

RECEIVED: December 11, 2021

REVISED: January 28, 2022

ACCEPTED: February 26, 2022

PUBLISHED: March 21, 2022

A Starry Byte — proton beam measurements of single event upsets and other radiation effects in ABCStar ASIC Versions 0 and 1 for the ITk strip tracker

M.J. Basso,^{a,*} J. Fernández-Tejero,^{b,c} B.J. Gallop,^d G. Greig,^b J.J. John,^e P.T. Keener,^f K. Krizka,^g P.V. Leitao,^e B. Norman,^h P.W. Phillips,^d L. Poley,^{b,c,*} C. Sawyer,^d T.L. Stack,^c S. Stucci,ⁱ D.A. Trischuk^j and M. Warren^k

^aDepartment of Physics, University of Toronto,
Saint George St., Toronto, Canada

^bDepartment of Physics, Simon Fraser University,
University Drive W, Burnaby, Canada

^cTRIUMF,
Wesbrook Mall, Vancouver, Canada

^dParticle Physics Department, STFC Rutherford Appleton Laboratory,
Harwell Science and Innovation Campus, Didcot, United Kingdom

^eExperimental Physics Department, CERN,
Geneva, Switzerland

^fDepartment of Physics and Astronomy, University of Pennsylvania,
South 33rd Street, Philadelphia, U.S.A.

^gLawrence Berkeley National Laboratory,
Cyclotron Road, Berkeley, U.S.A.

^hDepartment of Physics, Carleton University,
Colonel By Drive, Ottawa, Canada

ⁱBrookhaven National Laboratory,
Rochester Street, Upton, United States of America

^jDepartment of Physics, University of British Columbia,
Agricultural Road, Vancouver, Canada

^kDepartment of Physics and Astronomy, University College London,
Gower Street, London, United Kingdom

E-mail: mbasso@physics.utoronto.ca, lpoley@triumf.ca

*Corresponding author.

ABSTRACT: Single Event Effects (SEEs) — predominately bit-flips in electronics caused by particle interactions - are a major concern for ASICs operated in high radiation environments such as ABCStar ASICs, which are designed to be used in the future ATLAS ITk strip tracker. The chip design is therefore optimised to protect it from SEEs by implementing triplication techniques such as Triple Modular Redundancy (TMR).

In order to verify the radiation protection mechanisms of the chip design, the cross-section for Single Event Upsets (SEUs), a particular class of SEEs, is measured by exposing the chip to high-intensity particle beams while monitoring it for observed SEUs.

This study presents the setup, the performed measurements, and the results from SEU tests performed using the latest version of the ABCStar ASIC (ABCStar V1) using a 480 MeV proton beam.

KEYWORDS: Radiation-hard electronics; Front-end electronics for detector readout; Radiation damage to electronic components; Radiation-hard detectors

Contents

1	Introduction	1
2	ABCStar V1	2
3	Measurement setup	4
3.1	Setup during measurements	7
4	Measurements	9
5	Results from proton tests	11
5.1	Firmware TESTOUT counters	11
5.2	HPR reset and unlock rates	12
5.2.1	August results	13
5.2.2	December results	14
5.3	HPR and register read status bits	14
5.3.1	August results	16
5.3.2	December results	16
5.4	SEU cross-section for physics packet clusters	16
5.4.1	August results	19
5.4.2	December results	20
5.4.3	Combined results and estimated impact on operation	21
5.5	SEU cross-section for register read data	22
5.5.1	August results	22
5.5.2	December results	23
5.6	Current measurements	24
6	Complementary results from heavy ion measurements	26
7	Conclusion	29
A	Dose and fluence measurements from proton tests	30

1 Introduction

For the High-Luminosity Upgrade of the Large Hadron Collider [1], the ATLAS Inner Detector [2] will be replaced with the new ATLAS Inner Tracker (ITk) as part of the ATLAS Phase-II Upgrade [3]. The ITk will be an all-silicon tracker and consist of an inner pixel tracker [4] and an outer strip tracker [5].

Modules for the ITk strip tracker consist of silicon strip sensors diced from 6 inch wafers, onto which flexes, known as hybrids and powerboards, are glued [6]. Each sensor strip is read out by a single channel of an ATLAS Binary Chip (ABC) ASIC. ABC ASICs were developed in several iterations (ABCN-25, ABC130, ABCStar [7]) to support several generations of module prototyping programme [6, 8].

The latest version of the ABCStar ASIC, the ABCStar Version 1 (V1), was developed to be used for the construction of the ATLAS detector and designed to be able to operate for an accumulated ionising dose of up to 660 kGy(Si) (measured up to 700 kGy(Si)). A major consideration for the design revision of the ABCStar Version 0 (V0) was the protection against Single Event Effects (SEEs) by using triplication of registers, clocks, and resets. In order to verify the improved radiation protection mechanisms of the V1 ASIC and its readiness for production, extensive measurements were performed to quantify the cross-sections for Single Event Upsets (SEUs), a class of SEEs corresponding to bit-flips ($0 \rightarrow 1$ or $1 \rightarrow 0$), in two proton beam tests. These measurements, which are described in the following sections, comprise the first published results of SEU tests conducted on ABCStar V0 and V1 ASICs.

2 ABCStar V1

The ABCStar ASIC is fabricated using 130 nm CMOS technology. It is designed to read out signals from 256 sensor strips through binary readout channels. The ABCStar front-end channel is optimised for an input capacitance of 5 pF, corresponding to the length of sensor strips on long-strip barrel modules (i.e., 5 cm), at a 25 ns readout time.

After passing a discriminator, the binary outputs from each front-end channel are sampled and stored in the chip's L0Buffer (effectively a pipeline) for the duration of the programmed latency. The ABCStar architecture supports a multi-trigger data flow:

- Level-0 (L0A) trigger
A local version of the Level-0 Accept trigger (and associated "L0ID") used in the experiment. Event data is transferred from the beam crossing synchronous pipeline, the L0Buffer.
- Low Priority (LP) trigger
A first asynchronous read-out request with priority and low latency for fast processing of track information.
- Priority Request (PR) trigger
A second asynchronous read-out request (called PR trigger) intended for a global read-out.

When an L0A is received, the event at the end of the L0Buffer pipeline is copied to the EventBuffer memory. The L0A command includes an identification number (called a "L0tag", copied from the lowest 7 bits of the L0ID) that is used as the address of the event in the EventBuffer. An LP or PR command will supply a tag that is used to retrieve an event from the EventBuffer memory for processing by the cluster finder. The cluster finder creates clusters of channels where hits were detected, which are then formatted into packets together with the event identification (e.g. L0tag) within the read-out block. A command decoder block receives and distributes trigger signals internally,

as well as receiving configuration commands, register read-back instructions, and ASIC settings (e.g. mask and threshold settings). A full description of the ASIC can be found in reference [5].

Compared to the ABCStar V0, the SEU protection of the ABCStar V1 is completely revised. Instead of particular registers getting hard-coded protection, the entire project is run through a tool [9] that generates triplication for all flip-flops and combinatorial logic. Each of the triplicated instances runs from a different clock tree. The remaining parts that are not triplicated are the pipeline memories, the analog multiplexer block that reads ADCs, and the Hit Counter registers (used only for calibration/testing, these count hits per strip allowing occupancy histograms to be recorded on-chip and verification of the readout circuitry). An SEU counter is also included which increments when a corrective action in the triplication logic of specific configuration registers occurs. Figure 1 shows a block diagram of the ABCStar V1 ASICs indicating what is triplicated and what is not.

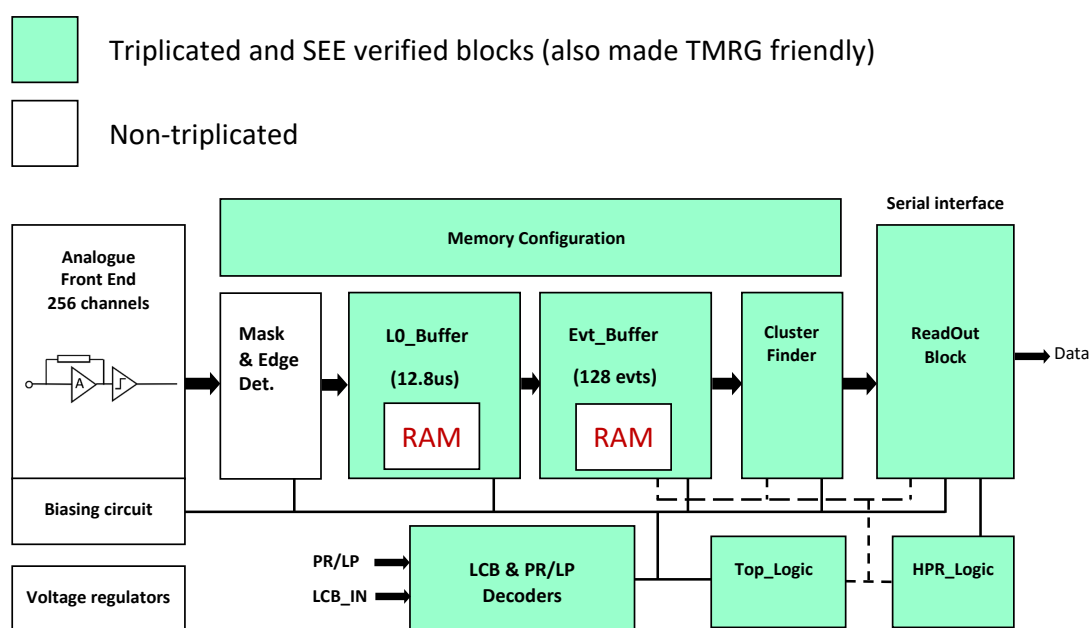


Figure 1. ABCStar V1 block diagram. Please note: “TMRG” := Triple Modular Redundancy Generator — for more details, see reference [9].

Further changes between ABCStar V0 and V1, which are relevant for the studies presented in this paper, include:

- SEU counter

An 8-bit SEU counter was added to the ABCStar V1, which increases when an SEU is encountered in any of the following registers: ADCS0, ADCS1, ADCS2, CREG0, or CREG1.¹ The counter does not indicate register errors, but the occurrence of an upset and a subsequent corrective action.

¹These registers store information related to ADC monitoring and digital configuration — details on all of the register types are given in table 3.

