The VMM3a User Guide

ATL-MUON-PUB-2022-002

George Iakovidis

Abstract

The VMM is a custom Application Specific Integrated Circuit (ASIC). It was designed for the New Small Wheel upgrade of the ATLAS Muon Spectrometer at CERN [1]. It is used in the front-end readout electronics of both the Micromegas and sTGC detectors. It was developed at Brookhaven National Laboratory (BNL) with its latest version, VMM3a, by the DG circuits on behalf of BNL. The L0 circuitry has been design by Sorin Martoiu (IFIN–HH, DFPE). The ASIC is fabricated in the 130 nm Global Foundries 8RF-DM process (former IBM 8RF-DM). The VMM3a has been proposed for many other experiments and application. It is used either packaged in a Ball Grid Array (BGA) with outline dimensions of $21 \times 21\,\mathrm{mm}^2$ and 1 mm pitch, or wire-bonded.

November 7, 2022 Version: 1.0

Contents

1 Introduction

The VMM3a [2] is composed of 64 linear front-end channels. A block diagram of one of the identical channels is shown in Figure 1. Each channel integrates a low-noise charge amplifier (CA) with adaptive feedback, test capacitor, and adjustable polarity (to process either positive or negative charge). The input MOSFET is a p-channel with gate area of $\mathrm{L} \times \mathrm{W} = 180\,\mathrm{nm} \times 10\,\mathrm{mm}$ (200 fingers, 50 µm each) biased at a drain current ID = 2 mA; this corresponds to an inversion coefficient IC ≈ 0.22, a transconductance $g_m \approx 44\,\mathrm{mS}$, and a gate capacitance $C_g \approx 11\,\mathrm{pF}$. The filter (shaper) is a third-order (one real pole and two complex conjugate poles) designed in delayed dissipative feedback (DDF) [3], has adjustable peaking time in four values (25, 50, 100, and 200 ns) and stabilized, band-gap referenced, baseline. The DDF architecture offers higher analog dynamic range, making possible a relatively high resolution at input capacitance much smaller than 200 pF. The gain is adjustable in eight values (0.5, 1, 3, 4.5, 6, 9, 12, 16 mV/fC).

Next to the shapers are the sub-hysteresis discriminators [4] with neighbor enabling logic, and individual threshold trimming, the peak detector, and the time detector. The sub-hysteresis function allows discrimination of pulses smaller than the hysteresis of the comparator circuit. The threshold is adjusted by a global 10-bit Digital to Analog Converter (DAC) and an individual channel 5-bit trimming DAC. The neighbor channel logic forces the measurements of channels neighboring a triggered one, even those channels did not exceed the set threshold. The neighbor logic extends also to the two neighboring chips through bidirectional IO. The peak detector measures the peak amplitude and stores it in an analog memory. The time detector measures the timing using a time-to-amplitude converter (TAC), i.e., a voltage ramp that starts either at threshold crossing or at the time of the peak and stops at a clock cycle of the BC clock. The TAC value is stored in an analog memory and the ramp duration is adjustable in four values (60 ns, 100 ns, 350 ns, 650 ns). The peak and time detectors are followed by a set of three low-power ADCs (a 6-bit, a 10-bit, and a 8-bit), characterized by a domino architecture [5] but of a new concept. These ADC are enabled depending on the selected mode of operation.

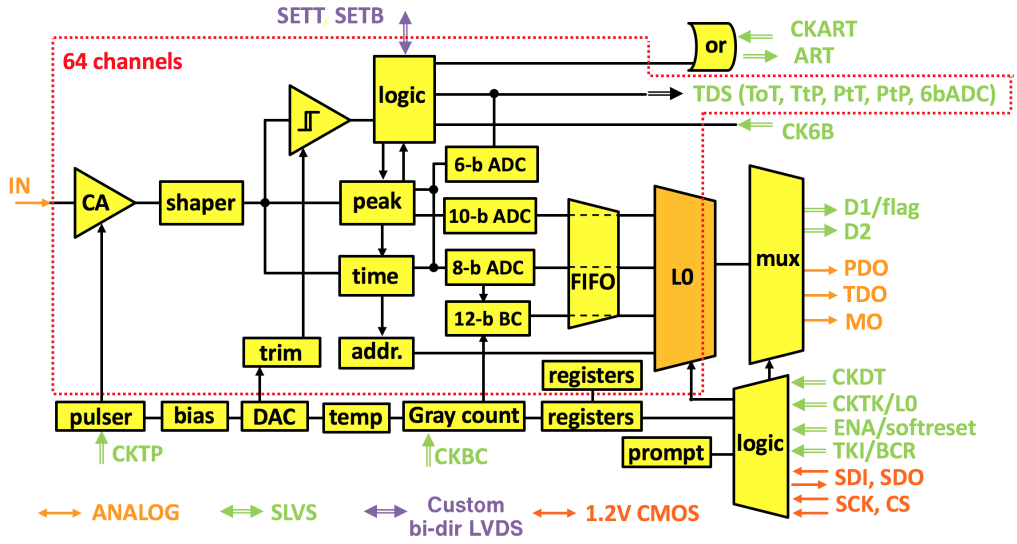

Figure 1: Architecture of VMM3a.

The VMM3a has three modes of operation, a two-phase analog mode, a continuous simultaneous read/write mode and the so-called L0 mode. The ASIC has four independent output data paths:

1. Multiplexed analog amplitude and timing with digital address.
2. Digitized (10-bit amplitude, 20-bit vernier time stamp) in a 2-bit (DDR readout) digital multiplexed mode in either trigger-less continuous mode or the Level-0 with the associated control logic.
3. Address in Real Time (ART) which provides the address of the first channel fired per BC clock.
4. Direct SLVS-400 outputs of all 64 channels in parallel in one of five selectable formats: time-over-threshold (ToT), threshold-to-peak (TtP), peak-to-threshold (PtT), and pulse-at-peak (PtP). Alternatively, it can provide an event flag (either asynchronous or synchronous with the external clock) followed by the 6-bit ADC of the event amplitude information.

The ASIC includes global and acquisition resets and an adjustable pulse generator connected to the injection capacitor of each channel, adjustable with a global 10-bit DAC, and triggered by an external clock. The band-gap reference circuit, a temperature sensor complete the basic features of the VMM. Finally, it integrates analog monitor capability to directly measure the global DACs, the band-gap reference, the temperature sensor, the analog baseline, the analog pulse, and the channel threshold (after trimming).

2 Physical Description

As mentioned above, the VMM ASIC is designed and fabricated in the 130nm Global Foundries 8RFDM process (former IBM 8RF-DM). It is produced on 8-inch wafers consisting of 113 usable ASICs. The VMM3a layout size is $15.308 \times 8.384\,\mathrm{mm}^2$ and the die size $15.308 \times 8.464\,\mathrm{mm}^2$.

The VMM3a is used either wire-bonded or packaged in a 400 ball, 1 mm pitch BGA. The die pinout of the VMM3a is shown in Figure 2[1]..

