BeamCal Instrumentation IC: Design, Implementation and Test Results

Angel Abusleme, *Member, IEEE,* Angelo Dragone, Gunther Haller, and Bruce A.Wooley, *Fellow, IEEE*

***Abstract*—The BeamCal detector, one of the calorimeters in the forward region of the International Linear Collider detector, will serve three purposes: ensure hermeticity of the detector for small polar angles, reduce the backscattering from pairs into the detector center, and provide a low-latency signal for beam diagnostics. The BeamCal specifications in terms of noise suppression, signal charge, pulse rate and occupancy pose unique challenges in the front-end and readout electronics design. The Bean – BeamCal Instrumentation IC – is a 32-channel front-end and readout IC that will address the BeamCal instrumentation requirements. By employing switched-capacitor filters and a slow reset-release technique, the Bean will process the signal charge at the International Linear Collider pulse rate. Each channel will have a 10-bit successive approximation analog-to-digital converter and digital memory for readout purposes. The Bean will also feature a fast feedback adder, capable of providing an 8-bit, low-latency output for beam diagnostic purposes. This work presents the design and characterization of the Bean prototype, a 3-channel IC that proves the principle of operation described.**

I. Introduction

THE BeamCal is one of the detectors planned for the International Linear Collider (ILC) very forward region [1]. The BeamCal will serve three purposes: ensure the hermeticity of the detector for very small polar angles, reduce the backscattering from pairs into the detector center, and provide a feedback signal for beam diagnostics. The input signal comes in pulse trains, where each input pulse represents the outcome of a bunch crossing. Fig. 1 shows the pulse train structure. Each pulse train has 2820 pulses, $308\,\mathrm{ns}$ apart, followed by a 199-ms idle period. Since 100% occupancy is expected, all events must be recorded and stored in the IC to be read out during the idle period. In addition to this per-channel, 10-bit output for science purposes, for each collision the BeamCal IC must provide a low-latency, 8-bit diagnostics output consisting of the sum of the outputs of all 32 channels in the IC.

The Bean signal path for each of the 32 channels must be able to cope with two different modes of operation: the *standard data taking* (SDT) mode, and the *detector calibration* (DCal) mode. According to the initial BeamCal specifications, the maximum input signals were to be about $36.9\,\mathrm{pC}$ in the SDT mode, and 50 times smaller in the DCal mode.

A. Abusleme was a PhD student at the Department of Electrical Engineering, Stanford University, Stanford, CA 94305. He is now with the Department of Electrical Engineering, Pontificia Universidad Católica de Chile (telephone: +56-2-354-4809, e-mail: angel@ing.puc.cl).

A. Dragone and G. Haller are with the SLAC National Accelerator Laboratory, Menlo Park, CA 94025 (e-mail: dragone@slac.stanford,edu, haller@slac.stanford.edu).

B. A. Wooley is with the Department of Electrical Engineering, Stanford University, Stanford, CA 94305 (e-mail: wooley@par.stanford.edu).

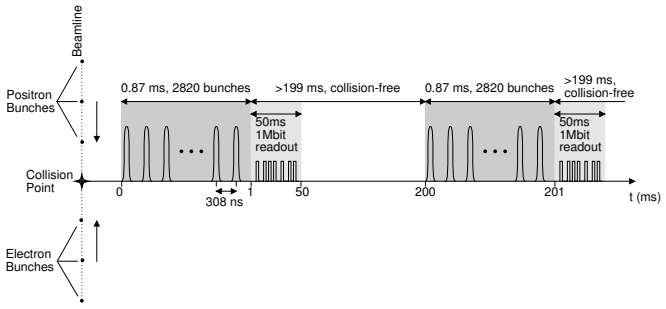

Fig. 1. Structure of the ILC pulse train.

This work introduces the Bean prototype, a scaled-down version of the Bean, intended as a proof-of-concept.

II. System-level design

The Bean prototype has three identical channels with 10-bit outputs for science purposes, and an additional 8-bit output with the sum of the outputs of all channels, for beam diagnostic purposes. Fig. 2 shows a simplified block diagram of the IC.

The Bean signal path has been designed taking into account noise considerations. The noise in a pulse processor can be modeled as an input-referred combination of a series noise generator and a parallel noise generator. The input-referred noise depends on the total capacitance at the input node, the detector leakage current (which affects the shot noise content), the charge-sensitive amplifier (CSA) input-referred noise, and the noise coefficients derived from the weighting function $W(t)$ [2]. With only $308\,\mathrm{ns}$ for signal processing, the series noise component [3] will be dominant in the Bean noise equation. This is because the series noise coefficient is proportional to the slope of $W(t)$, which needs to be large in order to define the weighting function shape in such a short time. Due to the series noise dominance in the Bean, a triangular-shaped weighting function, which minimizes the series noise, represents an adequate solution. Thus, the Bean signal path aims for a triangle-like weighting function when explicit noise-filtering is required.

