18.4 A Matrix-Multiplying ADC Implementing a Machine-Learning Classifier Directly with Data Conversion

Jintao Zhang, Zhuo Wang, Naveen Verma

Princeton University, Princeton, NJ

Embedded sensing systems conventionally perform A-to-D conversion followed by signal analysis. In many applications, the analysis of interest is inference (e.g., classification), but the sensor signals involved are too complex to model analytically. Machine learning is gaining prominence because it enables data-driven training of classifiers, overcoming the need for analytical models. This work presents: 1) an algorithmic formulation, where feature extraction and classification are combined into a single matrix, reducing the total multiplications needed, and 2) a matrix-multiplying ADC (MMADC) that enables multiplication of input samples by a programmable matrix. Thus, the MMADC combines feature extraction and classification with data conversion, mitigating the need for further computations. Two systems are demonstrated: an ECG-based cardiac-arrhythmia detector and an image-pixel-based gender detector.

In general, analog computation faces two key challenges: 1) energy scales severely with dynamic range, making multiplication particularly problematic; and 2) the ability to implement specific computations is limited by both the preciseness and the particular transfer functions achievable by circuits. Dynamic range in the MMADC is addressed by exploiting mixed digital and analog multiplication, described below. The need to implement precise transfer functions is overcome by employing a machine-learning algorithm, called Adaptive Boosting (AdaBoost) [1]. AdaBoost uses an ensemble of *weak* classifiers to form a *strong* classifier (in machine learning, a weak classifier is one that cannot be trained to fit arbitrary data, a strong classifier is one that can). AdaBoost has the benefit that very weak classifiers can be used. This opens the possibility of classifiers preferred in terms of circuit implementation (energy and hardware). The MMADC focuses on linear classifiers, which involve dot product (element-wise multiplication and accumulation) between a feature vector $\vec{u}$ (derived from *N*-sample input data $\vec{x}$) and a classifier vector $\vec{c}_i$. Algorithmically, this has the benefit shown in Fig. 18.4.1. Conventional strong classifiers apply non-linear operations, requiring $\vec{u}$ to be computed explicitly. However, if feature extraction is linear, as is often the case [e.g., the systems demonstrated use discrete wavelet transform (DWT) and principal component analysis (PCA)], then an ensemble of *K* linear classifier vectors $\vec{c}_1 \ldots \vec{c}_K$ can be combined with the feature-extraction matrix *F*. This gives a single *K*×*N* matrix *H*, not requiring explicit computation of $\vec{u}$. As we show, the number of weak classifiers *K* can be low, reducing the number of multiplications needed and making the system energy, when applying *H* via the MMADC, of same order as an ADC.

Figure 18.4.2 shows the MMADC, which implements multiplication in an 8b SAR ADC at negligible additional energy. To address analog dynamic-range limitations, the quantized multiplier M_{QUANT} is broken into three parts (provided through *MULT*): a sign bit M_S, a 4b unsigned analog multiplier M_A, and a 5b signed exponent (i.e., $M_{QUANT} = M_S \times M_A \times 2^{M_D}$). M_A is applied passively through attenuation in the CapDAC feedback path via the FB-Divider block. The FB-Divider implements feedback attenuation β from 1 to 16/(16+15)=16/31, corresponding to multiplication by values from 1 to 31/16 (in steps of 1/16). The ADC supports signed inputs with respect to a nominally mid-rail bias V_{MID}. Signed multiplication is thus supported by using M_S to pass through or negate the ADC code. Digital multiplication, based on the exponent M_D, is then performed on the resulting 8b code via barrel shifting, implementing multiplication by powers of 2 (2^{-16}, 2^{-15},…2^{14}, 2^{15}). With 5b M_D, shifting by up to 32b is supported, giving a 40b final output.

Figure 18.4.3 expresses the overall quantized multiplier value M_{QUANT} (with respect to an ideal, unquantized multiplier M_{IDEAL}). The proposed approach has two benefits. First, regardless of the overall multiplier, the analog input range is reduced by only the maximum value of M_A, namely a factor of 31/16≈2 (giving analog input range 0.35-0.85V). Second, through M_D, arbitrarily large multiplier values can be supported at the low hardware cost of additional bits in the barrel shifter. The resulting multiplication has quantization noise that scales with 2^{M_D} (i.e., with the multiplier value). As seen in the simulation of Fig. 18.4.3, traditional linear-quantized multiplication (shown for the 8-to-13b levels) causes the output SQNR to vary widely with the multiplier, resulting in inefficient use of dynamic range. This particularly degrades performance for lower-valued multipliers, which often occur with equal or substantially higher frequency in signal-processing applications (such as the two considered). On the other hand, the MMADC results in relatively constant SQNR over an arbitrarily large range of multipliers, achieving performance substantially greater than the 8b level, yet with passive 4b multiplication hardware (FB-Divider) and barrel shifting.

