Performance and Radiation Testing of a Low-Noise Switched Capacitor Array for the CMS Endcap Muon System.

Breedon, R.E., Holbrook, B., Ko, W., Mobley, D., Murray, P., Tripathi, M.

Department of Physics, University of California, Davis, CA 95616 USA
breedon@ucdhep.ucdavis.edu

Abstract

The 16-channel, 96-cell per channel switched capacitor array (SCA) ASIC developed at UC Davis for the cathode readout of the cathode strip chambers (CSC) in the CMS endcap muon system is ready for production.

For the final full-sized prototype, the Address Decoder was re-designed and LVDS receivers were incorporated into the chip package. Under precision testing, the chip exhibits excellent linearity within the 1V design range and very low cell-to-cell pedestal variation. Monitored samples of the production design were subjected to exposure to a 63.3 MeV proton beam. The performance of chips after exposures up to 100 krad was within tolerances of an unexposed part.[1]

I. INTRODUCTION

During the first level trigger latency period of approximately 128 bunch crossing (3.2 μs), signals from the front-end electronics of sub-detectors in CMS must be held in temporary storage before being passed
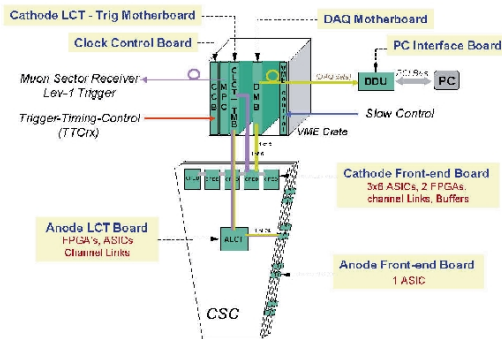

Figure 1: Endcap muon front-end electronics schematic

to the DAQ system or rejected. The cathode readout of the cathode strip chambers (CSC) of the endcap muon system employs analogue storage centred on a switched capacitor array.

The cathode strip chambers in each endcap are arranged in four stations that are interleaved among the steel disks of the flux return yoke. In order to provide for both charged track reconstruction and triggering in the forward region with its high charged particle fluxes, each chamber has six layers of radial cathode strips (for

the precise position measurement in the bending plane) and six layers of approximately azimuthal anode wires (for the coarse radial position measurement and timing for bunch crossing identification). Including the innermost CSC (ME1/1), there are more than 266,000 cathode channels, thus data reduction electronics are mounted directly on the detectors.

One cathode front-end board (CFEB) [1] reads out a tower of 16 adjacent strips from each of the six layers. The anode wires are ganged together in groups of ten to
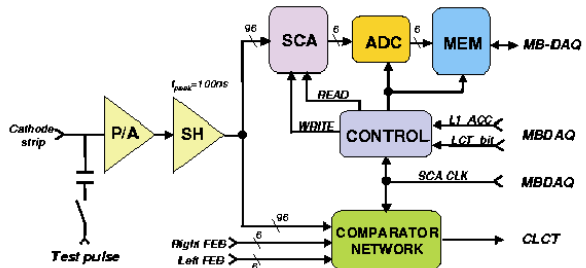

Figure 2: Block diagram of the cathode front-end board

twenty--depending on rapidity--before being connected to an anode front-end board (AFEB). The data from the front-end boards are sent to motherboards on an external VME crate (Figure 1) that provide the links to the main DAQ and the level-1 trigger of the experiment.

The anode wire signals are immediately discriminated to provide the bunch-crossing time stamp, so digital pipelining is used. However, full-wave sampling and analogue storage of the precise cathode measurement allows the charge on the cathode strips to be recorded accurately for cluster reconstruction. The switched capacitor array provides a higher level of control over pileup and baseline shift than other pipeline options.

II. CATHODE READOUT ELECTRONICS

The pulse from a cathode strip of a CSC is pre-amplified and shaped (peaking time=100 ns), then is split into two signals: one potentially to contribute to a trigger primitive (called a Local Charged Track, or LCT), the other to provide the strip charge measurement (Figure 2). The latter is sampled by the SCA at 20 MHz and the samples are stored awaiting the trigger decision. The eight samples of a pulse that are saved in the SCA enable the strip pulse to be reconstructed to better than 1% accuracy, as required for track fitting and identification of pileup events.

