lpGBT: Low-Power Radiation-Hard Multipurpose High-Speed Transceiver ASIC for High-Energy Physics Experiments

Paulo Moreira, Szymon Kulis, Sophie Baron, Stefan Biereigel, Eduardo Brandao De Souza Mendes, João P. Matos-Carvalho, Bram Faes, Miroslaw Firlej, Tomasz Fiutowski, Jose Fonseca, Rui Francisco, Datao Gong, Nour Guettouche, Ping Gui, *Senior Member, IEEE*, Di Guo, Daniel Hernandez Montesinos, Marek Idzik, Iraklis Kremastiotis, Thanushan Kugathasan, Pedro Leitao, Paul Leroux, *Senior Member, IEEE*, Julian Mendez, Jakub Moroń, Nuno Paulino, *Senior Member, IEEE*, David Porret, Jeffrey Prinzie, Adithya Pulli, Quan Sun, Krzysztof Świentek, Ken Wyllie, Dongxu Yang, Jingbo Ye, Tao Zhang, and Wei Zhou

***Abstract*— Commissioning of detector systems for the high-luminosity large Hadron collider (HL-LHC) is scheduled to take place between 2026 and 2028 at CERN. Application-specific integrated circuits (ASICs) for those systems have been in intense development over the past ten years. Some of those ASICs, as is the case of the low-power gigabit transceiver (lpGBT) described in this work, have now been produced in industrial quantities and have been fully qualified for operation in the HL-LHC environments that require, where the innermost detectors are concerned, radiation hardness over 1 MGy. The lpGBT is a multifunctional device, enabling data transmission between the off-detector and the on-detector systems. Data can be transmitted from the detector at 5.12 and 10.24 Gb/s and to the detector at 2.56 Gb/s. It implements data rate-configurable electrical links to communicate with the front-end ASICs and low-speed serial and parallel buses for experiment control. A set of analog functions for monitoring and control of the physics detectors is also included. This article describes the functionality and the architecture of the lpGBT ASIC and reports on its radiation hardness characterization.**

Received 29 August 2024; revised 29 October 2024; accepted 22 November 2024. Date of publication 4 December 2024; date of current version 21 January 2025. This work was supported in part by the Polish Ministry of Science and Higher Education under Contract 5179/H2020/2021/2 and in part by the European Union's Horizon 2020 Research and Innovation Program under Grant 101004761 (AIDAinnova). *(Corresponding author: Szymon Kulis.)*

Paulo Moreira, Szymon Kulis, Sophie Baron, Stefan Biereigel, Eduardo Brandao De Souza Mendes, Jose Fonseca, Rui Francisco, Nour Guettouche, Daniel Hernandez Montesinos, Iraklis Kremastiotis, Thanushan Kugathasan, Pedro Leitao, Julian Mendez, David Porret, Adithya Pulli, and Ken Wyllie are with the European Organization for Nuclear Research (CERN), 1211 Geneva, Switzerland (e-mail: Paulo.Moreira@cern.ch; Szymon.Kulis@cern.ch).

João P. Matos-Carvalho and Nuno Paulino are with UNINOVA-CTS and LASI, NOVA School of Science and Technology, NOVA University Lisbon, 2829-516 Caparica, Portugal.

Bram Faes, Paul Leroux, and Jeffrey Prinzie are with ESAT-ADVISE Research Laboratory, Katholieke Universiteit Leuven, 3000 Leuven, Belgium.

Miroslaw Firlej, Tomasz Fiutowski, Marek Idzik, Jakub Moroń, and Krzysztof Świentek are with the Faculty of Physics and Applied Computer Science, AGH University of Krakow, 30-059 Kraków, Poland.

Datao Gong, Ping Gui, Di Guo, Quan Sun, Dongxu Yang, Jingbo Ye, Tao Zhang, and Wei Zhou are with Physics Department, Southern Methodist University, Dallas, TX 75205 USA.

Color versions of one or more figures in this article are available at https://doi.org/10.1109/TNS.2024.3506753.

Digital Object Identifier 10.1109/TNS.2024.3506753

***Index Terms*— High-energy physics (HEP) communications, multigigabit transceiver, radiation-hard application-specific integrated circuit (ASIC).**

I. Introduction

TODAY'S particle physics experiments present recurring technological challenges. Among them are: clock and trigger information distribution, efficient data transfer from the detectors to the data acquisition (DAQ) systems, and the control and monitoring of environmental parameters inside the detector. With the backdrop of these challenges, the high-energy physics (HEP) community, drawing on the extensive experience in designing, building, and commissioning detector systems for the large Hadron collider (LHC), foresaw a trajectory where future systems would inevitably become more complex, harder to test and debug, and, concurrently, more expensive [1], [2]. Moreover, the growth in engineering resources within the community was not expected to keep pace with the escalating demands.

Reacting to this situation, the HEP community recognized the imperative to adopt common solutions, specifically for optical links and their integral components such as optoelectronics and application-specific integrated circuits (ASICs). This laid the foundation for the specification and development of the versatile optical link modules [3] and the gigabit transceiver (GBT) chipset [4], with the GBT X (GBTX) ASIC being the core transceiver chip [5]. Collaboratively developed across multiple HEP institutes, these projects provided a unified and versatile framework for addressing various services, including timing, trigger, and control (TTC) and DAQ, without the need for dedicated physical links to carry specific detector data types. That philosophy led to the development of link architecture as shown in Fig. 1, where both the downlink (off to on-detector) and the uplink (on to off-detector) are present with the downlink carrying detector control data (including detector configuration and calibration) and TTC.

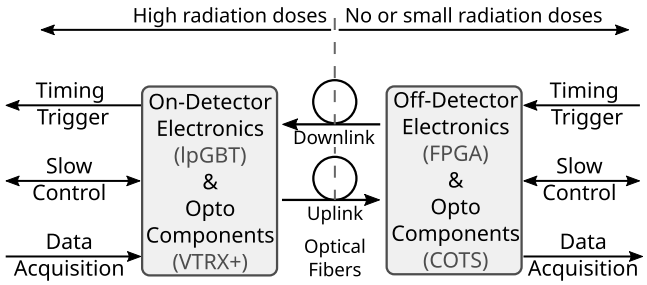

Fig. 1. LHC link architecture.

Embedded in the architecture is the transmission of the clock signal that, almost always, is not transmitted individually but is recovered from the downlink serial data stream. This clock is the LHC bunch crossing clock (40 MHz). The uplinks close the loop of the experiment's control systems and carry the DAQ data collected by the front-end detectors. The chipset development embodied the approach of designing for radiation tolerance, sharing physical resources, minimizing development costs, and efficiently utilizing engineering resources. Compatibility with commercial off-the-shelf (COTS) systems, especially field-programmable gate arrays (FPGAs), emerged as a primary consideration, ensuring seamless integration with prevailing technologies.

This article describes the low-power GBT (lpGBT), designed as the successor to the GBTX. Compared to the GBTX, the lpGBT not only enables higher data rates but also integrates additional features for controlling detectors, consolidating functions previously distributed across different chips. Notably, these advancements come with reduced power consumption and enhanced resilience to radiation, marking a significant stride in developing more efficient and robust tools for particle physics research. Anticipated for integration into experiments during the upgrades scheduled between 2026 and 2028, the lpGBT is strategically aligned to meet the specific requirements of A Toroidal LHC Apparatus (ATLAS) and Compact Muon Solenoid (CMS) detector systems [6], [7]. This article comprehensively presents the architecture of the lpGBT chip and gives circuit-level insights into selected components (see Section II) so that the choices adopted for the design methodology and optimization can be put into context. The key ASIC performance metrics are reviewed, and a summary of the characterization and radiation qualification results is presented in Section III.

II. Architecture and Features

The lpGBT [8] is a device designed to address the needs of high-luminosity LHC (HL-LHC) subdetectors considering various data transmission modes (simplex transmitter, simplex receiver, and transceiver), clock and timing distribution, experiment control, and monitoring. To cover those, the lpGBT functionality is highly configurable. This section provides, for completeness, a birds-eye view of the lpGBT architecture and its selected functionalities. For an in-depth description of the ASIC, refer to [8].

Physically, the ASIC is encapsulated in a 289-pin fine-pitch (0.5 mm) ball grid array (BGA) package with dimensions $9 \times 9 \times 1.24$ mm.

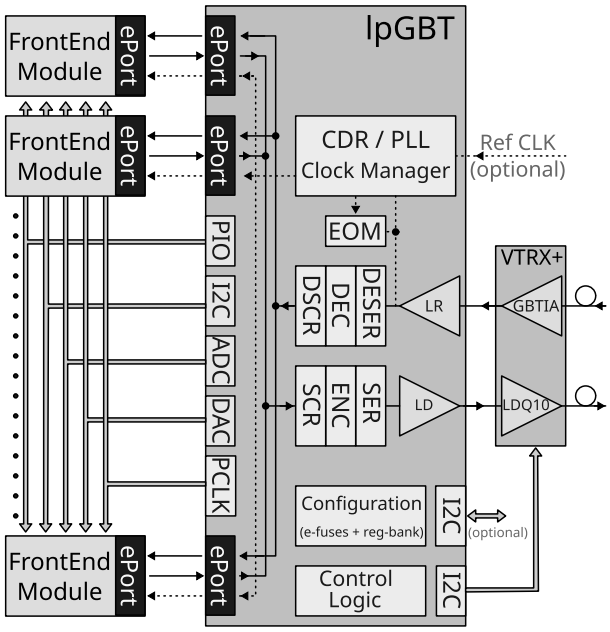

Fig. 2. lpGBT ASIC and system architecture.

A. Downlink Section

The architecture chosen for the ASIC and its data links interconnects is represented in Fig. 2. As shown, the lpGBT interfaces with the fiber-optic downlink (off-detector to on-detector direction) through a p-i-n-diode receiver [9], [10] embedded in the VTRX+ module [11]. The off-detector systems transmit data to the lpGBT at 2.56 Gb/s. The data transmitted to the ASIC can be targeted to the ASIC itself, to control its operation, and to the front-end modules/ASICs. In the latter case, these data can be either for experiment control or for detector timing (e.g., trigger commands). Because the lpGBT is an integral part of the synchronization of the detector systems in which it is embedded, it was designed as a fixed and deterministic latency device.[1]

In the lpGBT receiver section, a clock and data recovery (CDR)/phase-locked loop (PLL) circuit [12], [13], [14], in the clock manager block, locks to the incoming serial bit stream extracting a 5.12-GHz clock synchronous with the LHC machine clock (since the machine clock is used as the timing reference for the downlink transmitters). Note that the CDR/PLL circuit operates as a PLL, with an external reference clock provided to the ASIC, when the lpGBT is configured as a simplex transmitter.