- Clock disabling
The ABCStar V1 supports an option to disable one of the three triplicated clocks and thereby the mechanism for internally correcting SEUs occurring in only one of the three instances. Disabling one of the triplicated clocks effectively increases the cross-section by a factor of 2.
- Glitch filter
This filter delays a copy of the incoming signal by 1.3 ns and compares the copy with the original. If they do not match, the signal is not propagated further. In this way voltage spikes - glitches - of period less than 1.3 ns are not seen by the logic beyond this block. The filter is only applied to the configuration pads and to the 40 MHz clock input pads on the ABCStar V1 ASIC.
- Idle pattern
A programmable idle pattern was introduced to the ASIC, so that the default is not “0000” but repeating “0110” (like an 80 MHz clock).
- Programmable multiplexer
For testing purposes, a programmable multiplexer allows an external output pin (“TESTOUT”) to reflect one of a large number of internal signals inside the ASIC.
- Packet structure
A comparison of packet structures for the ABCStar V0 and V1 is shown in table 1.
- Status bits
An overview of the status bits in the ABCStar V0 and V1 is shown in table 2.
- Test mode
Used during chip testing to apply a set of test patterns to the pipeline in static or pulsed operation modes. In the static mode, the pipeline input continuously gets a set mask bit value.² In the pulse mode, there are three controllable options for the input.

Table 3 shows the different types of the ABCStar registers. The payload of register read packets (32 bits — see table 1) is the associated register’s contents.

3 Measurement setup

For the measurement setup, each ABCStar ASIC was mounted on a dedicated Single Chip Board (SCB) as shown in figure 2.

Each SCB is connected to one FMC-1701 printed circuit board (PCB), which is used to power the SCB and facilitate its readout. The cables used to power the SCB were flat ribbon cables, where the current carrying lines had been replaced by thicker wires soldered into the cable. This was done to prevent a voltage drop when the ASIC’s current increases due early absorbed dose (an effect called the Total Ionising Dose (TID) bump, see reference [6]). The chip under investigation was read out using twisted pair cables to minimise cross-talk. The FMC-1701 board was plugged into a

²N.B. this mode is what allowed for the all 0’s and all 1’s fills on the front-end of the chip, as described in section 5.4.1.

Table 1. Packet structure for the ABCStar V0 and V1.

ABCStar Packet	Physics Packet	V0 Register Read	V1 Register Read	V0 High Priority Register	V1 High Priority Register
3 start bits	3 start bits	3 start bits	3 start bits	3 start bits	3 start bits
Header (16 bits)	Type: 0001 (PR) / 0010 (LP) Flag (1 bit) L0tag (7 bits) BCID (4 bits)	Type: 0100 Register address (8 bits) 0000	Type: 0100 Register address (8 bits) SEU counter [3:0]	Type: 1101 Register address (0x3F) 0000	Type: 1101 HPR check bits: 111 “K2 pending” (1 bit) SEU counter [3:0]
Payload (48 bits)	Cluster 1 (12 bits)	Register contents (32 bits)	Register contents (32 bits)	HPR contents (32 bits)	HPR contents (32 bits)
	Cluster 2 (12 bits)				
	Cluster 3 (12 bits)				
	Cluster 4 (12 bits)				
1 stop bit	1 stop bit	1 stop bit	1 stop bit	1 stop bit	1 stop bit

Table 2. Status bits for ABCStar V0 and V1.

V0 status bits	bit	V1 status bits
chipID [3]	15	chipID [3]
chipID [2]	14	chipID [2]
chipID [1]	13	chipID [1]
chipID [0]	12	chipID [0]
0	11	OR of all SEU bits
BCIDFlag	10	Flag (1 bit)
PRFIFO full	9	PRFIFO almost full
PRFIFO empty	8	PRFIFO empty
LPFIFO full	7	LPFIFO almost full
LPFIFO empty	6	LPFIFO empty
RegFIFO overflow	5	RegFIFO overflow
RegFIFO full	4	RegFIFO almost full
RegFIFO empty	3	RegFIFO empty
ClusterFIFO overflow	2	ClusterFIFO overflow
ClusterFIFO full	1	ClusterFIFO almost full
ClusterFIFO empty	0	ClusterFIFO empty

Table 3. Brief descriptions of the different types of registers present in ABCStar ASICs. Each register is 32 bits deep. When referring to the number of registers of a given type, the bracketed quantity refers to the number present in V0 chips, if different from V1 chips.³

Register type	Description	Number of registers
SCReg	Special/control	1
ADCS	ADC monitoring	3 (6)
MaskInput	Mask for each channel	8
CREG	Digital config	2 (7)
STAT	Runtime statistics/counters	7 (5)
HPR	High priority status codes	1
TrimDAC	Trim for each channel	40
CalREG	Calibration mask	8
HitCountREG	Cluster info for each channel	64

Nexys Video FPGA board [10] with an Artix-7 FPGA [11], which was operated through a PC over a network connection.

During the tests, a constant voltage of 1.5 V was applied to each FMC-1701 which, in addition to powering the ASIC, provided a high resolution voltage and current measurement of the supplied voltage and the corresponding digital and analogue current.

³N.B. the ABCStar register map was compressed going from V0 to V1, hence the fewer V1 registers.




Figure 2. Picture of an ABCStar chip mounted on a Single Chip Board. The ASIC is powered and read out through wire bonds.

3.1 Setup during measurements

The Proton Irradiation Facility (PIF) [12] at TRIUMF in Vancouver, Canada provides protons with an energy of 480 MeV, which are able to traverse several layers of ASICs on PCBs without significant energy loss. Therefore, multiple boards can be mounted in series and tested. Two beam tests were conducted each using four boards in series. The first, conducted in August 2020, used three ABCStar V1 ASICs and one ABCStar V0 ASIC as a reference.⁴ The second, conducted in December 2020, used only V1 ASICs. Each SCB was connected to one FMC-1701 and one Nexys Video FPGA board, and each FMC-1701 was connected to an individual power supply channel providing 1.5 V analogue and digital voltage connected together. All Nexys Video FPGA boards were connected to one Gigabit network switch and read out by one computer.

Due to an inherent divergence of the proton beam, which increases the beam diameter while decreasing the proton flux per area accordingly, placing the ASICs behind each other exposes them to different dose rates and fluences. Figure 3 shows the beam profile measured at different distances from the end of the beam pipe to check the beam diameter.

Dosimetry was performed by measuring the beam flux at different distances to the beam pipe prior to the measurement and extrapolating the flux for each chip based on the particle counter system of the beam line. Tables 4 and 5 summarise the corresponding parameters for all boards mounted in the beam for the August and December beam tests, respectively.

⁴Dedicated measurements of the ABCStar V0 SEU rate were conducted during earlier beam tests; however, these results are not publicly available. Here, an ABCStar V0 ASIC was included only to quantify the relative improvement between both ASIC versions.

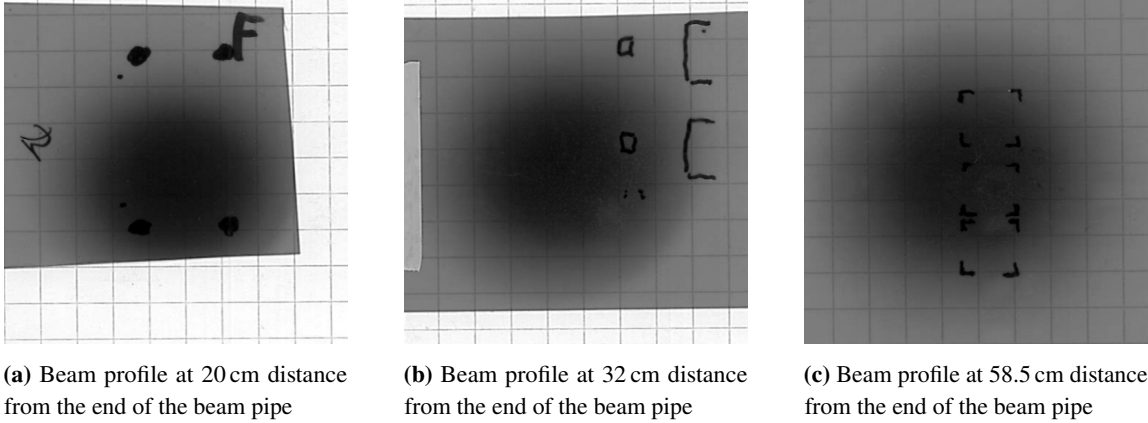


Figure 3. Beam profiles obtained using dosimetry foils, placed at various distances from the end of the beam pipe: the beam diameter increases with increasing distance from the end of the beam pipe. Each square in each grid is 5×5 mm.

A picture of the setup used in the August beam test is shown in figure 4. A similar setup was used for the December beam test.

Table 4. Doses and fluences accumulated by all boards under investigation as part of the August 2020 proton beam test. The parameters of each board are given as a function of their distance to the beam pipe.

Chip	Generation	Distance [cm]	Dose [Mrad]	Fluence [p/cm ²]
261	V1	18.8	4.30	1.18×10^{14}
002	V0	31.0	3.34	9.17×10^{13}
259	V1	44.0	2.67	7.33×10^{13}
267	V1	56.3	2.23	6.13×10^{13}

Table 5. Doses and fluences accumulated by all boards under investigation as part of the December 2020 proton beam test. The parameters of each board are given as a function of their distance to the beam pipe.

Chip	Generation	Distance [cm]	Dose [Mrad]	Fluence [p/cm ²]
259	V1	43.5	2.48	6.80×10^{13}
273	V1	46.0	2.38	6.55×10^{13}
267	V1	56.0	2.07	5.67×10^{13}
274	V1	58.5	2.00	5.49×10^{13}



Figure 4. A picture of the setup used in the December 2020 proton beam test. The beampipe is visible on the far right, and the four PCBs to its immediate left correspond to SCBs with ABCStar ASICs mounted. Each SCB is connected to flat and twisted ribbon cables for slow and fast signals, respectively. The red laser focused on the ABCStar ASICs indicates the position of the beam.

4 Measurements

During the beam tests, ABCStars were read out using the Inner Tracker Strips DAQ (ITSDAQ) software [6]. The Nexys Video FPGA board was loaded with ITSDAQ firmware (FW) [13] version vb48c for the August beam test and firmware version vb4b4 for the December beam test.⁵

In order to determine the SEU rates, the following data was monitored:

- register data (triplicated, for writable as well as read-only registers);
- physics data (non-triplicated event data);
- TESTOUT output counters.

The firmware TESTOUT counters are connected to the ABCStar TESTOUT debug output that can be switched to monitor a selection of signals inside the ASIC. Some of these may be affected by SEUs. To monitor this activity, three counters are available in the firmware. The first counts the number of times the TESTOUT signal changes from 0 to 1, the second counts the integrated period that the signal is asserted, and third is continuously running and can be used as a time reference. Firmware status information is read out as a block 16-bit words, with the TESTOUT counters wider than 16-bit assembled from multiple of these words:

⁵The important differences going from firmware version vb48c to vb4b4 include support for idle mode, the addition of TESTOUT counters, and the expansion of the timestamp register rollover size — it was increased from 32 bits to 48 bits (i.e., 107 s to ~2000 hours).

- Edge Counter: 32-bit.
Counts rising-edges in TESTOUT at 1.56 ns (1/640 MHz) resolution.⁶
- Period-high Counter: 32-bit.
Counts the period that TESTOUT is high in 1.56 ns steps.⁷
- Timecount Counter: 48-bit.
Free running counter, clocked by a 40 MHz clock, set to 0 at power-up/reset.

