10th ICALEPCS Int. Conf. on Accelerator & Large Expt. Physics Control Systems. Geneva, 10 - 14 Oct 2005, PO2.020-2 (2005)

THE LHC BEAM DUMPING SYSTEM TRIGGER SYNCHRONISATION AND DISTRIBUTION SYSTEM

A. Antoine, E. Carlier and N. Voumard
CERN, Geneva, Switzerland

ABSTRACT

Two LHC beam dumping systems (LBDS) will fast-extract the counter-rotating beams safely from the LHC collider during setting-up of the accelerator, at the end of a physics run and in case of emergencies. They consist of 15 fast pulsed magnets per ring for beam extraction from the accelerator combined with 10 fast pulsed magnets for horizontal and vertical beam dilution. Dump requests will come from 3 different sources: the machine protection system for emergency cases, the machine timing system for scheduled dumps or the LBDS itself in case of internal failures. These spontaneously issued dump requests will be synchronised with the 3 μs beam abort gap within a fail-safe trigger synchronisation unit (TSU) based on a digital phase lock loop (DPLL) locked on the beam revolution frequency with a maximum phase error of 40 ns. Afterwards, the synchronised trigger pulse will be distributed to the fast pulsed magnet high voltage generators through a redundant fault tolerant trigger distribution system based on the domino effect strategy. The time of flight of the beam through the different magnets as well as the electronic and high voltage generator turn-on delays will be individually compensated for each magnet through the trigger distribution cable length. This paper reviews the technical implementation of the LBDS trigger synchronisation and distribution system.

INTRODUCTION

The function of the LHC Beam Dumping System (LBDS) [1], actually under construction at CERN, will be to fast-extract the circulating beam in a loss-free way from each ring of the collider and to transport it to an external absorber positioned sufficiently far away to allow appropriate beam dilution. It will consist, per ring, of 15 extraction kickers, followed by 15 septum magnets, 10 dilution kickers, and finally an external dump, mounted in a separate cavern some hundred meters away.

A loss-free extraction requires a particle-free gap in the circulating beam during which the field of the extraction kicker magnets can rise to its nominal value. For a correct beam dump action, the 15 extraction kickers must all be fired at the same time with a precision better than 25 ns. The 2.8 μs kick rising edge must be synchronised with the 3 μs beam abort gap with a phase error of less than 40 ns. Incorrect synchronisation will lead to beam loss during the extraction process and thus to damage to the machine, to the experiments or to the beam dumping system itself.

The performance of the extraction kicker system is determined by three operational parameters: its state, its kick time and its kick strength [2]. To reflect this, its control architecture comprises three independent sub-systems, each one dedicated to the control of one specific parameter: the State Control and Surveillance System (SCSS), the Trigger Synchronisation and Distribution System (TSDS) and the Beam Energy Tracking System (BETS).

The TSDS will distribute the dump requests arriving from the client interface to the power triggers, after synchronisation with the beam abort gap and protect the machine against spontaneous firing in one of the pulse generators.

TSDS ARCHITECTURE

Dump requests will come from 3 different sources: the machine protection system for emergencies, the machine timing system for scheduled dumps or the LBDS itself in case of internal failures.

These spontaneously issued dump requests will be synchronised with the beam abort gap within the Trigger Synchronisation Units (TSU). Once synchronised, dump requests will be distributed through the Trigger Fan-Out units (TFO) to the Power Trigger Units (PTU) for firing of the extraction kicker pulse generators as shown in Figure 1. In addition, a redundant fault-tolerant Re-Triggering System (RTS) is foreseen to re-distribute, as fast as possible, a trigger request issued from a spontaneous firing of one generator to the remaining 14 generators.

Furthermore, any dump request will also send a dump trigger, delayed by more than one turn, via the RTS to all power trigger modules. This additional trigger path guarantees that at least an asynchronous beam dump is initiated, even when both TSUs fail.

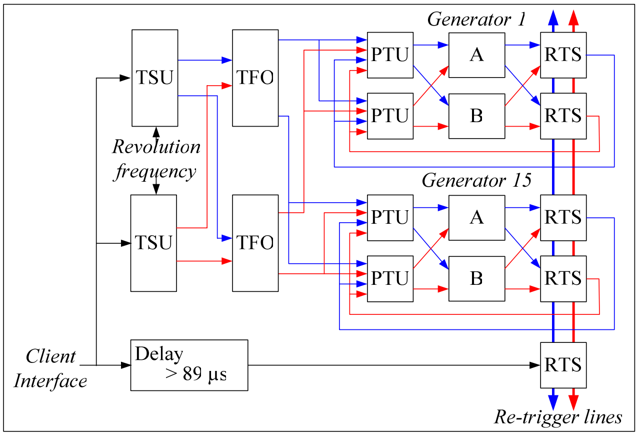

Figure 1: Trigger synchronisation and distribution

The TSDS will be based on a redundant fail-safe logic up to the TSUs and thereafter on a redundant fault tolerant system up to the high voltage generator PTU in order to avoid asynchronous beam dumps in case of failure within the TSDS itself.

