The DMILL readout chip for the CMS pixel detector

Wolfram Erdmann
Institute for Particle Physics
Eidgenössische Technische Hochschule Zürich
Zürich, SWITZERLAND

1 Introduction

The CMS pixel detector will be the innermost tracking detector of the CMS experiment at the Large Hadron Collider (LHC) at CERN [1]. Bunch crossings will occur in 25 ns intervals with several hundred charged particles produced inside the acceptance of the pixel detector at each crossing.

The readout chip must register the signals produced by the charged particles hitting the sensor and store the information during the latency of the first level trigger of the CMS experiment. The amount of charge collected in a pixel will be used to improve the accuracy of hit reconstruction in the CMS pixel detector. The pulse-height must be stored together with the pixel address and the correct bunch crossing association. Buffering and subsequent readout are only possible with zero-suppression because of the large number of channels involved.

Readout chip development was pursued in the radiation hard DMILL technology [2]. Device density and connectivity of this $0.8\ \mu m$ SOI BiCMOS process with two metal layers are limited compared to the more recent 0.25 μm processes with radiation tolerant layout. The column drain architecture of the PSI43 [3] is well suited to a layout in DMILL, because it permits a simple pixel cell and needs a small number of bus lines connecting pixels and periphery.

DMILL devices are expected to tolerate more than 10 Mrad or $10^{14}/cm^2$ neutrons. It is foreseen to replace layers of the CMS pixel detector when they have received an accumulated dose corresponding to 6×10^{14} charged particles per cm^2. The first part of this article will give a description of the PSI43 chip and its readout architecture pointing out sources of dead-time. Results from a high rate beam test conducted at PSI last summer with measurements of data-losses will be presented in the second part.

2 The PSI43 chip

The active area of the PSI43 chip has a size of about 8×8 mm^2 and contains 53 rows and 52 columns of pixels (fig. 1). The pixel cells are quadratic with a size of $150 \times 150\ \mu$m^2.

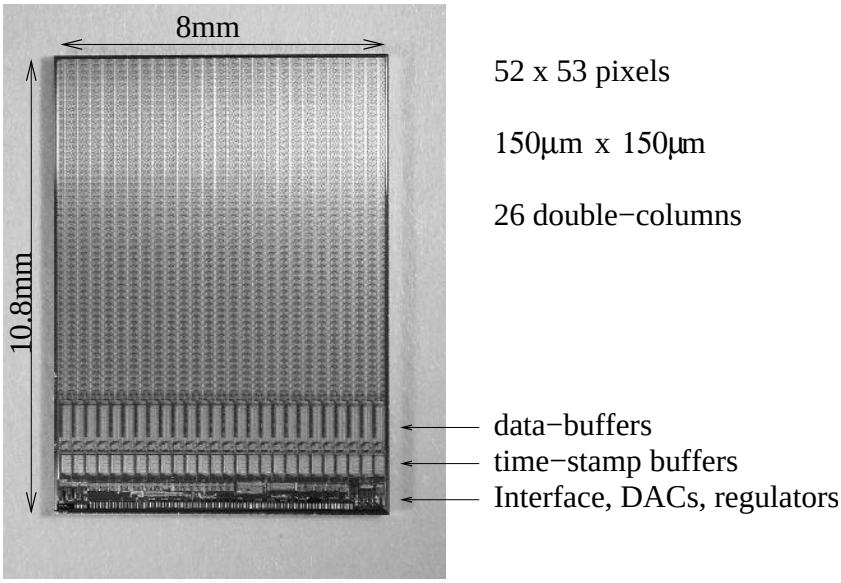

Figure 1: PSI43 readout chip in DMILL technology for the CMS pixel detector.

Two neighboring columns form a double-column with a shared vertical bus in between and a common periphery block containing data- and time-stamp buffers. The 26 double columns operate independently of each other.

The control- and interface block at the bottom of the chip receives and decodes the external control signals and distributes them together with time-stamp and trigger numbers in horizontal busses to the double-columns. It also contains the supply voltage regulators and 21 DACs for internal bias voltages. The DACs are loaded through a fast serial interface that is also used for configuring the pixels.

The height of the periphery is 2.8mm, which is dominated by the time-stamp and data buffers. In the modules of the pixel detector the periphery extends beyond the sensor allowing access to the row of wire bond-pads at the bottom of the chip.