Only the “LCB_Locked” signal was monitored. The LCB (L0A Command Bunch Counter Reset) input to the ABCStar carries L0A,⁸ Command⁹ and BCR¹⁰ signals. These are used to control the ASIC, including register reading and writing. The LCB is transferred in 16-bit frames at 160 Mb/s, to which the logic in the ABCStar must synchronise (or “lock”) to. Once a sufficient number of good frames are seen, the LCB decoder enters a “locked” state. The LCB signal is always active and encoded such that errors can easily be identified, and persistent errors will cause the ABCStar logic to de-assert “locked”. As such, this is either a good indicator of a serious failure in this part of the ASIC or that a reset has occurred.

Monitoring was grouped into several loops:

1. An inner loop filling and reading out the event data buffer;
2. A middle loop, during which:
 - the inner loop was executed, in sets of 10, in the case of the August beam test, or 100, in the case of the December beam test;
 - the register data was read out and writable registers were reset (the code was optimised to maximise integration time by resetting the register directly after reading out);
 - the monitoring data from the TESTOUT counters and the FMC-1701 was read out and stored.

Each step was logged with a timestamp for subsequent calculation of integration time and corresponding fluence.

3. An outer loop: in the case of the August beam test, the fill pattern was changed for sets of 10 middle loops (“fixed” versus BCID-dependent fills — see section 5.4.1 for more details).

Additionally, several options and running modes were implemented:

⁶The underlying logic operates on bytes at 80 MHz, so the maximum number of rising edges possible in a byte is 4 (N.B. the last bit of the previous byte is also used to check for an edge between the two). The max value is 0xfffffb+last byte edge-count. This means the max value is the range 0xfffffc-f.

⁷Similar to Edge Count, the maximum number of steps a byte can be high is 8, the max value is 0xffff7+last byte count. This means the max value is the range 0xffff8-f.

⁸L0A: Level 0 Accept: beam-synchronous pulse that transfers data from L0Buffer to the EventBuffer.

⁹Command: controls register reading and writing.

¹⁰BCR: Bunch Counter Reset: pulse synchronous with the machine orbit used to keep local counter aligned.

- Triplication enabled | disabled

The occurrence of SEUs in registers is rare due to their triplication. In order to verify that the setup was able to detect SEUs, a test mode was implemented where the third clock was disabled and the rate of observable SEUs increased. This mode was used for initial tests.

- Glitch filter mode enabled | disabled

A glitch filter was added to the configuration pads and to the 40 MHz clock input pads on the ABCStar V1 ASIC. The goal is to filter out glitches that could come from external sources to the chip. The glitch filter was activated for half of the data collection time in order to compare the SEU cross-sections with and without glitch filter enabled.

- Fill pattern fixed | BCID

Two modes of filling the ABCStar’s event data were tested: “Fixed”, where all hits were filled with 1’s or 0’s (alternating per outer loop), and “BCID”, where, during the inner loop, the LOBuffer was filled with all possible combinations of 4-strip clusters in the event buffer. While the “Fixed” pattern allowed for straightforward comparison of the cross-sections for $0 \rightarrow 1$ and $1 \rightarrow 0$ bit flip cross-sections and straightforward analysis, the BCID pattern data provided a more realistic approximation of real event data.

- Latency setting

The ABCStar is a pipelined device, where data is stored for some period of time awaiting a trigger. This setting allows the time in which data stays in the pipeline (in clock periods) to be changed, which means more SEUs might be integrated in this time.

A second mode of data collection was implemented for the December 2020 proton measurements called “Idle” mode: instead of gathering physics data, the ASIC was configured to send an “Idle” pattern - that is a repetitive pattern on the data output - of “0110”. In the absence of SEUs, data from the ASIC was expected to contain only the idle pattern and, at scheduled intervals, high-priority packets (HPRs). If unexpected packets were seen, as would happen if, for example, the ASIC saw a trigger, this would be known to be caused by SEEs.

5 Results from proton tests

The following sections describe the measurements performed as part of the proton beam tests at TRIUMF in August and December 2020. Where appropriate, the integrated fluence or dose for a given chip is calculated using the calibration curves and particle counts provided by the TRIUMF PIF, following the procedure described in appendix A.

5.1 Firmware TESTOUT counters

The results from the ABCStar’s TESTOUT counters for the December beam test are shown in figure 5. Only the “LCB_Locked” signal was monitored. N.B. the TESTOUT counter functionality was *not* included as part of the August beam test.

From figure 5a, no edges were counted — this indicates the LCB_Locked signal output never transitioned high-to-low or low-to-high. And from figure 5b, the LCB_Locked signal output is

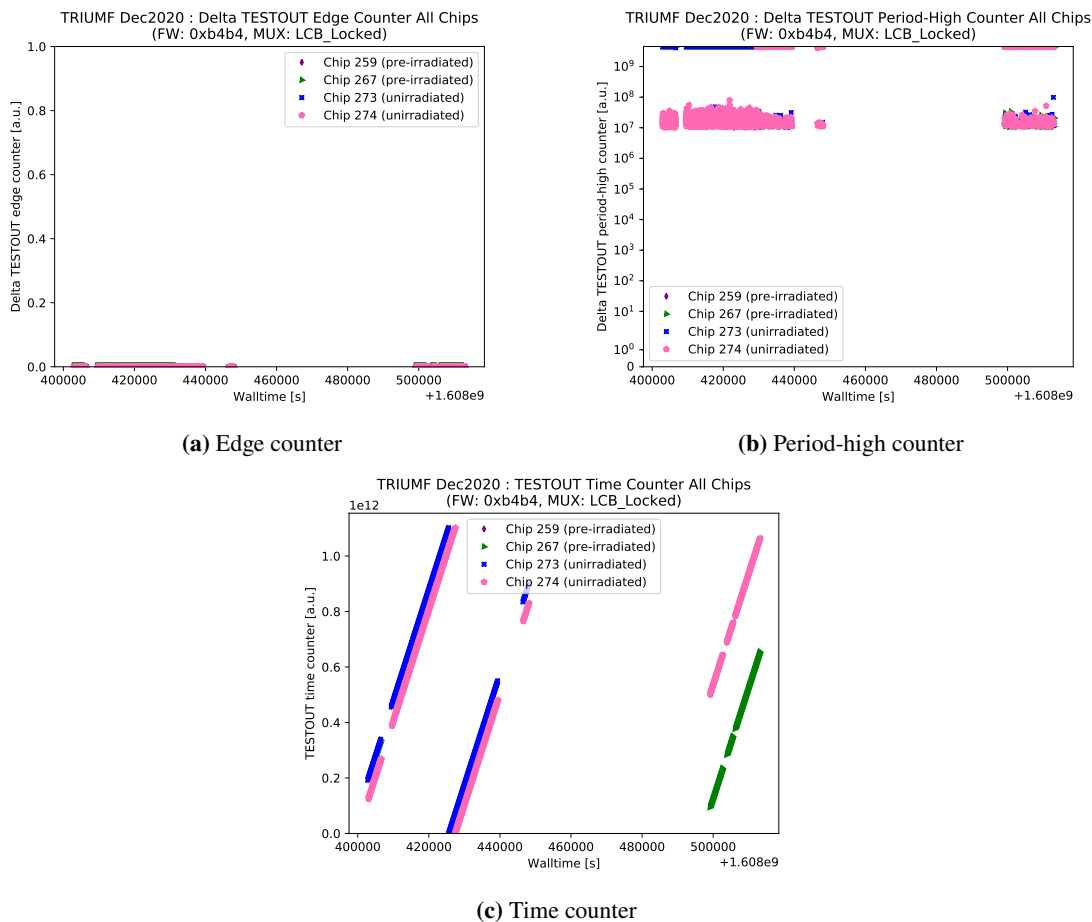


Figure 5. Firmware TESTOUT counters for the ABCStar V1 ASICs measured as part of the December 2020 beam test. “0xb4b4” refers to the firmware revision (vb4b4) used. N.B. the large time gap between some data points corresponds to the time between days 1 and 2 of the beam test. Additionally, the vertical scale for the period-high counter, (b), is logarithmic with a linear scale near 0.

always high, as the period-high counter is always at or near saturation following a counter reset. Altogether, these results indicate that *no* LCB unlocked events were measured for any of the ASICs over the duration of the beam test, indicating no serious failures.

From figure 5c, the time counted in firmware is proportional to the time counted in software and the slope of the linear relationship between firmware and software time is the same for all chips, as expected.

5.2 HPR reset and unlock rates

High Priority Register (HPR) packets are normal packets containing link status information that are sent with high priority at regular intervals. They are initiated with a higher priority than any other packet type. There are two rate measurements of interest made with respect to HPR packets: the rate of HPR resets and the rate of LCB unlocks.

During normal operation, the HPR packets are expected to be spaced apart in integer multiples of 40,000 bunch crossings (BCs), corresponding to integer multiples of 1 ms. If an HPR packet

occurs at another time after the previous HPR packet, then this would be indicative of a reset in the ASIC. However, if an HPR is scheduled to occur while another packet is already being read out (or a group of packets), the time when the HPR packet is read out may be modified by the time required to read out the conflicting packet (or packets). Therefore, the check for out-of-time packets is $|\Delta BC \bmod 40,000| > 67$ BCs, where 17 out of 67 BCs are for packet readout¹¹ and 50 out of 67 BCs are to account for further delays due to the readout of additional physics packets initiated by SEUs. For events in August with all 1's L0Buffer fills, if an HPR occurs during the LP or PR buffer scans, 17 BCs is modified to $16 \times 17 = 272$ BCs as there are 16 packets read from the LP or PR buffers per scan per LOA (as opposed to 1 packet for events with all 0's fills).

One of the status bits contained in the HPR is the “LCB_Locked” (bit 29). The number of LCB unlocks can therefore be determined by counting the number of times this bit is seen to switch from “1” to “0”. In this section as well as section 5.3, pattern matching for “good” HPR status bits (“0101001001” — these bits correspond to the “PRFIFO full/almost full” through “ClusterFIFO empty”, inclusive, status bits in table 2 followed by a stop bit followed by a “0000” or “0110” idle sequence) is used to guard against unlocked HPR packets which are actually mis-parsed, locked HPR packets.

For both the numbers of out-of-time and unlocked HPR packets, the rates are calculated by summing the integrated times between subsequent packets.

5.2.1 August results

The HPR packet results for the August 2020 beam test are summarised in table 6. With the exception of the BCID fill runs with the glitch filter enabled for Chips 261 and 267 (which see $O(0.5)$ packets/s), all of the V1 chips see $O(5)$ packets/s for all runs. In comparison, the V0 chip sees $O(1)$ packets/s for all runs.

