



Contents lists available at ScienceDirect

Nuclear Inst. and Methods in Physics Research, A

journal homepage: www.elsevier.com/locate/nima

Rate-capability of the VMM3a front-end in the RD51 Scalable Readout System

D. Pfeiffer^{a,b,c,*}, L. Scharenberg^{b,d,**}, P. Schwäbig^{d,***}, S. Alcock^a, F. Brunbauer^b, M.J. Christensen^e, K. Desch^d, K. Flöthner^{b,f}, F. Garcia^g, R. Hall-Wilton^{a,c}, M. Hracek^{b,h}, G. Iakovidisⁱ, D. Janssens^{b,j}, J. Kaminski^d, M. Lupberger^{d,f}, H. Müller^{b,d}, E. Oliveri^b, L. Ropelewski^b, A. Rusu^k, J. Samarati^{a,b}, M. van Stenis^b, A. Utrobicic^b, R. Veenhof^{b,l}

^a European Spallation Source ERIC (ESS), Box 176, SE-221 00 Lund, Sweden

^b European Organization for Nuclear Research (CERN), 1211 Geneva 23, Switzerland

^c University of Milano-Bicocca, Department of Physics, Piazza della Scienza 3, 20126 Milan, Italy

^d Physikalisches Institut, University of Bonn, Nußallee 12, 53115 Bonn, Germany

^e European Spallation Source ERIC (ESS), Data Management and Software Centre, Ole Maaløes Vej 3, 2200 Copenhagen, Denmark

^f Helmholtz-Institut für Strahlen- und Kernphysik, University of Bonn, Nußallee 14–16, 53115 Bonn, Germany

^g Helsinki Institute of Physics, P.O. Box 64, FI-00014 University of Helsinki, Finland

^h Institute of Experimental and Applied Physics, Czech Technical University in Prague, Husova 240/5, 110 00 Prague 1, Czech Republic

ⁱ Brookhaven National Laboratory, PO Box 5000, Upton, NY 11973, United States of America

^j Physics Department, Vrije Universiteit Brussel, Pleinlaan 2, 1050 Brussels, Belgium

^k SRS Technology, 30 Promenade des Artisans, 1217 Meyrin, Switzerland

^l Bursa Uludağ University, Görükle Kampusu, 16059 Niifler/Bursa, Turkey

ARTICLE INFO

Keywords:

VMM3a ASIC
Scalable Readout System (SRS)
Micro-Pattern Gaseous Detector (MPGD)
Readout electronics
Field Programmable Gate Array (FPGA)
X-ray imaging

ABSTRACT

The VMM3a is an Application Specific Integrated Circuit (ASIC), specifically developed for the readout of gaseous detectors. Originally developed within the ATLAS New Small Wheel (NSW) upgrade, it has been successfully integrated into the Scalable Readout System (SRS) of the RD51 collaboration. This allows, to use the VMM3a also in small laboratory set-ups and mid-scale experiments, which make use of Micro-Pattern Gaseous Detectors (MPGDs). As part of the integration of the VMM3a into the SRS, the readout and data transfer scheme was optimised to reach a high rate-capability of the entire readout system and profit from the VMM3a's high single-channel rate-capability of 3.6 Mhits/s. The optimisation focused mainly on the handling of the data output stream of the VMM3a, but also on the development of a trigger-logic between the front-end cards and the DAQ computer. In this article, two firmware implementations of the non-ATLAS continuous readout mode are presented, as well as the implementation of the trigger-logic. Afterwards, a short overview on X-ray imaging results is presented, to illustrate the high rate-capability from an application point-of-view.

1. Introduction

A common goal in the development of radiation detectors is to improve their rate-capability. Prominent examples driving the need for larger statistics, shorter data acquisition times and better handling of the detector occupancy are the high luminosity upgrades of the Large Hadron Collider (LHC) in high energy physics or the European Spallation Source (ESS) for neutron and material sciences. Alongside the development of new detector technologies coping with higher event rates comes the development of powerful readout electronics allowing to deal sufficiently with the new rate challenges.

This high-rate challenge was one of the reasons for the development of the VMM3a Application Specific Integrated Circuit (ASIC) within the ATLAS New Small Wheel Upgrade [1]. In parallel to the integration of the VMM3a into the ATLAS environment, the VMM3a was also integrated [2] into the RD51 [3] Scalable Readout System (SRS) [4], providing its features for small R&D setups and mid-sized experiments. Since the readout scheme of the VMM3a in the SRS differs from that used by ATLAS, the rate-capability of the VMM3a/SRS system was investigated. The results are presented in this article. In the main part (Section 2), the readout scheme of the VMM3a within the SRS is described. The focus is set on a description of the firmware, the

* Corresponding author at: European Spallation Source ERIC (ESS), Box 176, SE-221 00 Lund, Sweden

** Corresponding author at: European Organization for Nuclear Research (CERN), 1211 Geneva 23, Switzerland.

*** Corresponding author at: Physikalisches Institut, University of Bonn, Nußallee 12, 53115 Bonn, Germany

E-mail addresses: dorothea.pfeiffer@ess.eu (D. Pfeiffer), lucian.scharenberg@cern.ch (L. Scharenberg), schwaebig@physik.uni-bonn.de (P. Schwäbig).

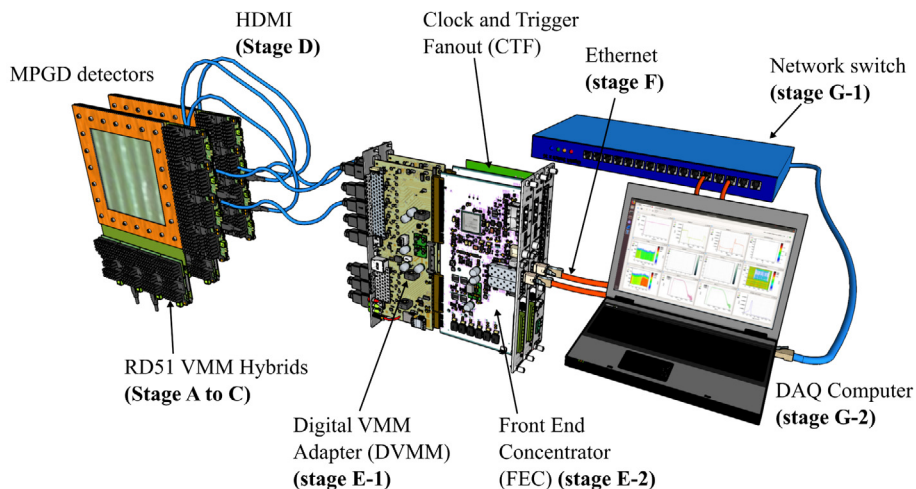
<https://doi.org/10.1016/j.nima.2022.166548>

Received 22 September 2021; Received in revised form 28 December 2021; Accepted 2 February 2022


Available online 17 March 2022

0168-9002/© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license

(<http://creativecommons.org/licenses/by/4.0/>).



(a) 3D drawing of the VMM3a/SRS stages.



(b) Photograph of an R&D set-up.

Fig. 1. Example of a typical SRS laboratory set-up. The RD51 VMM Hybrid contains two VMM3a ASICs (stage A) that are connected (stage B) to a Spartan-6 FPGA (stage C). Each hybrid is connected via HDMI (stage D) to the combination of DVMM card and FEC card (stage E) that here is shown outside the Minicrate 2k. The FEC is then connected via Ethernet (stage F) to the DAQ computer (stage G).

software to control and operate the system, and the data acquisition rates that can be achieved with well-defined test pulses. To confirm the test pulse measurements, and to demonstrate the rate performance under application conditions, X-ray measurements were carried out. The results are described in the second part (Section 3).

2. The VMM3a ASIC and the RD51 scalable readout system

The Scalable Readout System (SRS) is a versatile and multi-purpose readout system, developed by the RD51 collaboration for the readout of Micro-Pattern Gaseous Detectors (MPGDs). A typical configuration of the SRS is shown in Fig. 1.