The channel signal path starts with a dual-gain CSA, that is able to cope with the speed requirements of both modes of operation. The CSA output voltage after reset (baseline) is shifted by a precharger circuit (Cal), which injects a known charge into the front-end allowing to take advantage of the voltage range over which linearity is acceptable. Due to its finite bandwidth, the CSA alone is effective in filtering noise

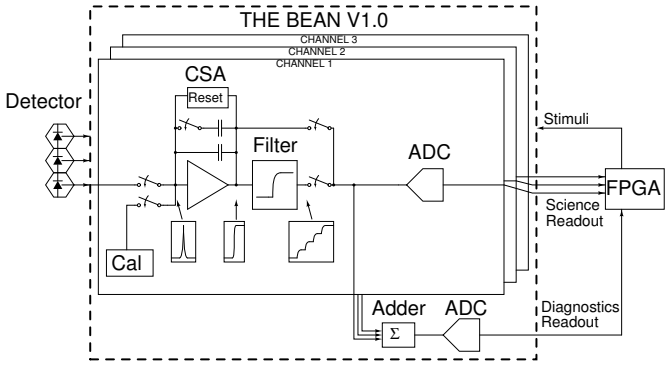

Fig. 2. The Bean prototype simplified block diagram.

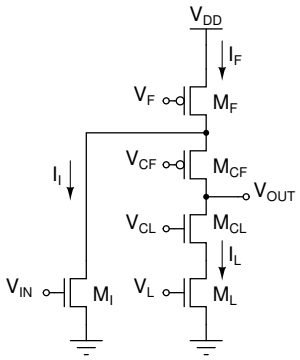

Fig. 3. Schematic of CSA amplifier. Bias circuits have been omitted.

in the SDT mode. However, an explicit filter is required for the DCal mode in order to minimize the series noise component. The filter is implemented using a switched-capacitor (SC) integrator that shapes the negative slope of the weighting function. The positive slope of the weighting function is shaped by a slow, programmable reset-release implemented in the CSA feedback network. Together, the filter and the slow reset-release produce a triangle-like weighting function. The last block in the signal path is a 10-bit successive approximation analog-to-digital converter (SAR ADC) that digitizes the output of each event. Future revisions of the Bean will include a digital memory array, which is currently implemented off-chip.

The beam diagnostics output of the IC is generated by a fully-differential SC adder, that sums the outputs of the three channels. The adder feeds an 8-bit ADC that produces a low-latency digital output.

In order to maximize the time allocated for filtering, a simple pipeline scheme has been implemented. The pipeline allows both the front-end and ADC to make full use of the $308\,\mathrm{ns}$ available between bunch crossings. When the front-end has finished processing the outcome of a bunch crossing, the ADC starts its operation. Then, the front-end is reset and waits for a new input, while the ADC is still converting the outcome of the previous input.

The budget for the Bean electronic noise has been set at $2.75 \cdot Q_n^2$, where Q_n^2 is the quantization noise power. This is equivalent to less than $0.5\,\mathrm{LSB}$ (RMS) of noise (standard deviation). The noise power is allocated as follows: Q_n^2 due to the CSA, $0.5 \cdot Q_n^2$ due to the filter, $0.25 \cdot Q_n^2$ due to the buffers, and Q_n^2 due to the ADC.

Early calculations and validating simulations of the SC filter showed that a sampling rate of $51.95\,\mathrm{MHz}$, equivalent to 16 sampling periods per collision, represents an adequate tradeoff between circuit complexity and performance. A higher sampling rate makes the design of the SC filter more challenging, whereas a lower rate has a reduced filtering capability as the aliasing increases the integrated noise. Therefore, the Bean is designed to process the input pulses at the 3.247-MHz ILC collision rate, with an internal clock 16 times faster. The 308-ns period between pulses constitutes one cycle, and the 16 clock periods within each cycle define subcycles and sampling periods of the switched-capacitor circuits.

III. Circuit design

The Bean prototype was designed for integration in a 180-nm mixed-signal technology. This section presents some of the circuit design details.

A. CSA

The CSA transfers the input charge Q_{in}, generated in the detector, into a capacitor of size C_F, producing an output voltage $V_o = Q_{in}/C_F$. This is done by means of capacitive negative feedback around a voltage amplifier [4], implemented in this case as a single-ended folded-cascode voltage amplifier with cascoded active load. Fig. 3 shows the amplifier schematic without the feedback network. This topology was chosen for its simplicity, linearity, low noise and excellent gain and bandwidth. Since the amplifier gain is finite, the closed-loop linearity is sensitive to the variations in the amplifier open-loop gain. In order to preserve the closed-loop linearity, the CSA output swing must be limited to the region where the variation in the amplifier open loop gain is sufficiently small. This is achieved by sizing the feedback capacitor at $45\,\mathrm{pF}$ in the SDT mode, and at $0.9\,\mathrm{pF}$ in the DCal mode. Thus, the output swing remains within $0.82\,\mathrm{V}$ from the baseline.

