31st Annual International Conference of the IEEE EMBS
Minneapolis, Minnesota, USA, September 2-6, 2009

Adaptive Resolution ADC Array for Neural Implant

Stephen O'Driscoll and Teresa H. Meng
Department of Electrical Engineering, Stanford University, Stanford, CA 94305 USA
stiofan@stanford.edu, thm@stanford.edu

***Abstract*—This paper describes an ADC array for an implantable prosthetic processor which digitizes neural signals sensed by a microelectrode array. The ADC array consists of 96 variable resolution ADC base cells. The base ADC has been implemented in $0.13\mu m$ CMOS as a $100kS/s$ SAR ADC whose resolution can be varied from 3 to 8-bits with corresponding power consumption of $0.23\mu W$ to $0.90\mu W$ achieving an ENOB of 7.8 at the 8-bit setting. The resolution of each ADC cell in the array is varied according to neural data content of the signal from the corresponding electrode. Resolution adaptation reduces power consumption by a factor of 2.3 whilst maintaining an effective 7.8-bit resolution across all channels.**

I. Introduction

Research groups have demonstrated that monkeys and humans can learn to move computer cursors and robotic arms to various target locations simply by activating neurons in the brain without actually performing the arm movements [1]. Clinically viable prosthesis will require much faster and more accurate movement commands which requires extracting as much information as possible from every neuron available. To this end an implantable prosthetic processor (IPP) that achieves high levels of prosthetic performance while minimizing power consumption, as shown in Fig. 1, has been proposed [2].

Neural signals are essentially voltage spike trains. The neural information is encoded in the variable timing between these spikes. Implanted electrodes are used to sense the neural signals. For a variety of clinical reasons electrodes cannot be oriented to sense individual neurons. Rather each electrode typically senses signals from more than one neuron. In order to extract the underlying information the signals sensed by each electrode must be decomposed into signals from individual neurons. To do this the signals are first converted to the digital domain and then "spike sorted". The sorted spikes are then decoded to identify intended movements. The IPP therefore comprises four major building blocks: a variable-resolution ADC array, a digital spike sorter [3], a maximum-likelihood neural decoder [4], and a wireless data and power transceiver. Low power is essential for both supply considerations and heat dissipation in the surrounding tissue.

This paper describes the adaptive resolution ADC array and the ADC cell including a CMOS implementation with measured results.

II. Adaptive Resolution ADC Array

ADC power consumption increases strongly with resolution, for example the capacitance and thus power consumption of the sub-DAC in a charge redistribution SAR ADC are $\propto 2^N$, where N is the number of output bits. In the IPP application the

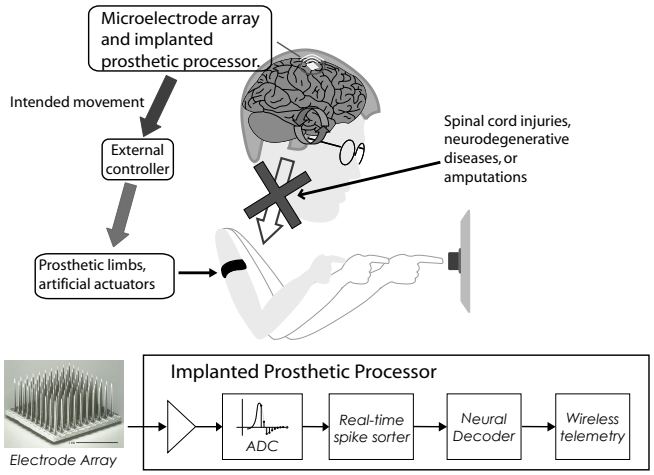

Fig. 1. Neuroprosthetic application and block diagram of implantable prosthesis processor.

voltages sensed by each electrode can contain signals or spikes from multiple neurons, exhibit different signal-to-noise ratios and relative signal strengths. Therefore the resolution required in digitizing the signal from each electrode is not uniform across all electrodes and energy can be saved by setting each ADC resolution no greater than required for that electrode.