The BGA device size is $21 \times 21\,\mathrm{mm}^2$. The ball assignment is shown in Figure 3. A detailed BGA pin list and their functions their are shown in Table 1

3 Power

The VMM is designed to operate at a nominal voltage of 1.2 V . It requires four different supplies in order to minimize the contribution to the Equivalent Noise Charge (ENC) of the digital and mixed analog-digital circuits. These four power supplies are:

- **Vddp:** Charge amplifier supply connected to the sources of the p-channel input MOSFETs
- **Vdd:** Powers all other analog circuits

[1]The pin assignment and layout of the VMM3a are identical to those of the VMM3 version.

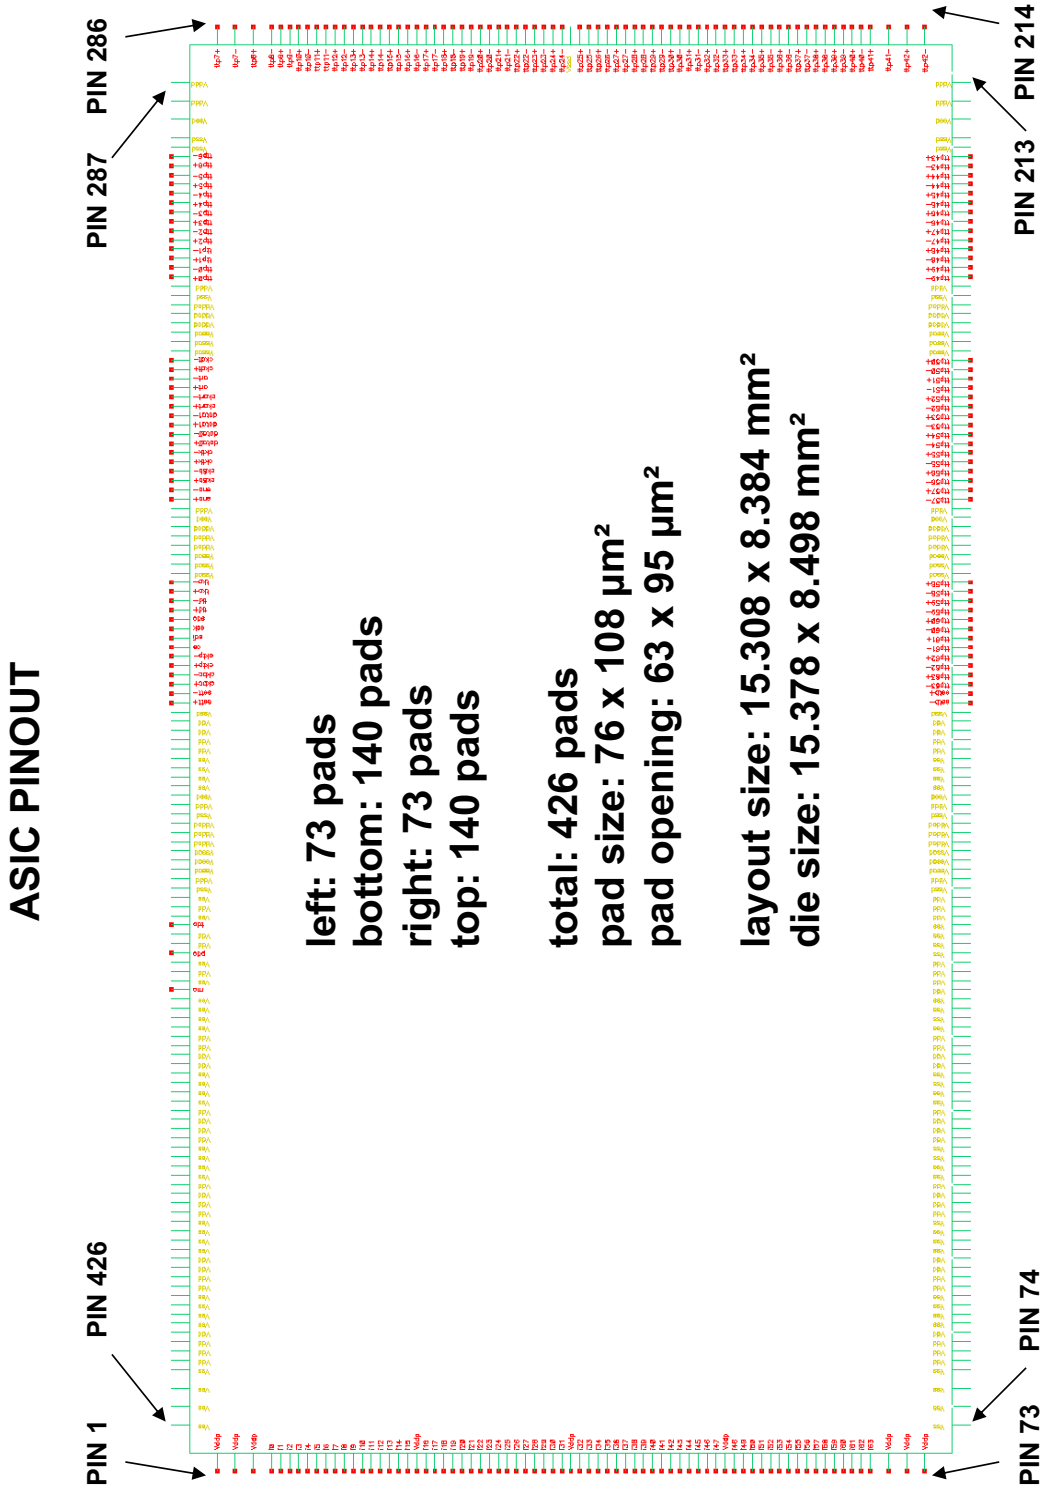

Figure 2: VMM3a pinout (top view)