The amplifier transistor-level design was optimized for power dissipation, and taking into account the bandwidth and noise performance constraints in both modes of operation. In the SDT mode the CSA settling goes through a slewing regime, limited by the folding transistor DC current. The non-linear settling behavior, however, does not affect the linearity since the CSA output is no further processed before being sampled and converted to digital. In the DCal mode, the CSA settling is linear, and therefore, linearity is not affected when the CSA output is processed by the filter.

SPICE simulations on the CSA show an open-loop gain of $76.5\,\mathrm{dB}$, a unity-gain bandwidth of $800\,\mathrm{MHz}$, and a closed-loop bandwidth of $19.1\,\mathrm{MHz}$ for the DCal mode.

The CSA feedback network, shown in Fig. 4, includes the two feedback capacitors (C_{Op} for the SDT mode and C_{Cal} for the DCal mode), two reset switches (M_{ROp} and M_{RCal}), and the mode-select switch (M_{OM}). The feedback network has three different functions: CSA operation, full reset, and slow reset-release. During CSA operation, the reset switches remain open. When a full reset is engaged, the

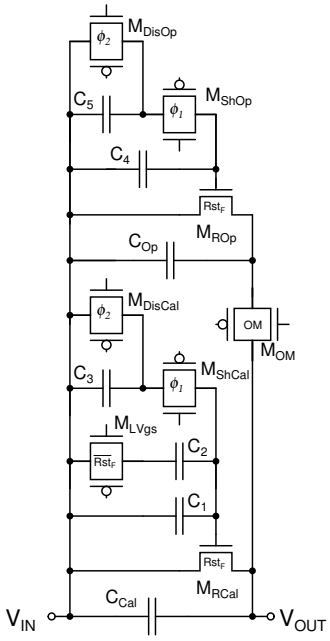

Fig. 4. Schematic of CSA feedback network.

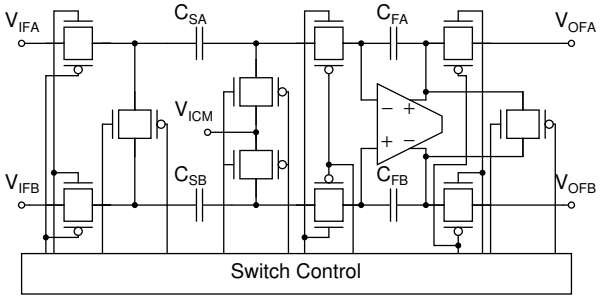

Fig. 5. Simplified schematic of the filter.

reset switches are directly driven by a reset control signal. During the reset-release operation, the gate-to-source voltage of the reset switches is gradually reduced in every clock cycle by switched-capacitors C_1 through C_5. The slow reset-release results in a reduced positive slope of the weighting function, which consequently mitigates the noise effects of a split doublet [2].

B. Filter

In order to produce the negative slope of the weighting function when in the DCal mode, a fully-differential SC integrator is used. The filter input corresponds to the buffered and level-shifted difference between the CSA output and a dummy CSA that produces an artificial, constant baseline voltage. Fig. 5 shows a simplified schematic of the filter, where some control logic and switches have been omitted.

The filter operates on a two-phase non-overlapping symmetric clock, having only $9.625\,\mathrm{ns}$ for settling. To conserve power, a class-AB operational transconductance amplifier (OTA) is used, capable of providing enough output current during the transients. A schematic of the two-stage OTA, based on the

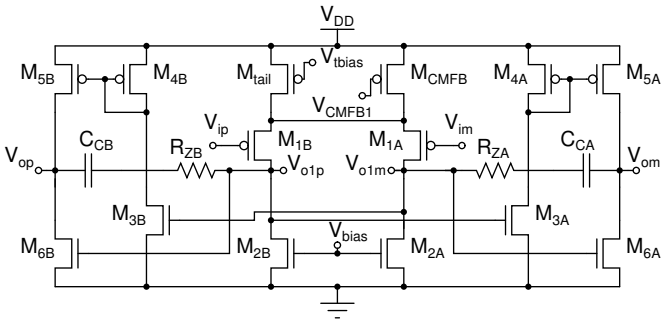

Fig. 6. Simplified schematic of the filter OTA. The common-mode feedback networks have been omitted.

topology presented in [5], is shown in Fig. 6. The first-stage quiescent current is set by the tail current source of the differential pair. The second-stage quiescent current is set by the common-mode output voltage of the first stage, a reference for which is established by a diode-connected transistor biased by a reference current source (not shown). The OTA compensation network, C_C and R_Z, has been designed to ensure stability during reset, which represents the worst-case feedback factor scenario.