Fig. 2 illustrates signal characteristics which are important to understand in order to develop methods which determine the required resolution. Fig. 2(a) shows a typical neural spike. In applications where only one signal source contributes to the sensed signal, SNR and dynamic range are sufficient to determine the required ADC resolution. Frequently as many as 6 neurons may contribute spikes to the signal sensed by a single electrode and the subsequent signal processing must be able to differentiate the spikes from each neuron. Fig. 2(b) shows the overlay of spikes from three different neurons sensed by a single electrode. This suggests that in order to be able to allow differentiation of these spikes, ADC resolution should be determined by considering relative signal amplitudes or signal-to-signal ratios (SSR) in addition to signal-to-noise ratios. If neural spikes were sorted solely according to signal amplitude then SSRs and SNRs would be sufficient to determine the required ADC resolution. Fig. 2(c) shows an overlay of spikes from two neurons which are very similar in amplitude, recorded by the same electrode. Required resolution estimation methods based solely on SNR and/or SSR would assume these spikes came from the same source,

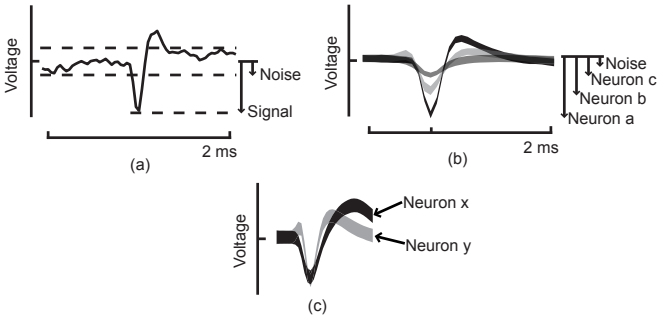

Fig. 2. (a) Signal and noise on a single spike signal (after HPF) (b) Multiple spikes with different magnitudes (c) Spikes with similar magnitude but different shape.

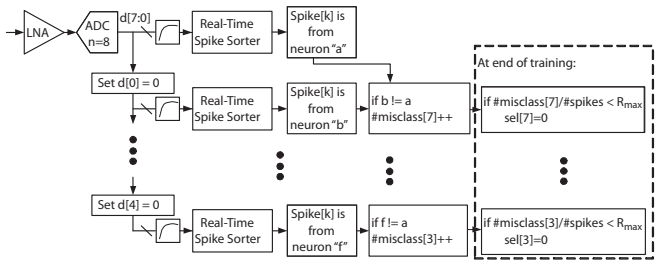

Fig. 3. Resolution Determination and Assignment Method.

so that a very low resolution would be deemed adequate and the subsequent spike sorting would be unable to distinguish these spikes. In order to allow differentiation of these spikes the resolution assignment criteria must be equivalent to the sorting criteria in the downstream signal processing. This is accomplished by using feedback from the real-time spike sorter to determine the resolution of each ADC cell.

The spike sorter classifies each spike as originating from a specific neuron. Let n_i be the resolution for the i^{th} ADC cell. The required ADC resolutions are estimated as shown in Fig. 3. Misclassification rates, $R(n_i) = \frac{\#\ misclassifications(n_i)}{\#\ spikes}$, are calculated for each possible resolution. The lowest n_i for which $R(n_i) < R_{max}$ is chosen to be the ADC's resolution until its next training cycle. Resolution estimation uses 5 extra spike sorts for each electrode in a time window equal to the training period of the spike sorter, which is 120 seconds every 12 hours. For a spike sorter power consumption of $1.4\mu W$ per channel [3], this gives a power overhead for resolution estimation of $20nW$ per channel, or 5% of the ADC power with optimally assigned resolutions. This method does not interrupt the throughput as 8-bit data is available throughout the calibration phase.

Fig. 4 shows an example resolution assignment for typical neural data with $R_{max} = 1\%$. The spike sorter itself has a spike misclassification rate of approximately 5% based on measured neural data, so choosing $R_{max} = 1\%$ does not significantly compromise performance.

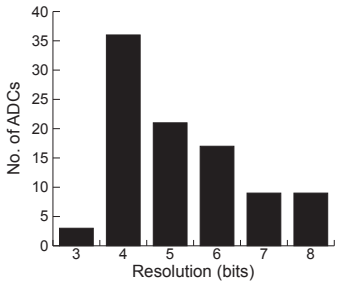

Fig. 4. Histogram of chosen ADC resolutions over 96 channels of synthetic neural data.