The 5.12-GHz clock is further divided in the clock manager circuit, making it available to the front-end modules as eClock signals at frequencies that are the power of two multiples of 40 MHz, ranging up to 1.28 GHz. The lpGBT implements 29 independently programmable clock outputs (part of the ePorts in Fig. 2). The phase of these clocks is fixed in relation to the LHC machine clock. To allow for timing calibration of the detector systems, the lpGBT provides four additional clock ports (phase-shifted clock (PCLK) in Fig. 2) that are

[1]That is, the delay and phase of the data received through the downlink in relation to the LHC main clock is guaranteed to be constant during operation and to have always the same value even upon a restart.

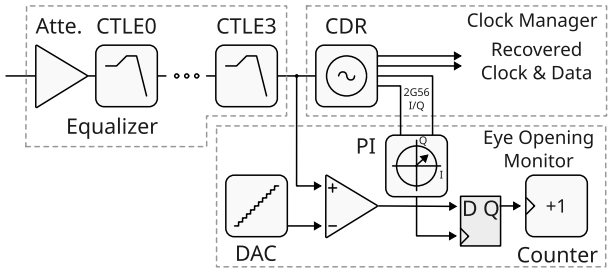

Fig. 3. CTLE and EOM schematic.

not only programmable in frequency but also in phase with a resolution of 48.8 ps over the full clock period, independent of the selected frequency [15].

The 5.12-GHz clock and all its submultiples, down to 40 MHz, are further used inside the ASIC for the operations of de-serialization, serialization, de-scrambling, scrambling, decoding, and encoding.

The downlink signal path starts with a passive attenuator, followed by four stages of a continuous-time linear equalizer (CTLE) illustrated schematically in Fig. 3. The purpose of the attenuator is to accommodate large-amplitude signals, up to 1 V_{pp}, to avoid saturation of the CTLE stages. Each CTLE stage provides flexibility in shaping the transfer function by adjusting the stage resistor and capacitor values, thereby influencing the transfer function zero locations within the frequency range of 80–750 MHz.

The eye-opening monitor (EOM) circuit enables on-chip eye-diagram measurements of the incoming 2.56-Gb/s downlink data (see Fig. 3). Controlled by the user, the EOM facilitates vertical (voltage) and horizontal (time) scans of the data stream, allowing the reconstruction of a sampled eye diagram. This functionality proves instrumental in monitoring the downlink signal quality and in assisting the configuration of the downlink receiver equalizer parameters. The equalizer output signal undergoes comparison with a reference voltage (V_{ref}), and the comparator's result is sampled at the rising edge of a clock synchronized with the incoming data. The sampling results control a ripple counter, accumulating statistics over a defined period. To eliminate dependence on the input signal common mode, the comparator incorporates a differential difference amplifier. The phase interpolator, a component of the circuit, receives the 2.56-GHz in-phase "I" and quadrature "Q" clocks from the clock manager block. These clock signals are used for phase interpolation, allowing for full phase rotation with an approximate resolution of 6.1 ps (64 possible phases).

The received serial data are de-serialized into a 64-bit word. This word is further de-interleaved and decoded using a forward error correction (FEC) code [16], [17] to correct for any errors that might have occurred during transmission due either to noise or to single-event effects (SEEs). The latter can occur in the chain of circuits preceding the decoder stage, including the external p-i-n-receiver. The decoder and following circuits use triple modular redundancy (TMR) to achieve robustness against SEEs.

The data payload of the 64-bit downlink word is organized into four bytes plus four bits. Each byte is addressed to a specific transmitter, called ePortTx, where it is serialized and driven to the front-ends using a set of differential electrical links, called eLinks. Each ePortTx can drive up to four eLinks depending on the chosen data rate. The number of eLinks driven per ePortTx scales down with the programmed data rate: one at 320 Mb/s, two at 160 Mb/s, or four at 80 Mb/s. The eLink data rate for each of the four ePortTx can be set independently. A maximum of 16 transmit eLinks can be established if all the ports are programmed to operate at 80 Mb/s.

The extra four bits of the downlink payload form the external control (EC) and internal control (IC) links. The two EC bits are associated with a dedicated ePortTx which operates at 80 Mb/s and drives a single eLink. This port is primarily intended for transmitting experiment control data, but it can be used for other purposes. The two IC bits are dedicated to the configuration of the ASIC itself using a custom protocol. Commands sent over the IC link extend across multiple frames (spanning 25 ns) and employ bit stuffing [18]. Consequently, the execution latency of data/commands is not fixed, making the IC channel asynchronous in nature.

B. Uplink Section

The lpGBT uplink (on-detector to off-detector direction) receives data from the front-ends through the eLinks associated with the receiver ports, called ePortRx, and forwards it to the off-detector systems via an optical fiber link driven by a vertical cavity surface emitting laser (VCSEL) driver ASIC [19], [20] in the VTRX+ module [11]. The lpGBT supports two uplink data rates, 5.12 Gb/s (5G mode) and 10.24 Gb/s (10G mode), and two FEC codes, FEC5 and FEC12, which will be discussed later. As a consequence, the available user bandwidth depends both on the chosen uplink data rate and the FEC code used.

Similar to the ePortTx, each ePortRx is associated with four eLinks, and the data rate of the eLinks (and thus the number of active eLinks) in each ePortRx can be set independently for each ePort. The maximum eLink data rate is 1.28 Gb/s, when the lpGBT transmitter operates at 10.24 Gb/s and 640 Mb/s for 5.12-Gb/s operation.

Each ePortRx deserializes and combines the data received over the associated eLinks into either a 16- or 32-bit word for transmission at 5.12 or 10.24 Gb/s, respectively. The lpGBT can trade off bandwidth for SEE protection offering either seven or six active ePortRx when the FEC5 or FEC12 codes are used. The maximum number of receiving eLinks that can be established is thus 28 when all the eLinks are operated at the minimum data rate and FEC5 encoding is used.

As is the case for the receiver section, the uplink payload field embeds two further logical links. Two EC bits are assigned to an ePortRx with a single eLink operating at 80 Mb/s to implement a detector control link. Two additional bits are assigned to the IC uplink. These bits are used by the ASIC to send replies to commands transmitted via the downlink IC link.

Data from ePortRx, EC, and IC links are scrambled (for dc balance), encoded, and interleaved (for SEE protection) before serialization and transmission over the uplink.

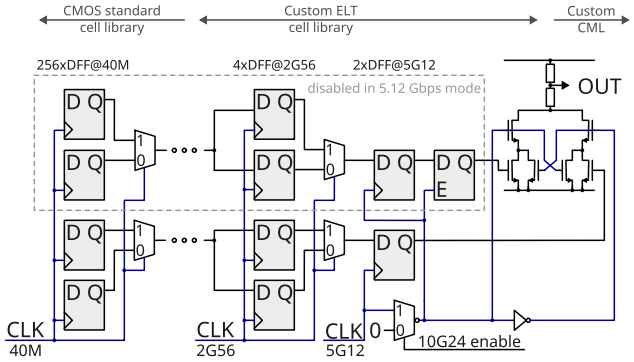

Fig. 4. Serializer schematic.

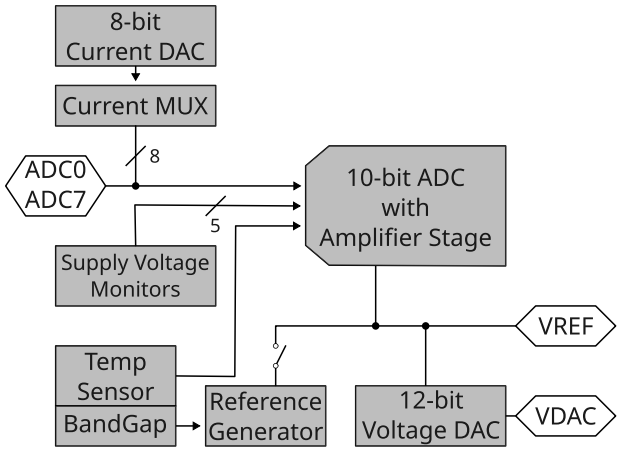

Fig. 5. Analog unit block diagram.

The lpGBT serializer, shown in Fig. 4, works at two data rates 5.12 and 10.24 Gb/s converting the parallel data (128 or 256 bits) at 40 MHz into a serial bit stream. The serializer implements a multistage scheme where the serialization is performed in ten stages and consumes only about 4% of the power of a classical shift-register serializer architecture. Since the clock frequency increases with the serialization level, it was possible to further optimize the power and timing margins at each level.

The serializer uses the half-rate architecture and omits the resampling stage at 10.24 GHz. The last multiplexing state is a custom-designed cell that doubles as a single-ended to differential converter block.

C. Experiment Control and Monitoring

The lpGBT implements a set of features dedicated to experiment control and environment monitoring [8].

1) Digital Functions: The ASIC includes digital functionality to aid in the implementation of the detector control functions.

- *General-Purpose I/O Ports:* 16 I/O pins, can be individually configured as inputs or outputs. When set as outputs, the I/O pins can be configured either as complementary metal oxide semiconductor (CMOS) drivers or pull-up/pull-down drivers (with the pull/down resistors internal to the circuit). All I/O transactions are synchronous with the internal system clock (40 MHz).
- *Interintegrated Circuit (I^2C) Mains:* Three I^2C [21] mains are implemented in the ASIC and their main features are a concurrent operation of the three channels, programmable data transfer rates (0.1, 0.2, 0.4, and 1.0 MHz), support for 10-bit addressing mode and the compatibility with both single-byte and multibyte I^2C read/write operations.
- *RSTOUTB Pin:* An output pin is implemented that delivers an active low reset pulse. The pulse is generated by the ASIC power-up state machine and its behavior is thus closely coupled to the initialization of the lpGBT itself.

2) Analog Functions: A set of analog functions is provided. An in-depth discussion can be found in [8] and [22] but a brief overview follows for completeness.

The analog unit, shown in Fig. 5, is composed of a 10-bit analog-to-digital converter (ADC), preceded by a programmable gain amplifier, a 12-bit voltage digital-to-analog converter (DAC), an 8-bit current DAC, and of a temperature sensor. These blocks are further supported by a 16-input analog multiplexer, an eight-pole current switch, and a reference voltage generator.