No unlocked packets were measured for any runs. A 95% upper confidence bound is calculated by assuming 3 unlocked packets in the rate calculation — these upper bounds are 1.67×10^{-5} packets/s and 5.04×10^{-6} packets/s for the V0 and V1 chips, respectively.

Altogether, 17 and $7\Delta t \neq 1$ ms packets were seen for the V0 and V1 chips, respectively. In general, $\Delta BC \bmod 40,000$ is large for the out-of-time packets — the distribution of $\Delta BC \bmod 40,000$ for out-of-time packets for all chips is shown in figure 6. As none of the chips showed a corresponding LCB unlock event or a register with unexpected contents, these are assumed to be related to the data-taking mode (with an extra reset as part of setup) rather than an unscheduled ASIC reset. Therefore, this extra reset was removed in the following run. We may still calculate the rate of $\Delta t \neq 1$ ms packets as:

$$\text{Rate}_{\Delta t \neq 1 \text{ ms}} = \frac{N_{\text{packets}, \Delta t \neq 1 \text{ ms}}}{\int dt}, \quad (5.1)$$

and the 95% confidence interval on the rate as:

$$\Delta_{\text{Rate}_{\Delta t \neq 1 \text{ ms}}} = 1.96 \times \frac{\sqrt{N_{\text{packets}, \Delta t \neq 1 \text{ ms}}}}{\int dt}, \quad (5.2)$$

¹¹There are 68 bits per packet and bits are read at a rate of 4 bits per BC, yielding 17 BCs to read the entire packet.

(assuming Poisson counting statistics). The factor of 1.96 comes from $\Phi(1 - \alpha/2) = \Phi(1 - 0.05/2) \approx 1.96$ where Φ is the quantile function of a standard normal distribution. The rates are calculated to be $(9 \pm 4) \times 10^{-5}$ packets/s and $(1.7 \pm 1.2) \times 10^{-5}$ packets/s for the V0 and V1 chips, respectively.

Table 6. HPR packets results for August 2020. Both the number of unlocked HPR packets (N_{unlocked}) and the number HPR packets with delta times not equal to 1 ms ($N_{\Delta t \neq 1 \text{ ms}}$) are shown.

Chip	Fill type	Glitch filter?	Clock disabled?	N_{unlocked}	$N_{\Delta t \neq 1 \text{ ms}}$	N_{packets}	Integrated time [s]
002	Fixed	No	No	0	17	171,557	179677.0
	BCID	No	No	0	1	65,787	15263.7
	BCID	Yes	No	0	0	66,211	15665.1
259	Fixed	No	No	0	0	96,858	18142.5
	Fixed	No	Yes	0	0	3,871	801.7
	Fixed	Yes	No	0	0	84,721	15930.2
	BCID	No	No	0	1	68,390	15825.9
	BCID	Yes	No	0	1	73,978	127189.0
261	Fixed	No	No	0	0	82,822	16263.3
	Fixed	No	Yes	0	0	7,102	1386.7
	Fixed	Yes	No	0	0	86,364	15949.9
	BCID	No	No	0	0	74,046	17007.5
	BCID	Yes	No	0	1	28,601	117843.0
267	Fixed	No	No	0	0	67,774	13509.9
	Fixed	No	Yes	0	1	7,283	1390.6
	Fixed	Yes	No	0	2	132,798	23811.3

5.2.2 December results

The HPR packet results for the December 2020 beam test are summarised in table 7. *No* unlocked packets or packets with delta times out of specification were measured — this is consistent with the FW counter results described in section 5.1: no ASIC resets were measured.

We can calculate 95% upper confidence bounds on the rate of LCB unlocked and $\Delta t \neq 1$ ms (altogether: “bad”) packets. Assuming the results are not dependent on the chip’s latency, upper bounds on the rate of bad HPR packets for the normal and idle running modes are 1.94×10^{-5} packets/s and 1.29×10^{-4} packets/s, respectively. Here, we have distinguished between the normal and idle running modes as the normal mode sees $O(100)$ packets/s while the idle mode sees $O(500)$ packets/s.

While the ASIC specifications do not place an explicit limit on these effects, the observed rates for both August and December were found to be of no concern for the operation of ABCStar ASICs in the detector.

5.3 HPR and register read status bits

The status bits for HPR and register read packets are described in table 2. For all HPR and register read packets in August and December, the status bits were checked for FIFO almost-full (“*FIFO_almost_full”) or, in the case of V0 chips, full statuses (“*FIFO_full”) as well as for register or cluster FIFO overflows, which could indicate a reset in the ASIC.




Figure 6. Histogram of delta times (in BCs) between subsequent HPR packets for the out-of-time packets measured in August 2020. The dotted red lines denote the domain of the modulo 40,000 operation (i.e., $\pm 20,000$ BCs).

Table 7. HPR packets results for December 2020. Both the number of unlocked HPR packets (N_{unlocked}) and the number HPR packets with delta times not equal to 1 ms ($N_{\Delta t \neq 1 \text{ ms}}$) are shown.

Chip	Idle mode?	Latency	N_{unlocked}	$N_{\Delta t \neq 1 \text{ ms}}$	N_{packets}	Integrated time [s]
259	No	20	0	0	3,323,805	33173.1
	No	503	0	0	1,395,467	12312.2
267	No	20	0	0	2,591,380	25863.8
	No	503	0	0	1,417,932	12313.0
273	No	20	0	0	3,543,084	34829.8
	No	503	0	0	1,418,156	12313.7
274	Yes	20	0	0	12,101,755	23227.2
	No	20	0	0	1,015,722	9888.1
	No	503	0	0	1,617,822	13964.5

5.3.1 August results

The status bit results for the August beam test are summarised in table 8. Chip 002, the V0, returned 159 and 228 PRFIFO_full and LPFIFO_full statuses, respectively, corresponding to 0.0007% and 0.001% of all packets measured. Chip 261, a V1, returned 18 LPFIFO_almost_full statuses for its BCID fill run, corresponding to 0.00007% of all packets measured. Chip 267, a V1, returned 150 ClusterFIFO_almost_full statuses for its clock-disabled run, corresponding to 0.0005% of all packets measured. Additionally, *no* FIFO overflow statuses were measured — this is consistent with the HPR out-of-time results described in section 5.2.1: no ASIC resets were measured.

An attentive reader may notice that the number of HPR and register read packets measured for the fixed LOBuffer fill runs do not sum to the number of expected packets given table 6 and (later, in section 5.5.1) table 12. This is because the fixed LOBuffer fill runs perform full LP/PR buffer scans, leading to a full readout of the HitCountREG registers (0x80 through 0xbf, inclusive). As a result, an additional 64 register read packets per buffer read per scan type are measured for these runs.

5.3.2 December results

The status bit results for the December beam test are summarised in table 9. *No* FIFO almost-full or overflow statuses were measured — this is consistent with the FW counter and HPR out-of-time results described in sections 5.1 and 5.2.2, respectively: no ASIC resets were measured.

5.4 SEU cross-section for physics packet clusters

Given the LOBuffer is filled in a predictable way, the physics packet clusters returned from the buffer have an expected structure — the expected structure can be compared to what was actually measured and the differences enumerated.

It is worth noting that there two contributions to the number of the SEUs measured during physics packet readout: the first contribution comes from SEUs accumulated while the data is stored in the LOBuffer and the second contribution comes from SEUs accumulated while the packets are built, processed, and sent. In the subsequent analysis, the two sources are *not* disentangled when calculating the SEU cross-sections.

The SEU cross-section for bit-flips in physics packet clusters, σ_{SEU} , is calculated as:

$$\sigma_{\text{SEU}} = \frac{n_{0 \rightarrow 1} + n_{1 \rightarrow 0}}{\int d\phi}, \quad (5.3)$$

where $n_{0 \rightarrow 1}$ and $n_{1 \rightarrow 0}$ are the number of $0 \rightarrow 1$ and $1 \rightarrow 0$ bit-flips in the measured clusters, respectively. The cross-section is normalised to the total integrated fluence, $\int d\phi$, accounting for integration time *only* between when the LOBuffer is written to and when it is read out (i.e., when the packet is received). The 95% confidence interval (assuming Poisson counting statistics) on cross-section, $\Delta_{\sigma_{\text{SEU}}}$, is calculated as:

$$\Delta_{\sigma_{\text{SEU}}} = 1.96 \times \frac{\sqrt{n_{0 \rightarrow 1} + n_{1 \rightarrow 0}}}{\int d\phi}. \quad (5.4)$$

In the case of zero measured SEUs, a 95% upper confidence bound is calculated by assuming 3 SEUs in the cross-section calculation.

Table 8. HPR and register read status bits for the August 2020 beam test. The status bit columns indicate the number of packets which had 1 instead of 0 for that status bit. For Chip 002, the V0, the “almost_full” FIFO statuses actually correspond to “full” statuses. N.B. a packet may count towards the yields in more than one status bit column.

Chip	Fill type	Glitch filter?	Clock disabled?	PRFIFO_almost_full	LPFIFO_almost_full	RegFIFO_overflow	RegFIFO_almost_full	ClusterFIFO_overflow	ClusterFIFO_almost_full	N_{packets}
002	Fixed	No	No	159	228	0	0	0	0	22,916,647
	BCID	No	No	0	0	0	0	0	0	179,107
	BCID	Yes	No	0	0	0	0	0	0	180,236
	Fixed	No	No	0	0	0	0	0	0	14,624,436
	Fixed	No	Yes	0	0	0	0	0	0	596,129
259	Fixed	Yes	No	0	0	0	0	0	0	12,374,915
	BCID	No	No	0	18	0	0	0	0	186,872
	BCID	Yes	No	0	0	0	0	0	0	202,725
	Fixed	No	No	0	0	0	0	0	0	12,438,698
	Fixed	No	Yes	0	0	0	0	0	0	1,092,867
261	Fixed	Yes	No	0	0	0	0	0	0	12,590,366
	BCID	No	No	0	0	0	0	0	0	202,560
	BCID	Yes	No	0	0	0	0	0	0	78,079
	Fixed	No	No	0	0	0	0	0	0	10,218,863
	Fixed	No	Yes	0	0	0	0	0	150	1,092,914
267	Fixed	Yes	No	0	0	0	0	0	0	19,300,187

Table 9. HPR and register read status bits for December 2020 beam test. The status bit columns indicate the number of packets which had 1 instead of 0 for that status bit. N.B. a packet may count towards the yields in more than one status bit column.