To implement matrix multiplication, an MMADC input sample must be multiplied by the *K* elements in each column of *H*. To realize this, input sampling is performed on a separate capacitor C_{SAMP}, rather than within the CapDAC (whose charge must be reset following each conversion). Thus, as shown in the waveforms of Fig. 18.4.2, an input sampled once (by *SAMP*) can be applied to the MMADC multiple times (enabled by *CNVRT*) with different multiplier values (M_S, M_A, and M_D). Row-wise accumulation of products is then performed by a digital adder (not implemented on the chip). Since the sampled input must be held on C_{samp} (~1pF) for the duration of *K* conversions, leakage is a concern. The sampling switch consists of a series of two CMOS transmission gates (as shown in Fig. 18.4.2), exploiting the stacked effect to reduce leakage [2]. For the analog input range, simulated sample error <0.5LSB is achieved across all values of M_A for *K* up to 50.

Figure 18.4.4 shows the CapDAC and FB-Divider. To enable bipolar inputs (and signed multiplication), both the output of the FB-Divider and the output of the CapDAC are switched to V_{MID} at the start of a conversion (i.e., *CNVRT*=0). At the same time, feedback attenuation β is set by C_{MULT}, by switching in composing unit capacitors. Nominally, the unit capacitors have a value C_M set to $C_{DAC}||C_C$. Using, $C_C \approx 380fF$, a C_M of modest value is required (~25fF). Regarding parasitic capacitances, thanks to switching to V_{MID}, $C_{P,DAC}$ does not degrade ADC linearity. $C_{P,MULT}$ leads to a constant term in the denominator of β, causing error in the analog multiplier (measured below). Further, separating input sampling from the CapDAC causes the voltage of critical bit decisions to vary with input, leading to non-linearity (measured below) due to variation of sampling-switch charge injection and comparator offset. However, the errors can be partially overcome by the AdaBoost algorithm, where iterative training of weak classifiers enables errors in all *i-1* iterations to be compensated during training of the i^{th} classification vector $\vec{c}_i$ [3].

The MMADC is prototyped in 130nm CMOS (photo in Fig. 18.4.7). Figure 18.4.5 shows the parameters and measured performance of two systems implemented using the MMADC. Patient ECG from [4] and image-pixel data from [5] are replayed by an arbitrary waveform generator and presented to the MMADC along with multiplier values for the matrix *H*, which is formed from the feature-extraction matrix *F* (DWT and PCA, respectively) and the classification vectors $\vec{c}_1 \ldots \vec{c}_K$, derived from off-line training. As is conventional, 10-fold validation is performed to divide the dataset for training and testing. The performance is also shown for a MATLAB-implemented support-vector machine (SVM) with a radial-basis function kernel (a widely-used non-linear strong classifier). The MMADC successfully achieves strong classification performance near an ideal SVM, with convergence in 5 and 1 weak classifiers, respectively. Comparing the MMADC implementations to conventional systems (as in Fig. 18.4.1), the total energy for the two applications (including data-conversion, multiplications, and additions) is thus reduced by 9.7× and 23×, respectively, for the parameters involved (*N*, *J*, *K*, number of support vectors).

Figure 18.4.6 summarizes the measurements. Conversions are performed at a rate up to 20kS/s (scalable down at constant energy). Slight deviation of the multiplier value is measured (~9%) due to $C_{P,MULT}$. The SNDR is measured to be 45.8dB (low freq.) and 38.9dB (near Nyquist), giving a FoM down to 209fJ/conv. step.

Acknowledgements:
Funding provided by SRC, NSF (CCF-1253670), and Systems on Nanoscale Information fabriCs (SONIC), one of six SRC STARnet centers, sponsored by MARCO and DARPA. IC fabrication is provided by MOSIS.

References:

[1] R. E. Schapire and F. Yoav, "Boosting: Foundations and Algorithms," *MIT Press*, 2012.