For calibration, the reference voltage is adjustable through an 8-bit register determining directly the accuracy of the ADC and voltage DAC. Calibration parameters for the analog unit are determined at production test time and are made available to the users and associated with the ASIC unique identifier set at testing time in the internal electrically fusible (e-Fuse) registers.

The analog multiplexer allows the ADC to measure the voltages of any of the eight analog pins and of the eight internal circuit monitoring points (e.g., internal supply voltage rails). The current switch enables the selective feeding of currents to any of the analog pins.

D. ASIC Initialization and Control

The lpGBT was designed to accommodate as many of the HL-LHC detectors/subdetectors requirements as technically feasible. As previously outlined, this requires a high degree of configurability. Moreover, for the systems in which it will be used (e.g., [23], [24]), the ASIC will be the connection bridge to the off-detector systems. It is, therefore, essential that it is capable of reliably auto-booting into a state where it is suitably configured for the system or securely establishing the link with off-detector electronics to allow a full configuration. To achieve this, the very basic modes of operation are set by a few hardwired pins. The hardwired setting of these pins determines, for example, the uplink data rate and FEC encoding as well as the type of boot sequence to be executed.

Depending on the chosen operation mode, fully configuring the ASIC might require programming up to 224 registers. The choice of basic operation mode (hardwired) determines how these registers will be initialized. The most basic boot sequence, *boot from ROM*, initializes only the strict minimum set of registers needed to establish communications with the off-detector system. The off-detector system can then complete the initialization using the IC channel.

At the other extreme, the ASIC can be fully configured at startup by reading the configuration from an internal e-fuse memory bank.

A third option, and the most flexible, is to hard-wire the ASIC to wait to be configured through its secondary I^2C interface (see Fig. 2). This of course assumes that an I^2C main is present in the system.

It should be noted that independent of the boot mode, all the configuration registers can either be accessed through the I^2C secondary port or the previously described IC channel and thus reconfigured during operation.

As mentioned, the lpGBT can be configured to operate as a simplex receiver or transmitter in which case the EC port can be used as an IC channel.

E. Special Features and Design for Testability

The lpGBT implements test features that were designed to validate the operation of the ASIC and the data transmission channels. Of particular relevance for the user are the data path test features. These can be used not only to test the ASIC itself but also as a tool to help during user system development and debugging since they allow standardized link test procedures at the system level. Among the most interesting features for the users are the loopbacks, test pattern generators, checkers, and the EOM (previously described in Section II-A).

Loopbacks allow the user to send data to the ASIC, having the data propagating down the ASIC to a selected depth of the ASIC data path and, at that point, looping the data back to the user. Loopbacks are provided both for the high-frequency ports (up and downlinks) and the ePorts. Due to the bandwidth asymmetry between the uplink and downlink directions, when using the loopback configuration, the downlink data bits are repeated on the uplink (up to four times, depending on the transmitter's data rate) [8].

Test pattern generators and checkers are also available both for high-frequency ports and ePorts. The generators allow the ASIC to generate known data patterns (e.g., pseudorandom and counting sequences) and count detected errors so that the user can verify their correct data transmission.

Finally, the ASIC contains a checker which, when enabled, tracks the number of corrections made by the receiver FEC decoder. It can be active during normal operation and read out through either the IC channel or the I^2C port. Although this is not a true bit error rate measurement, if read frequently, it can be used to monitor the quality of the data reception.

F. TID and SEE Mitigation Strategies

Due to the very high beam luminosity planned for the HL-LHC machine upgrade, the radiation levels for the innermost detectors in the LHC experiments are expected to exceed 10 MGy and 10^{16} 1 MeV n/cm^2 over the ten-year lifetime of the experiments [25]. Additionally, high particle fluxes will lead to SEE. These extremely high levels of total ionizing dose (TID), nonionizing energy loss, and particle fluxes, pose significant challenges to the electronics and optoelectronics components installed in the detectors. Designing for TID radiation tolerance and SEE mitigation were thus major specifications for the lpGBT.

Since the lpGBT will be physically close to the optoelectronics devices, it was decided to target a TID radiation hardness similar to what those devices can sustain [26]. The qualification target, with a safety factor of two, was thus set to 2-MGy TID. For SEE immunity, it is harder to set a target since the SEE rates are difficult to predict. However, as the lpGBT is considered a connection bridge to the off-detector systems, it was decided to make the ASIC free of functional interrupts for the range of expected linear energy transfer (LET) [25].

1) TID Effects Mitigation: TID effects are partially mitigated in the lpGBT by fabricating the ASIC in a commercial 65-nm CMOS technology extensively qualified for radiation tolerance by the HEP community [27], [28], [29], [30], [31] and by the use, when deemed necessary, of special radiation hardening layout and circuit techniques [32].

A range of circuit design techniques is used in the lpGBT to ensure robustness against TID. Most of the circuitry runs at the LHC bunch clock frequency (40 MHz). For circuits running at frequencies between 40 and 640 MHz, the intrinsic speed of the technology allows to accept a performance degradation of the logic even if the circuits are implemented with standard cells that were designed and layout without having radiation in mind [33], [34]. This was thus the technique adopted for the lpGBT circuits that operate with a clock frequency up to 640 MHz. For the digital circuit components operating with clock frequencies between 1.28 and 5.12 GHz, a custom digital library was developed based on an enclosed layout transistor (ELT), which was shown to have superior TID tolerance [33]. Finally, above 2.56 GHz, some circuits (in particular, the clock generator) are implemented in current-mode logic (CML) with ELT nMOS transistors. This offers the ultimate TID tolerance since it avoids the use of pMOS transistors altogether and uses polysilicon load resistors that are virtually insensitive to TID radiation.

2) SEE Mitigation: SEEs are a major concern for reliable data transmission in HEP experiments. The severity of an SEE depends primarily on the role of the circuit impacted by the SEE, the circuit cross section and the threshold LET. Single, or even multiple, bit errors in the data payload are typically not considered critical, and can often be dealt with easily. However, bit flip errors that occur, for example, in a state machine that monitors the circuit operation can lead to wrong decisions and result in functional interrupts lasting up to several milliseconds. The strategies used for SEE mitigation in the lpGBT are thus multiple and reflect the level of protection required for the specific function.

TMR for logic and clock trees is among the most robust SEE mitigation strategies but it incurs power consumption and speed penalties. Nonetheless, self-correcting TMR was the chosen technique for all state machines, logic, and configuration registers that operate at the LHC bunch-clock frequency. The TMR technique was introduced at design time with the help of the TMRG tool [35]. This was justified by the fact that running at 40 MHz, the power overhead was relatively small when compared with the power dissipation

in the high-frequency circuits. TMR was also selected for the high-frequency prescalars that are needed to generate all the frequencies down to 40 MHz starting from the PLL clock signal at 5.12 GHz. Given the speed penalty resulting from TMR, as discussed above, such circuits had to be implemented using custom logic libraries using either CMOS ELT or CML logic.

As already mentioned, the data paths between the FEC encoder and the uplink driver and, vice versa, between the downlink receiver and the decoder were not triplicated since the FEC encoding provides the error correction capability and, additionally, sporadic transmission errors were considered acceptable since they would not lead to system functional interrupts.

Three FEC approaches are used in the lpGBT. Most of the data bandwidth requirements are for the uplinks and this is reflected in the lpGBT by the adoption of either 5.12- or 10.24-Gb/s bandwidth for the uplink while the downlink operates at 2.56 Gb/s. The lower downlink bandwidth requirement also meant that a significant fraction of the bandwidth (effectively 37.5%) could be used to mitigate SEEs in the data path without imposing an excessive penalty on the user payload (56.25%), with the remainder of the bandwidth taken by the frame header. The Reed–Solomon [36] code RS(5, 3), which is a truncated version of the RS(7, 5), was chosen for the downlink. The outputs of four encoders are interleaved. This implementation choice allows transmitting a frame containing a 36-bit user payload (including the EC and ic field) while being able to correct any single random bit error in the frame and any burst of errors with a length of up to 10 bits (or up to 12 bits provided they are fully contained in the boundary of four symbols).

For the uplink, where the user bandwidth is at a premium, it was decided to allow the user to trade off error correction capabilities for additional bandwidth. Therefore, two FEC codes are implemented: FEC12 for high burst error protection and FEC5 for larger user payload. The FEC12 consists of RS(15, 13) interleaved three times for 5.12 Gb/s or six times for 10.24 Gb/s. The code offers the potential to protect up to 156 or 312 bits for three or six times interleaving. However, due to using a shortened RS code, the number of user bits protected is only 102 or 206, respectively. The FEC12 can correct up to 12 or 24 consecutive wrong bits for 5.12 or 10.24 Gb/s. Conversely, the FEC5 consists of an RS(31, 29) implementation with no interleaving for 5.12 Gb/s or two times interleaved for 10.24 Gb/s. Similar to the FEC12 case the code is truncated due to the frame size and the number of protected bits is 116 or 234. The FEC5 can correct up to five or ten consecutive wrong bits for 5.12 or 10.24 Gb/s, respectively. The summary of the frame format and the selected codes are presented in Table I.

Finally, potential sensitive circuits for long-lasting functional interrupts are the PLL or CDR circuits present in the transceivers. Although PLL and CDR circuits with digital architectures [37], [38] are becoming more common, and are thus amenable to TMR implementations, those circuits cannot be made fully SEE robust since they still partially rely on analog circuits, most commonly the oscillator. In the case of the lpGBT, an analog PLL/CDR is used thus precluding the use of circuit-wide TMR techniques. The approach was to search for SEE robust circuit architectures resulting in the adoption of an LC-voltage-controlled oscillator (VCO) that has demonstrated to be SEE robust and to have low jitter performance [12], [13], [39], [40]. Additionally, the TMR of selected components, like the phase frequency detector [41] and the feedback divider, was used.

TABLE I
SPECIFICATIONS FOR DIFFERENT BANDWIDTH CONFIGURATIONS

Field	Units	Uplink				Downlink
Bit Rate	Gbit s^{-1}	5.12		10.24		2.56
Max FEC Corr.	bits	5	12	5	12	12
Frame Size	bits	128		256		64
Header	bits	2				4
IC	bits	2				2
EC	bits	2				2
eLink Data	bits	112	96	224	192	32
FEC	bits	10	24	20	48	24
Efficiency	%	91	91	80	80	56

G. Verification

The lpGBT architecture and its implementation were verified before the submission. A verification environment based on the universal verification methodology (UVM) was developed to address the functional verification challenges arising from the architecture and versatility of the chip [42]. Furthermore, the UVM environment was specifically crafted to facilitate comprehensive verification of the robustness against SEE. Verification activities were conducted at various project stages, ranging from register-transfer level (RTL), throughout triplicated RTL, and all the way to a timing-annotated netlist of the design. The lpGBT test suite comprises a set of heavily randomized top-level tests and several standalone tests that target specific chip functionality. These tests simulate typical use cases of the lpGBT during boot-up and regular actions, mimicking intended use. Additionally, the suite exercises all recovery mechanisms, including watchdogs and timeouts, ensuring that the chip consistently returns to an operational state after prolonged error conditions. The thorough simulation of SEEs was a major part of the lpGBT verification effort.