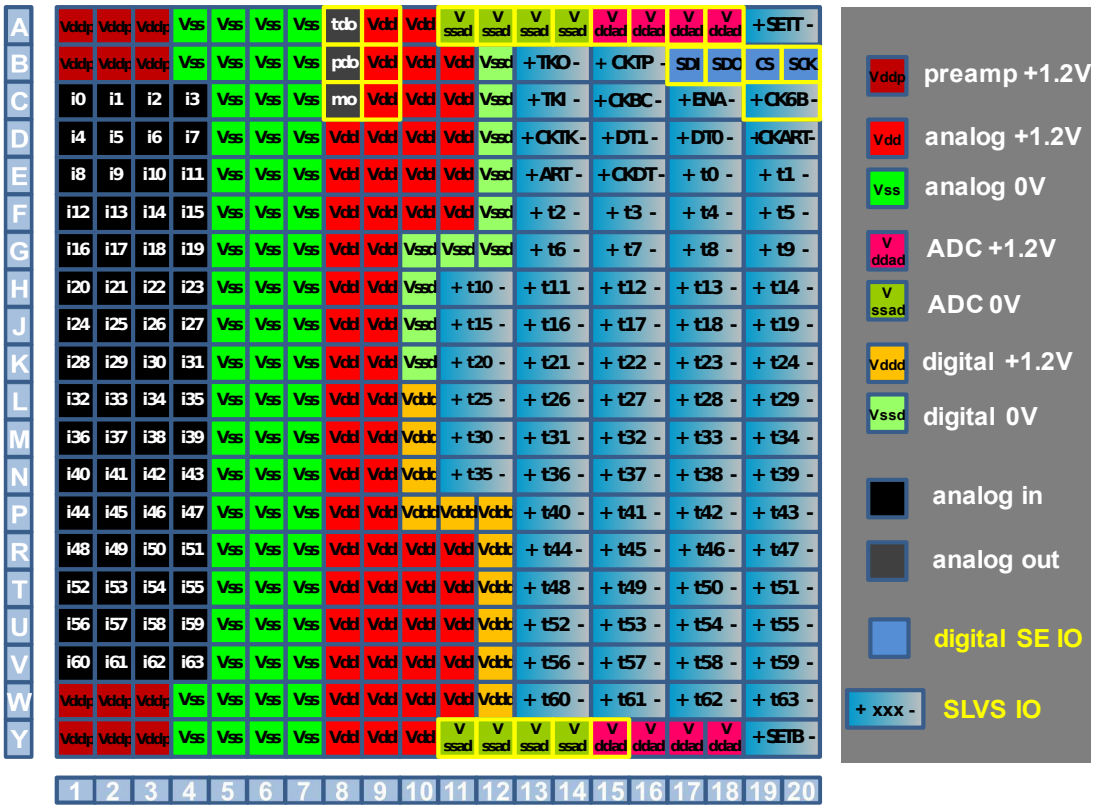

Figure 3: VMM3a pinout (top view)

Table 1: Pad/Pin Assignment/Function

BGA Ball Function	
Ball/Pin name	**Description-Comments**
Vdd,Vss	Analog supplies 1.2 V and grounds 0 V – 123 pins total, max current 400 mA
Vddad, Vssad	Mixed-signal (ADC) supplies 1.2 V and grounds 0 V – 16 pins, max current ∼200 mA
Vddd, Vssd	Digital supplies 1.2 V and grounds 0 V 22 pins
Vddp	Charge amplifier supplies 1.2 V 12 pins, max current ∼150 mA
i0-i63	Analog inputs ESD protected
mo	monitor multiplexed analog output
pdo	Peak Detector multiplexed output.
tdo	Time detector multiplexed analog output.
SETT	Ch 0 neighbor trigger, Custom LVDS - Bi-directional (chip-to-chip)
SETB	Ch 63 neighbor trigger, Custom LVDS - Bi-directional (chip-to-chip)
CKBC	Bunch Crossing clock, SLVS input, Advances 12-bit Gray-code BC counter
CKTP	Test Pulse Clock, SLVS input
SDI	Configuration SPI data input, 1.2 V CMOS
SDO	Configuration SPI data output, Tristated if chip is not selected , CMOS
CS	Configuration SPI chip select, CMOS
SCK	Configuration SPI clock, CMOS
TKI	BCR/OCR, Token input in analog mode. SLVS input
TKO	Token output, Used in analog mode only. SLVS output

ENA	Acquisition start/stop and provides acquisition reset at falling edge:
	• ENA high: acquisition is enabled
	  • internally enabled after 40 ns from ena high
	  • in two-phase (analog) mode is acquisition
	  • in continuous (digital) mode is acquisition and readout

CK6B	6-bit direct output clock, SLVS input
CKTK	Level-0 accept (L0). Token clock, SLVS input
DT0	First data line in digital DDR mode L0, Data 1 in continuous mode (flag and address in analog mode), SLVS output
DT1	Second data line (first bit) in digital DDR mode L0, Flag and Data 0 line in continuous mode, SLVS output
CKART	Address in Real Time (ART) clock, SLVS input
ART	ART output, SLVS output
CKDT	Data clock, SLVS input
ttp0–ttp63	Direct digital outputs, SLVS output

- **Vddad:** Mixed Analog–Digital (ADC)
- **Vddd:** Supplies the digital circuits and SLVS drivers

Table 2 summarizes the requirements and tolerances for the four supplies. The power dissipation depends on the selected functionality and mode of operation. It ranges from 500 mW to 800 mW. For example the SLVS outputs can be disabled when not needed.

Table 2: VMM Power Supply Requirements

Supply	**Voltage[V]**	**Ripple**	**Max Current [mA]**
Vddp	$1.2 \pm 5\%$	$< 10\,\mu$V rms , 1–10 MHz	150
Vdd	$1.2 \pm 5\%$	$< 100\,\mu$V rms , 1–10 MHz	400
Vddad	$1.2 \pm 5\%$	$< 100\,\mu$V rms , 1–10 MHz	200
Vddd	$1.2 \pm 5\%$	< 1 mV rms, 1–10 MHz	150

The requirement of $1.2 \pm 5\%$ must be respected especially in minus values otherwise the ASIC won't be stable for values larger than this.

4 Cooling

The VMM power dissipation depends on the features used with a maximum of ~ 1 W . Although not excessive, if the ASIC is enclosed in a Faraday cage in the high density environment or the expected age of the ASIC is high, cooling is highly recommended. A system with water as coolant is sufficient.

The IBM CMOS8RF Design Manual specifies the operating temperature range to be within the range of -55 °C to 125 °C. However device is designed to work up ∼65 °C. Moreover, the life time degrades rapidly at high temperatures. The case temperature should be kept below 50°C and preferably in the range 30–40 and should be verified and compared to the junction temperature provided by the VMM ASIC. The VMM includes a temperature sensor which can be read out by appropriately programming the monitor output and digitized externally (in configuration mode *scmx = 0, sm5–sm0 = 000100* (see Table 6)). The die temperature is approximately given by:

$$^{\circ}\mathrm{C} = \frac{725 - V_{\mathrm{sensor}}}{1.85}$$

where V_{sensor} is the temperature sensor reading in mV. The case temperature of a single-chip board was measured from turn ON until thermal equilibrium and was compared with the die temperature. The results are shown in Figure 4. The difference of about 2 °C is consistent with the typical junction to case thermal resistance for BGA devices of similar size, $\sim 1\,^{\circ}$C/W.

5 Input and Output

The input and output connections of the VMM are shown in Table 3. It describes the name of the signal and BGA ball address, the direction, the electrical specification and a short description.

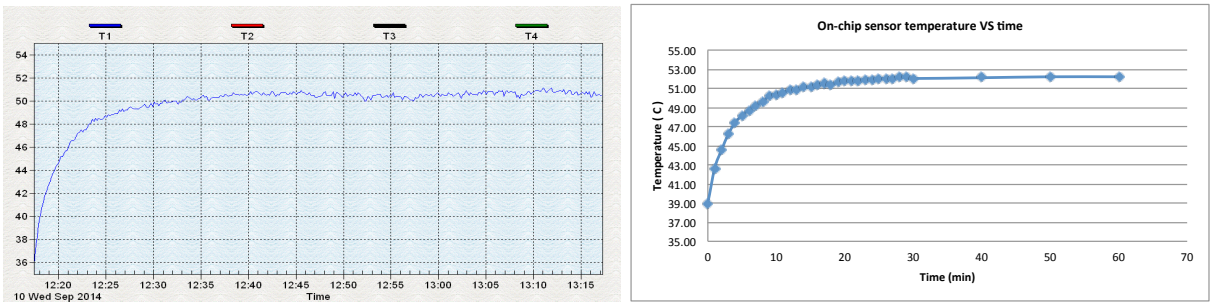

Figure 4: Left plot, the VMM case temperature. Right plot, the junction temperature as a function of time after turn-ON

The single-ended signals are 1.2 V CMOS. The SLVS signals conform to the SLVS-400 JEDEC standard.