Because of the space requirements of the DMILL devices and in order to optimize the chip yield, the transistor count was kept as low as possible. A pixel cell contains 125 devices, each double column periphery 3k and the control- and interface block contains 6k devices. The full chip uses about 430k transistors.

2.1 Readout architecture

The analog section of the pixel cell consists of a charge-sensitive preamplifier-shaper combination followed by a comparator and a sample-and-hold block. The preamplifier input is connected to a bump-bonding pad with a 15 μm wide opening. Sensor leakage current up to 10 nA is absorbed by the adjustable feedback resistor of the preamplifier. The comparator

threshold is set chip-wide with an 8-bit DAC and can be adjusted for each pixel with a 3-bit trim-DAC. Less than 25 ns time-walk at the comparator output for a 2500 electron input signal has been measured in un-irradiated test-structures with 40 μW power dissipation per pixel in the preamplifier and shaper.

When a pulse crosses the comparator threshold the pulse-height is latched on a storage capacitor after a fixed delay and the double-column periphery is notified. While the pixels operate independently of the bunch crossing clock, the periphery synchronizes the arriving hit information by latching the value of a bunch-crossing counter in the time-stamp buffer. The hit pixel remains inactive until its hit information is transferred to the periphery. The periphery prepares the column drain by setting up the precharged logic of a fast token scan [4] and freezing the hit pattern present in the double-column at that time. New hits arriving later in other pixels of the same double-column will be registered but will only be copied in a later column drain.

In parallel to the time-stamp entry, a marker bit is set in the data buffer, indicating the start of a new block of data that belongs to this time-stamp. The data-buffer can keep up to 23 hits. The pulse height and the analog coded address of the pixels are stored in 4 capacitors per entry. During a column drain the data-buffer is filled at a rate of 20 MHz until all pixels with hits have transferred their data. Empty pixels are skipped with 1.8 GHz. The time-stamp buffer can hold up to 8 time-stamps and is 8 bits wide allowing trigger latencies up to 255 bunch crossings. Both buffers are written in a circular manner and the oldest time-stamp and its hits are discarded after the trigger latency if not validated by a trigger.

Entries validated by a trigger must not be overwritten. In order to keep the control logic simple, the double-column stops data acquisition when a time-stamp is confirmed. The resulting inefficiency is tolerable because other double-columns that don't have data for this trigger are not affected and continue data acquisition. A 4-bit trigger number is latched in the double-column which allows keeping data sorted by triggers during readout.

The serial readout is controlled by a readout token that is passed from chip to chip and inside the chips from double-column to double-column. There must be exactly one token for every trigger and a token will only collect data that belongs to the same trigger number even when more triggers arrive before a chip is read out.

The PSI43 has a differential open collector output for analog data. The readout is clocked by the bunch crossing clock, optionally divided by two internally. Every readout chip starts sending a three cycle header when it receives a token. While the header is transmitted, the token is passed through the chip looking for a double-column waiting for this token. The length of the header is sufficient to go through all 26 double-columns if necessary and then pass on the token with the right timing for the next chip. The header starts with a full swing signal outside of the range used for pixel data ("ultra-black") followed by an arbitrary but well separated level ("black"). This sequence unambiguously identifies the beginning of a new chip in a sequence of analog levels. No further chip ID is present in the data stream. The third cycle of the header can contain a level representing the value of

the most recently programmed DAC. This is the only available read-back of configuration data.

When no double-column has data for a given token, the next chip will follow immediately. Otherwise, a six-cycle sequence is sent for every hit. Two values represent the analog coded double column address, three values identify the pixel inside a double column and the last value is the analog pulse height. The addresses are encoded using 5 discrete levels except for the last double column, where a sixth level is needed in the double column address.

2.2 Control and supplies

The PSI43 is configured and operated with a set of LVDS signals. Two signals carry a fast serial protocol for downloading the setting of 21 DACs, a control register, trim- and kill-bits. Running with a 40 MHz clock, a full configuration takes less than one millisecond. Random access to individual pixels is possible and frequent reloading is foreseen to correct misconfigurations caused by SEU. The address identifying the chip on the serial bus is defined by wire-bonds.