The TSUs will be housed within a LynxOS VME front-end powered with redundant power supplies while the TFOs will be based on dedicated hardware lodged in two separate crates.

Operating system and software capabilities available at the VME front-level will only be used for remote monitoring and post-mortem analysis of the system and will not enter within the operational aspect of the TSU.

TRIGGER SYNCHRONISATION UNIT

The TSU will re-phase the spontaneous dump requests with the circulating beam such that the rising edge of the magnetic field in the extraction kicker coincides with the beam abort gap in the circulating beam.

Locked to the LHC beam revolution frequency, two redundant TSUs produce continuously dump trigger pulse trains synchronised with the beam abort gap. The distribution of these pulse trains will be inhibited until a beam dump is requested. The pulses which pass the inhibit stage will then be sent via two redundant trigger fan-outs (TFO) to all the power trigger modules. In this way the first pulse after the reception of a dump request will synchronously trigger the system.

In case the beam revolution frequency will be lost during more than one turn, an internally synchronised direct digital synthesiser based on a numerically controlled oscillator and on a Digital Phase Lock Loop (DPLL), precisely locked on the beam revolution frequency, will issue a dump trigger which is still synchronous with the beam abort gap. In addition the two redundant TSUs will continuously cross-check their DPLL synchronisation and a phase discrepancy greater than 40 ns between the two TSUs will automatically issue a synchronous dump. This mechanism reduces the probability of unsynchronised dumps to almost zero. If the synchronisation of only one of the TSU fails, a synchronous dump trigger will be forced by the redundant system.

A TSU consists of four independent sub-systems as shown in Figure 2:

- the synchronisation sub-system consisting of a DPLL oscillator unit and a timing unit for compensation of the phase error between the beam revolution frequency and the beam abort gap position at the extraction kicker location;
- the dump requests client interface sub-system that consists of detection of the dump request issued by clients through redundant fail-safe 10MHz frequency detectors;
- the dump request management sub-system that re-phases the dump request with the beam abort gap and generates the synchronous beam dump trigger output pulse;

- the supervisory and diagnostic sub-system that cross-checks the correct operation of the TSU with the redundant unit and interfaces the board to the VME interface for surveillance of all operational parameters and post-mortem diagnostic.