SPICE simulations predict an OTA differential open-loop gain over $70\,\mathrm{dB}$, a phase margin of $83°$, and a crossover frequency of $143.5\,\mathrm{MHz}$ when connected in the integrator configuration. The bandwidth allows more than seven time constants for settling in each clock phase, reducing the settling error to less than $1\,\mathrm{LSB}$.

C. ADC

A 10-bit, fully-differential SAR ADC completes the signal path. Based on a charge-redistribution topology similar to that shown in [6], the ADC is designed to operate at a clock frequency equal to that of the SC filter, producing a conversion within 10 subcycles. The ADC uses two switched-capacitor DAC arrays, as shown in the simplified ADC schematic of Fig. 7, where the SAR logic has been omitted. In order to reduce the converter input capacitance, the capacitor array employs the smallest metal-insulator-metal (MIM) capacitors that can be fabricated in the process, sized at $16\,\mathrm{fF}$. In order to prevent the small capacitors mismatches from compromising the converter differential nonlinearity (DNL), the five most significant bits (MSBs) of the array were thermometer-coded [7].

The ADC comparator, shown in Fig. 8, comprises a preamplifier that reduces offset, noise and possible metastability problems, followed by a latched comparator. The comparator operates on a two-phase clock. During the reset phase (En input low), transistors M_{RA} and M_{RB} pull the comparator outputs to V_{DD}. During the compare phase (En input high), the reset transistors are open and the positive feedback through M_{4A} and M_{4B} swings the output nodes toward different rails according to the differential input of the comparator.

An alternative implementation of the ADC architecture, using smaller lateral-field metal-oxide-metal (MOM) capacitors, has been also fabricated as a test structure for independent

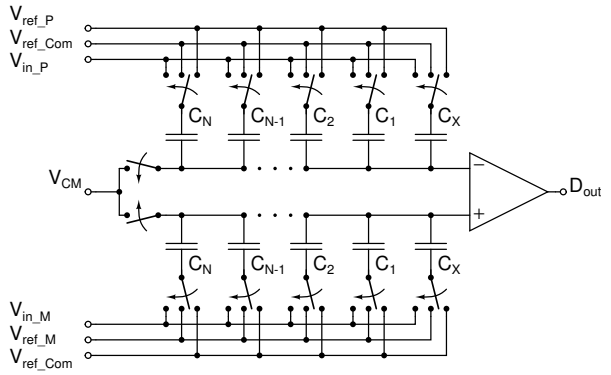

Fig. 7. Simplified schematic of the SAR ADC. The logic circuits have been omitted.

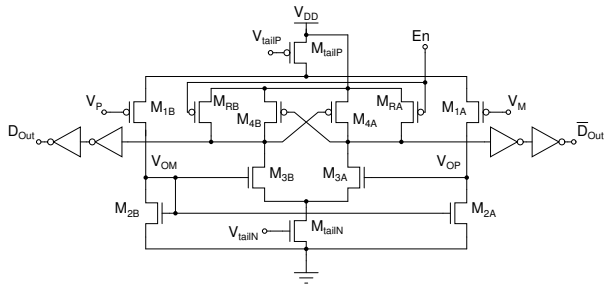

Fig. 8. ADC comparator schematic.

characterization. The unit MOM capacitor size was estimated at $2\,\mathrm{fF}$. This standalone ADC is being considered in future revisions of the Bean, due to its reduced input capacitance and feasibility of being integrated in a standard CMOS process.

D. Adder

As mentioned on Section II, the Bean must provide a dedicated low-latency output for beam diagnostics purposes. The output consists of the sum of the outputs of all channels. The addition is done in the sampled-data domain by using the fully-differential switched-capacitor adder shown in Fig. 9. The adder takes one subcycle to perform the addition, and its output is converted into digital using a dedicated ADC. The adder OTA is identical to that used in the filter. The series and feedback capacitors were sized to produce a gain of $1/3$ for each input. Thus, the adder full-scale output range is the same as the filter full-scale output range.

E. Signal buffer

Signal buffers are used at the outputs of the CSA and the CSA baseline generator to prevent the filter kickback noise from affecting the CSA operation. The buffer also shifts the baseline to produce a differential output spanning a more symmetric range. The buffer schematic, shown in Fig. 10, comprises a level shifter and a common-drain amplifier. The latter uses a cascoded current source and an auxiliary transistor M_2 to improve its linearity.

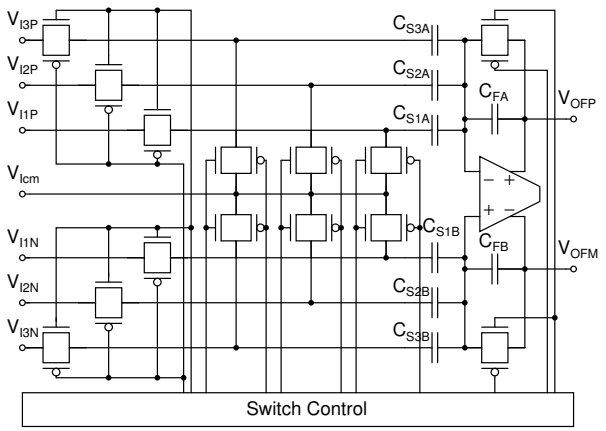

Fig. 9. Adder schematic.