The VMM3a ASIC on the hybrid (stage A) records the charge signals from the detector readout. It is connected (stage B) to the Spartan-6 FPGA on the hybrid (stage C) which serially reads the digital representation of these signals from the VMM3a channels. It creates hit data from this information, and then sends them out via an HDMI cable (stage D) to the combination of the Digital VMM adapter (DVMM) card and the Front End Concentrator (FEC) card (stage E). The DVMM card has eight HDMI ports, and can also provide power to the hybrids, either via HDMI cable or via a power cable plugged into its power outlet. From the FEC card, the data is sent in the form of UDP frames via Ethernet (stage F) to a DAQ computer (stage G).

Two different options exist to take data with more than one FEC. Up to 40 FECs can be connected to the Scalable Readout Unit (SRU), which

has one 10 Gbps Ethernet port that communicates with the network card of the DAQ computer. Alternatively, the Clock and Trigger Fanout (CTF) card can be used, which provides a common clock for eight FECs. For the measurements presented here, a CTF card and two FECs were used. The data from the two FECs were sent via the FEC 1 Gbps network connection to a 10 Gbps switch, and from there to a PC with a 10 Gbps network card.

In the following sections, a detailed bottom-up description of the VMM3a ASIC and the SRS is given. The focus lies on the integration of the VMM3a ASIC into the SRS and the firmware of the different readout stages. Further, the software tools needed to control the SRS and to acquire and analyse data are presented. For a more comprehensive technical description of all components, please refer to [4,5] (SRS), [1,6] (VMM) and [2] (VMM3a/SRS).

2.1. The VMM3a ASIC

The VMM3a is a 64-channel readout ASIC, developed by the Brookhaven National Laboratory for the ATLAS New Small Wheel Upgrade, specifically for the readout of gaseous detectors [1]. The analogue part of each VMM channel consists of a charge sensitive amplifier, a shaper, a discriminator and a peak finder. The gain of the amplifier can be adjusted within a range of 0.5 mV/fC to 16.0 mV/fC (8 discrete settings), and the peaking time of the shaper can be set to values between 25 ns and 200 ns (4 discrete settings). A test pulse

generator optionally generates an internal test pulse to inject charge into the charge sensitive amplifier. When the signal amplitude crosses the discriminator threshold and the signal peak has been found by the peak detector (PDO), the charge and time information of the signal is subsequently digitised. This digitisation or conversion takes 250 ns. The maximum rate-capability of a VMM3a channel amounts thus to 4 Mhits/s.

The peak height of the signal is obtained with a 10 bit (effectively ~ 8 bit) two-stage current-mode domino ADC [1]. The available time information consists of a 12 bit coarse time stamp and an 8 bit fine time stamp. The coarse timestamp, which is called BCID, is implemented as 12 bit Gray counter and counts clock cycles of the bunch crossing clock CKBC.¹ The BCID of the digitised signal is the last BC clock before the peak of the signal. Depending on the chosen BC clock frequency (40 MHz and 44.4 MHz in the presented measurements), the resolution of the BCID is 25 ns or 22.5 ns. The fine time stamp is provided by the time detector (TDO), which uses an 8 bit Time to Amplitude Converter (TAC) with adjustable slopes (60 ns to 650 ns in 4 discrete settings). The measured value represents the time between the peak of the signal and the falling edge of the next BC clock. Hence in the case of a 40 MHz BC clock and a TAC slope of 60 ns, the complete time of the peak can be calculated from the BCID and the TDO with the following formula:

$$t_{\text{chip}} = \text{BCID} \times 25 \text{ ns} + (1.5 \times 25 \text{ ns} - \text{TDO} \times \frac{60 \text{ ns}}{255}). \quad (1)$$

The VMM3a configuration data consists of 1728 bits and contains global and individual channel settings. The available per-channel settings are the enabling of the internal test pulse, the masking of the channel in case of malfunctioning, and finally baseline corrections for the PDO and the TDO. Global settings are gain, peaking time, TAC slope and polarity of the input signal. Further different modes can be chosen, which are not all relevant for the non-ATLAS continuous mode used by RD51. Settings that are useful are the neighbouring-logic (also called ‘neighbour trigger’), the sub-hysteresis discrimination and the Double Data Rate (DDR) mode for the data clock CKDT. Channels are normally only read out if the signal amplitude crossed the discriminator threshold. If the neighbouring-logic (NL) is enabled, the channels adjacent to channels over threshold are also read out, even when their signal stayed below the threshold (e.g. channel 8 is over threshold, with the neighbouring-logic also channels 7 and 9 are read). The DDR mode for the data clock CKDT means that new bits are available on the data lines $data_0$ and $data_1$ on the rising and the falling edge of the clock, whereas in single data rate (SDR) mode this happens only on the rising edge of CKDT. The timing diagram for the data readout of one channel is illustrated in Fig. 2.

2.2. The RD51 VMM3a hybrid

The RD51 VMM3a hybrid is a small 5 cm × 8 cm Printed Circuit Board (PCB) equipped with two VMM3a ASICs, providing thus 128 charge input channels. A 144 pin Hirose connector on the hybrid connects to the readout of a gaseous detector. From the Hirose connector, the 128 channels are routed to the channel inputs of the two

¹ The ATLAS experiment at CERN measures particles coming from collisions of proton bunches in LHC. Every 25 ns one proton bunch arrives at the interaction point, hence the bunch crossing frequency is 40 MHz. The name bunch crossing clock refers thus to the ATLAS situation, but the clock is in reality just a timer/counter and can have any chosen frequency. At a CKBC frequency of 160 MHz missing BCIDs were observed. The maximum CKBC frequency that works reliably lies therefore somewhere between 44 MHz and 160 MHz, but has not been determined by the authors. The lowest CKBC frequency is around 2.5 MHz. When the peak finder finds the peak, the ramp of the TDC starts. The next rising edge of CKBC arms the stop circuit and the next falling edge stops the timing ramp. The TAC slope has thus to be at least 1.5 times as long as the period of CKBC. For the longest available TAC slope of 650 ns, the lowest possible CKBC is thus $1.5/650 \text{ ns} = 2.31 \text{ MHz}$. Slower CKBC frequencies therefore mean worse time resolution.

Table 1
Data format of one 38 bit VMM3a hit.

Name	Length (bit)	Values
Data flag	1	always 1
Over-threshold flag	1	Over threshold: 1 Below threshold (requires NL on): 0
Channel number	6	0 to 63
ADC (PDO)	10	0 to 1023
TDC (TDO)	8	0 to 255
BCID (clock counter)	12	0 to 4095

VMM3a ASICs, with each channel having a spark protection circuit on the PCB. The other I/Os on the ASICs (clock inputs, configuration inputs, data outputs) are presently connected to a Xilinx Spartan-6 Field Programmable Gate Array (FPGA). The FPGA will be upgraded to Spartan-7 in the second half of 2021. The FPGA provides the necessary clocks for the VMM3a (bunch crossing clock CKBC, token clock CKTK, data clock CKDT, and the test pulse clock CKTP), and reads out the two data lines $data_0$ and $data_1$. Further, the FPGA takes care of sending the configuration data to the ASICs.

The format of one 38 bit long VMM3a hit (a hit is the data read out from one channel) is shown in Table 1. After reading out the hits from the two VMM3a, they are sent from the hybrid via the DVMM card to the RD51 Front-End Concentrator (FEC) card. The maximum possible readout rate that can be obtained for the complete VMM3a ASIC depends on the passing speed of the readout token, and the speed of the data transfer. In the following two sections, different firmware approaches to obtain high rates on the hybrid are presented.

The hybrid is connected via HDMI cable to the RD51 FEC. In the HDMI cable² four differential pairs are used for high-speed transmission, and two single wires for I2C. Via the clock differential pair, a base clock is sent from the FEC to the hybrid. The purpose of the configuration/trigger pair is to transmit configuration data and to send control commands like the start and stop of the acquisition from the FEC to the hybrid. For each of the two VMM3a ASICs on the hybrid, a dedicated differential pair is available for data transmission from hybrid to FEC. VMM3a data, as well as configuration/trigger data, are 8 b/10 b encoded.