Data acquisition is controlled by the 40 MHz bunch crossing clock and the first level trigger signal, which also encodes resets and calibration triggers.

The PSI43 needs four supply voltages. Internal regulators, controlled by configurable DACs, produce the voltages used inside the chip.

A $-5V$ supply feeds the digital sections. Separate regulators are used for data acquisition related parts and communication. The remaining voltages are used in the pixel cells only. The preamplifiers and shapers are operated from a $-2.5V$ supply regulated to about $-2V$. The regulators ensure that this very sensitive voltage is well controlled for every readout chip regardless of voltage drops in different length supply cables. The value of the analog voltage, and thus the preamplifier current, is adjusted with an 8-bit DAC. The supply voltage of the comparator determines the hit threshold and is also set by an 8-bit DAC.

The regulators provide good rejection of supply ripple up to 200 kHz. The regulated voltages are connected to wire-bond pads allowing external bypass capacitors to be connected that filter higher frequencies.

3 High rate testbeam

The PSI43 was operated in a high rate pion beam at PSI in Summer 2002. The main purpose of this test was to verify that the implemented readout architecture was able to cope with the particle rates encountered at the LHC.

The beam intensity was variable up to 30×10^6 tracks$/cm^2 s^{-1}$, which is about the expected intensity in the 4 cm layer of the CMS pixel detector at high luminosity operation

of the LHC.

3.1 Setup

The beam consisted mostly of pions and the beam momentum was 300 MeV/c. With a beam spot size of 10 mm$\times$20 mm (FWHM) it covered the whole readout chip. The PSI accelerator runs with a 50 MHz bunch structure. The PSI43 was operated with a synchronized 40 MHz clock. Triggers were only allowed in one out 5 cycles where the two clocks lined up. This ensured that triggered particles had the correct timing relative to the 40 MHz clock.

Two small plastic scintillators, $2 \times 2 \times 2$ mm^3, mounted upstream and downstream of the chip under test, provided particle triggering and beam intensity measurement. The scintillators were mounted inside air filled plastic cylinders wrapped in reflective mylar acting as lightguides to photo-multiplier tubes mounted on either side of each scintillator outside of the beam envelope. A coincidence of all four photomultipliers was required for a trigger.

At high beam intensity the rate of scintillator triggers is much higher than the first level trigger rate in CMS. The coincidence with a random signal provided by a radioactive source was used to reduce the trigger rate to 30 kHz or less.

Assemblies consisting of a single PSI43 chip bump-bonded to a matching 280 μm thick silicon sensor and mounted on a PCB were operated in the beam. Clean detection of beam particles is evident in the hit pattern found in triggered events that shows the shape of the scintillators (fig.2). The decoding of the analog coded pixel addresses was complicated by the fact that the levels apparently differed between double columns and to some extent also between pixels of the same double column due to transistor mismatching. Nevertheless the fraction of events reconstructed with a wrong pixel address was at the per-mille level.

No provisions were made for cooling and the speed of the chip turned out to be marginal for 40 MHz operation under those conditions. In particular the data buffer control did not function reliably sometimes leading to the readout of the wrong data. It was possible to circumvent these problems with an asymmetric 40 MHz clock that relaxes the most critical timing constraint.

The correct assignment of hits to bunch crossings is an important issue for LHC detectors. Ideally all events triggered by the scintillator at the testbeam should be found with the same chip setting for the trigger latency. The maximum efficiency observed at low beam intensity was 98% in one bunch crossing. About 2% were found in the following bunch crossing while all other latencies yielded hit rates compatible with random background.

3.2 Data loss

The readout architecture of the PSI43 is not entirely dead-time free. Given the limitations of the chip technology, data losses were accepted in places where avoiding them would