[2] S. Narendra, *et al.*, "Scaling of Stack Effect and its Application for Leakage Rreduction," *IEEE/ACM International Symp. on Low-Power Electronics and Design*, pp. 195-200, 2001.

[3] Z. Wang, *et al.*, "Error-Adaptive Classifier Boosting (EACB): Exploiting Data-Driven Training for Highly Fault-Tolerant Hardware," *IEEE International Conf. on Acoustics, Speech and Signal Proc.*, pp. 3884-3888, 2014.

[4] G. B. Moody and R. G. Mark, "The Impact of the MIT-BIH Arrhythmia Database," *IEEE Engineering in Medicine and Biology*, vol. 20, no. 3, pp. 45-50, 2001.

[5] J. P. Phillips, et al., "The FERET Evaluation Methodology for Face-Recognition Algorithms," *IEEE Trans. On Pattern Analysis and Machine Intelligence*, vol. 16, no. 5, pp. 295-306, 1998.

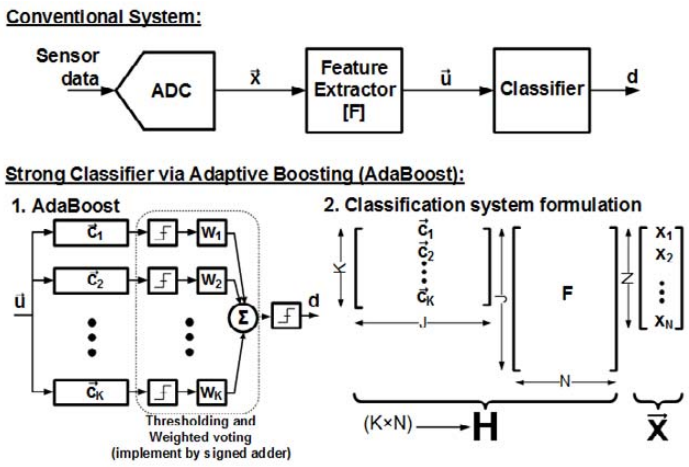

Figure 18.4.1: Compared to conventional system (top), AdaBoost enables ensemble of K linear classifier vectors $\vec{c}_i$ which can be combined with linear feature extraction F in the matrix H.

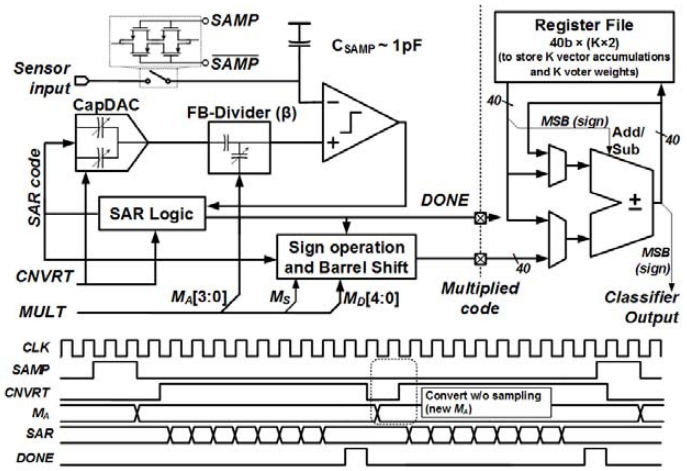

Figure 18.4.2: MMADC implements mixed analog/digital multiplication, with analog multiplication via feedback attenuation by FB-Divider and digital multiplication via barrel shifting.

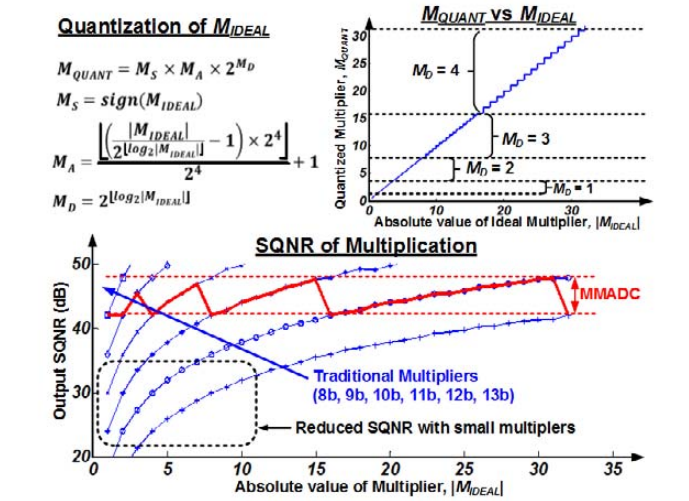

Figure 18.4.3: MMADC results in multiplier quantization error that scales with 2^{M_D} (and thus M_{QUANT}), giving relatively constant output SQNR over arbitrarily large range of multipliers.