All four differential pairs are connected to serialiser/deserialiser (SERDES) components in the FEC and the hybrid firmwares. In the hybrid firmware, the two data pairs use output serialiser/deserialiser (OSERDES) components, whereas the clock and configuration/trigger pairs are connected to input serialiser/deserialiser (ISERDES) components. The ISERDES of the configuration/trigger receiver is used in combination with an input/output delay (IODELAY) component. The delay of the IODELAY component is automatically adjusted so that each bit is read exactly at the centre of the data eye. For the correct word alignment of the 10 bit word, the bit slip feature of the ISERDES is used.

The timing diagram for the data readout from the VMM3a is shown in Fig. 3. With the help of a counter, the readout token clock CKTK is generated from the readout process clock with a pulse length of about 28 ns and a frequency of up to 9 MHz. The token clock is always running and continuously sending tokens. If data is available, the VMM3a pulls the data line $data_0$ high as a flag. The firmware detects this flag and starts the data clock CKDT that can have a frequency of up to 180 MHz. Although designed for CKDT frequencies up to 200 MHz [2], no higher CKDT frequencies than 180 MHz could be reached, due to the observation of occasional data corruption on some ASICs (see Section 2.2.1).

² The HDMI cable just serves as electrical connection between hybrid and FEC. The wires of the HDMI cable are not mapped according to the HDMI standard, and a custom made protocol different from the HDMI protocol is used for data transmission.




Fig. 2. Timing diagram of the VMM3a in the non-ATLAS continuous readout mode. The Peak Detector (PDO) starts the 10bit ADC for the charge information, and the voltage ramp (TAC slope) for the Time Detector (TDO). This voltage ramp is stopped by the next falling edge of the bunch crossing clock CKBC. The 8 bit value of the TDO represents thus the time between the peak and the subsequent BCID. To obtain the digitised data, a token is generated with the token clock CKTK, and passed around to channels with data. The channel that has the token is then bit-wise read out with the data clock CKDT on the two data lines $data_0$ and $data_1$.

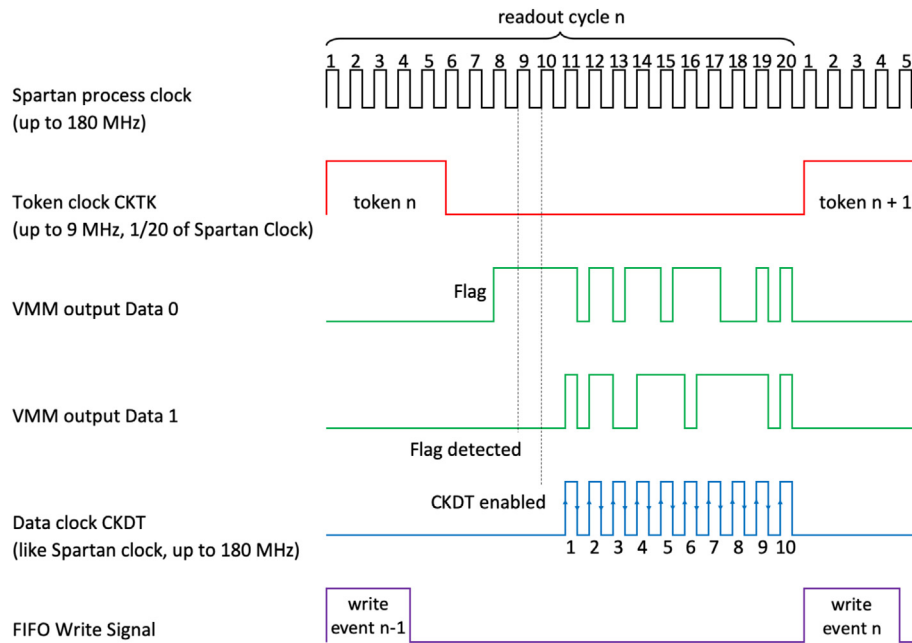


Fig. 3. The diagram explains the readout of the VMM3a data on the hybrid FPGA, and displays one full readout cycle (cycle n). In case hit data is available, it is read out in the time between two tokens. In the depicted DDR mode, the maximum readout frequency is 9 MHz, whereas in SDR mode it is 6 MHz. Since the time for token generation and readout of the flag remains unchanged, the maximum DDR readout frequency is not twice as large as the SDR frequency.

The data clock then pushes out the 38 bit VMM3a hit. Only after the entire hit has been transferred by the data clock, the channel is clear and can store a new hit.³ The diagram depicts the DDR readout

³ It should be mentioned that in the non-ATLAS continuous mode, which is used in the SRS integration of the VMM3a, the 4-hit deep FIFO mentioned in [1] does not work (it was identified by the authors together with the designer of the chip [7]). Due to a bug that occurred when the L0 mode was implemented in the 2016 revision of VMM3, the FIFO-read advancement does not work. This means the next three hits after a hit are stored, but cannot be read out. The FIFO is not used in the L0 mode used by ATLAS.

mode, with 19 bits coming out on each of the two data lines $data_0$ and $data_1$ on the rising and falling edge (when using SDR mode only on the rising edge) of CKDT. When the complete hit has been read out, the 19 bits from $data_0$ and 19 bits from $data_1$ are combined to a 38 bit hit, which is padded with two zeros to 40 bits. The 40 bit hit is first stored in a 1024 hit deep FIFO. This occurs at the same time as the sending of the next token. The hit data is then read out from the FIFO, encoded in 8 bit/10 bit format and subsequently sent to the FEC.

The following two chapters describe two different high-speed firmware implementations, the Bonn version and the ESS version. The structure of both hybrid firmware versions is shown in Fig. 4. The

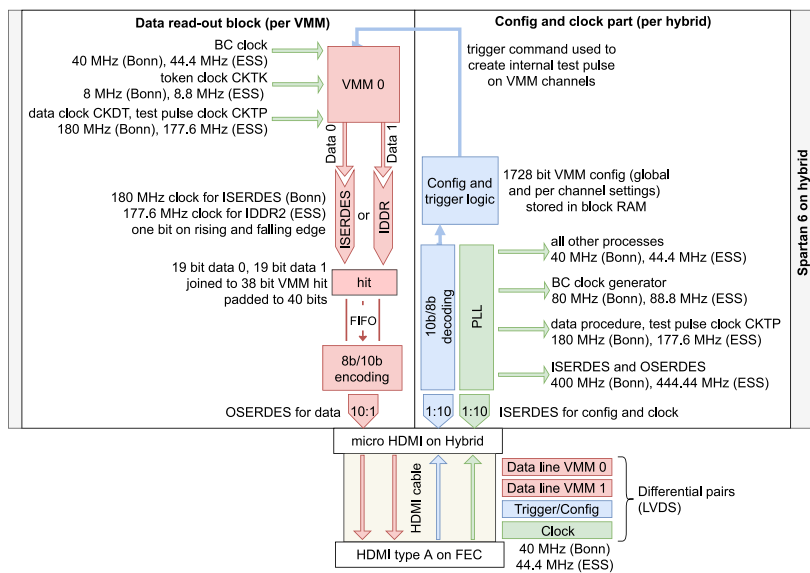


Fig. 4. Schematic drawing of the main functional components of the Spartan-6 firmware versions for the RD51 VMM3a hybrid. There are two data readout blocks on the hybrid (one per VMM3a), and a common configuration and clock part.

versions differ mainly in the choice of clock frequencies and the use of different components (ISERDES for Bonn, IDDR for ESS) to read the data from the VMM3a ASIC.

2.2.1. Bonn firmware for the RD51 VMM3a hybrid

The Bonn firmware of the hybrid was used for initial research of the high-speed readout capabilities of the VMM. It serves as a general-purpose firmware with many configuration possibilities. The base clock of the system is 40 MHz. Using two phase-locked loops (PLLs) all remaining necessary clocks are generated based on the 40 MHz clock. One of the PLLs is used for generating the clocks necessary for data transfers coming from the VMM while the other PLL is used for generating all remaining clocks needed by the firmware. To the latter belongs the bunch crossing clock CKBC which can be chosen from a list of seven different frequencies: 2.5 MHz, 5 MHz, 10 MHz, 20 MHz, 40 MHz, 80 MHz and 160 MHz. The incoming data from the VMM3a ASICs is received by ISERDES components with a deserialisation factor of five, reducing the clock frequency behind the IOB by the same factor. This way most of the logic of the firmware uses a lower clock which relaxes timing. To align the data an IODELAY component with fixed tap delay is used in front of the ISERDES.