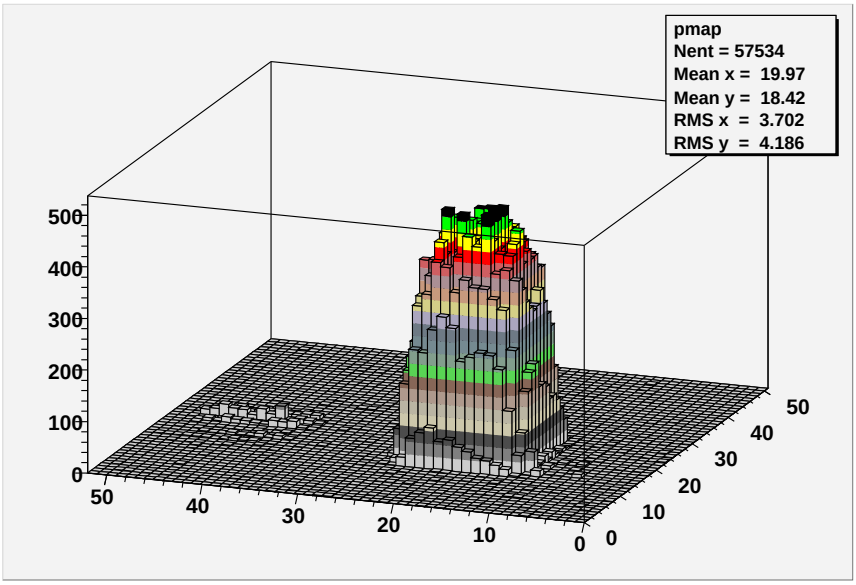

Figure 2: Hit map for events triggered with the scintillators. The 2 mm$\times$ 2 mm region containing hits reflects the size of the scintillators. A small fraction of the events has hits with misreconstructed row addresses.

have increased the complexity of the chip disproportionally.

Data losses occur at several stages in the readout chain (table 1), starting in the pixel cells which are unable to register new hits while waiting for a column drain ("pixel overwrite"). All other pixels of a double-column are sensitive during the column-drain and new time-stamps can be set with two important exceptions. The periphery can only process one pending column drain ("column busy") and during the clock cycle immediately following a hit no new time-stamp can be set in the same double-column ("CD-setup").

Data loss also occurs when either the time-stamp or the data-buffer are full temporarily ("TS-,DB-overflow"). While the column-drain related losses depend only on the hit rate, the buffer overflows are also influenced by the trigger latency. Larger buffers would have made the size of the periphery unacceptable for module construction. The dead-time caused by stopping a double-column after time-stamp validation ("DC-waiting") depends on latency and readout time but is usually small.

The dead-time during the column drain setup is expected to be the dominant source of data loss under test beam conditions. Other losses are smaller because the pixel multiplicity per track is lower than it will be in LHC conditions with non perpendicular tracks and charge sharing due to the magnetic field. The measured inefficiency for calibration signals confirms this expectation (fig.3). For this measurement calibration signals are injected at regular intervals regardless of scintillator triggers. In the absence of tracks passing through the assembly the injected signals are read out with 100% efficiency. With increasing beam

	LHC 4 cm layer $L = 10^{34} cm^2 s^{-1}$	LHC 7 cm layer $L = 10^{34} cm^2 s^{-1}$	PSI test beam
track density	22 MHz/cm^2	12 MHz/cm^2	30 MHz/cm^2
CD setup	3.90%	1.80%	1.90%
Column busy	0.95%	0.23%	0.15%
pixel overwrite	0.38%	0.16%	0.08%
DB overflow	0.12%	0.04%	0
TS overflow	2.80%	0.05%	0.11%
DC waiting	0.59%	0.21%	0.12%

Table 1: Simulation of data-loss in the PSI43 in different environments. A description of individual sources can be found in the text. The average hit multiplicity is close to one in the test beam and 4-5 at the LHC. The 4 cm layer was originally not intended for high luminosity operation. For low luminosity LHC running ($L = 10^{33} cm^2 s^{-1}$) the sum of losses is less than 2% at 4 cm radius.

intensity, the fraction of empty readouts rises proportionally to the beam intensity as expected from data-loss simulations (fig.3). The number of events where a hit is found in a randomly selected bunch crossing increases at almost the same rate. This is a direct measure of the probability that a hit occurs in the same double column in the bunch crossing preceding any triggered event. It is almost equal to the total data-loss observed with calibration signals and other sources of data-loss are indeed small in the test-beam environment.