III. EXPERIMENTAL RESULTS

This section presents the experimental methodology and the various results with emphasis on TID and SEE experiments.

A. Experimental Setup

The lpGBT was manufactured using a commercial 65-nm six-metal CMOS technology, and the resulting die was flip-chip assembled in a fine-pitch BGA package. To validate the chip's functionality and assess its performance, a specialized test setup was created based on an FPGA [43]. In this setup, a custom mezzanine card connects all the eLinks interfaces to an AMD Virtex 7 device, emulating a connection with a front-end device. A high-speed port interface employs two pairs of subminiature version A connector (SMA) connected to two FPGA transceivers. The eClocks and digital interfaces (general-purpose input/outputs (GPIOs) and I^2C mains) are also interfaced with the FPGA, which provides the necessary

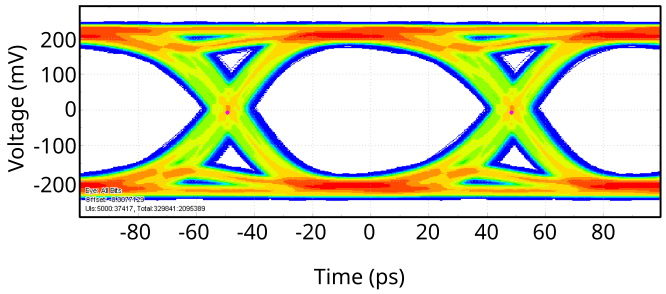

Fig. 6. Uplink eye diagram at the output of the lpGBT operating in 10.24-Gb/s mode captured using a Tektronix DPO73304SX real-time oscilloscope.

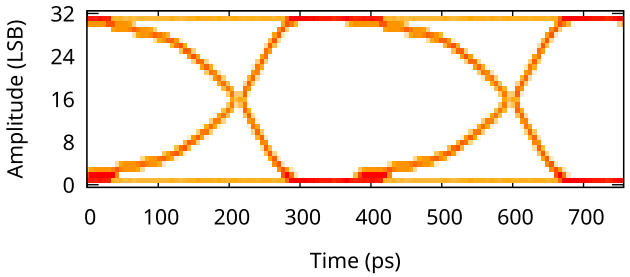

Fig. 7. Downlink eye diagram captured with the built-in EOM. The input electrical signal with an amplitude of 800 mV_{pp} was generated with the FPGA test system. The lpGBT internal attenuator was bypassed and no channel equalization was performed.

logic to exercise these functionalities. The analog channels in the ASIC, which cannot be directly monitored by the FPGA, are connected to ADCs or DACs. The lpGBT tester firmware is designed for the AMD Virtex 7 FPGA on the VC707 development kit, featuring an Ethernet interface for tester control via the IPBus protocol [44]. The lpGBT tester has proved effective in the characterization of the lpGBT ASIC prototypes and in conducting the SEE and X-ray test campaigns. Furthermore, it was also used for production testing [45] of the ASIC (200k devices).

B. Qualification and Characterization

Comprehensive testing of the chip's functionalities has been conducted to ensure robust operation and adherence to the specifications across a range of conditions. Tests were carried out for the specified power supply voltage range, encompassing variations from -10% to $+10\%$ from the nominal value, and over a temperature range spanning from -20 °C to 50 °C. The chosen temperature range reflects the environmental conditions expected in the HEP detectors and experimental halls.

1) Data Transmission: The data transmission capabilities of the chip were rigorously evaluated across all modes of operation and no errors were observed. An illustrative electrical eye diagram, for the chip operating at 10.24 Gb/s, is provided in Fig. 6, demonstrating the chip's high-speed data transmission performance.

Correspondingly, the link operation in the downlink direction is confirmed through the recording of an eye diagram utilizing the built-in EOM, visually represented in Fig. 7.

2) Power Consumption: The optimization of power consumption stood as one of the primary considerations throughout the design and implementation of the chip. Given its highly versatile structure accommodating multimode operation, presenting a single power consumption figure is challenging. The ASIC power consumption varies substantially depending on its configuration. In the minimal setup (after reset and with all features disabled), at the nominal power supply voltage of 1.2 V, the power consumption is approximately 150 mW. For most typical transceiver configurations, the power consumption is below 300 mW. As an example, an lpGBT working in the transceiver mode (5G FEC5), with 14 ePortRx enabled at 320 Mb/s, four ePortTx at 320 Mb/s, four ePortClock at 320 MHz, and two PSCLKs at 320 MHz consumes 275 mW. In its maximal configuration, which encompasses the activation of all features (although unlikely in practical applications), the power consumption can reach 600 mW.

3) Clock Performance: The clock manager circuit underwent comprehensive characterization in both PLL and CDR modes. With a power consumption below 34 mW, the circuit exhibited integrated jitter (100 Hz–100 MHz) levels of 1.3 and 1.6 ps in PLL and CDR modes, respectively [13], [14], [40]. Simultaneously, the phase-shifter demonstrated full functionality, providing adjustable phase shifts with 48.4-ps resolution for all output frequencies. The phase-shifter's linearity was confirmed with integral nonlinearity (INL) below 0.61 LSB and differential nonlinearity (DNL) below 0.44 LSB [15].

4) Analog Functions: The analog subsystem underwent a comprehensive characterization to evaluate its performance across various parameters. The reference voltage generator circuit produced the specified 1 V consistently across power supply voltage and temperature ranges. The programmable gain amplifier demonstrated reliable performance within its specified range and the ADC showcased excellent linearity, yielding a static effective number of bits (ENOB) of approximately 9.6 bits. Both the current and voltage DACs performed as designed, with INL values staying below 0.5 and 1 LSB for the current and voltage DACs, respectively. The temperature sensor exhibited linearity in the range from -20 °C to 50 °C, with a measured sensitivity of 2.42 mV $°C^{-1}$. A comprehensive set of results was reported in [8] and [22].

C. SEE Test Results

The SEE tests were conducted at the heavy-ion facility (HIF) at UCL in Belgium. The beam offered by the facility is 25 mm in diameter with homogeneity of $\pm 10\%$, sufficient for the die size of 4.5 $\times$ 4.5 mm. The test campaign consisted of two phases: an initial phase to make sure that there were no functional interrupts, with the chip remaining operational even under the most stringent irradiation conditions; a second phase dedicated to determining the cross section (σ) and threshold LET (LET_{th}) of various functional units in the ASIC.

To irradiate the lpGBT, the plastic package over-mold was removed and the die was thinned down to 50 μm which is below the ion penetration range, allowing the ionization in the sensitive volume. The irradiation was done through the back side due to the need to maintain the ASIC in its BGA package.

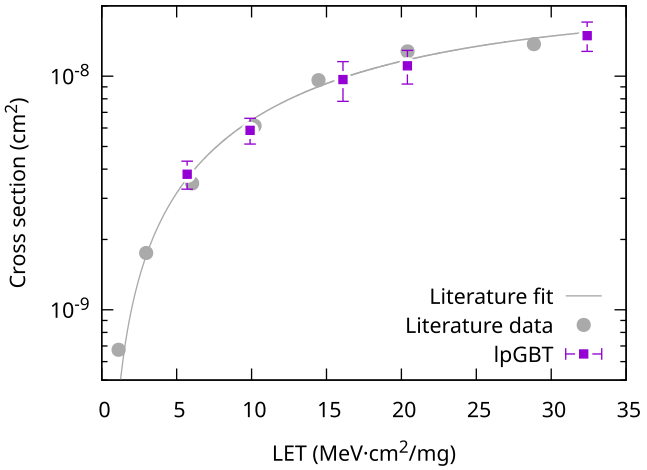

Fig. 8. Individual registers (non-TMR) cross section compared to results from [31].

The initial phase was run with Krypton (Kr) ions with LET of 32.4 MeV cm^2 mg^{-1} and the maximum flux provided at HIF 1.5×10^4 s^{-1} cm^{-2}. This phase extensively covered the power-up initialization and the I^{2}C communications with the ASIC. During that experiment 3×10^3 complete boot cycles were executed and more than 1×10^7 I^{2}C transactions were performed across an accumulated fluence of 5×10^7 cm^{-2}. No errors or functional interrupts were observed, resulting in a limit cross section of 2×10^{-8} cm^2 at an LET of 32.4 MeV cm^2 mg^{-1} for both processes. Throughout the full duration of the experiment, the lpGBT power consumption remained stable (no latch-up events were observed).

The next phase of the experiment was dedicated to the detailed characterization of the ASIC functions. The focus was on monitoring data transmission errors in both simplex modes and the transceiver mode using the different available uplink data rates and FEC codes. To optimize the use of time in the facility, the ePorts (ePortRx and ePortTx) were set to operate at a combination of data rates and operation modes. Moreover, during the experiment, the contents of all the configuration registers, the power consumption, and the output clock performance were regularly monitored. No TMR registers were found to have been corrupted by SEE.

1) ASIC Configuration: The absence of single-event upsets (SEUs) in the TMR registers confirmed the correct implementation of the TMR protection. It should be noted, however, that, as expected, the individual registers forming TMR register sets were in fact upset during the experiment. The lpGBT implements a global counter (TMR protected), which increments each time a TMR correction is made. This allows monitoring of the correction activity.

That counter was used to acquire statistics on the individual flip-flops cross section, which is represented in Fig. 8.

As a cross-check, the graph displays also data from a previous study [31] (gray points and corresponding line for Weibull fit), which characterized foundry-provided cells, which closely fits the data acquired during this study. To establish a reference point for this and the following SEE graphs, a scaled version of this curve will be used in each case. The scaling, equivalent number of D flip-flop (DFF) counted, will be such that the sensitive areas of the two circuits being compared (lpGBT circuit) and $N \times$ DFF are estimated to be the same.