The data transfer clock CKDT has a maximum frequency of 180 MHz which is lower than the design value of 200 MHz, as tests showed that frequencies above 180 MHz resulted in incomplete data transfers when using double data rate. Measurements of this can be seen in Fig. 5. The lower readout frequency showed to be a design problem of the VMM that was verified in simulation. Since all clocks for the data transfers from the VMM3a are generated by the same PLL from which the remaining firmware is independent, it is possible to reprogram the PLL using its Dynamic Reconfiguration Port (DRP) to chose transfer clock frequencies of 90 MHz, 45 MHz and 22.5 MHz instead.

When using data transfers with double data rate the token clock CKTK has a frequency of 8 MHz with a duty cycle of roughly 20% and is generated using a counter. The resulting token length of roughly 28 ns was tested to be necessary such that all tested VMM3a ASICs recognise the token.

The number of active channels has a significant influence on the highest possible data rate. This is due to an empty data transfer which occurs when the token has to be passed to a channel with lower channel number: when there is a data transfer from channel n but no channel with number $> n$ has data to transfer, the token must be passed down to a lower channel or has to stay at n . The latter is the case if no other

channel has any data to be transferred. The VMM3a will then require an additional token on CKTK during which it does not shift out any data, effectively reducing the data rate of the VMM. Having N active channels and a sufficiently high hit rate the total transfer time for all channels is given by $(N + 1) \times t_d$ with $t_d = 1/f_{CKTK}$. In order to improve the data rate, the firmware will increase the token clock CKTK and send the next token after a shorter time t_p if the VMM3a does not initiate a data transfer some time after the previous token. The highest expected hit rate considering the effects of the additional token and influence of active channel count and varying token clock is given by

$$\text{hit rate} = \frac{\text{active channels}}{\text{active channels} \times t_d + t_p} \times \frac{\text{acquisition window}}{4095}, \quad (2)$$

with the acquisition window being a configurable amount of counts by the Gray counter for which the VMM is enabled during a complete BCID cycle.

Measured hit rates in dependence on the number of active channels as well as the corresponding expected rate calculated using Eq. (2) are shown in Fig. 6. It can be seen that theoretical limit and measurement are in very good accordance, with the extrapolated maximum hit rate close to 7.2 Mhits/s. When looking at the rates for one active channel note that the upper limit for a single channel is 4 Mhits/s although in tests no VMM performed with rates above 3.6 Mhits/s. Since this upper limit varied between different VMM ASICs it is assumed that this is due to the ASIC itself. Tests with different waveforms and added time offsets between channels did not increase the measured hit rate.

An additional feature of the firmware is the auto alignment: for updates and changes of the firmware – especially those regarding the readout and clocking scheme – the data read from the VMM3a might become misaligned such that bits are missing at the beginning or end. Therefore the firmware offers the possibility to automatically align the 38 bit data word correctly.

2.2.2. ESS firmware for the RD51 VMM3a hybrid

The implementation of the data readout in the ESS firmware is depicted in Fig. 4. Via the clock differential pair, a 44.4 MHz base clock is sent from the FEC to the hybrid. The 44.4 MHz clock from the FEC is fed into a phase-locked loop, which creates all the clocks needed in the firmware.⁴ The output clocks are distributed via global

⁴ When the hybrid is used as part of the ESS readout and not the SRS, the basic clock frequency send to the hybrid will be 44.0265 MHz, half of the ESS facility clock frequency of 88.053 MHz.




Fig. 5. Maximum data transfer clock measurement, results for eight tested VMM3a ASICs and five different transfer clocks CKDT. The relative amount of valid data transfers (with a total of about 20 000 transfers per VMM3a) is shown for each channel.

clock buffers (BUFG) components. The general firmware logic uses a frequency of 44.4 MHz. The clock generator for the bunch crossing clock works with 88.8 MHz and generates a fixed 44.4 MHz CKBC. The data readout procedure and the test pulse clock generator all use 177.6 MHz. The fastest clock is the 444.4 MHz clock for the SERDES components.

The clock and configuration part of the firmware is common for both ASICs, whereas one data readout block has been implemented per VMM3a. The two data lines ($data_0$ and $data_1$) of the VMM3a are connected to input double data rate register (IDDR2) components.

Assuming the clock settings described above (the 44.4 MHz clock from the FEC is taken as basic clock for the hybrid), the maximum hit rate is 8.8 Mhits/s. This rate is perfectly matched to the speed of the SERDES. With a clock frequency of 444.4 MHz and 8 b/10 b encoding, the effective bit rate of the SERDES amounts to

$$\text{bit rate}_{\max} = 444.4 \text{ Mbps} \times 0.8 = 355.5 \text{ Mbps} . \quad (3)$$

For a 40 bit hit, the hit rate is hence

$$\text{hit rate}_{\max} = 355.5 \text{ MHz} / 40 \text{ bits} = 8.8 \text{ Mhits/s} . \quad (4)$$

This is close to the maximum of 9 Mhits/s, which can be read from a single VMM3a. This maximum rate would require a SERDES clock of 450 MHz and a CKDT of 180 MHz.

2.3. The RD51 SRS FEC

The combination of RD51 FEC and DVMM card can read out up to 8 RD51 VMM3a hybrids. The hybrids are connected via HDMI cables to the DVMM card, from which the signals are routed to a Xilinx Virtex 6 FPGA on the FEC. Fig. 7 explains the data path of the FEC firmware.

The differential pairs in the HDMI are described in chapter 2.2.2. The ESS version of the FEC firmware operates the four SERDES in DDR mode with a 222.2 MHz clock, and sends a 44.4 MHz clock to the

hybrids, whereas the Bonn version uses 200 MHz and sends a 40 MHz clock. The bit rate of each SERDES is thus 444.4 Mbps for the ESS version, and 400 Mbps for the Bonn version. Due to 8 b/10 b decoding, the effective incoming bit rate on each data line on the FEC is reduced by 20% to 355.5 Mbps (3) and 320 Mbps, respectively. After 8 b/10 b decoding, the hits from the VMM3a are then time-stamped on the FEC in the latency logic block, and subsequently as 48 bit hits written into a hit FIFO. For each hybrid connected to the FEC, the firmware provides two separate FIFOs (one per VMM3a). Whenever a hit is available, it is written with a 44.4 MHz write clock to the FIFO.

In the readout part of the firmware, there is one common FEC hit FIFO for all hits. A fair scheduler runs with 125 MHz and reads one hit at a time from all the non-empty VMM3a data FIFOs. Provided that it is not full, the hits are transferred to the FEC FIFO. To be sent out via UDP, the 48 bit hits have to be decomposed into six 8 bit chunks and written to the UDP FIFO. The 8 bits from the UDP FIFO are read with 125 MHz, as determined by the 1 Gbps Ethernet of the FEC. Since the maximum data rate per VMM3a is around 356 Mbps, with the theoretical maximum at 360 Mbps, the FEC can only support two VMM3a or one RD51 hybrid at maximum rates.

2.3.1. Latency logic on the FEC

The coarse time stamp in the hit data, the BCID, is a 12-bit counter and overflows every 4096 BC clock periods or 92.16 μs .⁵ The FEC generates a 42 bit FEC time stamp with 22.5 ns resolution, which only overflows every 1.2 days. To save bandwidth, the 42 bit timestamp is only send every 16 BCID overflow periods of 92.16 μs , i.e. every 1.47 ms for each VMM. This combination of 42 bit timestamp, 5 bit VMM-ID and a 1 bit flag (always 0) is called *marker*.

⁵ To simplify the description, in the following the 22.5 ns BC clock period of the ESS version is taken for the calculations. For the Bonn version, one simply has to replace the 22.5 ns with 25 ns.