The *sett, setb* signals which, are used for the inter-chip communication of the neighbor logic, are custom bi-directional LVDS signals.

6 Configuration Process

The ASIC can be put in configuration mode by having the *ena* signal low and the *cs* low. When in the configuration mode, the ASIC registers are accessible through the SPI clock *sck* and the data inputs *sdi*. The data transmitted are shifted at the falling edge of *sck* in groups of 96-bits and latched when the *cs* is high. The contents of the configuration register is available at the *sdo* output for daisy-chain configuration. Several ASICs can be assembled in a front end card and they can be individually configured. In this mode there are common data input, data output, clock buses and individual chip select pins. For this reason if a VMM is not selected its configuration data output, *sdo*, is tri-stated. The timing diagram of the configuration of up to 8×VMMs is shown in Figure 5.

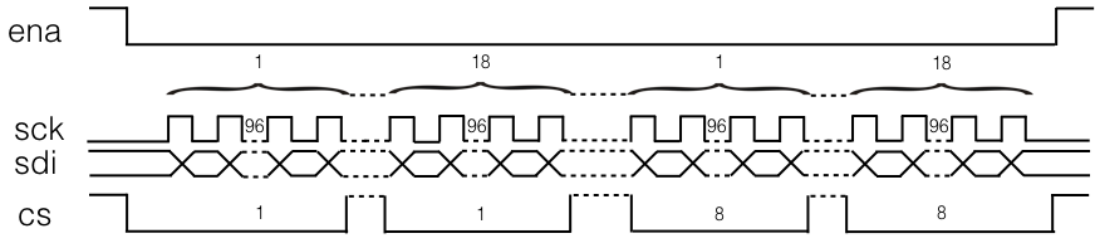

Figure 5: Configuration timing diagram of 8×VMMs.

The configuration is 18 × 96-bits of which 2 × 96-bits are the global registers and the rest are channel registers. Each of the 64 channels has a 24-bit configuration. The list of registers is shown in Tables 6,5. The sequence of registers to be written is global bank-1, channel registers, global bank-2. The first bit to write is the last bit of global bank-2 and the last bit to write is the global bank-1 first bit. The sequence is shown in Table 4.

Table 3: Input and Output Signals of the VMM3a.

Name, Position	**In, Out or I/O**	**Type of Signal or Max/Min**	**Description**
sett A19-20	I/O	Custom LVDS Bi-directional	Channel 0 force-neighbor signal
setb Y19-20	I/O	Custom LVDS Bi-directional	Channel 63 force-neighbor signal
ckbc (BCclk) C15-16	In	SLVS	Bunch crossing clock of 40 MHz / External trigger signal
cktp (Test Pulse) B15-16	In	SLVS	Test pulse clock
cktk (Level-0) D13-14	In	SLVS	Token clock / L0
ckdt (ROclk) E15-16	In	SLVS	Data clock
ckart (ARTclk) D19-20	In	SLVS	ART clock
sdi B17	In	CMOS	Configuration data input
sdo B18	Out	CMOS	Configuration data output
cs B19	In	CMOS	Chip Select, active low
sck B20	In	CMOS	Input SPI clock
t0-t63 E17-W20	Out	SLVS	Direct digital outputs
mo C9	Out	0-1 V	Analog output
tki (BCR/OCR) C13-14	In	SLVS	Token input (an. mode) / (BCR-OCR in L0) / acceptance window in cont. mode
tko B13-14	Out	SLVS	Token output (analog mode)
ena (ENA/Soft Reset) C17-18	In	SLVS	Acquisition start/stop
ck6b C19-20	In	SLVS	6-bit ADC Clock
art E13-14	Out	SLVS	Address in Real Time
data0 D15-16	Out	SLVS	data line
data1 D17-18	Out	SLVS	data line first bit, flag in cont.

Table 4: Sequence of Configuration Registers of the VMM3a.

Register Type	**Sequence**
global registers - first bank	sp sdp sbmx sbft sbfp sbfm slg sm5:sm0 scmx sfa sfam st1:0 sfm sg2:0 sng stot sttt ssh stc1:0 sdt9:0 sdp9:0 sc010b:sc110b sc08b:sc18b sc06b:sc26b s8b s6b s10b sdcks sdcka sdck6b sdrv stpp res00 res0 res1 res2 res3 slvs s32 stcr ssart srec stlc sbip srat sfrst slvsbc slvstp slvstk slvsdt slvsart slvstki slvsena slvs6b sL0enaV slh slxh stgc nu nu nu nu nu reset reset
channel register (64×)	sc sl st sth sm smx sd0:sd4 sz010b:sz410b sz08b:sz38b sz06b:sz26b nu
global registers - second bank	nu0:30 nskipm_i sL0cktest sL0dckinv sL0ckinv sL0ena truncate_i0:5 nskip_i0:6 window_i0:2 rollover_i0:11 l0offset_i0:11 offset_i0:11 nu0:7

Table 5: Channel Configuration Registers of the VMM.

Channel bits (defaults are 0)	**Description**
sc [0 1]	large sensor capacitance mode ([0] <∼200 pF , [1] >∼200 pF)
sl [0 1]	leakage current disable [0=enabled]
st [0 1]	300 fF test capacitor [1=enabled]
sth [0 1]	multiplies test capacitor by 10
sm [0 1]	mask enable [1=enabled]
sd0-sd4 [0:0 through 1:1]	trim threshold DAC, 1 mV step ([0:0] trim 0 V , [1:1] trim -29 mV)
smx [0 1]	channel monitor mode ([0] analog output, [1] trimmed threshold))
sz010b, sz110b, sz210b, sz310b, sz410b	10-bit ADC offset subtraction
sz08b, sz18b, sz28b, sz38b	8-bit ADC offset subtraction
sz06b, sz16b, sz26b	6-bit ADC offset subtraction

Table 6: Global Configuration Registers of the VMM3.