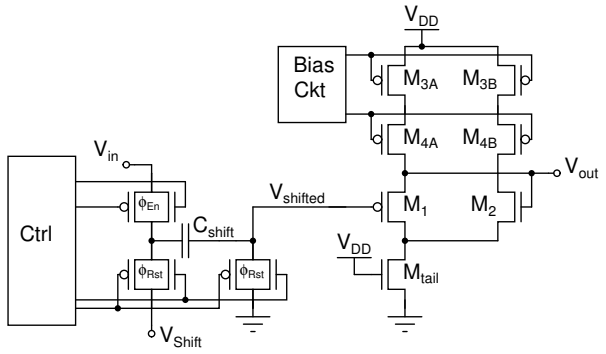

Fig. 10. Signal buffer schematic.

F. Rail-to-rail buffer

The rail-to-rail buffers in the Bean serve two purposes: they drive the MIM ADC input nodes and buffer some of the internal nodes allowing external probing. The buffer is based on a two-stage amplifier topology with rail-to-rail differential input and a Class-AB single-ended output stage [8], as shown in Fig. 11.

SPICE simulations show an open-loop DC gain of $101.6\,\mathrm{dB}$, a crossover frequency of $49\,\mathrm{MHz}$, and a phase margin of $80.3°$ measured with an 8-pF load.

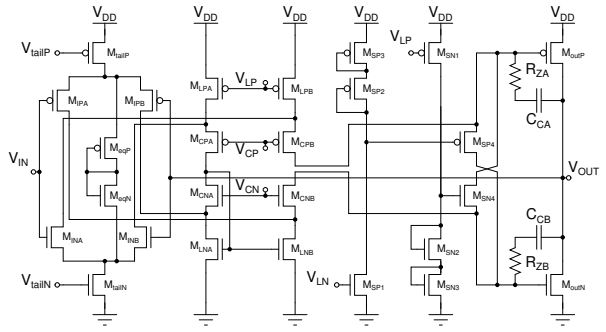

Fig. 11. Rail-to-rail buffer schematic.

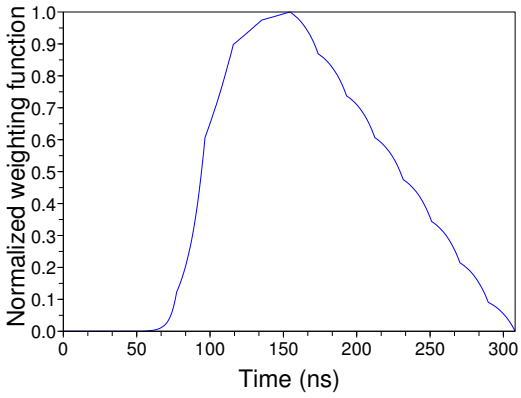

Fig. 12. Simulated weighting function in the DCal mode, for a switched-capacitor integrator and slow reset-release technique.

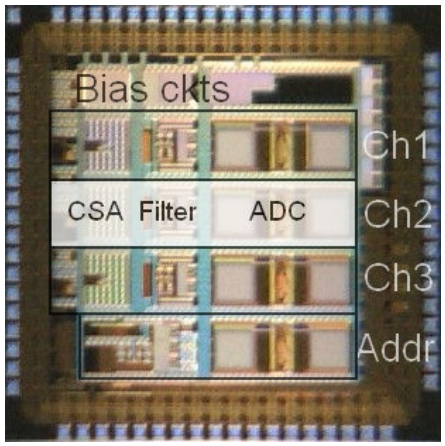

Fig. 13. Microphotography of the Bean prototype.

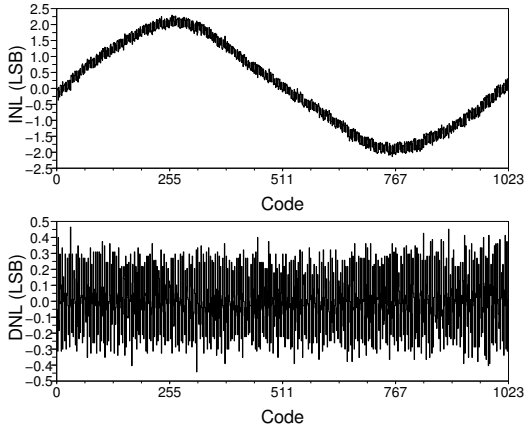

Fig. 14. MOM capacitor ADC linearity test results.