The eight samples stored per pulse include two or three (the number is programmable) leading samples before the pulse that are available to be averaged offline for an estimate of any baseline shift due to event pileup or any other electronic effect. In this way the "pedestal" is determined on an event-by-event basis, providing a high level of control over baseline shifts.

One CFEB handles sixteen cathode strip channels from each of the six layers of a CSC. Sixteen channels are handled by one SCA, so there are six SCA chips per CFEB. Fast-shaped pulses from one layer are combined in a comparator network that locates the centroid of an induced strip charge cluster with a resolution of half the strip width, and marks its time. The cluster positions from all layers in a tower are fed to the local charged track (LCT) processor, which searches for a coincidence of at least four out of six layers from a pre-determined set of acceptable patterns. The time, location and local angle of the LCT are passed via the cathode LCT motherboard as trigger primitives for the level-1 muon trigger.

Upon a delayed coincidence of the LCT associated with a pulse and the level-1 trigger accept, the stored voltage samplings in the SCA are digitised and stored in memory to await transfer to the main DAQ via the DAQ motherboard. The readout control chip controls writing to and reading from the SCA, digitisation and memory storage.

III.   THE SWITCHED CAPACITOR ARRAY

The SCA analogue storage chip is subdivided into 16 parallel channels of 96 sample-and-hold cells (capacitors) per channel. Approximately 64 of the cells on one channel are required to cover the level-1 latency period of the CMS global trigger, while the additional cells provide a readout buffer. Each of the 16 channels has a dedicated analogue input that is connected to a bus distributing the input signal to the 96 cells.

Each sample-and-hold cell consists of a pair of complementary CMOS transmission gates and a double-polysilicon capacitor. An externally supplied reference voltage is applied to the bottom plate of the capacitor. The voltage stored on each sample-and-hold capacitor then corresponds to the difference between the input signal from a given channel at the time its sample-and-hold switches are closed and the applied reference voltage. The reference voltage can be adjusted to shift the baseline to optimise level matching. During readout, each capacitor is placed in the feedback path of an op-amp with one op-amp per channel.

The SCA is an externally address analogue memory, with addresses generated by the readout controller chip. The cells may be addressed in any order, but on the CFEB addresses are always generated in blocks of eight in sequential order in grey code (see Ref. [1]). Charge can be stored simultaneously in one capacitor while being retrieved from a different capacitor in the same channel. This ensures deadtime free operation of the pipeline.

Switched capacitor arrays used in other high energy physics applications have often relied upon a set of calibration constants for each capacitor cell to achieve the desired charge measurement resolution. The cathode readout of the CMS endcap muon system (plus ME1/1, the innermost CSC and a Dubna responsibility) will have more than 25.5 million capacitor cells among 16,632 SCA chips. Thus from the start of the project, a major design goal has been for the cell-to-cell variation within a channel to be small enough to make such a calibration unnecessary.

A. SCA development

Beginning in 1994, a three-channel, 28-cell development prototype ASIC[2] went through five iterations in the HP 1.2 µm, double-metal process with linear capacitors using the university consortium MOSIS service at the University of Southern California. The first full-sized prototype was produced in November 1996. In 1998 the HP process (CMOS34) was transferred to American Microsystems, Inc. (AMI) of Pocatello, Idaho so the chip had to be partially redesigned following new design rules.

The final production prototype in the AMI standard CWL 0.8 µm process was received in September 1999. It bears the appellation SCA-4A. The last major design improvements include a redesigned address decoder so that no NAND gate has more than three inputs for more consistent throughput delay. In addition, to conserve space on the CFEB, LVDS receivers (which were originally in six separate chips per CFEB) were incorporated into the SCA ASIC. Both of these changes likely contributed to the reduced noise observed compared to the previous iteration (SCA-3B).

The power consumption of SCA-4A is low, about 30 mA per 16-channel chip.

B. SCA performance testing

For bench testing the prototypes, a Tektronix 9200 Data Acquisition System (DAS) is used. A custom designed test board with 12-bit ADC and DAC provides the interface to the DAS under computer control and readout. Under DAS control, typically a test pulse is written into the SCA under test at 60 ns intervals[2] and read out at approximately 120 ns intervals. It is important to note that at the sensitive level of testing required, the test board itself has its own characteristics (e.g. noise) that cannot always be isolated from those of the SCA under test. Thus many tests yield only an upper limit on SCA performance.