Chip	Idle mode?	Latency	PRFIFO_almost_full	LPFIFO_almost_full	RegFIFO_overflow	RegFIFO_almost_full	ClusterFIFO_overflow	ClusterFIFO_almost_full	N_{packets}
259	No	20	0	0	0	0	0	0	6,491,530
	No	503	0	0	0	0	0	0	2,616,791
267	No	20	0	0	0	0	0	0	4,955,352
	No	503	0	0	0	0	0	0	2,639,594
273	No	20	0	0	0	0	0	0	6,910,671
	No	503	0	0	0	0	0	0	2,643,497
274	Yes	20	0	0	0	0	0	0	13,701,314
	No	20	0	0	0	0	0	0	2,075,346
	No	503	0	0	0	0	0	0	3,043,142

5.4.1 August results

For the August beam test, the LOBuffer was filled with “fixed” patterns - all 0’s or all 1’s, alternating every 10 events - or with “BCID” patterns - patterns dependent on the BCID. Technical issues were encountered with the BCID fill runs and so physics packets from those runs are not analysed.

Per event, 128 LOA’s are sent in order to read out the entire 128-bit deep buffer. Per LOA, 4 reads of the current slice of the LOBuffer are performed. Per read, 1 packet is returned for all 0’s and 16 packets for all 1’s fills. Thus, per event, 128 groups of either 4 or 64 packets (with each group sharing a common LOID) are expected for fixed fill runs. Both LP and PR buffer scans are performed, doubling the number of expected physics packets per event.

Within the packets returned for a given LOA, only the last read (i.e., the last 1 or 16 packets in the case all 0’s or all 1’s fills, respectively) is used in the SEU cross-section calculation. This is because the first, second, and third reads are not independent of the fourth, which also has the longest integration time. In general, the first, second, and third reads were measured for posterity and only scrutinised in select cases.¹²

The measured and expected clusters are converted to their corresponding 256-bit strip data sequences, and the measured and expected sequences are compared to one another and the number of bit-flips are counted. The LP and PR buffer scans are combined when counting the number of bit-flips. The SEU cross-sections for bit-flips in physics packet clusters are summarised in table 10.

Table 10. A summary of the SEU cross-sections for bit-flips in physics packet clusters per chip per running mode for the August 2020 beam test.

Chip	Glitch filter?	Clock disabled?	$n_{0 \rightarrow 1}$	$n_{1 \rightarrow 0}$	$N_{\text{bits, total}}$	$\int d\phi$ [p/cm ²]	$\sigma_{\text{SEU}} \pm \Delta\sigma_{\text{SEU}}$ [cm ² /p]
002	No	No	33,699	34,546	15,250,176	1.86×10^{16}	$(3.69 \pm 0.03) \times 10^{-12}$
259	No	No	16,343	16,506	7,642,368	4.47×10^{15}	$(7.35 \pm 0.08) \times 10^{-12}$
	No	Yes	1,016	784	411,136	2.34×10^{14}	$(7.7 \pm 0.4) \times 10^{-12}$
	Yes	No	13,282	13,913	6,321,408	3.83×10^{15}	$(7.09 \pm 0.08) \times 10^{-12}$
261	No	No	18,679	18,853	8,543,488	6.67×10^{15}	$(5.63 \pm 0.05) \times 10^{-12}$
	No	Yes	2,233	1,971	930,304	6.38×10^{14}	$(6.6 \pm 0.2) \times 10^{-12}$
	Yes	No	18,848	19,643	8,687,360	6.20×10^{15}	$(6.21 \pm 0.06) \times 10^{-12}$
267	No	No	8,459	8,775	4,129,280	2.85×10^{15}	$(6.04 \pm 0.09) \times 10^{-12}$
	No	Yes	1,342	1,053	562,432	3.38×10^{14}	$(7.1 \pm 0.3) \times 10^{-12}$
	Yes	No	15,235	14,798	7,214,592	5.04×10^{15}	$(5.97 \pm 0.07) \times 10^{-12}$
All V1 chips	All modes	129,136	130,842	59,692,544	4.89×10^{16}	$(5.32 \pm 0.02) \times 10^{-12}$	

Within a particular chip, the different running scenarios (i.e., glitch filter enabled/disabled, clock enabled/disabled) result in cross-sections which are very consistent with one another. This behaviour is expected — the glitch filter counts glitches on the input clock which is only expected when the generation is part of the system under test, as is the case with a Hybrid Control Chip (HCC) [5]. Moreover, disabling of the clock only enhances the number of SEUs in triplicated registers (see section 5.5).

¹²For an example of this, in situations where the fourth read is truncated due to edge effects, the third read is analysed in place of the fourth. However, this is of no consequence with respect to the measured results.

Additionally, the cross-sections obtained for different chips are very consistent with one another, all $O(10^{-12})$ cm²/p. The V0 and V1 chips measure cross-sections of $(3.69 \pm 0.03) \times 10^{-12}$ cm²/p and $(5.32 \pm 0.02) \times 10^{-12}$ cm²/p, respectively. While the two cross-sections don't agree within statistical error, they are likely much more consistent when including systematic effects such as the uncertainty on the integrated fluence.¹³

Very rarely, the output data was truncated during readout as a result of issues encountered due to edge effects in the DAQ software — excluding these cases, no SEUs were measured as part of packet building, processing, and sending. This is expected, as these mechanisms utilise TMR protection.

5.4.2 December results

During the December 2020 beam test, the L0Buffer was repeatedly filled and read out (*only* LP buffer scans) in groups of 8 triggers, with the fill pattern set according to the BCID. To represent a more realistic running scenario, the L0Buffer was filled and immediately read out for each group of 8 triggers. The measured packets were grouped by closest-in-time (i.e., within 1000 BCs of one another). The integration time for the measured packets is assumed to be constant and given by:

$$\Delta t = (40 \text{ MHz})^{-1} \times (255 \times (7 - i) + \text{latency}), \quad (5.5)$$

where i is the index of the packet in the group of 8 ($i = 0, \dots, 7$) and the latency is either 20 or 503, depending on the run. N.B. the integration times used in December were *considerably* shorter than those used in August: $O(10)$ μ s/packet for December versus $O(10)$ s/packet for August. However, this difference is accounted for in the total integrated fluence.

With the BCID-dependent fills, the clusters returned in a packet are predictable based on the BCID. The BCID decrements by 1 from one packet to the next within a group (accounting for overflow, modulo 256). On group edges, the BCID increments by the number of triggers sent, 8, plus the number of measured packets in the current group, 8 (accounting for overflow, modulo 256). Packets are identified as “bad” if the BCID corresponding to the returned clusters does not match the value expected within the larger pattern or if the returned clusters do not correspond to any of the clusters possible with BCID fills.

The number of bit-flips is calculated in the same way as for the August results (i.e., by comparing measured and expected strip data sequences). The SEU cross-sections for bit-flips in physics packet clusters are summarised in table 11.

Summing all chips, the SEU cross-section for bit-flips in physics packet clusters is $(1.6 \pm 1.4) \times 10^{-11}$ cm²/p. This result is consistent with the cross-section measured in August for the V1 ASICs, $(5.32 \pm 0.02) \times 10^{-12}$ cm²/p. Additionally, for the run scenarios where no bit-flips were observed, the upper confidence bounds of $O(10^{-11})$ to $O(10^{-10})$ cm²/p are nearly consistent with the cross-sections for the runs where bit-flips were observed in both August and December.

¹³There are two sources of uncertainties affecting the integrated fluence: the uncertainty on the number of SEM counts, which has a linear effect, and the uncertainty on the distance of the SCBs from the end of the beampipe, which has a quadratic effect. The number of SEM counts comes with a Poisson counting uncertainty, which is expected to be small, and an uncertainty on the efficiency for measuring a particular SEM count, which is unknown. Focusing on the uncertainty on the distance, if we made the (*very*) conservative assumption of a 25% uncertainty, then this translates to a 50% uncertainty on the integrated fluence. Even still, this does not affect our conclusion concerning the effect of SEUs in physics packet clusters on the detector's performance.

Table 11. A summary of the SEU cross-sections for bit-flips in physics packet clusters per chip per running mode for the December 2020 beam test. Cross-sections with a prefixed asterisk (*) indicate 95% upper confidence bounds assuming 3 SEUs.

Chip	Latency	$n_{0 \rightarrow 1}$	$n_{1 \rightarrow 0}$	$N_{\text{bits,total}}$	$\int d\phi$ [p/cm ²]	$\sigma_{\text{SEU}} \pm \Delta\sigma_{\text{SEU}}$ [cm ² /p]
259	20	0	0	219,809,024	2.09×10^{10}	$*1.44 \times 10^{-10}$
	503	0	0	299,576,832	4.44×10^{10}	$*6.78 \times 10^{-11}$
267	20	1	0	218,816,256	1.73×10^{10}	$(0.6 \pm 1.1) \times 10^{-10}$
	503	1	0	299,680,256	3.70×10^{10}	$(3 \pm 5) \times 10^{-11}$
273	20	2	1	826,031,104	8.05×10^{10}	$(4 \pm 4) \times 10^{-11}$
	503	0	0	300,588,544	4.28×10^{10}	$*7.02 \times 10^{-11}$
274	20	0	0	259,894,528	2.25×10^{10}	$*1.34 \times 10^{-10}$
	503	0	0	349,664,768	4.41×10^{10}	$*6.81 \times 10^{-11}$
All chips	20 & 503	4	1	2,774,061,312	3.10×10^{11}	$(1.6 \pm 1.4) \times 10^{-11}$

As with the August results, DAQ-related issues sometimes caused only 7 triggers to be sent instead of 8. However, no SEUs were measured as part of packet building, processing, and sending, as expected.

5.4.3 Combined results and estimated impact on operation

Altogether, the combined cross-section for V1 chips using data from August and December is $(5.32 \pm 0.02) \times 10^{-12}$ cm²/p. This combined result is dominated by the August runs due to their much larger integrated fluence as compared to those in December. As the time in the pipeline is long for (i.e., $O(10)$ s) for the data measured in August and comparatively short (i.e., $O(10)$ μ s) for the data measured in December, the combined number of SEUs is dominated by those accumulated while the data is in the L0Buffer.