Fig. 6. The hit rate received from a single VMM as a function of the number of active channels for three different Bonn firmware versions (similar limitations apply to the ESS firmware, but are not shown here). The markers show the measurement and the lines the expected rate calculated based on the times t_d and t_p . Note that for a single active channel the measurements for the upper two curves are identical. The horizontal lines mark the rates in case the hit rate would be independent of the amount of active channels. Stimulus for all measurements is an externally generated, periodic sawtooth waveform.

In the time between markers, the 5 bit VMM-ID and a 5 bit overflow cycle or offset counter are added to the 38 bit VMM data. The DAQ can distinguish between markers and data by looking at the 1 bit flag. If the flag is zero it is a marker, if it is one it is a hit.

The hits from the VMM3a are read out and transmitted in a roughly time ordered way. The readout token inside the VMM3a is passed around between channels containing data in a round-robin like fashion, going from lower channel numbers to higher channel numbers. Therefore slight disturbances of the time order can occur, as illustrated with the following example: the ASIC registered three hits that occurred within one BC clock cycle. Channel 2 has a hit with BCID 0, whereas channel 1 and channel 63 have a hit with BCID 4095. After reading channel 1 (BCID 4095), the token will be passed to channel 2 (BCID 0), and subsequently to channel 63 (BCID 4095). For the hit in channel 2 with BCID 0, the BCID overflow already happened. The BCID in channel 1 and 63 has not overflowed yet, since the hit occurred slightly earlier. When the FEC generates the higher-order time stamp, the firmware has thus to consider that the hit with BCID 0 did not occur 4094 clock cycles earlier than the hits with BCID 4095, but instead slightly later.

To be able to make the decision to which overflow cycle a hit belongs, the FEC also counts BC clock cycles, similar to the BCID counter in the VMM3a. After receiving the acquisition start signal from slow control, the FEC starts incrementing its BC counter and then sends via the hybrid a soft reset command to the VMM3a. This soft reset command sets the BCID counter on the VMM3a to zero. After sending the soft reset, the FEC is ready to accept data from the hybrids. The time at which the FEC sends the soft reset is called *reset latency* and can be adjusted via the slow control.

This *reset latency* time is comprised of two components: the time that it takes the soft reset command to reach the VMM3a and reset the BCID, and the time that it takes for a single hit from the VMM3a to reach the FEC. If correctly set, a single hit with e.g. BCID 100 will arrive on the FEC at the moment where the FEC BC counter is also 100. The aim is thus to have the same values in BCID and FEC BC. With 2 m HDMI cables, the *reset latency* latency is of about 47 clock cycles. That means the soft reset is sent when the FEC BC counter has a value of $4096 - 47 = 4049$. After setting the reset latency to the correct value, the FEC BC counter is now identical to the BCID of a single hit upon arrival on the FEC. Fig. 8 shows data from one channel on one VMM acquired with the correct latency settings.

A small jitter of a few clock cycles can nevertheless occur. This parameter is called *latency jitter* and is usually set to around 4 clock cycles. The third important parameter is the *maximum latency*. If all 64 channels are activated at the same time, there is a substantial difference in data arrival time on the FEC between the first and the last channel that is read out. On the hybrid one hit can be read out every 5 clock cycles of the BC clock or every 112.5 ns. So when reading out 64 hits that have been digitised at the same time, the last hit will arrive with a maximum latency of 7.2 μ s.

Fig. 9 explains how the latency logic is implemented. In addition to counting BC cycles, the FEC also counts the overflows of its BC counter. After 16 overflow cycles, a new marker is generated and written into each VMM FIFO. The overflow counter is then reset from 15 to zero. Upon arrival on the FEC, the BCID of a hit is compared with the FEC BC counter. Based on the result of the comparison, the FEC firmware can decide to which overflow cycle a hit belongs. A hit can belong either to the present overflow cycle, the previous overflow cycle or violate the latency conditions. In the easiest and most common case, the hit belongs to the present overflow cycle. The present overflow counter is then added to the data. If the hit belongs to the previous overflow cycle, the overflow counter is decremented by one. Therefore overflow counters with values between -1 and $+15$ can be added to the valid data. Invalid hits that violate the latency conditions are marked with an overflow counter of -16 . With a correctly configured system and fully operational hardware, invalid hits do not occur.

The latency logic makes it possible to continuously take data without applying an acceptance window and without rejecting data. If measuring e.g. X-rays or white noise, the time distribution of the acquired data is flat. In the DAQ or the data analysis, the correct FEC timestamp is always calculated from the last marker that has been sent. The overflow periods are added to the marker (or subtracted from the marker if the overflow count is -1) according to the formula:

$$t_{\text{FEC}} = t_{\text{marker}} \times 22.5 \text{ ns} + \text{overflow counter} \times 4096 \times 22.5 \text{ ns} . \quad (5)$$

To obtain the complete hit time stamp, the FEC timestamp (Eq. (5)) and the VMM3a timestamp (Eq. (1)) have to be added together:

$$t_{\text{hit}} = t_{\text{FEC}} + t_{\text{VMM3a}} . \quad (6)$$

2.4. The DAQ PC

The RD51 slow control tool *vmmdds* configures FEC, hybrid and VMM3a [8]. The configuration data is sent in UDP frames via Ethernet cable from the PC to the FEC. The FEC can determine the destination of the configuration based on the different UDP ports of the configuration frame. For the data acquisition, two solutions are available. The *ESS DAQ* [9] provides event reconstruction and online monitoring but has limited rate-capabilities when the full non-reconstructed hit data is written to disk (so-called debug mode). Alternatively, the *tcpdump* [10] utility can directly read the data at the network adapter. The network packages are then written to disk, but online monitoring and visualisation are not possible. Thus, the solution for high-rate data-taking is to use *tcpdump* for the data acquisition, and the *ESS DAQ* only for the online monitoring. This combination allows to obtain the