Due to a low cross section of multiple circuits and the limited experimental time available at the test facility, it was impossible to gather sufficient statistics at lower LET. The limited statistics are reflected in the error bars associated with all measurement points. For the cases where no events were observed, the data point is omitted, but an error bar indicating the upper limit is included. Therefore, the same curve, or a scaled version of it, is used in the following cross section graphs as an indication of the anticipated cross section profile. The parameters of this fit were also used to better constrain the threshold and the shape of the Weibull fits required for extrapolation of the results to proton environments.

2) Clock Phase Stability: Previous research had identified PLL unlocks as a major contributor to prolonged periods of noncorrectable errors [46]. Given the critical role of the clocking subsystem in ensuring the reliability of data transmission, emphasis was placed on evaluating its functional radiation tolerance in both CDR and PLL modes. During the tests, no occurrences of PLL unlocks (phase or frequency errors leading to cycle slips) were detected in either operation mode, even when subjected to ionizing radiation with LET values as high as 32.4 MeV cm^2 mg^{-1}. This validates the effort put into the SEE-robust implementation of the clock generator. The value obtained for the limit cross section is below 2.5×10^{-8} and 1.7×10^{-7} cm^2 in the CDR and PLL modes, respectively.

The lpGBT acts at the clock reference for most of the HL-LHC subdetectors and it was thus important to evaluate the eClock jitter performance in radiation environments.

It has been observed that, induced by single-event transient (SET), the phase of the clock or data signals generated by the lpGBT may experience temporary deviations from their nominal values, particularly when operating in simplex transmitter mode with an external 40-MHz reference clock. Ionizing radiation interactions in the lpGBT clock generator can induce a temporary shift in phase for all output signals, affecting internal and external clocks, as well as clock-synchronous input and output signals like eLinks. This phase error was found to accumulate gradually over a few hundred reference clock periods due to a frequency error in the on-chip VCO. The worst case accumulated phase error observed in the heavy-ion experiments is around 400 ps [47], [48]. The lpGBT PLL eventually corrects the phase error and the phase returns to the lock value. The phase transients consistently have the same sign, with the on-chip VCO frequency temporarily increasing relative to the off-chip reference clock. The characteristics of the phase transients depend on the particle energy, its LET, and the PLL loop configuration, resulting in a distribution of response amplitudes. These phase transients are due to SETs in the inductor region of the LC oscillator, as detailed in [47] and [48].[2]

[2]The authors call attention to the fact that the explanation, proposed in the reference, for the mechanism inducing the phase-shifts was not confirmed by further investigations. Nonetheless, the observations reported are accurate.

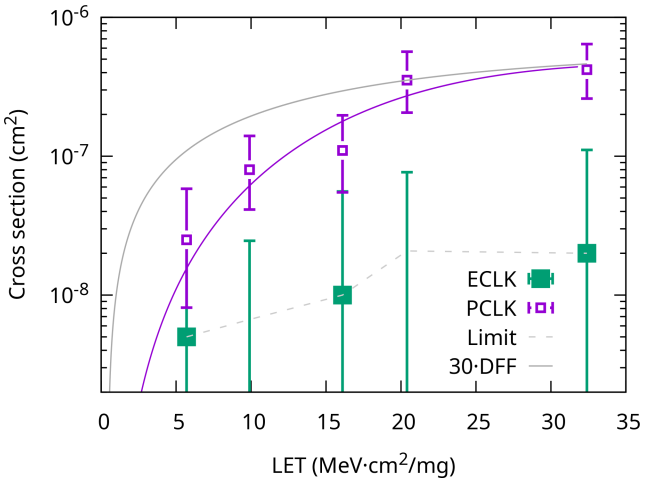

Fig. 9. Clock glitches cross section for phase-shifter clocks (PCLK) and eClocks (ECLK).

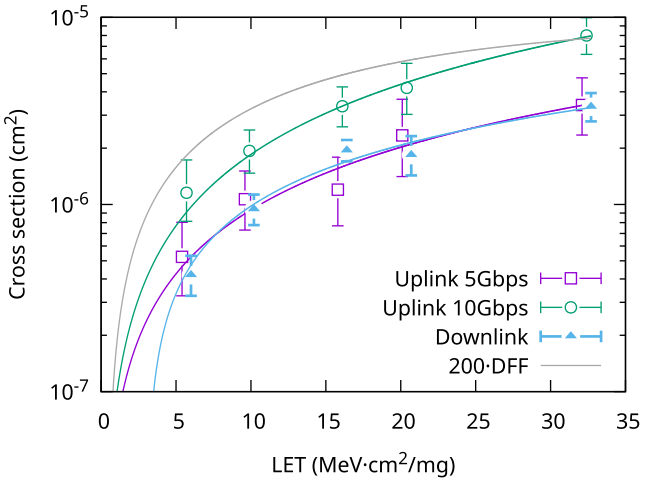

Fig. 10. Receiver and transmitter correctable errors cross sections.

The phase error buildup is either absent or smaller when the ASIC operates in the CDR mode, since the CDR loop bandwidth is much higher, which allows much faster correction of the stimulated frequency error. Those observations led to the recommendation of using the lpGBT in the CDR mode for timing applications. Although the jitter variance is slightly higher in the CDR mode, the probability of phase deviations exceeding 50 ps (peak-to-peak) is substantially reduced in the CDR mode.

Finally, the experiment searched for clock phase glitches that were defined as phase errors exceeding 100 ps and that could last for multiple clock cycles. Of the events detected, all were contained within a single 25-ns clock cycle.

The cross section of the clock glitches is presented in Fig. 9. The data show that the phase-shifter clocks are more susceptible to SEE than eClocks with only one event being observed for the eClocks during the entire run. The noticeably higher cross section observed for the phase-shifter clocks is attributed to the sensitivity in the full-custom part and the fact that the highest frequency digital stages of the phase-shifter circuit are not triplicated [15]. The higher threshold, when compared with the reference DFF, can be explained by the use of larger devices (e.g., ELT) in the phase-shifter circuit.

3) Data-Path Errors: As the lpGBT primarily functions as a data transmission ASIC, it is crucial to assess its susceptibility to data transmission errors caused by SEEs. Due to the lpGBT's architecture, these SEE-induced data transmission errors can be classified into two categories: correctable and noncorrectable errors.

Correctable errors are caused by SEEs that occur, for the transmitter, on the data path after the encoder and do not exceed the error correction capability of the uplink FEC code. Similarly for the receiver, these occur in the data path before the decoder and do not exceed the downlink FEC error correction capability. Correctable errors do not impair data transmission resulting in error-free communications. They can, however, be detected by monitoring the activity of the FEC decoders.

Noncorrectable errors are caused by SEEs that can occur anywhere in the data path and either exceed the error correction capability of the FEC code or occur in the non-TMR parts of the data path. Examples of blocks with non-TMR are the very front-end logic of the ePorts and the high-speed serializer and de-serializer circuits.

The test system allowed to feed or receive data to or from the lpGBT through the ePorts and the high-speed links. At any instant, it was possible to compare the received data to that fed to the lpGBT and thus count the number of errors occurring during data transmission. The cross section of the data transmission errors is calculated for each LET as the observed number of errors divided by the run fluence.

The cross section for correctable errors is collected together for the transmitter and the receiver in Fig. 10. It should be noted that the results presented for the transmitter were obtained in the FEC12 mode. However, the FEC mode has almost no influence on the results, and, therefore, the FEC5 results were omitted for clarity. It can be observed that the cross section for the transmitter at 5.12 Gb/s is roughly the same as that of the receiver while that of the transmitter at 10.24 Gb/s is twice as high as the transmitter at 5.12 Gb/s. This correlates very well with the frame size and therefore the number of flip-flops used in the serializer (5.12 and 10.24 Gb/s) and the de-serializer (the de-serializer uses additional latches to generate a phase shift). The results confirm that correctable errors originate in the unprotected part of the high-speed serializer and de-serializer, that is after the FEC encoder, for the transmitter, and before the FEC decoder for the receiver.

The cross section for noncorrectable SEE-induced downlink errors is presented in Fig. 11. All bit errors combined are marked with squares, per group eLink single-bit errors are marked with circles, frame corruptions are marked with triangles and the scaled DFF cross section (from Fig. 8) is added for reference.

Only two single-bit errors were observed during the full experimental run, one erroneous bit for each of the two lowest values of LET out of 16 active channels. In those cases, the lpGBT was set up to run the eLinks at 160 Mb/s in the mirror function configuration. In that configuration, pairs of eLinks are set to transmit the same data. For each single bit error

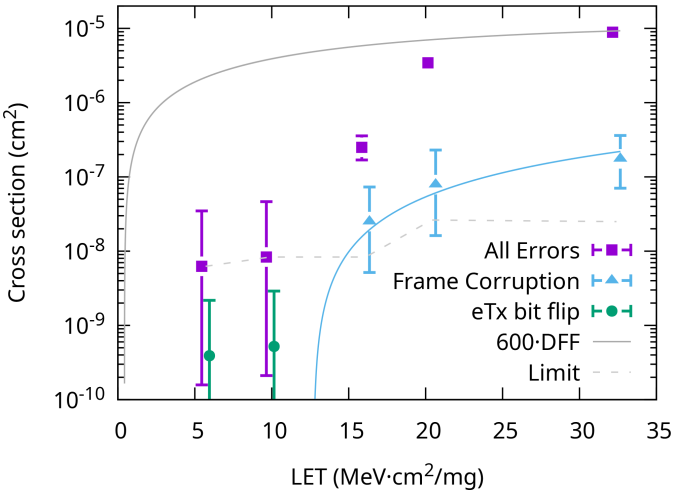

Fig. 11. Downlink noncorrectable bit errors cross section.

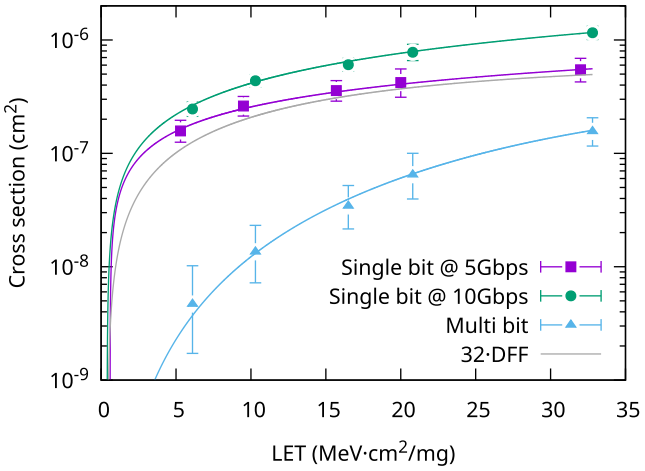

Fig. 12. ePortRxGroup errors cross section.