We can estimate the number of hit errors per second during typical ITk operation at the HL-LHC. ATLAS simulations show worse-case expected fluences of $O(10^{-3})$ hadrons/cm²/ pp collision (where pp collision refers to a proton-proton collision). Assuming a pileup of $\langle\mu\rangle = 200$ and given the rate of BCs is 40 MHz (i.e., the number of pp collisions per second),¹⁴ the hadron flux, Φ_{hadrons} , is given by:

$$\begin{aligned}
 \Phi_{\text{hadrons}} &= O(10^{-3}) \text{ hadrons/cm}^2/pp \text{ collision} \times 200 \\
 &\quad \times (40 \times 10^6) \text{ } pp \text{ collisions/s} \\
 &= O(10^7) \text{ hadrons/cm}^2/\text{s}.
 \end{aligned}
 \tag{5.6}$$

The data acquisition scheme in December is a more realistic running scenario (i.e., the event buffer is filled and rapidly read-out). For December, the average time in the pipeline for data is $O(10)$ μ s. Additional simulations show the worst case readout rates in the ITk are $10/9 \approx 1.1$ packets/event/ABCStar, where an event is defined as a *readout* event. Given equation (5.6) and

¹⁴The collision rate of 40 MHz represents a worst-case scenario, as approximately only 2,700 in 3,500 BCs are filled.

the above information, the rate of hit errors due to SEUs in physics packet clusters, Rate_{SEU} , is given by:

$$\begin{aligned} \text{Rate}_{\text{SEU}} &= \left(\Phi_{\text{hadrons}} \times \frac{\langle \text{time in pipeline} \rangle}{\text{packet}} \right) \times \frac{\text{packets}}{\text{event}} \times \sigma_{\text{SEU}} \\ &= (\mathcal{O}(10^7) \text{ hadrons/cm}^2/\text{s} \times \mathcal{O}(10^{-5}) \text{ s/packet}) \\ &\quad \times 1.1 \text{ packets/event} \times \mathcal{O}(10^{-12}) \text{ cm}^2/\text{p}/\text{ABCStar} \\ &= \mathcal{O}(10^{-10}) \text{ errors/event}/\text{ABCStar}. \end{aligned} \tag{5.7}$$

Here, the bracketed factors correspond to the effective integrated fluence per packet. So we expect $\mathcal{O}(10^{-10})$ hit errors per readout event per ABCStar. Given there are $\sim 230,000$ ABCStar ASICs in the ITk strip tracker¹⁵ and assuming a trigger rate of 1 MHz, this equates to $\mathcal{O}(20)$ hit errors/s during normal operation. This error rate is *significantly* below the noise occupancy due thermal noise expected for the detector (i.e., 10^{-2} to 10^{-1} hit errors per readout event per ABCStar) and so is of no practical concern.

5.5 SEU cross-section for register read data

A measurement of the SEU cross-section is calculated using the register reads performed during both proton beam tests. When reading a register, the returned packet is required to have a structure consistent with that described in table 1.

As SEUs are identified by comparing an expected payload to a measured one, the only register types considered are those whose content is well defined: SCReg, ADCS, MaskInput, CREG, STAT (*only* the fuse ID register, 0x32), TrimDAC, and CalREG (as a reminder, the different register types are enumerated in table 3). The HitCountREG registers, which summarise front-end cluster information, are (implicitly) analysed as part of the physics data cross-section analysis (see section 5.4). The HPR register, which contains the information returned by HPR packets, is reserved for the HPR packet analysis (see section 5.2).

The SEU cross-section for bit-flips in the register read data, σ_{SEU} , and its associated 95% confidence interval, $\Delta_{\sigma_{\text{SEU}}}$, are calculated using equations (5.3) and (5.4), respectively. The integrated fluence is only counted between the times when a register is written to and when it is read out.

5.5.1 August results

In August, the ADCS registers' contents are static in time. This is also true for the fuse ID STAT register. The MaskInput and TrimDAC registers are filled with all 1's for 10 events followed by all 0's for 10 events (where an event consists of a full read out of all registers), and this pattern repeats for the duration of all runs. The CalREG registers are filled with all 0's for 10 events followed by all 1's for 10 events (this is similarly true for the HitCountREG registers for fixed fill runs). For the V0 chip, the CREG registers are static in time; for the V1 chips, only the CREG register 0x20 is static in time (the other CREG register, 0x21, is dynamic with unobvious patterns). As a result, only 0x20 is included as part of the CREG registers for V1 chips. It's important to note that the SCReg was *not* measured as part of the August beam test due a mistake in the measurement software. In

¹⁵The number of ASICs also represents a worst-case scenario, as only a subset of all ASICs will be read out during a collision.

total, 70 and 61 registers were included in the SEU cross-section calculation for the V0 and V1 chips, respectively.

Figure 7 shows the SEU cross-sections for bit-flips in register read data as a function of register type, running mode, and chip. Table 12 summarises the inclusive cross-sections for each chip.




Figure 7. The SEU cross-sections for bit-flips in register read data for the August 2020 beam test. Visible data points as well as their errors are calculated using equations (5.3) and (5.4). Data points which are not visible on the scale of the graph correspond to 95% upper confidence bounds assuming 3 SEUs. “Nominal”, “BCID” (BCID fills), “GF” (glitch filter enabled), “BCID + GF” (BCID fills with glitch filter enabled), and “Clock disabled” refer to the different running scenarios. “INCLUSIVE” is inclusive in the other register types. N.B. the V0 chip, 002, was operated only in the nominal running mode.

The effect of the register triplication is striking: while the V0 ASICs saw a nonzero number of SEUs,¹⁶ the V1 ASICs saw *zero* SEUs for $O(10^7)$ bits measured. By disabling the clock of the triplicated registers in the V1 ASICs, an enhancement of the SEU cross-section is measured, as expected.¹⁷ Additionally, there is no difference in the behaviour between the fixed and BCID fill running modes as well as the glitch filter disabled and enabled running modes. This is expected, as the fixed and BCID fill running modes affect the clusters measured on front-end of the chip, and the glitch filter counters glitches on the 40 MHz clock input pad, which, for the purposes of the SEU beam test, was an ideal source.

5.5.2 December results

In December, the SReg, ADCS, MaskInput, CREG, STAT (fuse ID only), TrimDAC, and CalREG registers are all static in time. Additionally, with the exception of the fuse ID register, all of the

¹⁶The asymmetry in the number of $0 \rightarrow 1$ (16) and $1 \rightarrow 0$ (378) bit flips measured from the V0 register read packets is a result of a vulnerability in the V0 ASICs which was corrected in the V1 ASICs.

¹⁷The V1 register read SEU cross-section with the clock disabled is not necessarily comparable to the V0 register read SEU cross-section. Rather, this measurement is simply a check to verify that the triplication is indeed working in the V1 ASICs.

Table 12. A summary of the inclusive SEU cross-sections for bit-flips in register read data per chip per running mode for the August 2020 beam test. Cross-sections with a prefixed asterisk (*) indicate 95% upper confidence bounds assuming 3 SEUs.

Chip	Fill type	Glitch filter?	Clock disabled?	$n_{0 \rightarrow 1}$	$n_{1 \rightarrow 0}$	$N_{\text{bits,total}}$	$\int d\phi$ [p/cm ²]	$\sigma_{\text{SEU}} \pm \Delta\sigma_{\text{SEU}}$ [cm ² /p]
002	Fixed	No	No	16	378	3,518,912	3.19×10^{15}	$(1.24 \pm 0.12) \times 10^{-13}$
	Fixed	No	No	0	0	1,715,456	1.02×10^{15}	* 2.93×10^{-15}
	BCID	No	No	0	0	1,661,856	1.03×10^{15}	* 2.91×10^{-15}
259	Fixed	Yes	No	0	0	1,452,128	8.87×10^{14}	* 3.39×10^{-15}
	BCID	Yes	No	0	0	1,669,376	1.06×10^{15}	* 2.83×10^{-15}
	Fixed	No	Yes	3	69	68,128	5.32×10^{13}	$(1.4 \pm 0.3) \times 10^{-12}$
	Fixed	No	No	0	0	1,458,240	1.49×10^{15}	* 2.01×10^{-15}
	BCID	No	No	0	0	1,735,744	1.77×10^{15}	* 1.70×10^{-15}
261	Fixed	Yes	No	0	0	1,477,696	1.44×10^{15}	* 2.09×10^{-15}
	BCID	Yes	No	0	0	1,888,000	1.81×10^{15}	* 1.66×10^{-15}
	Fixed	No	Yes	12	166	126,624	1.44×10^{14}	$(1.23 \pm 0.18) \times 10^{-12}$
	Fixed	No	No	0	0	1,212,608	7.97×10^{14}	* 3.78×10^{-15}
	BCID	No	No	0	0	1,884,128	9.19×10^{14}	* 3.27×10^{-15}
267	Fixed	Yes	No	0	0	2,264,416	1.16×10^{15}	* 2.59×10^{-15}
	BCID	Yes	No	0	0	725,504	4.36×10^{14}	* 6.87×10^{-15}
	Fixed	No	Yes	3	79	124,864	7.88×10^{13}	$(1.0 \pm 0.2) \times 10^{-12}$
All V1 chips	All modes except clock disabled			0	0	19,145,152	1.39×10^{16}	* 2.16×10^{-16}

above registers are written to *once* at the start of a run. Accordingly, the SEU calculation verifies that contents of each register is constant with time and integrates the fluence from the initial write time of all registers to their last respective read time. In total, 63 registers were included in the SEU cross-section calculation.

Figure 8 shows the SEU cross-sections for bit-flips in register read data as a function of register type, running mode, and chip. Table 13 summarises the inclusive cross-sections for each chip.

As with the V1 measurements from August, *zero* SEUs were observed for $\mathcal{O}(10^7)$ bits measured. The upper limits for each register type are very consistent between the two beam test campaigns. This is not necessarily expected for upper limits in general, but it indicates a comparable amount of data was obtained by each campaign, which is expected. Combining the inclusive data for all V1 chips (excluding the runs with the clock disabled) in August and December, we have 95% upper confidence bound of 1.17×10^{-16} cm²/p on the SEU cross-section for bit-flips in register read data. Altogether, the results indicate that the triplication implemented in registers of V1 ABCStar ASICs offers excellent protection against SEUs in those registers. This performance is well beyond the needs of operation.

5.6 Current measurements

The analogue and digital currents measured by the FMC-1701 on the SCB for all chips from the August and December beam tests are shown in figures 9 and 10, respectively.




Figure 8. The SEU cross-sections for bit-flips in register read data for the December 2020 beam test. Visible data points as well as their errors are calculated using equations (5.3) and (5.4). Data points which are not visible on the scale of the graph correspond to 95% upper confidence bounds assuming 3 SEUs. “Idle mode”, “Latency 20”, and “Latency 503” refer to the different running scenarios. “INCLUSIVE” is inclusive in the other register types. N.B. for Chip 274’s idle mode runs, register 0x00 (“SCREG”) was not measured.

Table 13. A summary of the inclusive SEU cross-sections for bit-flips in register read data per chip per running mode for the December 2020 beam test. Cross-sections with a prefixed asterisk (*) indicate 95% upper confidence bounds assuming 3 SEUs.

