Energy-Efficient Signal Processing via Algorithmic Noise-Tolerance *

Rajamohana Hegde and Naresh R. Shanbhag
ECE Department/Coordinated Science Laboratory
University of Illinois, Urbana, IL 61801.
email:[rhegde, shanbhag]@uivlsi.csl.uiuc.edu

Abstract

In this paper, we propose a framework for low-energy digital signal processing (DSP) where the supply voltage is scaled beyond the critical voltage required to match the critical path delay to the throughput. This deliberate introduction of input-dependent errors leads to degradation in the algorithmic performance, which is compensated for via *algorithmic noise-tolerance* (ANT) schemes. The resulting setup that comprises of the DSP architecture operating at sub-critical voltage and the error control scheme is referred to as *soft DSP*. It is shown that technology scaling renders the proposed scheme more effective as the delay penalty suffered due to voltage scaling reduces due to short channel effects. The effectiveness of the proposed scheme is also enhanced when arithmetic units with a higher "delay-imbalance" are employed. A prediction based error-control scheme is proposed to enhance the performance of the filtering algorithm in presence of errors due to *soft computations*. For a frequency selective filter, it is shown that the proposed scheme provides $60\% - 81\%$ reduction in energy dissipation for filter bandwidths up to 0.5π (where 2π corresponds to the sampling frequency f_s) *over that achieved via conventional voltage scaling*, with a maximum of $0.5dB$ degradation in the output signal-to-noise ratio (SNR_o). It is also shown that the proposed algorithmic noise-tolerance schemes can be used to improve the performance of DSP algorithms in presence of bit-error rates of upto 10^{-3} due to deep submicron (DSM) noise.

1 Introduction

Energy-efficient VLSI circuit design is of great interest given the proliferation of mobile computing devices, the need to reduce packaging cost, the desire to improve reliability, and extend operational life of VLSI systems. Scaling of CMOS technology has made possible substantial reduction in energy dissipation and hence has lead to the proliferation of low cost VLSI systems with increasingly high levels of integration. At a given technology, reduction in energy dissipation has also been made possible due to energy-efficient design techniques at all possible levels of design hierarchy. Schemes at the lower levels of the design process such as the logic [1] and circuit levels [2] are usually application independent. At the algorithmic and architectural levels, features that are specific to a class of applications are exploited to develop application specific energy reduction techniques [3, 4]. Voltage scaling [3] is an effective means of achieving reduction in energy dissipation as a reduction in supply voltage by a factor K, reduces the dominant capacitive component of energy dissipation by a factor K^2 [5]. However, the extent of voltage scaling [5] is limited by the critical path delay of the architecture and the throughput requirements of the application.

A typical DSP system is designed in such a way that the critical path delay [10] T_{cp} (defined as the *worst case delay over all possible input patterns*), should be less than or equal to the sample period T_s, i.e., $T_{cp} \leq T_s$. Hence, given T_s, the DSP system is designed such that, at the rated supply voltage V_{dd}, the *delay condition*, $T_{cp} \leq T_s$, is satisfied. The relationship between V_{dd} and circuit delay τ_d is given by [10],

$$\tau_d = \frac{C_L V_{dd}}{\beta(V_{dd} - V_t)^\alpha}, \tag{1}$$

where C_L is the load capacitance, α is the velocity saturation index, β is the gate transconductance, and V_t is the device threshold voltage. We refer to the voltage at which $T_{cp} = T_s$ as the *critical supply voltage* $V_{dd-crit}$ of a given architecture. Note that violating the delay condition by reducing V_{dd} beyond $V_{dd-crit}$, i.e. setting

$$V_{dd} = K_v V_{dd-crit}, \tag{2}$$

where $0 < K_v < 1$, leads to erroneous output when the critical path is excited. In conventional voltage scaling, $V_{dd-crit}$ is seen as a lower bound on the supply voltage for a given architecture and throughput. Reduction in V_{dd} (without violating the delay condition) can be achieved by reducing the critical path delay of the VLSI implementation via architectural transformations such as pipelining and parallel processing [5]. We propose operating the DSP architecture at voltages lower than $V_{dd-crit}$. Such operation leads to errors in the system output when the critical paths and other longer paths are excited. Hence, the resulting computations are referred to as *soft computations*.