G. Weighting function simulation

The Bean weighting functions were obtained by means of SPICE simulations executed for the full channel. Each point of the weighting function $W(t_i)$ was obtained as the outcome of a transient simulation, measured at the ADC output, for an input occurring at time t_i. Fig. 12 shows the simulated weighting function in the DCal mode. The positive slope is due to the slow reset-release technique, whereas the negative slope is due to the SC integrator action, smoothed by the limited bandwidth of the CSA. The resulting shape is in agreement with the expectation of generating a quasi-triangular weighting function.

IV. Implementation

In the Bean prototype layout, special care has been taken to mitigate the effects of electrical and process-related non-idealities. Shielding, common centroid layout, extensive ground planes and separate supplies and grounds are among the measures taken.

Fig. 13 shows a Bean prototype die microphotograph. The Bean prototype measures $2.4\,\mathrm{mm} \times 2.4\,\mathrm{mm}$, has 72 pads, $23\,\mathrm{k}$ transistors (some parallel-connected) and $12\,\mathrm{k}$ capacitors. It has three channels, with a vertical pitch of $360\,\mu\mathrm{m}$ that includes the surrounding power bus. The layout was designed so that additional channels can be abutted to the existing layout, using the buses available and reducing the number of dedicated connections.

V. ADC test results

Both converters, the MIMCaps ADC with 16-fF unit capacitors and the MOMCaps ADC with 2-fF unit capacitors, were tested for integral nonlinearity (INL) and differential nonlinearity (DNL). The test setup consisted of a custom PCB with 16-bit DACs for differential input generation, and low-dropout linear voltage regulators with bypass capacitors for the power supply and DAC reference voltages. The DACs were driven by a Spartan3e FPGA on a Digilent Basys2 evaluation board, running at $50\,\mathrm{MHz}$. The FPGA was programmed to generate a ramp, to measure the output of the ADC to the input ramp, and to transfer the output to a PC via USB connection. On each run, $16\,384$ conversions were done, stored at the FPGA, and then transferred to the PC for further processing.

Fig. 14 shows the measured INL and DNL of MOMCaps ADC, whereas Fig. 15 shows the measured INL and DNL of the MIMCaps ADC. The MIMCaps ADC exhibits less nonlinearity than the MOMCaps ADC, which can be explained due to the relative capacitance mismatch. Both DNL plots show good uniformity and no missing codes. Both INL plots show non-ideal, cubic-like shapes that can be explained due to copper dishing effects on the capacitor arrays. In the MIMCaps ADC, the effect is less pronounced because only fringe capacitance is affected by copper dishing. The effect on the INL can be minimized by re-designing the interconnections of the array, which is currently insensitive only to linear (non-radial) gradients. From the measurements, the unit MOM capacitance mismatch was estimated to be 8.1%.

The ADC power consumption in both converters was measured to be $245\,\mu\mathrm{W}$ when operating at full speed, and without considering bias circuits.

VI. The Bean prototype test results

The Bean prototype was tested for functionality, linearity, cross-talk, bandwidth, weighting function, noise, and operation

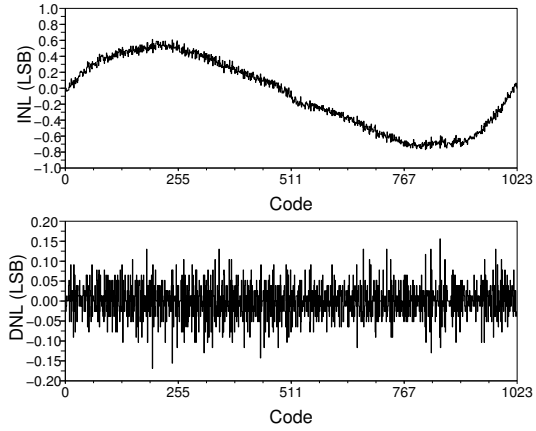

Fig. 15. MIM capacitor ADC linearity test results.

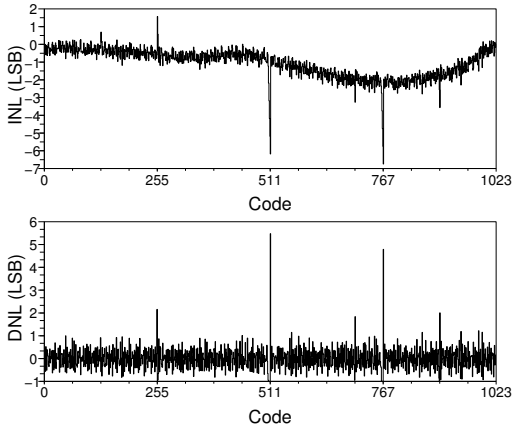

Fig. 16. The Bean prototype linearity test results, SDT mode.

of the fast feedback adder. The test setup consists of a custom PCB with voltage regulators, analog buffers, DACs for reference generation, and DAC-driven charge injectors for stimuli generation. The PCB was driven by an FPGA, programmed to produce a set of stimuli in all of the Bean channels, record 4096 values of the output of all channels, and transfer them to the PC for further processing.