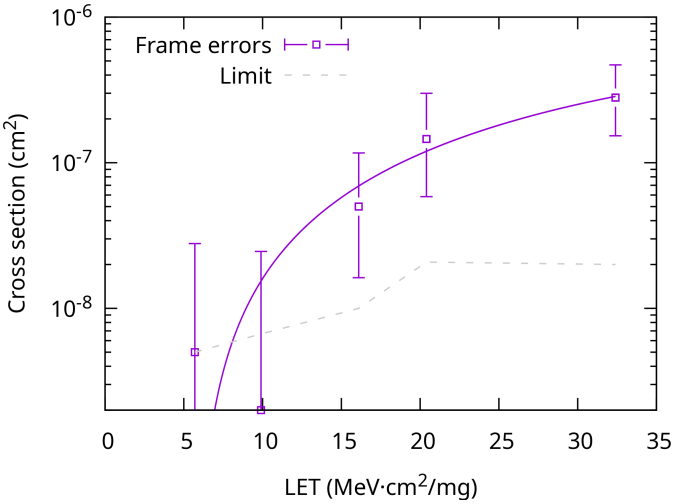

Fig. 13. Uplink frame errors cross section.

detected only one eLink in the pair was observed to have corrupted data. This means that the errors have occurred in the nontriplicated portion of the ePortTx, that is, the output driver. No errors have thus been detected on the ePortTx serializers that are fully TMR protected.

Frame errors are defined as errors that upset multiple bits in a frame and are detected as multiple eLinks having, simultaneously, corrupted data. As mentioned before, these occur when SEEs impact the receiver in a non-TMR part of the circuit (in this case, the de-serializer) and the disturbance exceeds the error correction capability of the FEC code. Given the high LET for which these events were detected (above 15 MeV cm^2 mg^{-1}) and the circuit architecture, it is likely that these are SETs on the clock signals of the de-serializer, which are not protected by TMR. Only a few events were detected during the full run of the experiment, resulting in low statistics.

In a parallel investigation of the downlink sensitivities, the SEE-induced uplink errors were scrutinized. No statistically significant difference was detected in the error cross section between transmitter and transceiver modes. Consequently, in the subsequent analysis, errors were integrated across all modes to enhance statistics. The intricacies of the observed errors, coupled with the presence of multiple modes of operation (specifically, variations in the uplink data rate), preclude the presentation of a comprehensive total cross section. Instead, the presentation of a selected set of detailed plots follows, to capture the diverse manifestations of SEE-induced errors in the uplink.

The cross section for SEE-induced data errors in the ePortRxGroup is presented in Fig. 12.

The cross section of single-bit errors for 10.24 Gb/s (green circles) is roughly twice as large as that for 5.12 Gb/s (purple squares). Both the choice of FEC code (FEC5/FEC12) and the transceiver mode (TX/TRX) have no (within the experimental uncertainty) influence on the cross section. This points to the single-bit error SEEs occurring in the ePortRx deserializer block. That is to be contrasted with the downlink case where no single bit SEE was attributed to the ePortTx. This observation is consistent with the fact that the data path ePortTx circuits are implemented using TMR while the ePortRx are not (due to the higher operation frequency, the TMR would result in a higher power penalty in the ePortRx than for the ePortTx). This was a design decision since it was considered more important to protect the downlink control and timing data than the physics data sent to the off-detector systems. Multibit errors (triangles) are errors that occur in multiple bits associated with the same ePortRx. Their cross section is approximately ten times lower than that of the single-bit errors, with a higher threshold. Again, there is no clear dependency with the FEC code used and they are likely due to SEEs on the clock buffers within a deserializer block of ePortRxGroup. It is important to note that no persistent errors were identified in the data stream, confirming the efficacy of the implemented SEE mitigation in control state machines.

Another class of uplink errors, presented in Fig. 13, are frame errors. These errors are flagged when multiple bits corresponding to multiple ePortRx are detected to be corrupted simultaneously. These are thought to occur in the high-speed serializer due to an SET on one of the clocks, leading to multibit upsets overwhelming the correction capability of the FEC code. The cross section for these events is minimized by using large fan-out buffers. Moreover, it should be noted

TABLE II
SEE ESTIMATES FOR THE HL-LHC ENVIRONMENT

	Effect	20 MeV Hadrons Cross Section (cm^2)	CMS tracker rate (Hz)
Conf.	Reset	Not observed for an LET of 32.4 MeV cm^2 mg^{-1}	
	Corruption	and a fluence of 5×10^7 cm^{-2}	
	TMR correction	9.2×10^{-11}	1.4×10^{-3}
Clock	PLL unlock	Not observed for an LET of 32.4 MeV cm^2 mg^{-1}	
	CDR unlock	and a fluence of 5×10^7 cm^{-2}	
	SET phase-shifter	2.7×10^{-14}	4.3×10^{-7}
	SET eClock	3.8×10^{-15}	6.0×10^{-8}
Downlink	Correctable error	8.6×10^{-13}	1.4×10^{-5}
	Bit error (80 Mbit s^{-1})	Not available	
	Bit error (160 Mbit s^{-1})	Not available	
	Bit error (320 Mbit s^{-1})	Not available	
	Word error	Not observed	
	Frame corruption	5.0×10^{-15}	8.0×10^{-8}
Uplink	Correctable error (5G)	2.2×10^{-12}	3.5×10^{-5}
	Correctable error (10G)	3.8×10^{-12}	6.0×10^{-5}
	Bit error (160 Mbit s^{-1})	1.8×10^{-13}	2.9×10^{-6}
	Bit error (320 Mbit s^{-1})	3.6×10^{-13}	5.7×10^{-6}
	Bit error (640 Mbit s^{-1})	7.3×10^{-13}	1.2×10^{-5}
	Bit error (1.28 Gbit s^{-1})	1.2×10^{-12}	1.9×10^{-5}
	Word error	1.0×10^{-14}	1.6×10^{-7}
	Frame corruption	6.5×10^{-15}	1.0×10^{-7}

that the clocks originate from the fully triplicated divider and, therefore, no long-lasting effects are expected. This approach enabled the optimization for the maximum timing margin required to account for anticipated TID-induced transconductance degradation.

4) SEE Estimates for the HL-LHC Environments: The lpGBT chip will be used across various subdetectors at the LHC, with the radiation field characteristics—particle types, energy, and flux—strongly dependent on each subdetector's position relative to the interaction point. Due to the optical component limitations discussed earlier, the lpGBT chips will not be operated in the innermost detector regions. Nonetheless, they will be exposed to significant radiation levels in many systems. The following analysis addresses the worst case scenario, exemplified by the CMS Outer Tracker [49]. In this subdetector, during 13-TeV p-p collisions, the innermost modules will be exposed to the highest particle fluxes, potentially reaching 1.6×10^7 $cm^{-2}\,s^{-1}$ for hadrons with energy exceeding 20 MeV [50]. Given that the SEU rates in CMS are dominated by such particles, as shown by Huhtinen and Faccio [25], the SEE rate evaluation adopts the approach proposed in this work. The estimated SEU cross section per lpGBT for hadrons above 20 MeV and the corresponding SEU rate in the hottest region of the detector are summarized in Table II.

The anticipated operational stability is a key feature of the lpGBT, with no expected occurrences of resets, configuration corruptions, or clock generator unlocks in both CDR and PLL modes. Infrequent glitches are projected, occurring approximately once every 27 days for the phase-shifter clocks and every 193 days for the eClocks.

The potential for downlink frame corruption is estimated to be of the order of one every 145 days. While the data collected for single-bit errors lacks sufficient information for extrapolation, the absence of events at higher LET levels suggests a cross section below that obtained for frame corruptions. Uplink bit corruptions are expected to be predominantly influenced by SEUs in the deserializer, with errors anticipated every several tens of hours, resulting in a bit error ratio of the order of 1×10^{-14}.

TABLE III
TID SAMPLES AND RESULTS

Sample	Mode	Die Temp. (°C)	Total Dose (MGy)	Threshold Dose for Uplink FEC Corrections (MGy)	Threshold Dose for Uplink Data Errors (MGy)
1	10G FEC5	30	3.4	1.4	2.1
2	10G FEC5	30	4.0	2.0	1.9
3	5G FEC5	30	3.5	2.7	3.1
4	5G FEC5	30	4.0	2.85	3.2
5	10G FEC5	-5	5.6	3.85	3.2
6	10G FEC5	0	6.2	4.15	3.25
7	10G FEC5	30	0.05	-	-
8	10G FEC5	30	0.05	-	-

The results obtained from the heavy-ion test campaign were not directly utilized to predict the magnitude and frequency of phase jumps expected in the clock generator. This decision was primarily driven by the novelty of the observed effects. Instead, a dedicated proton campaign was organized to comprehensively investigate this phenomenon. Through this proton experiment, the maximum observed phase excursion remained below 40 ps, and the corresponding cross section was estimated to be of the order of 1×10^{-13} cm^2 as reported in previous works [47], [48]. This cross section implies that, on average, one event of such magnitude is expected to occur once per week per ASIC in the CMS tracker environment, while events of smaller amplitude will occur more frequently.

D. TID Results

X-ray irradiation was performed on eight samples of the lpGBT under the conditions summarized in Table III. As displayed, irradiation was performed at several die temperatures and up to several TID doses, followed by different annealing procedures. The results, which will be discussed next, are presented for the preproduction devices. Furthermore, ten devices, randomly selected among the full production series, were irradiated following the same irradiation procedure with no significant differences found among the devices.

The results discussed below concentrate on the digital and data transmission functionality and they either try to establish a limit for the TID dose the ASIC can sustain while being fault-free or address a specific HL-LHC experiment environment. One needs, however, to consider that the TID degradation of the CMOS process used for lpGBT manufacturing is highly dependent on dose rate, die temperature, and bias conditions. The large variability of lpGBT use cases (expected lifetime dose, operational temperature, bias conditions) makes a thorough qualification of the device virtually impossible. The investigation conducted by Faccio et al. [28] revealed a substantial shift in threshold voltage for pMOS transistors subjected to high-temperature annealing under bias conditions. Consequently, it is strongly recommended to avoid exposure

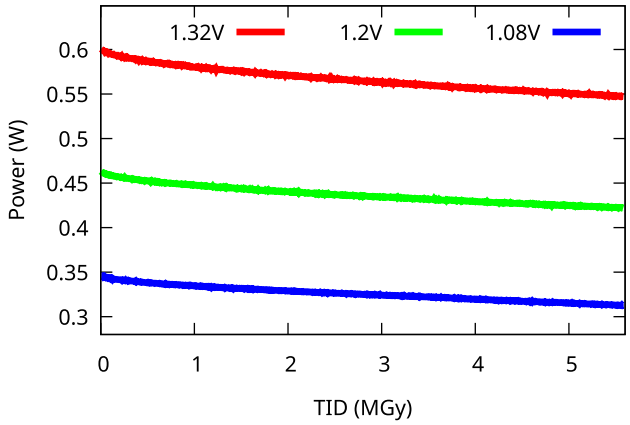

Fig. 14. ASIC power consumption versus TID for various power supply levels.