III. Circuit Design

The operation and design of fixed resolution charge redistribution ADCs has been discussed in the literature [5], and is modified to give a variable resolution converter as shown in Fig. 5. The successive approximation register (SAR) of [6] is used with the addition of ten 2-bit muxes and some switches to power down stages at lower resolution settings. Logic is implemented using I/O transistors to minimize leakage.

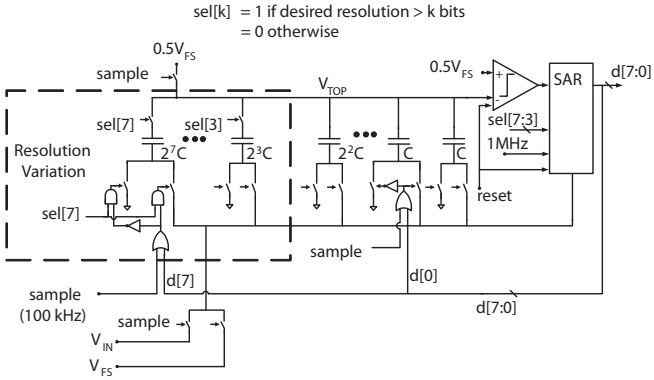

Fig. 5. Variable Resolution SAR ADC cell.

A. Variable Resolution Capacitor Array

In Fig. 5 the larger capacitors are switched out for lower resolution. Switches are placed on both sides of the capacitors. The parasitic capacitances at both the top and bottom plates are appreciable and significant power would be dissipated in charging and discharging them without the double switches degrading the efficiency at lower resolutions.

A $20fF$ MiM cap was chosen as the unit capacitance, C in Fig. 5, as it is low enough to achieve the required resolution and minimize power consumption whilst high enough to guard against process variations with some margin. The unit capacitors are $4\mu m \times 4\mu m$ in area. MiM caps were used as they are inherently linear. In the variable resolution structure the top plates of the capacitors are not connected to a single node; they must each connect to MOS devices. Technology layout constraints on interconnect and vias close to MiM caps and the use of top and bottom plate switches prevent use of a single array capacitance as is conventionally used in SAR ADCs.

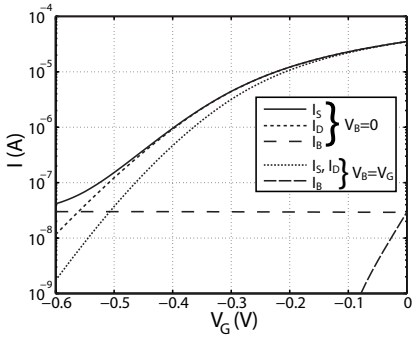

Fig. 6. Leakage Currents of min size NMOS; $V_D = 0.6V,\ V_S = -0.6V$.

Instead some separation is required between each capacitor. Nevertheless a common centroid structure is achieved for the entire array and each capacitor is made from unit capacitors surrounded by a ring of dummy units.

B. Reference Switch and Top Plate Switches

The switch connecting $0.5V_{FS}$ to V_{Top} is the *reference switch*. The voltage of the top plate of the array capacitance at the end of the hold phase is given by

$$V_{Top} = 0.5V_{FS} - V_{IN} - \frac{I_{Leak}t_{Leak}}{2^N C} \qquad (1)$$

where I_{Leak} is the leakage current from the top plate capacitance and t_{Leak} is the time between the hold phase and the first switch of the redistribution phase. Assuming that V_{FS}=1.2V, the top plate voltage varies from $0.6V$ to $-0.6V$ as V_{IN} varies from $0V$ to $1.2V$. If the reference switch is realized using a pass gate switch in which the applied gate and substrate voltages are $0V$ or V_{DD} then for V_{IN}=1.2V, $V_{GS,NMOS}$=0.6V during *hold* mode, resulting in substanial leakage to the top plate, i.e. I_{Leak}=$-35\mu A$ as shown in Fig. 6. This causes large signal-dependent error. Therefore the NMOS gate is driven with $V_{gate,Low} \leq -0.6V$ to turn the switch off. I_{Leak} then becomes $-41nA$, dominated by leakage across the $P\text{-}bulk/N^+$ diode which has a forward bias of $0.6V$. Given that $t_{Leak} \approx 0.75\mu s$, the resulting error in the top plate voltage is therefore given by

$$\Delta V_{Top} = \frac{I_{Leak}t_{Leak}}{2^N C} \geq 6.0mV = 1.28LSB \qquad (2)$$

which is still an order of magnitude or more too large. I_{Leak} is reduced to $-1.7nA$ by using deep n-well isolation and tying the bulk to $\leq -0.6V$, resulting in a satisfactory $\Delta V_{Top} = 0.25mV = 0.05LSB$.