We propose *algorithmic noise-tolerance* (ANT) to compensate for degradation in the system output due to errors from soft computations. ANT refers to algorithmic error-control schemes derived from the knowledge of the system transfer function, input and output signal statistics. The resulting framework is illustrated in Figure 1. The setup that comprises of the DSP architecture operating at a sub-critical supply voltage (lower than $V_{dd-crit}$)

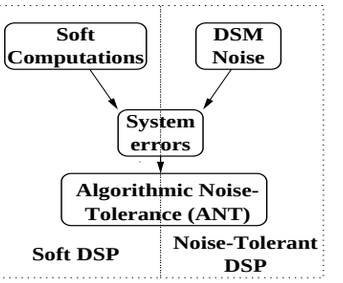

Figure 1: The proposed soft and noise tolerant DSP framework.

*This work was supported by NSF CAREER award MIP-9623737.

and the low-complexity error-control scheme is referred to as *soft DSP*. The goal of soft DSP is to achieve substantial energy-savings while meeting the algorithmic performance specifications. Note that the effectiveness of the proposed scheme depends on the error frequency. We show that the phenomenon of velocity saturation in short-channel devices (feature size is less than $0.5\mu m$) favours low-power operation via soft DSP. It is also shown that the ANT schemes can be employed to restore degradation in algorithmic performance due to deep submicron noise (DSM) noise [7] induced errors. We refer to the resulting setup as *noise-tolerant DSP* as shown in Figure 1.

The rest of this paper is organized as follows. In section 2, the proposed notion of soft DSP is introduced with a motivational example. In section 3, a low complexity prediction-based algorithm is developed to detect and mitigate the effect of soft errors on the performance of the digital filtering algorithm. In section 4, we study the energy savings due to the proposed approach in the context of frequency selective filtering. We also study the performance of algorithmic noise-tolerance schemes in presence of random errors in the system output due to DSM noise. Finally, in sections 5 conclusions and scope for future work on this topic are presented.

2 Energy Savings via Soft DSP

In this section, we illustrate the relationship between energy savings due to the proposed approach and the resulting degradation in performance due to errors in the system output. It is shown that error frequency due to soft computations is a function of the path delay distribution of the DSP block architecture and a new multiply-accumulate (MAC) architecture that improves the effectiveness of the proposed scheme for the filtering algorithm is presented.

2.1 Motivational Example

Consider the 5-bit adder shown in Figure 2(a), where the input operands are 00101 and 01011. Assuming that $T_{FA} = 3ns$, the critical path delay of this adder is $15ns$. Note that the time taken to compute the output corresponding to the two operands is also $15ns$. Let $T_s = 15ns$. If the supply voltage is now reduced such that $T_{FA} = 5ns$, the adder output at the end of the sample period will be 01000 as shown in Figure 2(a). Hence, the numerical value of the adder output will be 8 instead of 16. If the inputs do not excite the longer paths (e.g. 00001 and 00010), then the adder provides correct outputs.

In the absence of errors, the algorithmic performance of a filter transfer function $H(z)$ (shown in Figure 2(b)) is measured in terms of the output SNR given by

$$SNR_o = 20\log(\sigma_s/\sigma_w), \tag{3}$$

where σ_s^2 and σ_w^2 are the signal and noise variance, respectively. The output in this case can be expressed as

$$y(n) = s(n) + w(n), \tag{4}$$

where $s(n)$ is the desired signal and $w(n)$ is the signal noise. The filter output in presence of errors due to soft computations can be expressed as

$$\hat{y}(n) = y(n) + err(n), \tag{5}$$

where $err(n)$ is the error introduced in the output sample at the n^{th} instant. The output SNR using soft computations is given by

$$\hat{SNR}_o = 20\log(\sigma_s/\sigma_{w+c}), \tag{6}$$

where σ_{w+c} is the total noise power. Hence, errors in the system output lead to degradation in performance in terms of SNR_o. The

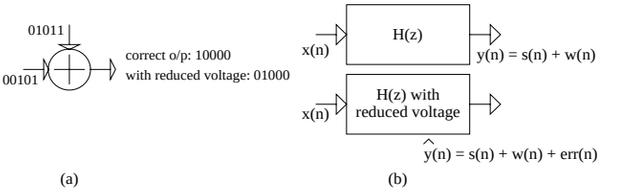

Figure 2: Effect of errors on performance of DSP algorithms

errors from to soft computations occur in the most significant bits (MSBs) due to longer path delays. This leads us to conclude that the error-control schemes should be effective in capturing MSB errors in order to result in substantial energy savings with marginal performance degradation.