Global bits (defaults are 0)	**Description**
sp	input charge polarity ([0] negative, [1] positive)
sdp	disable-at-peak
sbmx	routes analog monitor to PDO output
sbft [0 1], sbfp [0 1], sbfm [0 1]	analog output buffers, [1] enable (TDO, PDO, MO)
slg	leakage current disable ([0] enabled)
sm5-sm0, scmx	monitor multiplexing. • Common monitor: scmx, sm5-sm0 [0 000001 to 000100], pulser DAC (after pulser switch), threshold DAC, band-gap reference, temperature sensor) • channel monitor: scmx, sm5-sm0 [1 000000 to 111111], channels 0 to 63
sfa [0 1], sfam [0 1]	ART enable (sfa [1]) and mode (sfam [0] timing at threshold, [1] timing at peak)
st1,st0 [00 01 10 11]	peaktime (200, 100, 50, 25 ns)
sfm [0 1]	enables full-mirror (AC) and high-leakage operation (enables SLH)
sg2,sg1,sg0 [000:111]	gain (0.5, 1, 3, 4.5, 6, 9, 12, 16 mV/fC)
sng	neighbor (channel and chip) triggering enable
stot [0 1]	timing outputs control 1 (s6b must be disabled) • stpp,stot[00,01,10,11]: TtP,ToT,PtP,PtT • TtP: threshold-to-peak • ToT: time-over-threshold • PtP: pulse-at-peak (10ns) (not available with s10b) • PtT: peak-to-threshold (not available with s10b)
sttt [0 1]	enables direct-output logic (both timing and s6b)
ssh [0 1]	enables sub-hysteresis discrimination
stc1,stc0 [00 01 10 11]	TAC slope adjustment (60, 100, 350, 650 ns)
sdt9-sdt0 [0:0 through 1:1]	coarse threshold DAC
sdp9-sdp0 [0:0 through 1:1]	test pulse DAC
sc010b,sc110b	10-bit ADC conv. time (increase subtracts 60 ns)
sc08b,sc18b	8-bit ADC conv. time (increase subtracts 60 ns)
sc06b, sc16b, sc26b	6-bit ADC conversion time
s8b	8-bit ADC conversion mode
s6b	enables 6-bit ADC (requires sttt enabled)
s10b	enables high resolution ADCs (10/8-bit ADC enable)
sdcks	dual clock edge serialized data enable
sdcka	dual clock edge serialized ART enable
sdck6b	dual clock edge serialized 6-bit enable
sdrv	tristates analog outputs with token, used in analog mode
stpp [0 1]	timing outputs control 2
slvs	enables direct output IOs
stcr	enables auto-reset (at the end of the ramp, if no stop occurs)
ssart	enables ART flag synchronization (trail to next trail)

Table 6 (continued)

Global bits (defaults are 0)	Description
s32	skips channels 16-47 and makes 15 and 48 neighbors
stlc	enables mild tail cancellation (when enabled, overrides sbip)
srec	enables fast recovery from high charge
sbip	enables bipolar shape
srat	enables timing ramp at threshold
sfrst	enables fast reset at 6-b completion
slvsbc	enable slvs 100 Ω termination on ckbc
slvstp	enable slvs 100 Ω termination on cktp
slvstk	enable slvs 100 Ω termination on cktk
slvsdt	enable slvs 100 Ω termination on ckdt
slvsart	enable slvs 100 Ω termination on ckart
slvstki	enable slvs 100 Ω termination on cktki
slvsena	enable slvs 100 Ω termination on ckena
slvs6b	enable slvs 100 Ω termination on ck6b
sL0enaV	disable mixed signal functions when L0 enabled
reset reset	Hard reset when both high
sL0ena	enable L0 core / reset core & gate clk if 0
l0offset_i0:11	L0 BC offset
offset_i0:11	Channel tagging BC offset
rollover_i0:11	Channel tagging BC rollover
window_i0:2	Size of trigger window
truncate_i0:5	Max hits per L0
nskip_i0:6	Number of L0 triggers to skip on overflow
sL0cktest	enable clocks when L0 core disabled (test)
sL0ckinv	invert BCCLK
sL0dckinv	invert DCK
nskipm_i	magic number on BCID - 0xFE8
slh, slxh	increases bias current at input node from nominal 1nA to 15nA or 300nA respectively
stgc	extreme charge handling compensation

7 Readout

The VMM3a has three modes of operation, a two-phase analog mode, a continuous simultaneous read/write mode and the so-called L0 mode. It also features the ART output and the direct digital output per channel which can be used in parallel with the three readout modes.

7.1 Two-Phase Analog Mode

In two-phase (analog) mode (bit *s10* low), which is the mode originally implemented in the VMM1, the ASIC operates in two separate phases: acquisition with *ena* high and readout

with *ena* low. During the acquisition phase the events are processed and stored in the analog memories of the peak and time detectors. As soon as a first event is processed, a flag is raised at the digital output *data1*. Once the acquisition is complete the ASIC can be switched to the readout phase and the readout proceeds injecting a token at the token input *tki*. The first set of amplitude and time voltages is made available at the analog outputs *pdo* and *tdo*. Analog buffers can be enabled using the bits *sbfp* and *sbft*. The address of the channel is serialized and made available at the output *data1* using six data clocks. The next channel is read out by advancing the token with the token clock. The token is sparse, passed only among those channels with valid events. If, after the token clock occurs, the *data0* goes low, the readout is complete and the token is routed to the output *tko* for the readout of the next chip. This allows a daisy-chained readout with a single token input. Figure 6 shows the complete timing diagram of the analog mode operation. This mode was tested in VMM3a and found to be working correctly [6].

It should be noted that the two fast readout paths, the 64-channel direct digital outputs as well as the ART stream are active in this mode as well.

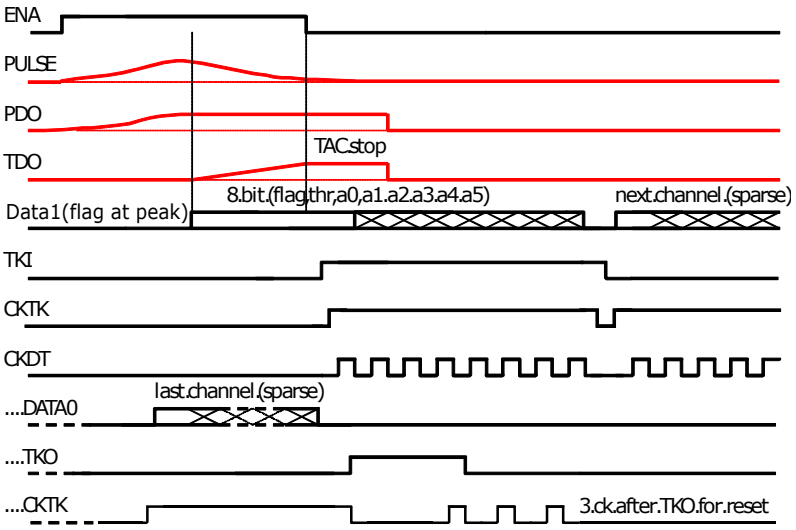

Figure 6: Data Readout with PDO,TDO and external ADC (2-phase mode).