A. Linearity

Fig. 16 shows the Bean prototype linearity test result in the SDT mode of operation, whereas Fig. 17 shows the linearity test result in the DCal mode. Both were obtained for input charges spanning the full-scale ranges. The nonlinearity visible in the top portion of the range is due to the finite CSA open-loop gain, and is within the specifications. The spikes in the INL, and the missing codes in the DNL, are an effect of the parasitic inductance in the long reference voltage traces that affect the ADC operation when the most significant bits transition. This effect can be mitigated by buffering the reference voltages internally.

B. Crosstalk

Crosstalk tests aim to determine the gain from the input of a channel, namely the aggressor channel, to the output of another

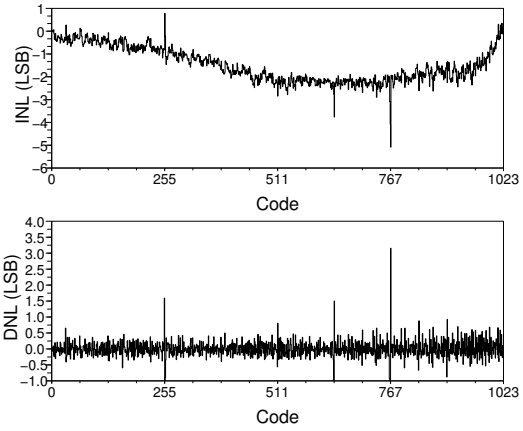

Fig. 17. The Bean prototype linearity test results, DCal mode.

channel, namely the victim channel. Crosstalk was measured in both modes of operation by ramping the aggressor channel input (channel 1) while maintaining the victim channels inputs constant (channels 2 and 3). In the SDT mode, the worst-case crosstalk gain was found to be 1.4% for both victim channels. In the DCal mode, the worst-case crosstalk gain was measured at 1.65% for the adjacent victim channel, and 1.3% for the non-adjacent victim channel. From the similarity of the measured crosstalk gains between adjacent and non-adjacent channels, it is concluded that most of the IC crosstalk in both modes of operation is the result of indirect channel-to-channel coupling, such as power supply and reference coupling. The crosstalk gains can be reduced in future revisions by increasing the number of power supply pins and buffering the references locally.

C. Bandwidth

Bandwidth tests intend to determine the residual effect of an input on the output of subsequent cycles. This is done by injecting an input charge and measuring the digital output on subsequent cycles. The results measured show no evidence of memory effect in either mode of operation. This result validates the effectiveness of the reset between cycles.

D. Weighting function

The weighting function of a pulse processor is like the signature of the signal path, since it reveals some details on how the signal path processes the input pulses. Calculations over the weighting function allow computation of the transfer functions for the input-referred series and parallel noise components.

The signal path was thoroughly simulated using SPICE. The weighting function results are presented in Fig. 18, for eight different reset-release schemes. The simulated weighting functions exhibit a low derivative in the negative slope, which supports the effectiveness of switched-capacitor filters in reducing the series noise component. The plot also shows that the slow reset-release schemes reduce the series noise, when compared to a fast reset-release condition, since the slow reset-release schemes exhibit a lower slope. Even though the

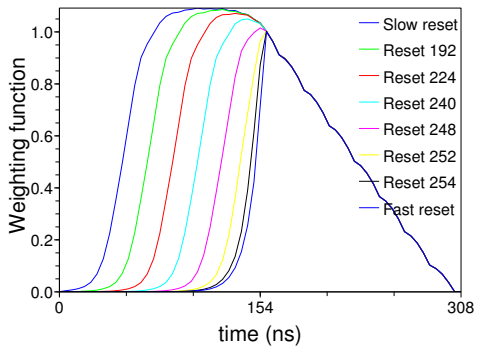

Fig. 18. The Bean prototype simulated weighting functions for different reset schemes, DCal mode.

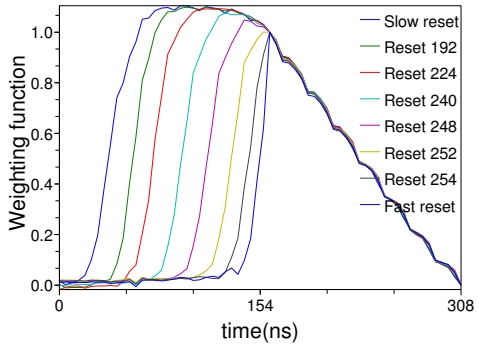

Fig. 19. The Bean prototype measured weighting functions for different reset schemes, DCal mode.

noise reduction is less than that achieved by correlated double sampling (CDS), the slow reset-release represents a simple alternative to reduce the effect of noise due to split doublets.