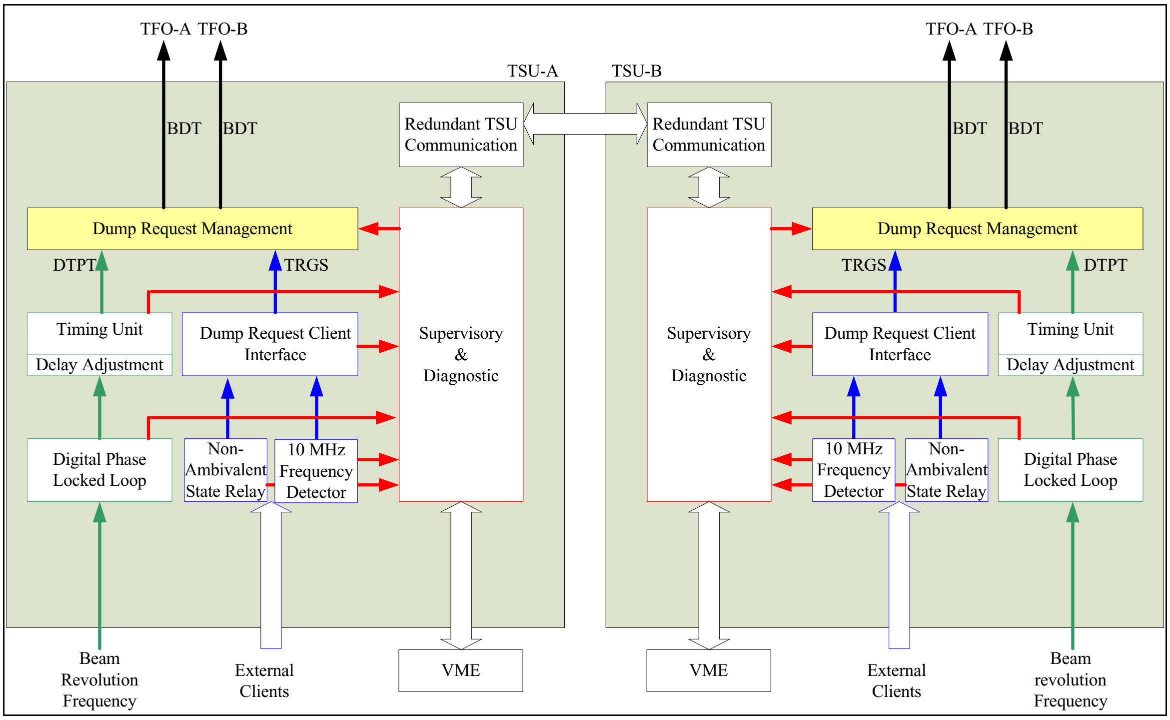

Figure 2: Trigger synchronisation unit architecture

The different TSU functionalities have been modelled in VHDL and implemented within a XILINX VIRTEX-2 FPGA. However, the functions with safety constraints have been kept outside the Xilinx FPGA for reliability reasons and implemented through traditional non-programmable logic components.

Synchronisation

Within the synchronisation sub-system, a DPLL is locked on the beam revolution frequency signal and provides a Dump Trigger Pulse Trains (DTPT) synchronised with the beam abort gap.

The dump trigger pulse train is then delayed by an internal 16 bit 10 ns resolution counter for re-phasing with the beam abort gap at the extraction kicker position and for compensation of the different trigger signal propagation delays. For safety reasons, the counter delay is hard-coded within the board through a DIP switch.

The principal advantages to use a DPLL within the synchronisation sub-system are:

- continuous adaptation of the output frequency to the input frequency avoiding long-term deviations of the internal reference clock;
- good damping properties preventing considerable oscillations and eliminating the noise and the jitter of the input signal while following its slow variation;
- insensitivity to an input signal loss once locked resulting in a correct frequency and phase stability of the dump trigger pulse train for a period of at least 400 µs.

Client interface

The client interface collects and records all dump requests issued from the different TSDS clients. Two different types of interface are supported:

- a non-ambivalent state relays issued by a industrial compliant SIL3 programmable logic controller;
- a 10 MHz frequency signal.

The detection of a change of the relay position or of an interruption of the 10 MHz frequency sets the Trigger Request Gate Signal (TRGS) that is sent to the dump request management sub-system. The detection of the interruption of the 10 MHz frequency will be based on two different techniques working in redundancy. One technique will be based on a resettable counter timer and the other one on digitial re-triggerable monostable.

Dump request management

The dump request management sub-system is responsible to generate the Beam Dump Trigger pulse (BDT) that will be distributed to the extraction kicker high voltage generators by the TFOs.

Within the dump request management sub-system, the DTPT generated by the synchronisation sub-system is continuously inhibited until a beam dump request has been detected by the client interface sub-system.

A simple logical AND function between the DTPT and the TRGS is implemented within the dump request management sub-system. The first pulse in the DTPT coming after the set of the TRGS will be sent out of the TSU to the TFO unit.

This signal will trigger the extraction kicker power trigger unit synchronously with the beam abort gap.

After emission of the beam dump trigger pulse, the dump request management sub-system inhibits the LHC beam permit signal until it has been re-armed trough the VME interface.

Supervisory and diagnostic

The supervisory and diagnostic sub-system is connected to all other sub-systems of a TSU. It supervises all TSU vital internal functions and has the following major tasks:

- the continuous check of the phase relation between the incoming beam revolution frequency and the DPLL output pulse of one TSU;
- the continuous check of the synchronisation of the DPLL output pulses between the two redundant TSUs;
- the surveillance of all internal dynamic parameters of one TSU;
- the recording of the internal registers when a BDT signal is issued for post-mortem analysis;
- the interfacing of the board to the VME bus for remote control and monitoring;
- the inhibition of the LHC beam permit loop until both redundant TSUs are armed and synchronised.

When an internal failure is detected within one of the sub-systems or an inconsistency is detected between the two redundant TSU, an internal dump request will be directly sent to the dump request management sub-system.

Table 1 summarises the possible synchronisation failure scenarios and the possible actions in order to still guarantee a synchronous BDT signal. A phase difference above 40 ns either between the beam revolution frequency and the DPLL output pulses or between the two redundant DPLLs sets the TSU to a faulty state.