The bulks of the NMOS sides of the top plate switches are tied to $\leq -0.6V$ to prevent leakage to the substrate when the sel_i for that switch is low. Since the drain of that NMOS is then floating there is negligible drain to source leakage and no need to drive that gate with $-0.6V$.

Fig. 7 shows a simple single stage charge pump used to generate a voltage of $|V_{TP}| - V_{DD} \approx -0.7V$ [7]. The generated voltage is applied to the bulk of the NMOS transistors in the top plate switches and the reference switch. The gate

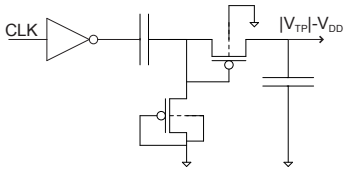

Fig. 7. Charge pump to generate negative voltage.

voltage for the NMOS in the reference switch is generated by passing the *sample* signal through a pair of inverters, the second of which has $V_{DD} = 1.2V$ and V_{SS} connected to the charge pump output.

C. Comparator

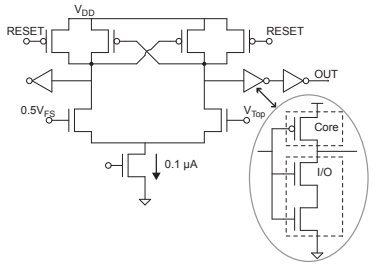

Fig. 8. Comparator schematic

The comparator is realized using a resettable latch baised to weak inversion, as shown in Fig. 8. Low comparator power requires low load capacitance, which is dominated by the comparator's intrinsic capacitance. Therefore minimizing comparator power dissipation demands low WL. Conversely, low V_T mismatch requires high WL. To reduce power consumption the latch itself is not pushed to give full output swing, rather it swings from 0.8 to 1.2V and the output is buffered. The first buffer uses a single core PMOS and a stack of two I/O NMOS to move the switching point close to the midrange of the signal swing, thus minimizing leakage current.

IV. Measurement Results

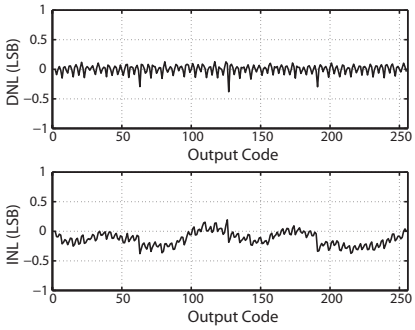

Fig. 9. Measured differential and integral non-linearties versus output code at 8-bit resolution setting.

Fig. 9 shows measured INL and DNL at 8 bit resolution. Both are well within $\pm 0.5LSB$. Jumps in the plots are visible

at codes 63, 127 and 191 corresponding to mismatch in the capacitance of the routing to the 2^6C and 2^7C capacitances.

Fig. 10 shows an FFT of measured digital output at 8-bit resolution for a 1 kHz sinusoidal input based on 8192 samples. The 3^{rd}, 7^{th}, and 9^{th} harmonics are clearly visible. The corresponding ADC performance metrics are ENOB=7.8 bits, SNDR=48.6dB, SFDR=60.0dB, and THD=-56.5dB.

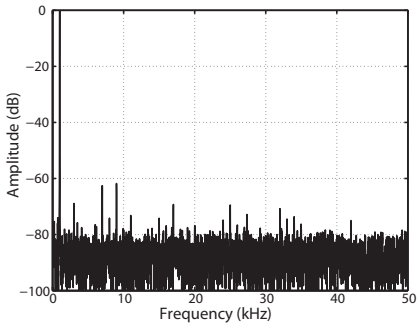

Fig. 10. FFT of measured digital output for 1 kHz input at 8-bit setting.