Chip	Idle mode?	Latency	$n_{0 \rightarrow 1}$	$n_{1 \rightarrow 0}$	$N_{\text{bits,total}}$	$\int d\phi$ [p/cm ²]	$\sigma_{\text{SEU}} \pm \Delta\sigma_{\text{SEU}}$ [cm ² /p]
259	No	20	0	0	7,638,624	2.48×10^{15}	* 1.21×10^{-15}
	No	503	0	0	2,941,344	8.76×10^{14}	* 3.42×10^{-15}
267	No	20	0	0	5,695,200	1.60×10^{15}	* 1.88×10^{-15}
	No	503	0	0	2,941,344	7.30×10^{14}	* 4.11×10^{-15}
273	No	20	0	0	8,116,416	2.50×10^{15}	* 1.20×10^{-15}
	No	503	0	0	2,949,408	8.42×10^{14}	* 3.57×10^{-15}
274	Yes	20	0	0	23,675,072	1.37×10^{15}	* 2.19×10^{-15}
	No	20	0	0	2,556,288	6.04×10^{14}	* 4.95×10^{-15}
	No	503	0	0	3,429,216	8.07×10^{14}	* 3.72×10^{-15}
All chips	All modes		0	0	59,942,912	1.18×10^{16}	* 2.54×10^{-16}

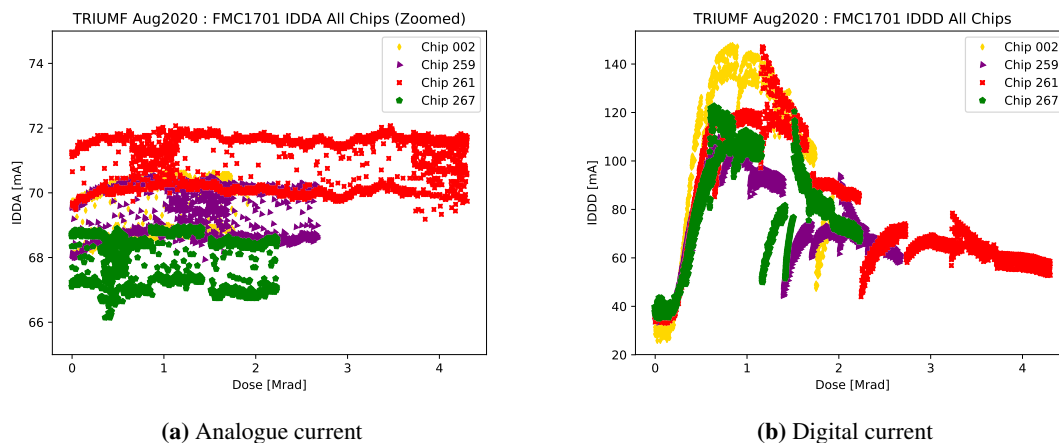


Figure 9. The (a) analogue and (b) digital currents drawn by each ABCStar ASIC as measured by the FMC-1701 during August 2020 beam test. The currents are shown as a function of dose. The analogue currents for the yellow data points lie nearly underneath the purple data points. The piece-wise behaviour of the digital current curves is a result of stopping/starting data collection runs. N.B. the band-like structures in the analogue currents are a result of switching between the different LOBuffer fill configurations during the beam test.

The analogue currents remain constant at approximately 70 mA, regardless of the radiation dose delivered to the ASIC. This is consistent with what was observed for ABC130 ASICs [6]. The digital currents see an initial increase in current, reaching a maximum of 140 mA at 1 Mrad before decreasing to 60 mA at higher doses. This TID bump has been studied extensively in ATLAS ASICs (see for instance section 4.10 of reference [6]). Pre-irradiation digital currents of 40 mA followed by a near 100% increase in current near 1 Mrad before returning to nominal levels is consistent with what was observed for ABC130 ASICs. Given that the ABCStar ASICs were manufactured with the same technology as the ABC130 ASICs, this indicates the TID bump to be a stable feature. It should be noted that the digital current increase is roughly comparable for the ABCStar V0 and V1 ASICs.

During the December beam test, where two new ABCStar V1 ASICs (273 and 274) were irradiated alongside two previously tested ASICs (259 and 267), a comparable current measurement shows a similar increase of the digital current for the new ASICs, as shown in figure 10b). The two previously irradiated ASICs show a digital current increase much lower than the unirradiated ASICs, confirming the effect of pre-irradiation on the TID bump observed for ABC130 ASICs [6].

6 Complementary results from heavy ion measurements

Proton interactions within the ASIC volume cause the protons to recoil, and the deposited energy depends on their recoil angle. Therefore, despite a uniform beam energy of 480 MeV, the observed SEUs correspond to an integrated deposited energy spectrum. In order to calculate the expected rate of SEUs within the detector based on the expected particle energies, complementary measurements were performed for a single ABCStar V1 ASIC using heavy ions (an overview of the ion energies used here is given in table 14) at the cyclotron facility at Louvain-la-Neuve, Belgium [14], where the deposited energy corresponded to the ion energy.

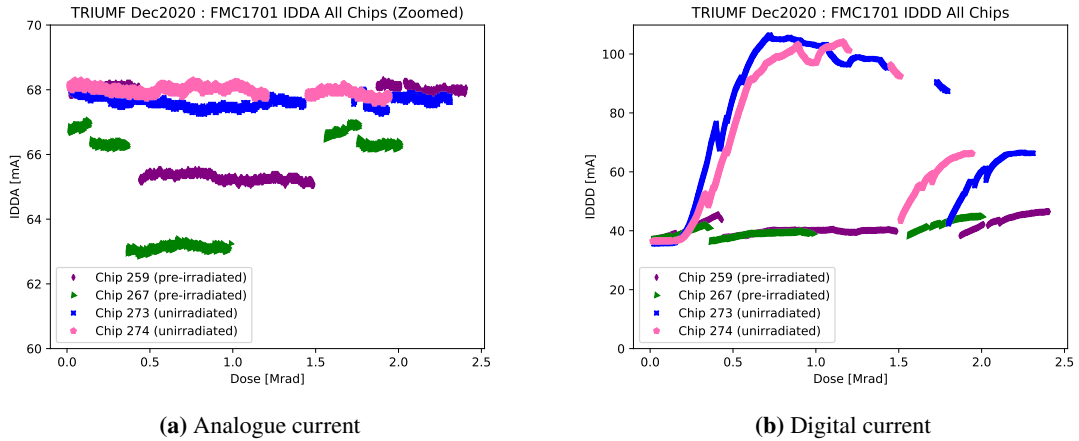


Figure 10. The (a) analogue and (b) digital currents drawn by each ABCStar ASIC as measured by the FMC-1701 during December 2020 beam test. The currents are shown as a function of dose. The piece-wise behaviour of the digital current curves is a result of stopping/starting data collection runs. Two new ASICs were tested and compared with two previously irradiated ASICs, confirming that the total current increase during irradiation is lower for pre-irradiated ASICs of the ABCStar V1 generation.

Table 14. Ions, tilt angles, and corresponding Linear Energy Transfers. The effective LET, LET_{eff} , is given by equation (6.1).

Ion	Energy [MeV]	LET at 0° [MeV/(mg/cm ²)]	Extra tilt angle [$^\circ$]	LET_{eff} [MeV/(mg/cm ²)]
$^{13}\text{C}^{4+}$	131	1.3	67	3.3
$^{22}\text{Ne}^{7+}$	238	3.3	55	5.8
$^{27}\text{Al}^{8+}$	250	5.7	55	9.9
$^{36}\text{Ar}^{11+}$	353	9.9	52	16.1
$^{53}\text{Cr}^{16+}$	505	16.1	38	20.4
$^{103}\text{Rh}^{31+}$	957	46.1	–	–
$^{124}\text{Xe}^{35+}$	995	62.5	–	–

While a full description of the performed tests and results go beyond the scope of this paper, the key result of these measurements showing the energy-dependent cross-sections, based on the same measurement setup as the December proton beam test, are shown below.

The wrong packet and fast command (“Fcmd”) bug cross-sections as a function of the effective linear energy transfer (LET_{eff}) are shown in figure 11. The “Fcmd” bug was identified after chip submission and is due to a small piece of combinatorial logic not being properly triplicated. The sensitive area in the chip of the poorly triplicated logic is very small but it has the nefarious effect of possibly issuing a global chip reset or a BC counter reset if 2 SEUs occur within a few clock cycles of one another. The wrong packet cross-section is given only for physics packets and defines a wrong packet as one with at least one wrong bit in any of the clusters (i.e., the data) — the packet itself is still required to have a structure consistent with that in table 1. N.B. the effective LET may be

calculated from the LET at 0° (LET_{0°), where the ions traverse the ASIC normal to its surface, using:

$$\text{LET}_{\text{eff}} = \frac{\text{LET}_{0^\circ}}{\cos \theta}, \quad (6.1)$$

where θ is the angle from the normal.

A fit has been performed to the resulting wrong packet data, both with and without an angle, as well as to the fast command bug cross-sections using a Weibull function:

$$\sigma(\text{LET}_{\text{eff}}) = \sigma_0 \times \left(1 - \exp \left(- \left[\frac{\text{LET}_{\text{eff}} - L_0}{W} \right]^S \right) \right), \quad (6.2)$$

where σ_0 is the cross-section at saturation, L_0 is the threshold in effective LET, and W and S are parameters of the Weibull function. For the wrong packet cross-section, the fitted parameters are:

$$\begin{aligned} \sigma_0 &= 1 \times 10^{-12} \text{ cm}^2/\text{particle} \\ L_0 &= 0.1 \text{ MeV}/(\text{mg}/\text{cm}^2) \\ W &= 22 \text{ MeV}/(\text{mg}/\text{cm}^2) \\ S &= 1.5 \end{aligned}, \quad (6.3)$$

and for the fast command bug cross-section, the fitted parameters are:

$$\begin{aligned} \sigma_0 &= 5 \times 10^{-14} \text{ cm}^2/\text{particle} \\ L_0 &= 0.1 \text{ MeV}/(\text{mg}/\text{cm}^2) \\ W &= 30 \text{ MeV}/(\text{mg}/\text{cm}^2) \\ S &= 1.2 \end{aligned}. \quad (6.4)$$

These fits are also shown in figure 11.

The resulting Weibull parameters may be used to obtain the cross-sections in the LHC outer tracker environment, which are found to be $4 \times 10^{-19} \text{ cm}^2/\text{proton}$ and $2.4 \times 10^{-20} \text{ cm}^2/\text{proton}$ for the wrong packet and fast command bug cross-sections, respectively.

As noted in figure 11, the wrong packet cross-section is normalised to the total number of physics packets. Given each ASIC from the proton beam test in December (which had a similar readout routine as Louvain) measured $\mathcal{O}(5 \times 10^8)/48 \sim \mathcal{O}(10^7)$ packets and only ran for 25% as long for a single ASIC, this yields a total number of packets¹⁸ $N_{\text{packets}} \sim \mathcal{O}(10^6)$. Additionally, wrong packets tend to have at most one bit flip, and so $N_{\text{packets}}/n_{\text{bit-flips}} \sim \mathcal{O}(1)$. Multiplying these factors onto the wrong packet cross-section yields a cross-section $\mathcal{O}(10^{-13}) \text{ cm}^2/\text{proton}$. This number is very consistent with the SEU cross-sections for bit-flips in physics packet clusters measured from the proton beam tests described in section 5.4, serving as an excellent consistency check.