Among the parameters tested at UC Davis are the following: gain, linearity (optimised over a 0-1V range), pedestal variation (capacitor cell-to-cell), noise, cross-talk and capacitor droop rate. The effect of the latter two parameters was so negligible that the tests were not repeated on the latest prototype. (Results for an earlier prototype may be found in Ref. [3].)

---

[2] In CMS the SCA will operate with a 50 ns Write interval, however, the test setup only allows either 60 ns or 40 ns intervals.
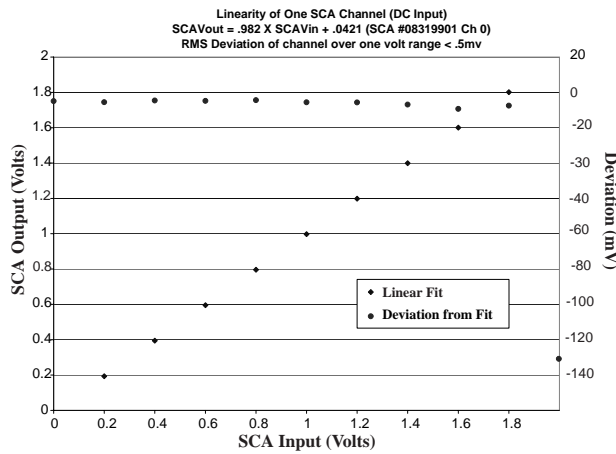
Figure 3: Linearity measurements on one channel of an SCA

In actual operation, each channel of the SCA will be a link in a readout chain connected to a CSC strip that will be calibrated individually. To assess the linearity of the SCA, the deviation from a linear fit is measured (Figure 3). The figure shows an example of the linearity measurement on one channel of a sample SCA. A constant level was written into the channel over 0-2V effective input voltage in 0.2V steps. The SCA output is shown on the left scale of the figure plotted against the input voltage. A linear fit was made over the 0-1V and the deviation of the measured output levels from this fit is plotted on the right scale. The RMS deviation from linearity over 0-1V is less than 0.5 mV.
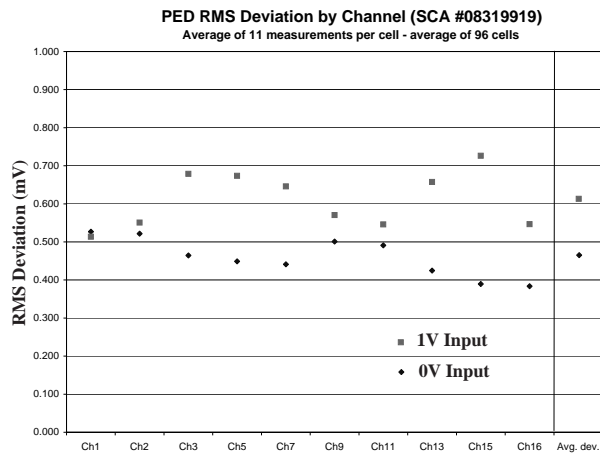

Figure 4: RMS pedestal variation for selected channels of an SCA and their averages

The design goal for pedestal variation among the individual capacitors of an SCA channel is for the average RMS variation to be <0.5 mV. A constant voltage level is written into and read from each of the 96 cells in a channel eleven times. The RMS deviation from the average output of all 1056 measurements is calculated for each of the 16 channels (Figure 4). The figure demonstrates a slight dependence on input voltage: The average deviation for an effective input of 0V is 0.46 mV while for 1V it is 0.62 mV, as shown in the right column of the figure. Note that these figures represent upper limits as they include an unknown contribution from the test board.

Uncorrelated noise levels (Figure 5) exhibit no dependence on input voltage. The average for input voltages both 0V and 1V is 0.5 mV.