7.2 Continuous Mode

In this mode the peak and time detectors convert the voltages into currents that are routed to the 10-bit ADC and 8-bit ADC respectively. The 10-bit ADC provides a high resolution A/D conversion of the peak amplitude in a conversion time of about 200 ns from the occurrence of the peak. The conversion time and baseline (150 mV pre-subtracted) are adjustable using the global bit set *sc10b* (the conversion time is a 200 ns base plus a 60 ns increment for the MSB and LSB phases, set by the *sc10b* bits) and the channel bit set *sz10b* respectively. The 8-bit ADC provides the A/D conversion of the timing (measured using the TAC) from the time of the peak or the threshold to a stop signal. The TAC stop signal occurs at a next clock cycle of a shared 12-bit Gray-code counter which is incremented using the external clock signal BC. The counter value at the TAC stop time is latched into a local 12-bit memory, thus providing a total of 20-bit deep timestamp with a nanosecond resolution. The conversion time and baseline (zeroing) are adjustable using the global bit set *sc8b* (the conversion time is a 100 ns base plus

a 60 ns increment for the MSB and LSB phases, set by the *sc8b* bits) and the channel bit set *sz8b* respectively. The channel is reset once both the 8-bit and 10-bit conversions are complete (or earlier if the *sfrst* is high) and the digital values are latched in digital memories. Thus, in continuous (digital) mode a total of 38-bits are generated for each event. The first bit is used as a readout flag, the second is the threshold crossing indicator (allows discrimination between above-threshold and neighbor events). Next is a 6-bits word for the channel address, followed by 10-bits associated with the peak amplitude, and 20-bits associated with the timing. The 38-bit word is stored in a derandomizing FIFO (there is one such FIFO per channel) and it is read out using a token-passing scheme where the token is passed first-come first-serve only among those FIFOs that contain valid events. The first token is internally generated as needed (*tki* becomes an acceptance window, if high at least for one *ckbc* then the event is processed otherwise discarded, *tko* IO is not used) and advanced with the token clock. The data in the FIFOs is thus sequentially multiplexed to the two digital outputs *data0* and *data1*. The second output *data1* is also used as a flag, indicating that events need to be read out from the chip. The external electronics release a sync signal using the token clock as well (i.e. the token clock provides both advancement and data output synchronization), after which the 38-bit data is shifted out in parallel to the *data0* and *data1* outputs using either 19 clock cycles or 19 clock edges of the external data clock, depending on the global bit *sdcks*. The timing diagram relevant to this mode is shown in Figure 7.

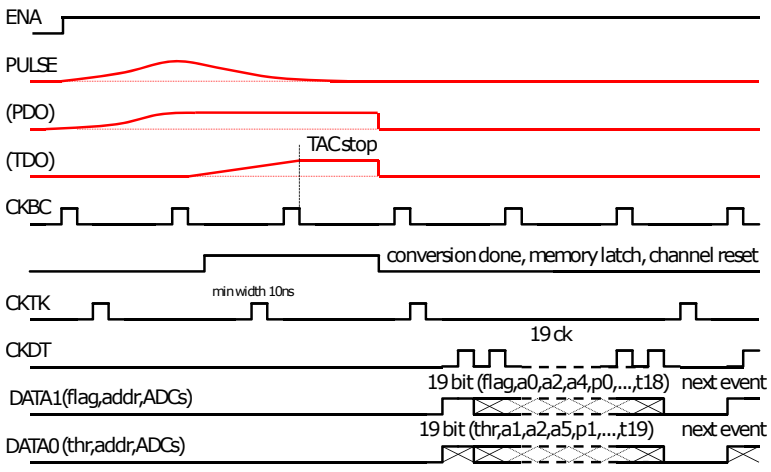

Figure 7: Data Readout with ADCs (continuous mode, 1 bit/ck). This mode can be configured to work in DDR with 2 bit/ck.

7.3 L0 Readout

This mode was specifically designed to be compatible with the ATLAS readout scheme adding the ability to buffer the data in deeper memory while enabling the chip to select the data based on an external trigger. The ASIC functions the same way like the continuous mode but the readout is different. The data are registered in a 64-deep "latency" FIFO where all hits can be buffered. Each channel has a selector circuit based on the arrival of the external trigger signal which finds events within the timing window configured through the *window i0:2* register (maximum size is eight clock cycles). A timestamp generated by a bunch-crossing counter is stored in the

latency FIFO together with the digitized hit data. At every clock cycle, the timestamp of the first hit available for reading in the FIFO is compared with the timestamp corresponding to the trigger signal, corrected by a configurable latency offset configured through the *l0offset i0:11* register. Once there is a match within the timing window, the data are copied to the back-end "channel" FIFO, but it is not immediately deleted from the latency FIFO, because it may be required by a subsequent trigger which may have its time window overlapping with the first one. The hit is only deleted if its timestamp is equal with the oldest timestamp in the current trigger time window, therefore there can be no other trigger requiring it. This condition is observed regardless of whether there is an active trigger or not, such that old hits are flushed automatically from the latency FIFO. If the a latency FIFO does not contain a hit matching the trigger condition, a "null" hit is inserted into the corresponding back-end FIFO, identified by a dedicated bit. This makes the back-end FIFOs behave as a unitary big FIFO such that the control pointers and signals can be shared. The bunch crossing information corresponding to each triggered event is stored in a separate FIFO (identified as "L0 BCID FIFO"), while the back-end channel FIFOs contain only the relative timestamp with respect to the trigger, within the trigger time window. This reduces the required bit-width of the event information. The trigger timestamp FIFO is double the size of the channel latency FIFO such that, in case the latter overflows, the bunch crossing information remains valid. In case that the bunch crossing FIFO overflows as well, then the VMM3a sends out data with a so-called "magic bc"(0xFE8) which is a value of the bunch crossing higher than the maximum allowed value in the ATLAS experiment. This way, the back-end readout electronics system is aware of the overflow. The maximum BC value allowed is configure through the *rollover i0:11* register while an offset can be set for the channels configured through the *offset i0:11* register. The overall buffering schema is shown on Figure 8.

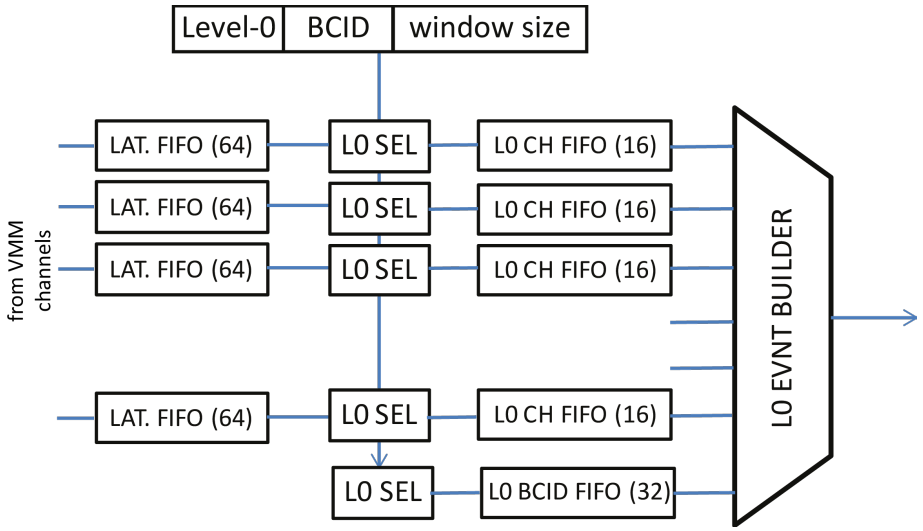

Figure 8: VMM3a internal buffers for Level-0 event output. The formats of the data at points A, B, C and D are shown in the figures following.