The fast command bug cross-section is *very* small — we can estimate the expected number of fast command bugs in a given year. The hadron flux calculation follows the same prescription as equation (5.6). With 31,536,000 s/year (the detector will not run the entire year, so this represents a

¹⁸This approximate number is confirmed from actual data taking: each point for the wrong packet cross-sections presented in figure 11 includes a normalisation of 10^6 – 10^7 packets.

worst-case scenario) and $\sim 230,000$ ABCStar ASICs in the ITk, the rate of fast command bugs is:

$$\begin{aligned} \text{Rate}_{\text{Fcmd bugs}} &= (230,000) \text{ ASICs} \times 31,536,000 \text{ s/year} \\ &\quad \times \mathcal{O}(10^7) \text{ hadrons/cm}^2/\text{s} \times (2.4 \times 10^{-20}) \text{ cm}^2/\text{p}/\text{ASIC} \\ &= \mathcal{O}(2) \text{ Fcmd bugs/year} . \end{aligned} \quad (6.5)$$

The rate of fast command bugs is also *very* small. Moreover, even when the fast commands are ASIC resets (as what was measured by the cross-section measurement), reconfiguring a single segment takes 6 ms, and so the data loss due to the fast commands bugs may be concluded to be negligible.




Figure 11. The heavy ion cross-section results from Louvain. The cross-sections for wrong packets (normalised to the total number of packets, with or without an angle) and fast command bugs as well as fits to those cross-sections are shown. The results are plotted as a function of the effective LET, corrected for the angle of the penetrating ions.

7 Conclusion

We have characterised the performance ABCStar ASICs Versions 0 and 1 when irradiated with protons at TRIUMF, Canada. This performance includes measurements of the HPR reset and LCB unlock rates (including those reported in the firmware), the physics packet and register read SEU cross-sections, and the digital currents drawn by the chips.

No HPR reset or LCB unlock events were measured for any of the chips included as part of the proton beam tests at TRIUMF. The SEU cross-sections for bit-flips in physics packet clusters were measured to be $(3.69 \pm 0.03) \times 10^{-12} \text{ cm}^2/\text{p}$ and $(5.32 \pm 0.02) \times 10^{-12} \text{ cm}^2/\text{p}$ for V0 and V1 chips, respectively. The corresponding error rate is $O(10^{-10})$ errors/event/ABCStar during normal operation, which is significantly smaller than the noise occupancy due to thermal noise. The SEU cross-section for bit-flips in register read data was measured to be $(1.24 \pm 0.12) \times 10^{-13} \text{ cm}^2/\text{p}$ for the V0 chip. No SEUs were measured in register read data for the V1 chips, corresponding to a 95% upper confidence bound of $1.17 \times 10^{-16} \text{ cm}^2/\text{p}$ on the cross-section — this result validates the excellent protection conferred by the TMR implemented for V1 registers. Additionally, results obtained with and without the glitch filter enabled were found to be in agreement with each other. As the glitch filter counters glitches on the 40 MHz clock input pad and because this input pad was an ideal source for both SEU beam test campaigns, these results are expected.

Measurements of the digital current drawn by the ABCStar ASICs has confirmed the validity of pre-irradiating the chips to guard against high current loads during runtime. This behaviour is consistent with what was observed for the prototype ABC130 ASICs, demonstrating it to be a stable feature of the technology used to manufacture both the ABC130 and ABCStar ASICs.

Complementary measurements have been performed for ABCStar ASICs irradiated with heavy ions in the cyclotron facility at Louvain-la-Neuve, Belgium. The cross-section for physics packets with wrong clusters obtained from heavy ion measurements were found to agree with the magnitude of the cross-section for bit-flips in physics packet clusters obtained from proton beam measurements. Additionally, the cross-section for fast command bugs obtained from heavy ion measurements was shown to have a negligible impact on detector operation and data taking.

Acknowledgments

This work was supported by the Canada Foundation for Innovation and the Natural Sciences and Engineering Research Council of Canada as well as the Alexander von Humboldt Foundation. We acknowledge the financial support by the Federal Ministry of Education and Research of Germany.

A Dose and fluence measurements from proton tests

Using the calibrations provided by the TRIUMF PIF, the fluence and dose delivered per Secondary Emission Monitor (SEM) counts as a function of the radial distance from the beampipe was fit. The calibration data and resulting fits as well as the integrated SEM as a function of time is shown in figure 12. The total integrated dose or fluence at a time t since the start of data taking, $\phi(t)$, is calculated as:

$$\phi(t) = \text{Calibration}(R) \times \text{SEM}(t), \quad (\text{A.1})$$

where $\text{Calibration}(R)$ is the value of the calibration curve for a chip at radial distance R from the beampipe and $\text{SEM}(t)$ is the integrated SEM.

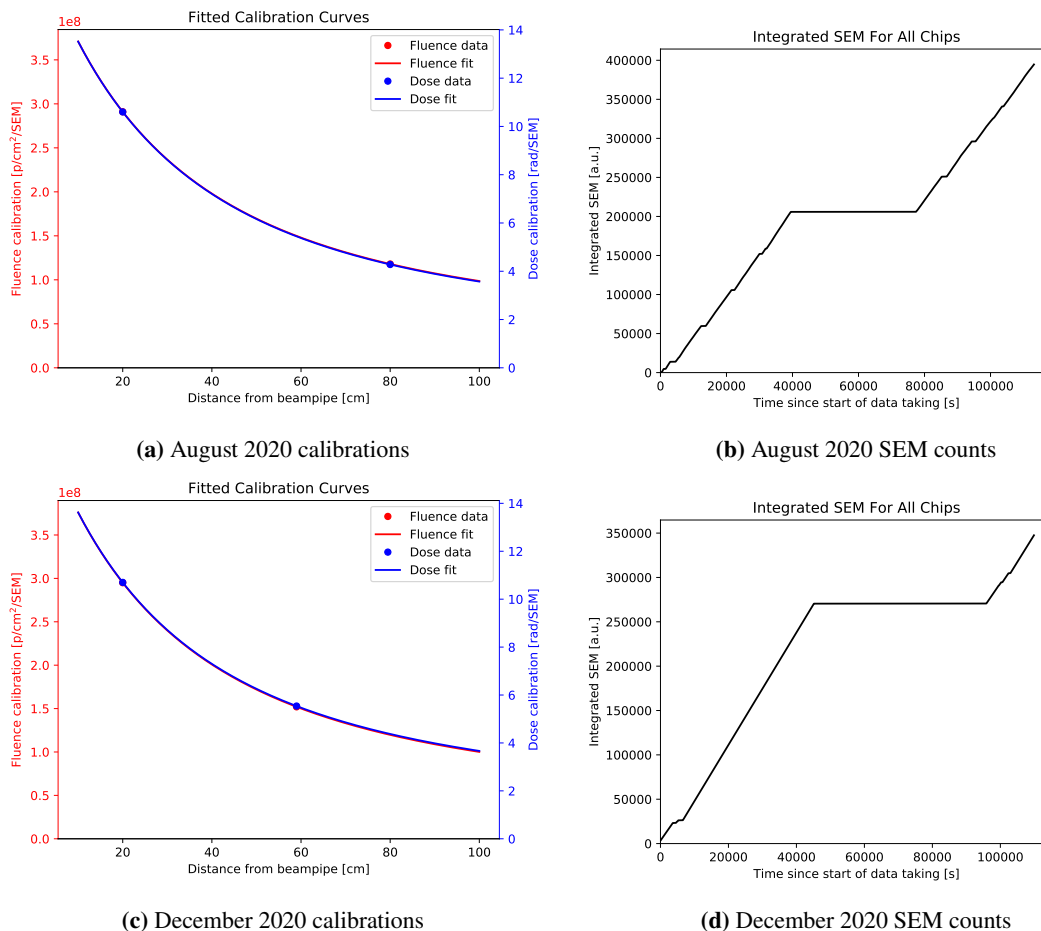


Figure 12. The results of the TRIUMF PIF irradiations. (a) and (c) show the fluence and dose delivered per SEM as a function of radial distance from the beampipe for the August and December 2020 proton beam tests, respectively. The red data points and curves lie underneath the blue data points and curves. (b) and (d) show the integrated SEM for all chips as a function of time since the start of data taking for the August and December 2020 proton beam tests, respectively.

References

- [1] I. Béjar Alonso et al., *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*, Tech. Rep., [CERN-2020-010](#), CERN, Geneva (2020).
- [2] ATLAS collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, [2008 JINST 3 S08003](#).
- [3] C. ATLAS, *Letter of Intent for the Phase-II Upgrade of the ATLAS Experiment*, Tech. Rep., [CERN-LHCC-2012-022](#), [LHCC-I-023](#), CERN, Geneva (2012).
- [4] ATLAS collaboration, *Technical Design Report for the ATLAS Inner Tracker Pixel Detector*, Tech. Rep., [CERN-LHCC-2017-021](#), [ATLAS-TDR-030](#), CERN, Geneva (2017).
- [5] ATLAS collaboration, *Technical Design Report for the ATLAS Inner Tracker Strip Detector*, Tech. Rep., [CERN-LHCC-2017-005](#), [ATLAS-TDR-025](#), CERN, Geneva (Apr, 2017).

- [6] ATLAS collaboration, *The ABC130 barrel module prototyping programme for the ATLAS strip tracker*, 2020 *JINST* **15** P09004 [[arXiv:2009.03197](https://arxiv.org/abs/2009.03197)].
- [7] K.J.R. Cormier, B. Gallop, J. Kaplon, Z. Li, M. Noy, P.W. Phillips et al., *Development of the front end amplifier circuit for the ATLAS ITk silicon strip detector*, 2021 *JINST* **16** P07061.
- [8] S. Kuehn et al., *Prototyping of hybrids and modules for the forward silicon strip tracking detector for the ATLAS Phase-II upgrade*, 2017 *JINST* **12** P05015.
- [9] S. Kulis, *Single Event Effects mitigation with TMRG tool*, 2017 *JINST* **12** C01082.
- [10] Digilent Inc., *Nexys Video FPGA development board documentation*, <https://digilent.com/reference/programmable-logic/nexys-video/reference-manual>.
- [11] Xilinx Inc., *Artix-7 FPGA board documentation*, <https://www.xilinx.com/products/silicon-devices/fpga/artix-7.html>.
- [12] E.W. Blackmore, *Operation of the TRIUMF (20-500 MeV) proton irradiation facility*, in *IEEE Radiation Effects Data Workshop. Workshop Record. Held in conjunction with IEEE Nuclear and Space Radiation Effects Conference*, Reno, NV, U.S.A., 24–28 July 2000, pp. 1–5.
- [13] M. Warren, *ITSDAQ Firmware*, <https://www.hep.ucl.ac.uk/~warren/upgrade/firmware/>.
- [14] L. Standaert, N. Postiau and M. Loiselet, *UCL irradiation facilities status*, in proceedings of the 17th *European Conference on Radiation and Its Effects on Components and Systems*, Geneva, Switzerland, 2–6 October 2017, pp. 1–3.