The standard figure of merit is given by Eq. (3) below. Our 40 fJ/conv-step is better than previous work [8], [9].

$$FOM = \frac{Power}{2^{ENOB} f_s} = 40fJ/conversion\ step. \quad (3)$$

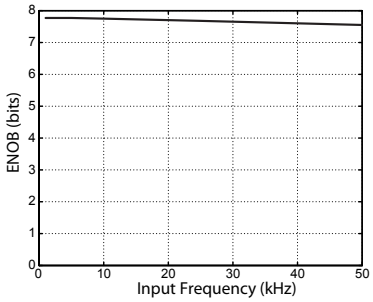

Fig. 11. Effective number of bits versus frequency at 8-bit setting.

Fig. 12 shows that the ADC cell power consumption increases strongly with resolution. At low resolutions comparator power dominates. If the resolutions of 96 channels are assigned according to Fig. 4 then the total power dissipation is $36.9\mu W$. Future work could implement a comparator whose power consumption scales with resolution which would increase the energy savings gained by resolution adaption further.

V. Conclusion

The ADC cell occupies $0.07mm^2$ in $0.13\mu m$ CMOS. Power consumption ranges from $0.23\mu W$ at 3-bit to $0.90\mu W$ at 8-bit resolution. The sampling rate is 100kS/s. The 96 element ADC array dissipates $37\mu W$ for a typical neural data set corresponding to a power saving of 2.3 $\times$ due to adaptive resolution. This extremely low level of power consumption was attained through integrating signal analysis, system architecture, and thorough analog design.

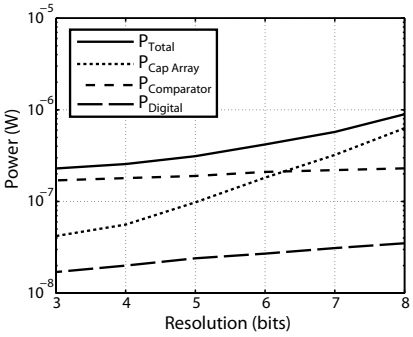

Fig. 12. Power consumption of base ADC cell versus resolution.

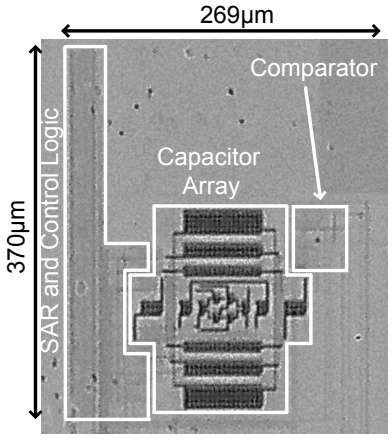

Fig. 13. Micrograph of ADC cell on die, area=$0.07mm^2$.

Acknowledgment

Stephen O'Driscoll is supported by the Lu Stanford Graduate Fellowship. The authors acknowledge the support of the Focus Center for Circuit & System Solutions (C2S2), one of five research centers funded under the Focus Center Research Program, a Semiconductor Research Corporation Program, and the generous fabrication support of Taiwan Semiconductor Manufacturing Corporation.

References

[1] G. Santhanam et al. A high performance neurally-controlled cursor positioning system. *Proc. of IEEE EMBS ICNE*, pages 494–500, 2005.

[2] S. O'Driscoll et al. Neurons to silicon: Implantable prosthesis processor. *Proc. of ISSCC*, pages 552–553, 2006.

[3] Z. Zumsteg et al. Power feasibility of implantable digital sorting circuits for neural prosthetic systems. *IEEE Trans. on NSRE*, 13:272–279, 2005.

[4] B. Yu et al. Improving neural prosthetic system performance by combining plan and peri-movement activity. *Proc. of IEEE EMBC*, pages 4516–4519, 2004.

[5] J. McCreary and P. Gray. All MOS charge-redistribution ADC techniques. *IEEE Journal of Solid State Ciruits*, 10:371–379, 1975.

[6] T. O. Anderson. Optimum control logic for successive approximation ADCs. *Computer Design*, 11:81–84, 1972.

[7] R. Baker et al. *CMOS Circuit Design, Layout, and Simulation*. Wiley-IEEE Press, 1998.

[8] M. Scott et al. An ultralow-energy ADC for smart dust. *IEEE JSSC*, 38:1123–1129, 2003.

[9] Y. K. Chang et al. A 8-bit 500-ks/s low power sar adc for bio-medical applications. *Proc. ASSCC*, pages 228–231, 2007.