The "Event Builder" circuit reads the triggered hit data corresponding to an event from all channels of the back-end FIFO starting with the trigger timestamp FIFO along with the obit

counter and the parity bit. If the current trigger caused the L0 Ch FIFOs to become full, the overflow bit, V, in the BCID item is set and further writing to the L0 Ch FIFOs is suspended for the configured number of triggers to be skipped through the *nskip i0:6*. This allows the FIFOs to recover. After the header, the hit data of each channel in consecutive order are appended, while channels containing null hits are skipped. The hit data contain a parity bit signalling if there was an error, a truncate bit which is triggered once the event size contains more hits that the configured through the *truncate i0:5* register, the channel 6-bit data, the charge 10-bit data, the time 8-bit data, a flag showing if this was a neighbouring channel and the relative BCID calculated based on the distance to the event BCID. The data format is shown in Figure 9.

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
header	V	P	orbit		BCID (12 bits)																											
hit data	1	P	R	T	Channel (6 bits)						Charge PDO (10 bit)										Time TDO (8 bits)								N	rel BCID		

Figure 9: Format of the data sent to the ROC ASIC (Point D in Figure 8) Zero or more hits may follow the header.

The data transfer from VMM3a is achieved via two serial lines running at 160 MHz with DDR giving a total bandwidth of 640 Mb/s. The data are encoded using the 8b/10b protocol which leads to an effective bandwidth of 560 Mb/s. The overall buffer scheme is shown in Figure 8. The maximum hit rate per channel is 4 MHz driven by the 250 ns digitisation dead-time and therefore the maximum latency guaranteeing that no data loss is occurred is 16 µs. The clock-out of the data is shown in Figure 10.

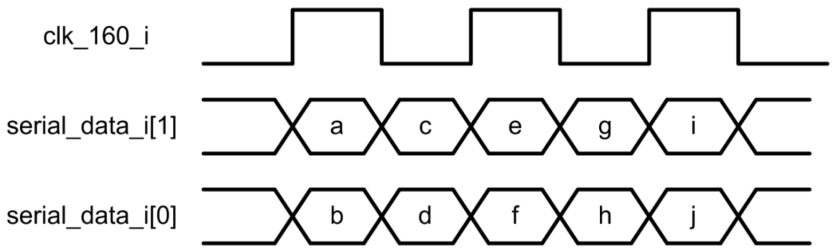

Figure 10: Data as clocked out from the VMM.

7.4 Fast direct outputs

The VMM, in addition to the information recorded at Level-0, provides fast outputs that can be used to get quick signals out of the detector.

Address in Real time (ART) The ART is the address with the earliest registered hit within a BC clock. This mode is enabled with the bit *sfa*. Either at the pulse threshold crossing (bit *sfam* low) or at the pulse peak (bit *sfam* high) a flag is released at the art output. The flag is followed by the serialized address of the event. Also in this case the address is released either at each clock cycle or at each clock edge of the external ART clock, depending on the global bit *sdcka*. It takes 2 CKART clock cycles (13ṅs at 160 MHz) to reset the ART circuit after the release of the last ART bit while the *ssart* register enables the ART flag synchronization with the CKART. Figure 11 shows the ART readout in 1-bit per clock while DDR is also possible.

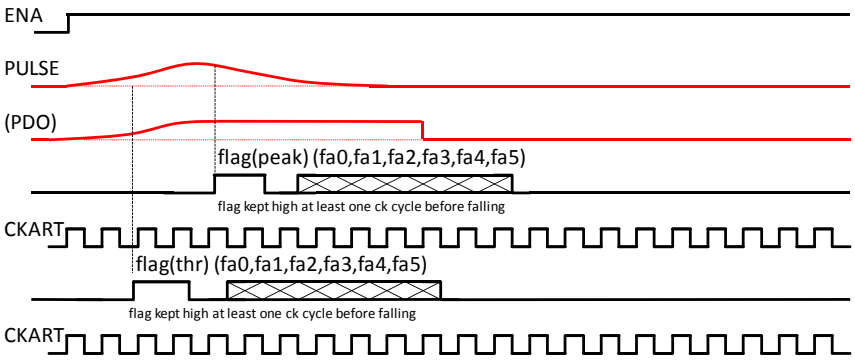

Figure 11: ART data as clocked out from the VMM (1 bit/cl).

Direct Digital Outputs Each of the 64 channels feature a dedicated digital output *ttp0-ttp63*. The outputs are activated when *sttt* is high and provide one of four different timing pulses: time-over-threshold (ToT), threshold-to-peak (TtP), peak-to-threshold (PtT), or a 10 ns pulse occurring at peak (PtP), and can be set using the global bits *stot* and *stpp*. The channel self resets at the end of the timing pulse, thus providing continuous and independent operation of all 64-channels. Alternatively, if the bit *s6b* is set high, the peak detector converts the voltage into a current that is routed to the 6-bit ADC. The 6-bit ADC provides a low-resolution A/D conversion of the peak amplitude in a conversion time of about 25 ns from the peak time. The conversion time and the baseline (zeroing) are adjustable using the global bit set *sc6b* (the conversion time is the number of data clocks set by the *sc6b* bits) and the channel bit set *sz6b* respectively. The serialized 6-bit data is made available at the channel output immediately after an event flag which occurs at the peak time. The flag is lowered at the next clock cycle of the data clock, and the 6-bit ADC data is shifted out after that, either at each clock cycle or at each clock edge of the data clock depending on the global bit *sdck6b*. The channel reset occurs after the last bit has been shifted out. The dead time because of the 10-bit ADC latency is ~ 200 ns after the peak. However it is possible to interrupt the 10-bit ADC conversion when the signal drops below threshold (the earliest the channel can be reset) or when the 6-bit conversion finishes thus providing a lower resolution peak value with the minimum dead time by enabling the bit *sfrst*. The readout of the data is shown in Figure 12.

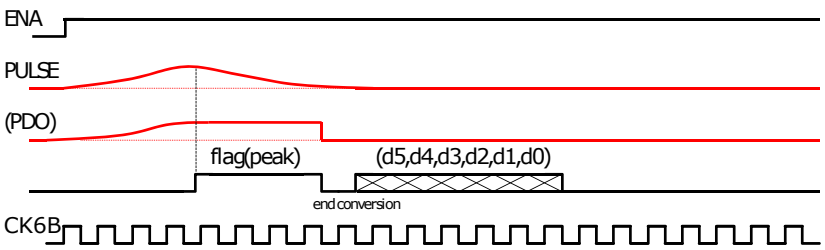

Figure 12: The 6-bit direct data as clocked out from the VMM (1 bit/cl). DDR is also possible.