Table 1: Synchronisation Failure Scenarios and Possible Actions

TSU-A	TSU-B	State	TSU-A Action	TSU-B Action
Faulty	Faulty	Revolution Frequency Faulty	Generate DTP Signal after 5 cycles	Generate DTP Signal after 5 cycles
Faulty	OK	Synchronisation TSU-A Faulty	Inhibit Output Stage	Generate DTP Signal
OK	Faulty	Synchronisation TSU-B Faulty	Generate DTP Signal	Inhibit Output Stage
OK	OK	All OK	No action	No action

In case of loss of synchronisation, the LHC beam permit signal will be inhibited until a re-synchronisation of the DPLL with the beam revolution frequency has been possible. After each generation of a DTP signal, the LHC beam permit signal will also be inhibited until a successful re-arming of the TSUs has been performed.

TRIGGER FAN-OUT AND DISTRIBUTION

The dump request distribution will be based on a redundant chain of stages using the “domino effect” strategy to trigger the next stage in the chain. The energy required to propagate the request from the TSU up to the extraction kicker high voltage generators will be pre-stored within capacitors at each stage of the triggering chain. This energy will be used to trigger the next stage of the chain and its level is checked before a beam permit signal is issued. An active high logic has been used for the trigger distribution chain in order to reduce the effect of an accidental cable disconnection or a bad contact in the chain.

Two redundant TFOs will be used to distribute the dump request coming from the TSUs to the PTUs. Each TFO is divided into three main circuits: the Trigger Fan-Out Receiver (TFOR), the Trigger Fan-Out Transceiver (TFOT) and a +/- 15V power supply to power the crate. Each crate will contain one TFOR and six TFOT with a capacity to drive up to 48 PTU 50 Ohms inputs. Four additional outputs are foreseen for inhibiting the BETS or for the acquisition of the UTC dump trigger time.

The TFOR is responsible for acquiring the two redundant signals coming from each of the two TSUs and to distribute them safely to the six TFOTs.

Each of the TFOTs is responsible for transmitting the dump request signal coming from the TFOR to the eight PTUs. The driver of each output will be based on a blocking oscillator circuit providing galvanic insulation through a high frequency pulse transformer. The pulse is transmitted using a charged capacitor pulled to ground via the primary winding of the transformer by two redundant switching transistors. As the pulse is about 1 μs wide and the time between two trigger pulses is more than 1 second there is no time constraint to recharge the capacitor for the next dump request and the capacitor can be re-charged with a slow ramping profile in order not to overload the power supply at the moment of the dump request due to the low output impedance of 50 Ohms.

The trigger distribution between the output of the TFO and the input of the PTU involves only signal propagation delays. Time of flight of the circulating beam through the magnets as well as electronic and high voltage turn-on delays will be compensated for each kicker magnet individually by fine adjusting the trigger distribution cable length.

RE-TRIGGERING CIRCUIT

The main task of the RTS is to re-distribute, as fast as possible, a trigger request issued from a spontaneous firing of one generator to the remaining 14 generators.

A redundant chained input/output system has been chosen for the RTS. Each pulse generator has 5 re-trigger source sensors with enough powering capabilities to trigger all the PTUs of the remaining other 14 high voltage generators. Due to the architecture of the system an avalanche mechanism is started after a detection of a spontaneous firing.

Since there is no stored energy in the system itself it is difficult to create spurious triggers; neither a disconnected cable nor a defective trigger source could cause triggers.

The typical reaction time between the detection of a spontaneous firing of one pulse generator and the re-triggering of the 14 other generators is 750 ns.

Any dump request detected by the TSUs will send, without any synchronisation, a trigger pulse, delayed by more than one turn to all the PTUs via the re-trigger system. This second trigger path guarantees that at least an asynchronous beam dump is initiated, even when both TSUs are in a faulty state. It should be noted that in case of a spontaneous firing of one of the generators no synchronisation is possible.

CONCLUSION

The LBDS must remove both beams safely from the LHC to prevent any damage to the machine. A complete redundant path for the entire triggering system from the synchronisation unit up to the triggering of the generators has been designed to provide the required degree of safety. The high availability of the selected architecture and its redundant structure minimises the risk of an asynchronous beam dump.

ACKNOWLEDGEMENTS

The authors would like to thank Martin Rampl for a lot of good input provided by his thesis and the development of the first TSU prototype that has guided our technical choices.

REFERENCES

[1] O. Brüning et al., “LHC Design Report, Volume 1, The LHC Main Ring”, CERN-2004-0003, pp. 441-464, Geneva, CERN, June 2004.

[2] E. Carlier et al., “Design Aspect related to the Reliability of the Control Architecture of the LHC Beam Dump Kicker Systems”, LHC Project Report 679, Geneva, CERN, October 2003 and in Proceedings of ICALEPCS 2003, Gyeongju, Korea.