Recent SCA prototypes have also been tested in a variety of other environments. At Ohio State University where the cathode front-end board is being produced, the chip was tested on CFEB prototypes with excellent results. 36 SCA chips populated the six cathode boards used to fully instrument full-size prototype cathode strip chambers for beam tests at CERN in 1998 and 1999. Some of these tests took place at the Gamma Irradiation Facility (GIF), where a radioactive source is combined with a muon beam to simulate the expected endcap muon background environment at the LHC. Pedestal analysis of GIF beam test data shows results similar to bench tests of the SCA (average of 0.64 mV RMS variation of cells within a channel [4]). Finally, the SCA is part of cosmic ray tests of CSC prototypes at Fermilab. In all these cases the SCA has performed well within specifications.
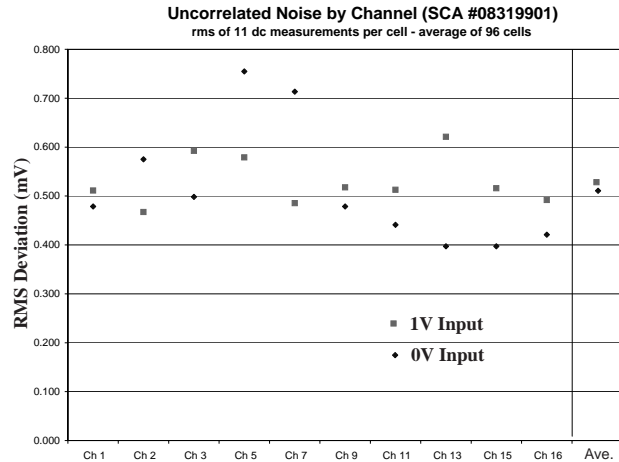

Figure 5: Uncorrelated noise measurements of selected channels of an SCA and their averages

C. Radiation testing

As was mandated by the CMS administration, all electronics must be certified against the radiation exposure that might be anticipated during many years of LHC running. Because the SCA, which is not produced in a radiation-hard process, has no on-chip memory, it needs to be tested only for total ionisation dose (TID) and single event latch-up (SEL). Expected radiation levels in CMS in the LHC environment have been estimated using the simulation program FLUKA by M. Huhtinen at CERN. Assuming 10 years of LHC running at full luminosity ($10^{34}/cm^2/s$) the largest estimated exposure to any part of the endcap muon system was still rather low: TID=1.8 krad/10 LHCy on the 50 cm of ME1/1 closest to the beam line. For the remainder of ME1/1 and the rest of the endcap muon system, no part is to receive more than 0.8 krad/10 LHCy. A factor three safety factor was chosen by the endcap muon electronics community, so all EMU electronics must be able to survive an exposure of 6 mrad.

Figure 6: Proton beam profile at Crocker Nuclear Lab

Samples of SCA-4A were exposed to the 63.3 MeV proton beam[3] in the Radiation Effects Area at the Crocker Nuclear Laboratory (CNL) on the campus of UC Davis in March 2000. The CNL's 76-inch isochronous cyclotron can provide a beam of protons with a current from 2 pA ($\sim$3.4 x 10^5 p/cm^2) up to 100 nA. Originally developed to perform radiation effects measurements on photonic and electronic devices, especially for aerospace applications, the Radiation Effects Area has hosted the radiation testing of nearly all the electronics that will be mounted on or near the endcap muon system.[4]

The CNL proton beam profile (Figure 6) is almost flat over a radius of $\sim$35 mm, with 97% of maximum available at 20 mm and 92% at 30mm. This allows uniform radiation of even large chip carriers. The sample SCA (size 14 mm x 20 mm) was mounted in a Zero Insertion Force (ZIF) socket centred on the beam that could be adjusted over a range of angles.

A number of SCA chips were tested under a variety of conditions:

- both powered and un-powered;
- both ceramic (development) and plastic (production) chip packages;
- at three angles with respect to the beam (5°, 45° and 90°);
- at several radiation exposure levels from a few krad up to a total single exposure of 1 Mrad.

During exposure and for part of the annealing time that followed, the chip under test was monitored with a PC that recorded the following parameters:

- total power supply current draw (to test for latch-up);
- output quiescent level (multiplexer baseline output voltage);
- transistor threshold levels (V_{bias1}, V_{bias2}).

Even at exposure levels up to 100 krad, the performance of chips bench tested after irradiation was within tolerances of an unexposed part.