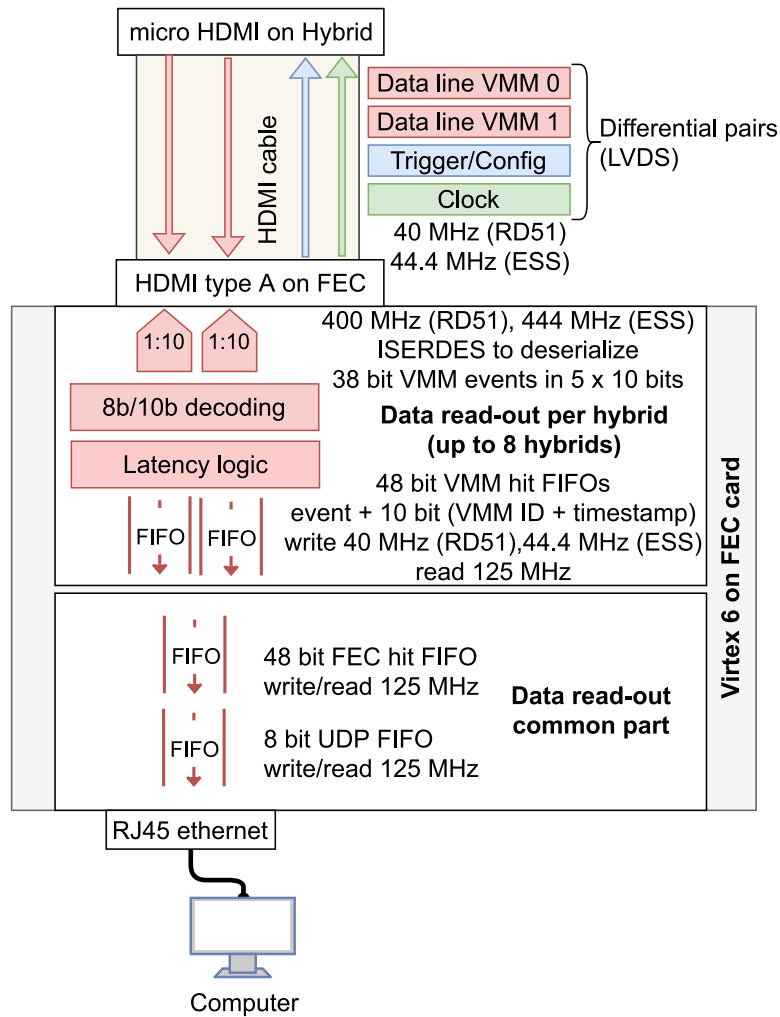


Fig. 7. Schematic drawing of the data path in the Virtex 6 firmware of the RD51 SRS FEC.

```

> Hit: 1, offset: 4, vmmID: 2, ch: 0, bcid: 100, trg: 99, latency: -1
> Hit: 2, offset: 5, vmmID: 2, ch: 0, bcid: 100, trg: 99, latency: -1
> Hit: 3, offset: 6, vmmID: 2, ch: 0, bcid: 100, trg: 100, latency: 0
> Hit: 4, offset: 7, vmmID: 2, ch: 0, bcid: 100, trg: 101, latency: 1
> Hit: 5, offset: 8, vmmID: 2, ch: 0, bcid: 100, trg: 101, latency: 1
> Hit: 6, offset: 9, vmmID: 2, ch: 0, bcid: 100, trg: 99, latency: -1
> Hit: 7, offset: 10, vmmID: 2, ch: 0, bcid: 100, trg: 100, latency: 0
> Hit: 8, offset: 11, vmmID: 2, ch: 0, bcid: 100, trg: 100, latency: 0
> Hit: 9, offset: 12, vmmID: 2, ch: 0, bcid: 100, trg: 101, latency: 1
> Hit: 10, offset: 13, vmmID: 2, ch: 0, bcid: 100, trg: 99, latency: -1
> Hit: 11, offset: 14, vmmID: 2, ch: 0, bcid: 100, trg: 100, latency: 0
> Hit: 12, offset: 15, vmmID: 2, ch: 0, bcid: 100, trg: 100, latency: 0
> Marker: 1, VMM ID 2, SRS timestamp: 81530880
> Marker: 2, VMM ID 3, SRS timestamp: 81530880

```

Fig. 8. Channel 0 on VMM 2 is pulsed once every 4096 BC clock cycles at BC clock 100. The Wireshark trace shows the hit data with the offset period and the BCID. The field *trg* represents the value of the FEC BC counter. The difference between the *trg* and the BCID is the latency, which has values here between -1 and 1 . After all data from offset 15 is transmitted, a marker with the 42-bit SRS time is sent for every active VMM.

highest possible rates with VMM3a and SRS. It should be noted that when using standard PC operating systems for high speed UDP network communication several techniques need to be applied to minimise the processing overhead and reduce avoidable packet loss [11]. The offline data analysis *vmm-sdat* [12] is able to cope with the data from *tcpdump* and *ESS DAQ*. It clusters the hit data, time matches the clusters in between the readout planes and reconstructs thus the interaction point of the particle in the gas volume of the detector. This way the rate of the detected particles can be calculated. The bit rates on the

other hand were directly measured on the network adapter with the Linux command-line tool *ifstat*. With the DAQ PC⁶ as last stage of the readout scheme, the maximum rates per readout stage, which could be achieved, are summarised in Table 2.

⁶ For the measurements (e.g. the ones presented in Section 3), a typical 2.5-inch SATA (6 Gbps) SSD combined with a 10 Gbps network card was used. No specific optimisation/modification of the software or hardware components was applied.

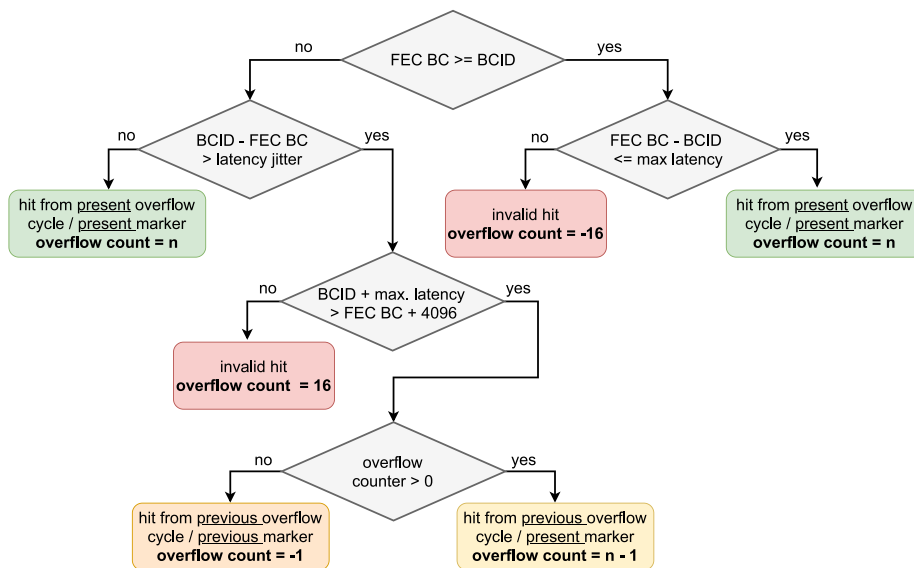


Fig. 9. The latency logic flow chart explains how the FEC firmware determines the correct higher order time stamp for the hits that arrive on the FEC. The BCID of the hit is compared with the FEC BC counter. Depending on the parameters latency jitter and maximum latency, an overflow counter is added to the hits. In the most common case the hit belongs to the present overflow cycle (cases in green boxes). If the FEC BC is a lot smaller than the BCID (FEC BC counter plus 4096 is smaller than the BCID plus the maximum latency), then the FEC BC counter has already overflowed while data of the previous overflow cycle are still read out. The previous overflow cycle can either belong to the present marker (yellow box), or be the last overflow cycle of the previous marker (orange box). In case the latency conditions are violated, the hit is marked as invalid. With correctly set parameters and working hardware this case does not occur.

Table 2

Rate-capability of the various stages of VMM3a/SRS in the current implementation of the non-ATLAS continuous readout mode. The values given in number of hits/active channels per second (Mhits/s) and data bandwidth in megabits per second (Mbps). The rate at the DAQ computer depends on the number n of network ports on the switch and the bandwidth limit of the network card (e.g. 10 Gbps).

Readout stage		Maximum rate	
Unit	Quantity	(Mhits/s)	(Mbps)
(A)	VMM3a channel Rate per channel	3.6	144
(B)	VMM3a to Spartan-6 Rate per VMM3a	8.8	355.5
(C)	Spartan-6 Rate per HDMI-SerDes	17.7	711.1
(D)	HDMI Rate per Hybrid	17.7	888.8
(E)	DVMM and FEC Rate per DVMM card	142.2	7111.1
(F)	Gigabit Ethernet Rate per FEC	20.8	1000
(G)	DAQ computer Rate per Switch Port	$20.8 \times n$	$1000 \times n$

3. Measurements with X-rays

For the rate measurements in the previous descriptions, test pulses have been used, due to their constant and well-known input rate. In the following, the results of rate measurements under application conditions are presented. First, the experimental set-up is described (Section 3.1), followed by the measurements on the rate limitations (Section 3.2). As last part, a few examples of high-rate X-ray imaging are presented, demonstrating the capabilities of VMM3a/SRS (Section 3.3).

3.1. Experimental methods

For the measurements, a COMPASS-like triple-GEM detector [13], which was operated at a gain of around 10^4 , was uniformly irradiated with X-rays. The anode featured an $x - y$ -strip readout with 256 strips

of $400 \mu\text{m}$ pitch in each direction and an active area of $10 \times 10 \text{ cm}^2$. As X-ray source, a copper target X-ray tube (Ital Structures Compact 3K5 X-ray Generator) was used. The acceleration voltage was set to 20 kV, while the X-ray tube current was varied.

The electronics set-up was chosen in a way that allowed to achieve the highest readout rates. For this, only a single hybrid was connected to one FEC, with two FECs in total being used; one for the readout of the x -strips of the detector and one for the readout of the y -strips of the detector. Hence, the actively read out detector area was only $5 \times 5 \text{ cm}^2$, with the rest of the area being shielded to absorb the other X-ray photons.