8 Radiation Tolerance and SEU

The VMM ASIC was designed targeting the expected radiation levels of the NSW upgrade. It is expected to be exposed to a total ionization dose (TID) of ∼100 krad according to the simulations done [1]. It is also expected to be operated in harsh neutron environment of $\sim 10^{12}\, n/\mathrm{cm}^2$/year of the ATLAS operation.

Design techniques are applied to mitigate issues that may affect the operation of the ASIC under the above targeted conditions. Although TID may degrade the performance of the ASIC, the VMM3a was tested for TID tolerance in the ^{60}Co source irradiation facility at BNL for the expected radiation and no performance degradation was noticed. Single event upsets (SEU) though become increasingly more serious as the CMOS technology feature size decreases because of the smaller capacitance in the storage elements that need smaller energy depositions in order to flip their state. In the VMM3a there are several types of elements that require SEU protection. The places protected are listed in Table 7. It uses established techniques like Dual Interlocked storage Cells (DICE) and Triple Module Redundancy (TMR) while in the long storage elements like the latency FIFO, an upset is just flagged once detected and the FIFO resets.

The DICE uses redundancy to significantly reduce susceptibility to an upset. The D flip flops based on the dual interlocked cell latches have redundant storage nodes and restore the cell's original state when an SEU error is introduced in a single node [7]. The scheme fails if multiple nodes are upset but this is far less likely, especially at the rather modest neutron levels of $\sim 10^{12}\, n/\mathrm{cm}^2$/year at luminosity $7 \times 10^{34}\, \mathrm{s}^{-1}\mathrm{cm}^{-2}$.

The L0 circuit is implemented using standard logic cells. Within the SEU mitigation policy applied for this block, data loss is accepted given that corrupted data can be identified and discarded, while for the control circuits, higher degree of protection is used such that synchronism is preserved as much as possible.

A parity bit is computed for each hit data which is registered in the latency FIFO on each channel. This bit travels throughout the entire readout system together with the hit data and is checked at the end of the chain. Since this bit is generated before the relative timestamp is computed, the validation can only be done after the initial timing information is reconstructed. If, finally, a parity error is detected, data cannot be reconstructed and the hit is discarded. Similarly, the trigger timestamp information is protected by a dedicated parity bit.

The latency FIFOs are implemented as shift-registers, as this topology requires a single pointer and multiplexer, reducing the area and routing requirement at the expense of a little increase in power consumption. Another positive consequence is that the number of failure points is reduced. The FIFO pointers are protected by a parity bit which is checked at each clock cycle. If a mismatch occurs, the channel FIFO is reset. The data is therefore lost, but any possible deadlock caused by pointing to the wrong data cell is avoided.

As discussed earlier, the back-end L0 FIFOs form essentially a big unitary FIFO therefore their controls signals are shared. Also in this case the number of potential failure points is greatly reduced. The back-end FIFO control circuits and the Event Builder state machines are implemented using Triple Modular Redundancy.

At the measured upset cross sections of order $10^{-14}\mathrm{cm}^2$/bit the upset probability is of order one tenth of an upset per VMM3a per year in the ATLAS operation.

Table 7: Protection schema in VMM3a

Block	Type of protection
Global configuration and channel registers	DICE
VMM State Machine	TMR
Bunch Crossing Counter	TMR
L0 FIFO Control	TMR
L0 Event Builder	TMR
L0 Accept register, NSkip Circuit	TMR
Latency FIFO	Parity on pointer, FIFO reset on parity error

Appendices

The overall connection scheme of the VMM in the NSW is shown in Figure 13.

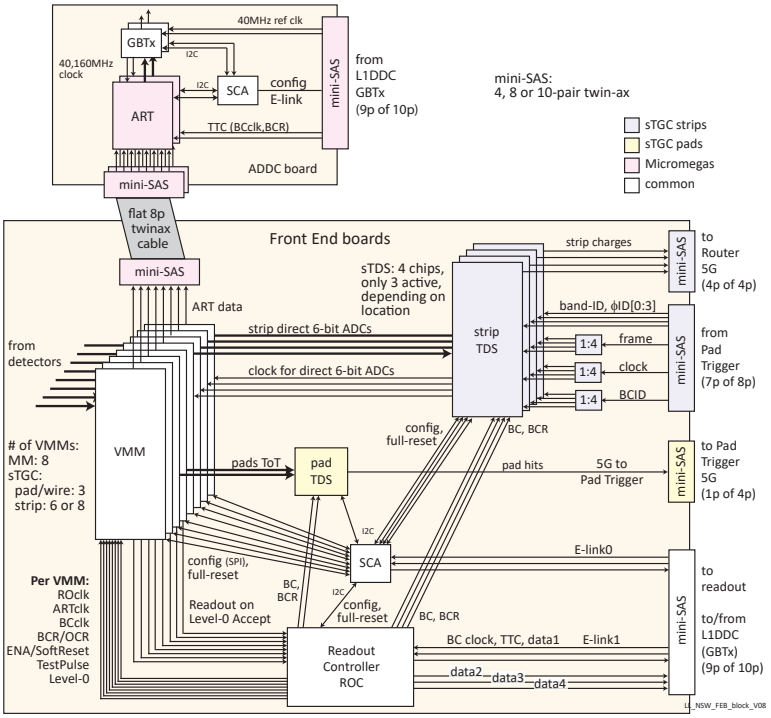

Figure 13: Overall connection diagram of the VMM in NSW, courtesy of L. Levinson

Conventions and Glossary

References

[1] CERN, *New Small Wheel Technical Design Report*, Geneva, Jun 2013, no. CERN-LHCC-2013-006. ATLAS-TDR-020.

[2] G. De Geronimo, G. Iakovidis, S. Martoiu, and V. Polychronakos, "The vmm3a asic," *IEEE Transactions on Nuclear Science*, pp. 1–1, 2022.

[3] G. De Geronimo and S. Li, "Shaper design in cmos for high dynamic range," *Nuclear Science, IEEE Transactions on*, vol. 58, no. 5, pp. 2382–2390, Oct 2011.

[4] G. De Geronimo, J. Fried, S. Li, J. Metcalfe, N. Nambiar, E. Vernon, and V. Polychronakos, "Vmm1 - an asic for micropattern detectors," *Nuclear Science, IEEE Transactions on*, vol. 60, no. 3, pp. 2314–2321, June 2013.

[5] G. De Geronimo, J. Fried, G. Smith, B. Yu, E. Vernon, C. Britton, W. Bryan, L. Clonts, and S. Frank, "Asic for small angle neutron scattering experiments at the sns," *Nuclear Science, IEEE Transactions on*, vol. 54, no. 3, pp. 541–548, June 2007.

[6] T. Alexopoulos, C. Bakalis, G. de Geronimo, and V. Polychronakos, "Characterization of the vmm front-end asic for high-resolution applications," in *2020 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, 2020, pp. 1–3.

[7] F. Faccio *et al.*, "Total dose and single event effects (SEE) in a 0.25 µm CMOS technology," in *Electronics for LHC experiments. Proceedings, 4th Workshop, Rome, Italy, September 21-25, 1998*, 1998.