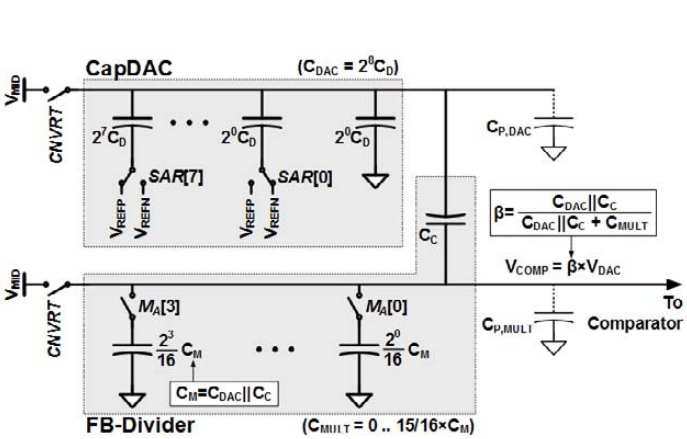

Figure 18.4.4: SAR feedback path, composed of CapDAC and FB-divider, wherein FB-Divider applies switchable attenuation factor β through capacitive division implemented with C_{MULT}.

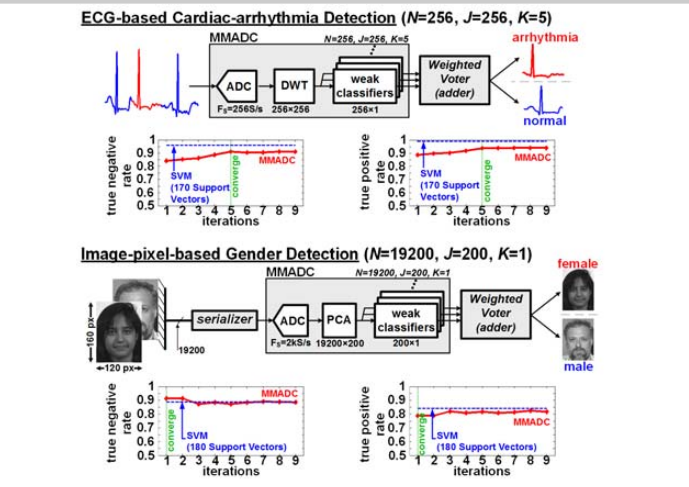

Figure 18.4.5: Two MMADC systems achieve classification performance near that of an SVM with ~9.7× and ~23× energy reduction vs. conventional systems (assuming 3.8pJ for 13b digital multiplication and 0.95pJ for 40b digital addition in 130nm CMOS).

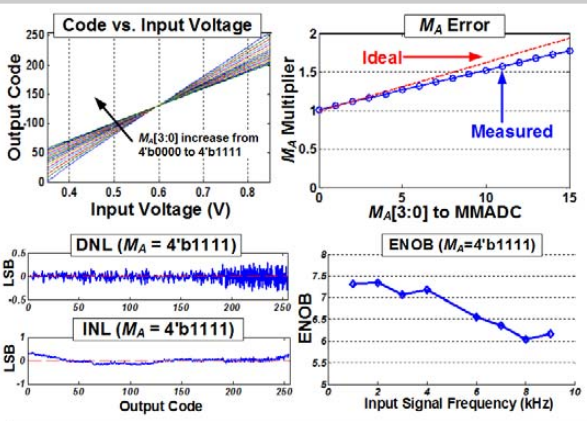

Technology	130 nm CMOS	Power @ 20 kHz	Digital: 376.8 nW Analog: 286.8 nW
Supply Voltage	1.2 V Analog 1.2 V Digital	DNL, INL	DNL:0.30 (LSB); INL:0.32 (LSB)
Resolution	8 bits	SNDR/ENOB	45.8dB/7.31b (w/ f_{in}=1kHz)
Input Range	350mV - 850mV		38.9dB/6.17b (w/ f_{in}=9kHz)
Sampling Rate	Up to 20 kHz	Achieved FOM	209fJ/conv. step (w/ f_{in}=1kHz)

Figure 18.4.6: Measurement summary of prototype MMADC.

18