In an attempt to irradiate a chip to the point of failure, one SCA sample was exposed to a total of 1 Mrad. The exposure was accomplished in steps of TID levels of 75, 225 and 1000 krads. After each step, the chip was bench tested (Figure 7). The input signal was a triangular 1V pulse, timed so that it was sampled at 0V by cell number 1 of the channel under test, peaking at the middle, and coming back to 0V by cell number 96. The output was compared to the output of a "standard" channel on a non-radiated chip, and the difference is shown in the figure. As the exposure level increased there was an unimportant overall level shift, but no other effect was observed until the maximum exposure level of 1 Mrad. The curve in the set of measurements associated with the 1 Mrad exposure indicates a change in the gain of no more than 2 mV. There was also evidence of a slight increase in noise.

IV.   CONCLUSIONS AND PRODUCTION PLANS

Innovative design features applied over years of development have resulted in an SCA design with excellent linearity and very low noise. Beam test data analysis is also consistent with the conclusion that the pedestal variation is so small that an elaborate cell-to-cell calibration will not be necessary. Radiation testing of the SCA has also shown its design to be remarkably robust against radiation levels far above what is expected during years of LHC operation at full luminosity. Specifically, no SEL (latch-up) was observed, and even at 1 Mrad exposure the SCA continues to work with only about twice the low noise level and a slight degradation in linearity.

With the passing of the final hurdle of radiation tolerance certification, the procurement of 25,000 SCA chips is underway. The wafers will be produced by AMI. The vendor will also arrange packaging in injection-moulded plastic and is also responsible for the chip production yield. AMI has been supplied with a description of the tests they are to perform on each chip and a list of test vectors. In this way AMI will deliver only working and tested parts so a yield estimation does not have to be included in the production order.

After making the masks, but before mass production begins, AMI will provide a number of engineering samples of the SCA for local testing and approval. Randomly chosen chips will also be tested as production progresses.

---

[3] Note that as dE/dx at 63.3 MeV is about six times that at minimum ionising energies, 2 krads can be provided by the CNL proton beam in a matter of seconds.

[4] For example, see the paper by T.Y. Ling in these Proceedings.

V. References

[1] Proceedings of the Second Workshop on Electronics for LHC Experiments, Balatonfüred, Hungary, September 23-27, 1996, CERN/LHCC 96-39, p. 309.

[2] Proceedings of the First Workshop on Electronics for LHC Experiments, Lisbon, September 11-15, 1995, CERN/LHCC 95-56, p. 204.

[3] Proceedings of the Third Workshop on Electronics for LHC Experiments, London, September 22-26, 1997, CERN/LHCC 97-60, p. 324.
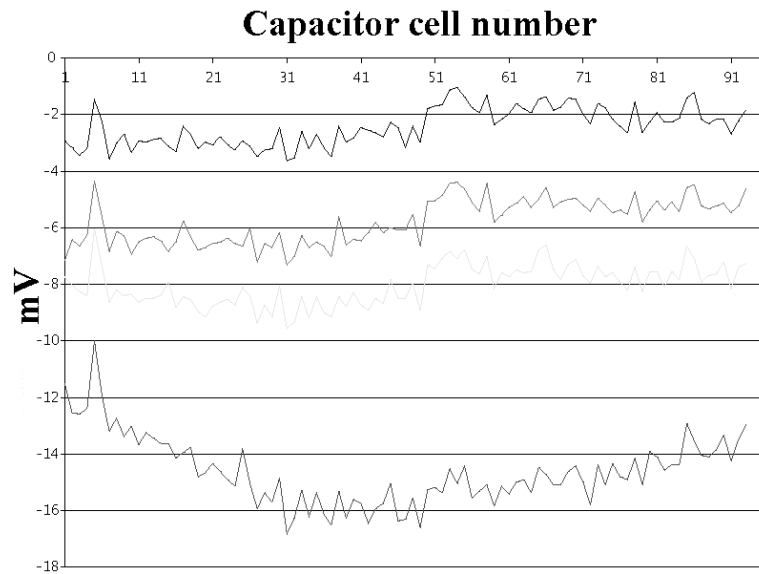
[4] CMS Internal Note IN-1999/049.

Figure 7: Change in the output of one SCA channel by cell at four increasing radiation doses, (from top to bottom) 0 krad, 75 krad, 225 krad and 1000 krad. The left axis is the difference of the output of each cell in mV from a chosen "standard" channel on a non-irradiated chip.