The inefficiency for beam particles, measured by looking for pixel hits in events triggered by the scintillators, is also shown in figure 3. In contrast to the calibration signals, a 2% inefficiency is observed at very low beam intensity. As shown before, data-losses in the readout do not occur in the limit of vanishing beam intensity and the inefficiency must be attributed to other sources such as time-walk or false triggers. As in the case of calibration signals a linear increase of the inefficiency with rising beam intensity is observed, however, the slope is larger than expected by almost a factor of two. The reason for this additional data-loss is unclear. The measurement with calibration signals was carried out with a large injected signal (14k e) and a high threshold (6k e) in order to have a well defined situation that tests the readout architecture. A much lower threshold of 3k was needed to get maximum efficiency for beam particles and it has been observed that also the slope measured with calibration signals becomes larger when the threshold is lowered.

3.3 Single Event Upset

A small fraction of the charged particles going through the readout chip can lead to charge depositions in sensitive nodes of the circuits large enough to cause a memory cell to change state. This has been observed in DMILL test-structures [6].

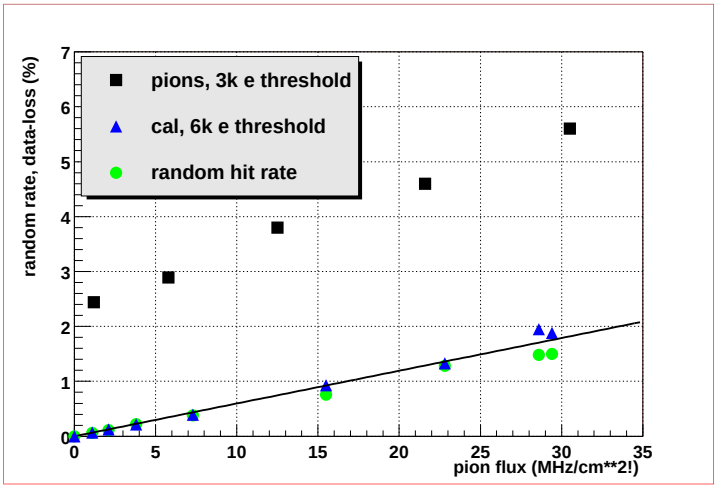

Figure 3: Data-loss measured with calibration signals and beam particles. The data-loss for calibration signals is almost equal to the random hit rate per double column, which is expected when the CD-setup is the dominant source of dead-time. The straight line shows the expected hit rate based on intensity and double-column area.

A measurable rate of such effects is expected to occur in the full readout chip when exposed to the high rate test beam. Since the PSI43 does not permit the readback of any internal registers, this was measured by setting up the readout chip for data taking but disabling all 2756 pixels. A single event upset in the relevant cell can then be detected by the fact that the pixel will start to report hits. After $6\frac{1}{2}h$ irradiation with a total of 4.2×10^{11} pions per cm^2, 17 pixels were turned on. The corresponding SEU cross section of a single cell is $(1.5 \pm 0.4)10^{-14} cm^2$, in good agreement with test structure results. Reloading pixels in intervals of a few seconds will be sufficient to keep the fraction of misconfigured pixels at the per-mille level even in 4 cm layer. With a millisecond download time this can be done with negligible dead-time.

4 Summary

The PSI43 is the first full readout chip for the CMS pixel detector. It has been fabricated in the radiation hard DMILL technology and is fully functional. Single chips with sensors were operated in a high rate test beam at PSI with track densities comparable to those expected at the LHC. Architecture inherent data-losses were found to be in agreement with simulations. The overall efficiency for beam particles was found to drop from 98% to 94% for highest beam intensities which is lower than expected from data-losses alone.

References

[1] The CMS collaboration, *CMS Tracker Technical Design Report*, CERN/LHCC 98-6.

[2] M. Dentan *et al.*, IEEE Trans. Nucl. Sci. **43** (1996), p. 1763.

[3] R. Baur, R. Horisberger, R. Schnyder, M. Lechner, B. Meier, in *Proceedings of the Fourth Workshop on Electronics for the LHC*, CERN/LHCC 98-36.
R. Baur [CMS Pixel Collaboration], Nucl. Instrum. Meth. A **465**, 159 (2000).

[4] B. Meier, Diploma Thesis, ETH-Zürich, April 1998.

[5] D. Kotliński, Nucl. Instrum. Meth. A **477**, 446 (2002).

[6] M. Barbero, R. Baur, K. Gabathuler, R. Horisberger, CMS-Note 2000/036.