 means that 2006 will be a crucial year for the Beam Interlock System.

The Beam Interlock System has been developed to be a fast and dependable communications system which is expected to fulfil the requirements of both speed and safety set forth in the design specification. These strict requirements are reached by exploiting simple technologies in a sensible way to ensure that such a large and distributed system can respond rapidly, and be safe, forming a solid backbone for the operation of the LHC Machine Protection System.

REFERENCES

[1] B. Puccio et al., “Beam Interlocks for LHC and SPS”, ICALEPCS’2003, Gyeongiu, Korea, October 2003.

[2] F. Bordry et al., CERN - LHC Project Report 521, December 2001.

[3] V. Kain et al., “Damage Levels-Comparison of Experiment with Simulation”, CERN-LHC Project Workshop ‘Chamonix XIV’, January 2005.

[4] M. Werner et al., A Fast Magnet Current Change Monitor for Machine Protection in HERA and the LHC, these proceedings

[5] E. Carlier et al., “Design Aspect Related to the Reliability of the Control Architecture of the LHC Beam Dump Kicker Systems”, ICALEPCS’2003, Gyeongiu, Korea, October 2003.

[6] International Electrotechnical Commission IEC, “Functional Safety of Electrical-Electronic Programmable-Electronic Safety Related Systems”, IEC 61508 International Standard, 1998.

[7] J. Uythoven at al.. “Will we ever get the Green Light for LHC Operation”, CERN- LHC Project Workshop ‘Chamonix XIV’, January 2005.

[8] A. Vergara, “Reliability of the Quench Protection System for LHC Superconducting Elements”, CERN-Thesis 2004-019, 2004.

[9] J. Goldie, National Semiconductor, “Application Note 1057”, 1996.

[10] B. Puccio et al., Interlocking strategy between the LHC and its injector, these proceedings