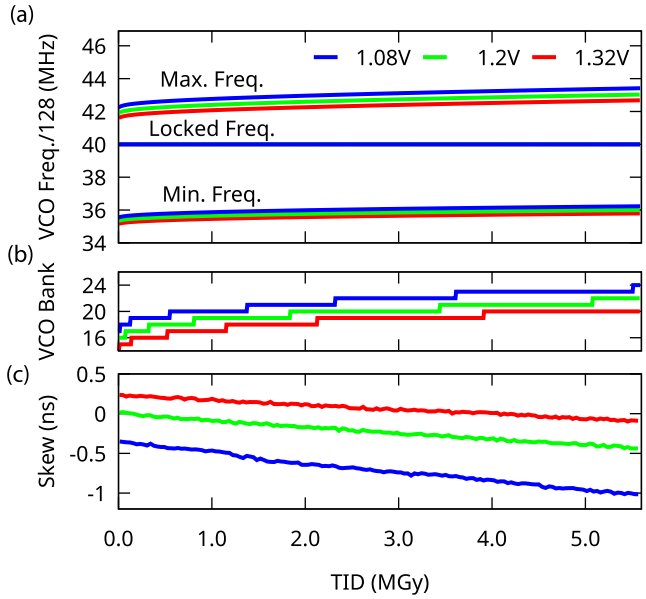

Fig. 15. (a) VCO tuning range. (b) Automatically selected VCO bank. (c) eClock0 skew normalized to prerad value versus TID.

of the ASICs developed in this technology, including the lpGBT, to elevated temperatures under bias after irradiation.

Similar to the heavy-ion irradiation, the plastic package over-mold was removed exposing the back side of the die. However, in this case, no thinning of the die was required but a dose correction factor was applied to account for the die thickness. An X-ray source at CERN was used operating at the maximum recommended power with the ASICs irradiated at the relatively high dose rate of 25 kGy h^{-1}. Constant monitoring ensured the ASIC's correct operation during irradiation. The chip's performance, assessed at power supply voltages ranging from 1.08 to 1.32 V, underwent scrutiny across approximately 500 metrics per test. These metrics encompassed vital parameters such as power consumption, bit error ratio, phase stability, PLL performance, digital functions and interfaces, e-fuse readout, and characterization of analog functions. This thorough evaluation provided a comprehensive understanding of the chip's behavior under irradiation and different supply voltages.

1) ASIC Configuration: In all cases, and up to the maximum dose used, no failures were detected in the configuration memory (including its checksum verification mechanism), the EC and IC channels, the I^2C interfaces, and the power-up state machine (chip booting procedure).

However, the e-fuses showed failures which, in some cases, revealed themselves early on during the irradiation. As a consequence, the use of the e-fuses to store the lpGBT configuration is not recommended.

2) Power Consumption: Fig. 14 shows the evolution of the power consumption with TID at the three test supply voltages.

It can be observed that for all the cases the power consumption decreases with TID. This effect is, however, mild when compared with the impact the supply voltage has on power consumption. The decrease in power consumption is consistent with the degradation of g_m that individual transistors are known to have with TID [29], [30] which, in the lpGBT case, mostly impacts the power consumption of the digital gates, I/O drivers, receivers and the CML circuits. This hypothesis is also coherent with the decrease observed in the oscillation frequency of the four process monitor ring oscillators included in the lpGBT.

3) Clocking: The typical performance of the clocking system is illustrated in Fig. 15.

During irradiation, a minor shift in the free-running VCO frequency was observed. Notably, the VCO tuning range consistently remained well-centered around the LHC reference frequency of 40 MHz, ensuring proper PLL lock. This shift is automatically compensated for by the capacitor bank selection process, occurring automatically at each power-up [see Fig. 15(b)]. Simultaneously, the stability of the phase for eClocks was closely monitored, as depicted in Fig. 15(c). Gradual radiation-induced skew changes were observed, attributed to clock buffers distributing the signal across the chip but lying outside the PLL feedback loop. These gradual changes are expected to be periodically calibrated in the final experiment. The effectiveness of the proposed source-switched charge pump [51] was validated, as no abrupt jumps caused by VCO bank selection were detected during the experiment.

4) Data Transmission Errors: The data transmission in both the uplink and downlink directions was meticulously monitored throughout the experiments. To maximize efficiency in the facility, the ePorts (ePortRx and ePortTx) were configured to operate at various data rates, coupled with different phase selection modes per ePortRx. The first instances of data transmission errors in the uplink were identified at a low power supply voltage of 1.08 V. At moderately elevated temperatures (around 0 °C), exceeding those expected in radiation-challenging environments like trackers, errors manifested at approximately 3.2 MGy (refer to Table III). Notably, at the nominal power supply voltage of 1.2 V, errors were observed over 1 MGy later. The chips irradiated at higher temperatures exhibited earlier failures, aligning with the transistor-level characterization findings [29]. Chips operating at lower data rates (5.12 Gb/s) tolerated higher doses than those at higher rates (10.24 Gb/s). No significant disparities

were observed between FEC5 and FEC12 modes or PLL and CDR modes. Importantly, in none of the experiments were any data transmission errors detected in the downlink direction.

5) Analog Functions: Comprehensive monitoring of all analog blocks was conducted throughout the irradiation process. Detailed results regarding the irradiation of analog functions have been extensively discussed in prior publications [8], [22] and will not be reiterated here for brevity. For completeness, it is noteworthy that the radiation-induced reference voltage shift remains below 18 mV (1.8%) across all irradiation samples. The ADC and DACs maintained full functionality, with most characteristic changes directly attributable to the aforementioned reference voltage shift. However, the temperature sensor exhibited heightened sensitivity to radiation, rendering it unreliable for its intended purpose—monitoring the ASIC's operating temperature. Consequently, its usage is recommended solely for detecting transient or trends in temperature changes induced, for instance, by a malfunction in the cooling system.

IV. Conclusion

In this work, the lpGBT, a low-power radiation-hard multipurpose high-speed transceiver ASIC for HEP experiments, was presented together with the analysis of the data collected during the qualifying irradiation test campaigns. The experimental data were further used to estimate the expected upset rates in the CMS tracker operating at HL-LHC. The absence of configuration upsets and PLL unlock events during testing, alongside the projected low bit error rate for the considered HEP detector environments, validate the strategies implemented for SEE mitigation.

The TID tolerance of the device, based on the experimental data, is estimated to be 2 MGy and exceeds the design target (1 MGy). Approximately 270,000 devices have now been produced and will be installed in HL-LHC detector systems, scheduled to take place between 2026 and 2028 at CERN. Quality control of the radiation-hardness of the production devices was done on a sample basis by running TID tests on samples extracted from each wafer lot.

The lpGBT is primarily a communications ASIC but also incorporates analog and digital functions targeted at probing the detector environment and enabling experiment control. Because of its versatility, it has become a fundamental building block in the construction of all HL-LHC detectors.

Acknowledgment

The authors extend their sincere gratitude to Miroslav Havranek and Tomas Benka from Czech Technical University, Prague, for contributing to the DAC macro cell. They also acknowledge the valuable input of Stefano Michelis from CERN, Geneva, in the development of the bandgap. Additionally, they thank Alessandro Caratelli, also from CERN, Geneva, for providing the RTL implementation of the I^2C Main.

References

[1] E. H. M. Heijne, "Particle physics experiments: From photography to integrated circuits," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 1055, Oct. 2023, Art. no. 168466.

[2] P. Moreira and S. Kulis, "Radiation-hard ASICs for data transmission and clock distribution in high energy physics," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 1053, Aug. 2023, Art. no. 168364. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0168900223003546

[3] J. Troska et al., "The versatile transceiver proof of concept," in *Proc. Topical Workshop Electron. Part. Phys.*, Paris, France, 2009, pp. 347–351. [Online]. Available: https://cds.cern.ch/record/1235837

[4] P. Moreira et al., "The GBT project," in *Proc. Top. Workshop Electron. Part. Phys.*, Paris, France, Sep. 2009, pp. 342–346. [Online]. Available: https://cds.cern.ch/record/1235836

[5] M. Paulo. (2021). *GBTX Manual; V0.18*. CERN, Switzerland. [Online]. Available: https://cds.cern.ch/record/2809057

[6] C. Atlas, "Letter of intent for the phase-II upgrade of the ATLAS experiment," CERN, Geneva, Switzerland, Tech. Rep. CERN-LHCC-2012-022, 2012. [Online]. Available: https://cds.cern.ch/record/1502664/files/LHCC-I-023.pdf?version=7

[7] D. Contardo, M. Klute, J. Mans, L. Silvestris, and J. Butler, "Technical proposal for the phase-II upgrade of the CMS detector," Univ. Minnesota, Geneva, Switzerland, Tech. Rep. CERN-LHCC-2015-010 Detectors and Experimental Techniques, 2015. [Online]. Available: https://cds.cern.ch/record/2020886

[8] P. Moreira et al. (2022). *LpGBT Documentation: Release*. CERN, Geneva, Switzerland. [Online]. Available: https://cds.cern.ch/record/2809058

[9] M. Menouni, P. Gui, and P. Moreira, "The GBTIA, a 5 Gbit/s radiation-hard optical receiver for the SLHC upgrades," in *Proc. Topical Workshop Electron. Part. Phys.* Geneva, Switzerland: CERN, 2009, pp. 326–330. [Online]. Available: http://cds.cern.ch/record/1235833

[10] M. Menouni et al., "The lpGBTIA, a 2.5 Gbps radiation-tolerant optical receiver using InGaAs photodetector," *PoS*, vol. 2019, p. 30, Sep. 2019. [Online]. Available: https://cds.cern.ch/record/2724953

[11] J. Troska et al., "The VTRx+, an optical link module for data transmission at HL-LHC," *PoS*, vol. 17, p. 48, Sep. 2017. [Online]. Available: https://pos.sissa.it/313/048

[12] J. Prinzie, J. Christiansen, P. Moreira, M. Steyaert, and P. Leroux, "A 2.56-GHz SEU radiation hard LC-tank VCO for high-speed communication links in 65-nm CMOS technology," *IEEE Trans. Nucl. Sci.*, vol. 65, no. 1, pp. 407–412, Jan. 2018.