However, most measurements have been performed with the neighbouring-logic of the VMM enabled, which slightly reduces the rate-capability. As explained in Section 2.1, the neighbouring-logic (NL) allows to read out the adjacent channels below threshold, if the triggering channel surpassed the threshold. This feature was enabled because it is beneficial for the position reconstruction and thus X-ray imaging applications [14].

3.2. Results of the rate limitation measurements

The measurement results are shown in Figs. 10 to 12. In Fig. 10 the measured hit rate, the number of recorded hits (active channels) per second, is plotted against the X-ray tube current. A larger X-ray tube current means that more X-ray photons are generated and thus more hits are expected.⁷ It can be seen that the measured hit rate saturates

⁷ The expected hit and cluster rates are illustrated as dashed lines in Figs. 10 and 11. The rates are estimated based on the known linear relation between photon flux and X-ray tube current in the selected current region for the specific X-ray tube. The linear behaviour is confirmed with a fit for the first three data points in the low X-ray tube current region, where both firmware versions lead to similar cluster rates and no cluster losses (plane clusters) are observed. In addition, for the data point at 0.6 mA X-ray tube current, the same rate was measured with an oscilloscope, while reading the induced signals from the bottom of the last GEM in the detector. Based on this, the linear cluster rate behaviour was extrapolated from the first three data points towards larger X-ray tube currents. To get the hit rate, the average cluster size of about 5 strips per readout plane is used (measured at 0.6 mA).




Fig. 10. Measured hit rate (number of active channels in the full set-up per second) depending on the X-ray tube current for the two firmware types.




Fig. 11. Measured cluster rate depending on the X-ray tube current for the two firmware types. Both, the cluster rates in each readout plane, as well as the detector cluster rate (where a cluster in the x -plane could be matched with a cluster in the y -plane), are shown.

at high X-ray currents, indicating a bandwidth limitation. The cause for the saturation was identified to be the token passing of the VMM3a itself in combination with the randomly⁸ occurring X-ray interactions (see also Section 2.2.1).

In Figs. 11 and 12 the effect of the saturation of the measured hit rate on the cluster reconstruction is shown; a cluster contains all hits that belong to the interaction of a single X-ray photon in the detector. It should be noted that some of the observed effects (see especially Fig. 12) are related to the x - y -strip readout⁹ and they may be different for other readout geometries. Due to the saturation in the hit rate, also a saturation in the cluster rate is expected. This is observed for the clusters in each individual readout plane (see Fig. 11). However, for the detector cluster rate (where a cluster in the x -plane could be matched with a cluster in the y -plane), a decrease after reaching a maximum is observed. The saturation of the measured cluster plane rate can be entirely explained with the bandwidth limits, while the decrease in the measured cluster detector rate is a combination of the bandwidth limit and the cluster reconstruction. For its explanation, the cluster reconstruction procedure is recalled. As mentioned in Section 2.4, the clusters are first reconstructed in each readout-plane individually. For this, the hits in a certain time interval (here $\Delta t_{\text{hits}} \leq 150$ ns) are considered as a cluster in time. Then the geometrical component is added,

⁸ Random in position and time, but still uniformly distributed.

⁹ Here with a 1:1 mapping, meaning that the strip number of the detector is the same as the readout channel number of the electronics.




Fig. 12. Fraction of matched clusters in x - and y -direction depending on the X-ray tube current for the two firmware types.




Fig. 13. Measured spectra from the copper target X-ray tube for various X-ray tube currents.

meaning that the hits in the time cluster occur on adjacent strips (in the present analysis, one missing strip was allowed). Hence, a saturation in the hit rate leads to a saturation of the measured cluster rate in the readout plane. Afterwards, the timestamps of the clusters found in each plane (calculated via centre-of-gravity from the individual hits, with the charge as weight) are matched in time again (here $\Delta t_{\text{plane}} \leq 150$ ns). If now, due to bandwidth limitations, data (meaning clusters) are lost/not processed; for example the hits of a cluster in the x -plane can still be reconstructed to a cluster, while the corresponding hits in the y -plane are lost due to the bandwidth limits and thus no combined detector cluster can be identified. This matching procedure gets less efficient, the larger the actual X-ray interaction rate in the detector gets, because it means that more hits can be lost. This explains the decrease in the rate of reconstructed clusters and is illustrated in Fig. 12, where the fraction of the matched clusters for each plane is shown.

In addition to the quantitative effects (fewer reconstructed clusters due to the reduced cluster finding efficiency), the bandwidth limitations also affect the quality of the reconstructed data. Fig. 13 displays the X-ray spectra for each current setting of the X-ray tube. The quality of the spectra decreases significantly with increasing hit rates and is particularly bad, as soon as the bandwidth saturation (≥ 8 mA X-ray current) is reached. The bandwidth limitation originates from the readout of the VMM3a on the hybrid level.¹⁰ The readout of each channel is achieved by ‘looping’ over all 64 VMM channels and reading one channel after another that has the active token flag for stored data; as mentioned in Section 2.2, a VMM channel can only register a new hit after the previous hit has been read out. This can lead to the loss of charge information required to reconstruct the full cluster, as illustrated in Fig. 14. To illustrate the effect of this VMM3a readout

¹⁰ A bandwidth problem of the FEC can be excluded. A hit rate of maximally 17.5 Mhits/s per FEC is well below the maximum of 20.8 Mhits/s derived from the gigabit Ethernet.




Fig. 14. Illustration of two clusters A and B that have a geometrical overlap of Δx . The time difference between the occurrence of the clusters is Δt . The active readout channels are read out token-by-token, with the readout time T_n for the token on channel n .




Fig. 15. Distributions of the time differences between two consecutive hits for a certain number of Activated Channels Per Hybrid (ACPH) and different X-ray tube currents. The different axis scales between (a) and (b) should be noted.

limitation, two charge clouds/clusters are shown. The two clusters overlap geometrically and are close in time. Each cluster will generate hits on several readout channels. The rate limitation occurs, if the time difference $\Delta t = t_B - t_A$ between the occurrence of two clusters A and B is now smaller than the time difference between the readout of the channels $\Delta t' = T_{n+1} - T_n$. Some of the hits in the first cluster A (in Fig. 14 it is channels $n+4$ and $n+5$) have not been read out when the second cluster B occurs. In this case, the information that is acquired by the VMM channels $n+4$ and $n+5$ cannot be stored in the channel buffer and is simply lost. The more channels are activated, the more likely this type of data loss is to occur due to the token passing scheme and the serial nature of the data transfer.




Fig. 16. Image of a pen containing 17×10^6 clusters. The full data set contains 50×10^6 clusters, that have been recorded in 30 seconds.

This explanation is confirmed by the distributions, shown in Fig. 15. There, the distributions of the time difference between two consecutive hits on a single channel are shown, for two X-ray tube currents (meaning different interaction rates) and for various numbers of Activated Channels Per Hybrid (ACPH). The time it takes to read out one channel and jump to the next one with the raised token flag is constant. Thus, the more channels are active, the more time it takes to loop over all channels with an active token and read them out. For a low X-ray current (2 mA), the shape of the distributions remains the same, independent of the number of channels that could be read out. By turning the NL off, only the number of hits changes as expected (less active strips per cluster), but not the overall behaviour. At higher X-ray tube currents (10 mA), in the bandwidth limitation, the effect of the token passing can be seen. The lower the number of channels which can be read out (ACPH), the quicker one ‘readout loop’ of the token is finished and the more hits can be read out in the same amount of time. This can be seen, as the peak of the time difference distributions moves towards smaller values for smaller ACPHs, but also for the same number of ACPH, when the NL is turned off.

3.3. X-ray imaging examples

As last part of the X-ray studies, a few measurements highlighting the high-rate imaging capabilities of VMM3a/SRS are shown. The X-ray tube settings were slightly changed compared to the previous measurements: the acceleration voltage was set to 18 kV, while the tube current was kept constant at 5 mA.