Fig. 19 shows the same series of weighting functions presented earlier, now measured experimentally. Series noise coefficients were computed for different weighting functions, and show that the combination of filter and slow reset-release scheme attenuate the signal path noise by 39%. The measurements are in agreement with the simulations, and validate the signal path design for series noise reduction.

E. Noise

Noise was measured in both modes of operation using the histogram method [9]. In the SDT mode, the RMS noise has an average of 0.6 LSB, which is within the specifications, and can be explained mostly due to the ADC noise. In the DCal mode, the RMS noise measurements were scaled to take into account the additional noise measured due to the input capacitance in the test setup. The estimated RMS noise in the DCal mode is 1.41 LSB. Further analysis and validating tests showed that the filter OTA is responsible for 1.35 LSB of the RMS noise in the DCal mode. This noise contribution can be reduced by means of an improved OTA design.

Noise measurements show that the combination of SC filter and slow reset-release scheme effectively reduces the series noise contribution. This result supports the use of SC filters in front-ends for particle physics experiments.

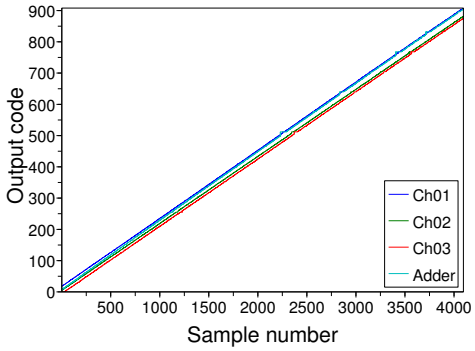

Fig. 20. Adder test results.

F. Adder

The adder operation was tested by injecting known input signals in the three input channels and measuring the adder output. Fig. 20 shows the outputs of the three channels and the adder, when the three channels inputs are ramps that span nearly the full-scale range. The discontinuity visible in the plots around codes 256, 512 and others is due to the effect of the inductance in the ADC reference nodes, and can be corrected by buffering the references internally.

In different tests, the adder gain from the different channels were measured to be 0.345, 0.344 and 0.329, close to the design value of 0.333. The adder digital output is available in less than 350 ns from the input pulse, and can be used for beam tuning and diagnostics purposes.

VII. Conclusion

An instrumentation ASIC prototype for particle physics experiments capable of processing the output of a detector with 100% occupancy has been successfully designed, integrated and tested. The Bean prototype demonstrates the use of modern CMOS techniques – switched-capacitor circuits – in the instrumentation of a particle physics experiment. Although switched-capacitor circuits produce fundamentally suboptimal results in pulse processors, due to a staircase-like weighting function, the time constant precision makes it an attractive option for instrumentation ICs designed for deep submicron processes.

References

[1] H. Abramowicz *et al.*, "Instrumentation of the very forward region of a linear collider detector," *IEEE Trans. Nucl. Sci.*, vol. 51, no. 6, pp. 2983–2989, 2004.

[2] E. Gatti and P. Manfredi, "Processing the signals form solid-state detectors in elementary particle physics," *Nuovo Cimento*, vol. 9, no. 1, pp. 2–145, 1986.

[3] G. De Geronimo and P. O'Connor, "MOSFET optimization in deep submicron technology for charge amplifiers," *IEEE Trans. Nucl. Sci.*, vol. 52, no. 6, pp. 3223–3232, 2005.

[4] W. Snoeys *et al.*, "Integrated Circuits for Particle Physics Experiments," *IEEE J. Solid-State Circuits*, vol. 35, no. 12, pp. 2018–2030, Dec. 2000.

[5] S. Rabii, "Design of Low-Voltage Low-Power Sigma-Delta Modulators," Ph.D. dissertation, Dept. of Elect. Eng., Stanford Univ., 1998.

[6] J. McCreary and P. Gray, "All-MOS charge redistribution analog-to-digital conversion techniques. I," *IEEE J. Solid-State Circuits*, vol. 10, no. 6, pp. 371–379, 1975.

[7] F. Kuttner, "A 1.2V 10b 20MSample/s non-binary successive approximation ADC in $0.13\ \mu$m CMOS," in *Dig. Tech. Papers, IEEE Int. Solid-State Circuits Conf., ISSCC, 2002*, vol. 1, pp. 176 –177 vol.1.

[8] R. Hogervorst *et al.*, "A compact power-efficient 3 V CMOS rail-to-rail input/output operational amplifier for VLSI cell libraries," *IEEE J. Solid-State Circuits*, vol. 29, no. 12, pp. 1505 –1513, Dec. 1994.

[9] S. Ruscak and L. Singer, "Using Histogram Techniques to Measure A/D Converter Noise," *Analog Dialogue*, vol. 29-2, 1995. [Online]. Available: http://www.analog.com/library/analogDialogue/archives/29-2/cnvtrnoise.html