[13] J. Prinzie et al., "A low noise fault tolerant radiation hardened 2.56 Gbps clock-data recovery circuit with high speed feed forward correction in 65 nm CMOS," in *Proc. IEEE 10th Latin Amer. Symp. Circuits Syst. (LASCAS)*, Feb. 2019, pp. 63–66.

[14] S. Biereigel et al., "The lpGBT PLL and CDR architecture, performance and SEE robustness," *PoS*, vol. 2019, p. 34, Sep. 2019. [Online]. Available: https://cds.cern.ch/record/2724954

[15] S. Kulis et al., "A high-resolution, wide-range, radiation-hard clock phase-shifter in a 65 nm CMOS technology," in *Proc. 26th Int. Conf. Mixed Design Integr. Circuits Syst.*, 2019, pp. 147–150.

[16] S. Lin and D. J. Costello, *Error Control Coding*. Upper Saddle River, NJ, USA: Prentice-Hall, 2004.

[17] A. Marchioro and P. Moreira, "An error-correcting line code for a hep rad-hard multi-gigabit optical link," in *Proc. 12th Workshop Electron. LHC Future Exp.*, 2007, p. 5. [Online]. Available: https://cds.cern.ch/record/1027475

[18] *Information Technology—Telecommunications and Information Exchange Between Systems—High-Level Data Link Control (HDLC) Procedures*, Standard ISO/IEC JTC 1/SC 6, Telecommunications and information exchange between system, Jul. 2002.

[19] Z. Zeng et al., "A compact low-power driver array for VCSELs in 65-nm CMOS technology," *IEEE Trans. Nucl. Sci.*, vol. 64, no. 6, pp. 1599–1604, Jun. 2017.

[20] Z. Zeng, T. Zhang, G. Wang, P. Gui, S. Kulis, and P. Moreira, "LDQ10: A compact ultra low-power radiation-hard 4 × 10 Gb/s driver array," *J. Instrum.*, vol. 12, no. 2, Feb. 2017, Art. no. P02020, doi: 10.1088/1748-0221/12/02/p02020.

[21] (Oct. 2021). *I2C-Bus Specification and User Manual*. NXP Semiconductors Std. Rev. 7.0. [Online]. Available: https://www.nxp.com/docs/en/user-guide/UM10204.pdf

[22] M. Firlej et al., "An lpGBT subsystem for environmental monitoring of experiments," *J. Instrum.*, vol. 18, no. 6, Jun. 2023, Art. no. P06008, doi: 10.1088/1748-0221/18/06/p06008.

[23] CMS Collaboration. (2017). *The Phase-2 Upgrade of the CMS Tracker*. CERN, Geneva, Switzerland. [Online]. Available: https://cds.cern.ch/record/2272264

[24] ATLAS Collaboration. (2017). *Technical Design Report for the ATLAS Inner Tracker Pixel Detector*. CERN, Geneva, Switzerland. [Online]. Available: https://cds.cern.ch/record/2285585

[25] M. Huhtinen and F. Faccio, "Computational method to estimate single event upset rates in an accelerator environment," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 450, no. 1, pp. 155–172, Aug. 2000. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0168900200001558

[26] F. Vasey et al., "The versatile link common project: Feasibility report," *J. Instrum.*, vol. 7, no. 1, Jan. 2012, Art. no. C01075, doi: 10.1088/1748-0221/7/01/c01075.

[27] F. Faccio et al., "Radiation-induced short channel (RISCE) and narrow channel (RINCE) effects in 65 and 130 nm MOSFETs," *IEEE Trans. Nucl. Sci.*, vol. 62, no. 6, pp. 2933–2940, Dec. 2015.

[28] F. Faccio et al., "Influence of LDD spacers and H^+ transport on the total-ionizing-dose response of 65-nm MOSFETs irradiated to ultra-high doses," *IEEE Trans. Nucl. Sci.*, vol. 65, no. 1, pp. 164–174, Jan. 2018.

[29] G. Borghello et al., "Ionizing radiation damage in 65 nm CMOS technology: Influence of geometry, bias and temperature at ultra-high doses," *Microelectron. Rel.*, vol. 116, Jan. 2021, Art. no. 114016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0026271420309033

[30] G. Borghello et al., "Effects of bias and temperature on the dose-rate sensitivity of 65-nm CMOS transistors," *IEEE Trans. Nucl. Sci.*, vol. 68, no. 5, pp. 573–580, May 2021.

[31] S. Bonacini et al., "Characterization of a commercial 65 nm CMOS technology for SLHC applications," *J. Instrum.*, vol. 7, no. 1, Jan. 2012, Art. no. P01015, doi: 10.1088/1748-0221/7/01/p01015.

[32] W. J. Snoeys, T. A. P. Gutierrez, and G. Anelli, "A new NMOS layout structure for radiation tolerance," in *Proc. IEEE Nucl. Sci. Symp. Conf. Rec.*, vol. 2, Nov. 2001, pp. 822–826.

[33] L. M. J. Casas et al., "Study of total ionizing dose effects in 65 nm digital circuits with the DRAD digital RADiation test chip," in *Proc. 17th Eur. Conf. Radiat. Its Effects Compon. Syst. (RADECS)*, 2017, pp. 1–6.

[34] L. M. J. Casas et al., "Characterization of radiation effects in 65 nm digital circuits with the DRAD digital radiation test chip," *J. Instrum.*, vol. 12, no. 2, Feb. 2017, Art. no. C02039, doi: 10.1088/1748-0221/12/02/c02039.

[35] S. Kulis, "Single event effects mitigation with TMRG tool," *J. Instrum.*, vol. 12, no. 1, Jan. 2017, Art. no. C01082, doi: 10.1088/1748-0221/12/01/c01082.

[36] I. S. Reed and G. Solomon, "Polynomial codes over certain finite fields," *J. Soc. Ind. Appl. Math.*, vol. 8, no. 2, pp. 300–304, Jun. 1960, doi: 10.1137/0108018.

[37] S. Biereigel, S. Kulis, P. Leroux, P. Moreira, and J. Prinzie, "Radiation-tolerant digitally controlled ring oscillator in 65-nm CMOS," *IEEE Trans. Nucl. Sci.*, vol. 69, no. 1, pp. 17–25, Jan. 2022.

[38] S. Biereigel, S. Kulis, P. Moreira, A. Kölpin, P. Leroux, and J. Prinzie, "Radiation-tolerant all-digital PLL/CDR with varactorless LC DCO in 65 nm CMOS," *Electronics*, vol. 10, no. 22, p. 2741, Nov. 2021. [Online]. Available: https://www.mdpi.com/2079-9292/10/22/2741

[39] J. Prinzie, J. Christiansen, P. Moreira, M. Steyaert, and P. Leroux, "Comparison of a 65 nm CMOS ring- and LC-oscillator based PLL in terms of TID and SEU sensitivity," *IEEE Trans. Nucl. Sci.*, vol. 64, no. 1, pp. 245–252, Jan. 2017.

[40] S. Biereigel et al., "A low noise fault tolerant radiation hardened 2.56 Gbps clock-data recovery circuit with high speed feed forward correction in 65 nm CMOS," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 67, no. 5, pp. 1438–1446, May 2020.

[41] J. Prinzie, M. Steyaert, P. Leroux, J. Christiansen, and P. Moreira, "A single-event upset robust, 2.2 GHz to 3.2 GHz, 345 fs jitter PLL with triple-modular redundant phase detector in 65 nm CMOS," in *Proc. IEEE Asian Solid-State Circuits Conf. (A-SSCC)*, Nov. 2016, pp. 285–288.

[42] A. Pulli, I. Kremastiotis, and S. Kulis, "Verification methodology of a multi-mode radiation-hard high-speed transceiver ASIC," *J. Instrum.*, vol. 17, no. 3, Mar. 2022, Art. no. C03008. [Online]. Available: https://cds.cern.ch/record/2803756

[43] J. M. Mendez et al., "LpGBT Tester: An FPGA based test system for the lpGBT ASIC," in *Proc. Topical Workshop Electron. Part. Phys. (TWEPP)*, Santiago de Compostela, Spain, 2019, p. 59. [Online]. Available: https://doi.org/10.22323/1.343.0059

[44] C. G. Larrea et al., "IPbus: A flexible Ethernet-based control system for xTCA hardware," *J. Instrum.*, vol. 10, no. 2, Feb. 2015, Art. no. C02019, doi: 10.1088/1748-0221/10/02/c02019.

[45] N. Guettouche et al., "The lpGBT production testing system," *J. Instrum.*, vol. 17, no. 3, Mar. 2022, Art. no. C03040, doi: 10.1088/1748-0221/17/03/c03040.

[46] P. Leitao et al., "Test bench development for the radiation hard GBTX ASIC," *J. Instrum.*, vol. 10, no. 1, Jan. 2015, Art. no. C01038, doi: 10.1088/1748-0221/10/01/c01038.

[47] S. Biereigel, S. Kulis, P. Leroux, P. Moreira, A. Kolpin, and J. Prinzie, "Single-event effect responses of integrated planar inductors in 65-nm CMOS," *IEEE Trans. Nucl. Sci.*, vol. 68, no. 11, pp. 2587–2597, Nov. 2021.

[48] S. Biereigel, "Radiation-tolerant all-digital clock generators for high energy physics," Ph.D. dissertation, BTU Cottbus-Senftenberg & KU Leuven, Aug. 2022. [Online]. Available: https://cds.cern.ch/record/2834952

[49] (2017). *The Phase-2 Upgrade of the CMS Tracker*. CERN, Geneva, Switzerland. [Online]. Available: https://cds.cern.ch/record/2272264

[50] CMS Collaboration. (2023). *Radiation Simulation Plotting Tool V.2.0*. CERN. [Online]. Available: https://cms-project-fluka-flux-map.web.cern.ch/

[51] J. Prinzie, S. Biereigel, S. Kulis, P. Leitao, and P. Moreira, "Source switched charge-pump PLLs for high-dose radiation environments," *IEEE Trans. Nucl. Sci.*, vol. 70, no. 4, pp. 590–595, Apr. 2023.