In Fig. 16 an X-ray image of a pen is shown, containing 17×10^6 clusters. The full data set, due to the larger active area, consists of 50×10^6 reconstructed clusters, which have been recorded in 30 seconds, corresponding to a measured interaction rate of 1.7 MHz. The measured hit rate was on average 20.8 Mhits/s. Similar to the image of the pen, an X-ray image of a small mammal was taken (Fig. 17). This image was recorded with similar hit and measured cluster rates as the one of the pen. However, the acquisition time was slightly longer with 3 minutes, allowing to record 277×10^6 clusters for the image.

Due to the high rate-capability with the self-triggered readout and the individual event reconstruction, it is also possible to perform imaging of dynamic processes. To illustrate this, two examples are shown in Fig. 18. Fig. 18a displays the opening of the X-ray tube’s shutter. The length of each frame was here chosen to be 2 ms. In Fig. 18b, the rotating blades of a fan are shown, again with a frame length of 2 ms. Each frame of the fan contains around 4000 clusters. To increase the absorption of the X-ray photons, the blades have been covered with copper tape. Further, the X-ray tube voltage was reduced to 16 kV in order to decrease the Bremsstrahlung fraction in the spectrum.






Fig. 17. Image of a dead mammal. The data set for this image contains 277×10^6 clusters, that have been recorded in 180 seconds.



(a) Frames of the recorded opening of the X-ray tube's shutter.



(b) Frames of the rotating blades of a fan.

Fig. 18. Examples of the continuous data stream, which is sliced into frames, allowing to reconstruct dynamic processes.

4. Conclusion and outlook

In the past years, the VMM3a ASIC, which was specifically designed to read out gaseous detectors of the ATLAS New Small Wheel, has been successfully integrated into the RD51 Scalable Readout System. Two VMM3a ASICs are combined together with a Spartan-6 FPGA on a front-end board, the RD51 VMM hybrid. As part of the integration process, the non-ATLAS continuous self-triggered readout scheme between the VMMS and the FPGA, as well as the readout scheme towards the Front-End Concentrator (FEC) card, were optimised for high rates. In this article, two firmware implementations for the Spartan-6 FPGA have been presented. Additionally, the optimised firmware on the FEC is described. It was shown, that readout rates of up to 8.8 Mhits/s per VMM and 20.8 Mhits/s per FEC can be achieved. The results were then confirmed by using the electronics for high-rate X-ray-imaging studies,

where in an optimised set-up, interaction rates of around 2 MHz were measured.

The SRS and specifically the RD51 VMM3a are thus suitable for various applications that require high data rates. Future studies like the measurement of fast charge-up effects in detectors in high irradiation environments, or the multi-channel investigation of space charge effects are now possible. In the coming productions of the RD51 hybrid, the Spartan-6 FPGA on the RD51 VMM hybrid will be replaced with a Spartan-7 FPGA. A working prototype of the Spartan-7 version already exists. This hardware upgrade will also allow the investigation and implementation of other VMM readout modes, which are e.g. described here [15]. Implementations of even faster readout rates of about 10 Mhits/s can be envisaged.

CRedit authorship contribution statement

D. Pfeiffer: Conceptualisation, Formal analysis, Investigation, Methodology, Software, Writing – original draft. **L. Scharenberg:** Conceptualisation, Formal analysis, Investigation, Methodology, Writing – original draft. **P. Schwäbig:** Conceptualisation, Formal analysis, Investigation, Methodology, Software, Writing – original draft. **S. Alcock:** Resources, Software, Writing – review & editing. **F. Brunbauer:** Writing – review & editing. **M.J. Christensen:** Resources, Software, Writing – review & editing. **K. Desch:** Conceptualisation, Project administration, Writing – review & editing. **K. Flöthner:** Writing – review & editing. **F. Garcia:** Writing – review & editing. **R. Hall-Wilton:** Project administration, Writing – review & editing. **M. Hracek:** Resources, Software, Writing – review & editing. **G. Iakovidis:** Validation, Writing – review & editing. **D. Janssens:** Writing – review & editing. **J. Kaminski:** Resources, Supervision, Writing – review & editing. **M. Lupberger:** Resources, Supervision, Writing – review & editing. **H. Muller:** Conceptualisation, Resources, Writing – review & editing. **E. Oliveri:** Conceptualisation, Methodology, Supervision, Writing – review & editing. **L. Ropelewski:** Conceptualisation, Project administration, Writing – review & editing. **A. Rusu:** Resources, Writing – review & editing. **J. Samarati:** Writing – review & editing. **M. van Stenis:** Resources, Writing – review & editing. **A. Utrobicic:** Writing – review & editing. **R. Veenhof:** Writing – review & editing.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Alexandru Rusu (SRS Technology) is a licensee within the framework of a CERN Spin-Off, who currently produces the hardware components and who supported the development, assembly and testing of the system described in this article.

Acknowledgements

Parts of this work have been sponsored by the Wolfgang Gentner Programme of the German Federal Ministry of Education and Research (grant no. 05E18CHA).

Parts of this work have received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 846674 as well as from the BMBF (Germany) under 05K19PD1.

References

- [1] G. Iakovidis, VMM3a, an ASIC for tracking detectors, J. Phys. Conf. Ser. 1498 (2020) 012051, <http://dx.doi.org/10.1088/1742-6596/1498/1/012051>.
- [2] M. Lupberger, et al., Implementation of the VMM ASIC in the scalable readout system, Nucl. Instrum. Methods Phys. Res. A 903 (2018) 91–98, <http://dx.doi.org/10.1016/j.nima.2018.06.046>.

- [3] RD51 Collaboration, Development of Micro-Pattern Gas Detectors Technologies, URL <http://rd51-public.web.cern.ch/rd51-public/>.
- [4] S. Martoiu, et al., Development of the scalable readout system for micro-pattern gas detectors and other applications, *J. Instrum.* 8 (2013) C03015, <http://dx.doi.org/10.1088/1748-0221/8/03/c03015>.
- [5] S. Martoiu, et al., Front-end electronics for the scalable readout system of RD51, in: 2011 IEEE Nuclear Science Symposium Conference Record, 2011, pp. 2036–2038, <http://dx.doi.org/10.1109/NSSMIC.2011.6154414>.
- [6] G. De Geronimo, et al., VMM1 - An ASIC for micropattern detectors, in: 2012 IEEE Nuclear Science Symposium and Medical Imaging Conference Record (NSS/MIC), 2012, pp. 633–639, <http://dx.doi.org/10.1109/NSSMIC.2012.6551184>.
- [7] G. De Geronimo, personal communication.
- [8] M. Guth, et al., vmm3a - VMM3a Slow Control and Calibration Software, URL <https://gitlab.cern.ch/rd51-slow-control/vmm3a>.
- [9] M. Shetty, et al., ess-dm3c/essdaq: brightness-software-deliverable-MS35, <http://dx.doi.org/10.5281/zenodo.1298367>.
- [10] The Tcpdump Group, Tcpdump, URL <https://www.tcpdump.org/>.
- [11] M. Christensen, T. Richter, Achieving reliable UDP transmission at 10 Gb/s using BSD socket for data acquisition systems, *J. Instrum.* 15 (2020) T09005, <http://dx.doi.org/10.1088/1748-0221/15/09/t09005>.
- [12] D. Pfeiffer, et al., vmm-sdat - VMM3a/SRS Data Analysis Tool: Analysis software for VMM3a data, recorded with the SRS as PCAP or HDF5 files (GdGEM pipeline of the EFU). URL <http://github.com/ess-dm3c/vmm-sdat>.
- [13] C. Altunbas, et al., Construction, test and commissioning of the triple-gem tracking detector for compass, *Nucl. Instrum. Methods Phys. Res. A* 490 (2002) 177–203, [http://dx.doi.org/10.1016/S0168-9002\(02\)00910-5](http://dx.doi.org/10.1016/S0168-9002(02)00910-5).
- [14] L. Scharenberg, X-ray Imaging with gaseous detectors using the VMM3a and the SRS, *Nucl. Instrum. Methods Phys. Res. A* 1011 (2021) 165576, <http://dx.doi.org/10.1016/j.nima.2021.165576>.
- [15] T. Alexopoulos, et al., The VMM readout system, *Nucl. Instrum. Methods Phys. Res. A* 955 (2020) 163306, <http://dx.doi.org/10.1016/j.nima.2019.163306>.