2.2 Path Delay Distribution of Adders

In this subsection, we study the frequency of excitation of critical paths in the context of a ripple-carry adder. For an N-bit ripple-carry adder, the total number of possible input combinations is $2^N \times 2^N = 4^N$. Of these, some combinations such as $x = 01010101$ and $y = 10101010$ (N is assumed to be 8) are evaluated in just T_{FA} time units. Other combinations such as $x = 11111111$, $y = 00000001$, and $x = 11111111$, $y = 11111111$ excite the critical path requiring $8T_{FA}$ time units. The path delay histogram of an 8-bit two's complement ripple-carry adder where the input operands x and y are generated randomly with Gaussian distribution(shown in Figures 3(i) and (ii)), is shown in Figure 3(iii). Also shown in Figure 3(iii), is the histogram of path delays for an unsigned 8-bit adder for the same operands. A bias of $+128$ is added to both the operands to make them unsigned positive numbers. It can be seen that, for unsigned numbers, the fraction of inputs that excite the critical paths is significantly less. In this paper, we propose a new MAC architecture, that employs an unsigned array multiplier to improve the effectiveness of the proposed soft DSP scheme. ANT-based system.

2.3 MAC Architecture for soft DSP

The output $y(n)$ of an N-tap filter is given by

$$y(n) = \sum_{k=0}^{N-1} h(k)x(n-k), \tag{7}$$

where $h(k)$ denotes the k^{th} coefficient of the filter, $x(n-k)$ denotes the input at $n-k^{th}$ instant, and N is the order of the filter.

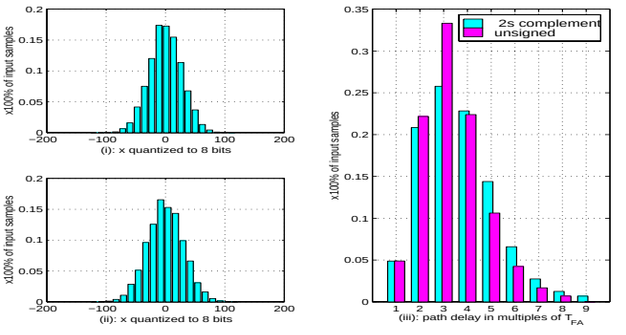

Figure 3: Path delay histograms of 8-bit ripple carry adder

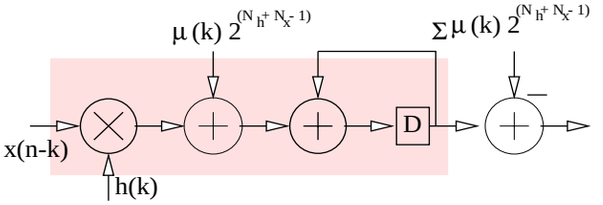

Figure 4: proposed MAC architecture

Typically, two's complement representation is used in representing the filter coefficients and the signal. However, as shown earlier, unsigned magnitude representation offers the advantage that a smaller fraction of inputs excite the critical path. Note that signed-magnitude representation has been employed in the past [3, 12], to reduce transition activity in correlators for wireless applications. The proposed MAC structure (refered to as the sign-magnitude architecture (SMA)), shown in Figure 4, employs signed magnitude representation and unsigned multiplier and adders. In this structure, the magnitude and sign of the product $h(k)x(n-k)$ are computed separately. If the product is negative, a bias term is added to make it positive before it is applied to the adder. Hence, we get

$$y'(n) = \sum_{k=0}^{N-1} \left(h(k)x(n-k) + \mu(k)2^{(N_h+N_x-1)} \right) \quad (8)$$

where,

$$\mu(k) = 0 \text{ if } h(k)x(n-k) \text{ is +ve,}$$
$$= 1 \text{ if } h(k)x(n-k) \text{ is -ve,}$$

and N_h and N_x are the number of bits in the representation of $h(k)$ and $x(n)$. An additional adder (operating at sample rate) is employed to subtract the bias term $\sum_{k=0}^{N-1} \mu(k)2^{(N_h+N_x-1)}$ from $y'(n)$ to obtain $y(n)$. Note that the multiplier in the SMA is smaller than that in the traditional structure due to signed magnitude representation. This leads to additional reduction in energy dissipation of the overall structure that compensates for the overhead of the additional adder.

3 Algorithmic Noise-Tolerance for Digital Filtering

In this section, we present an algorithmic error-control scheme for digital filtering in order to reduce the impact of errors on the algorithmic performance. The proposed scheme is shown in Figure 5, where the filter output is fed to an error-control block that detects errors in the filter output and reduces their effect on system performance. The output of the error control block is denoted by $y_o(n)$, and the goal of this approach is to obtain $y_o(n) \approx y(n)$, where $y(n)$ denotes the filter output in absence of errors. The term *noisy filter* represents a soft implementation of the digital filter or in presence of other noise inducing phenomena such as deep submicron noise. We assume that the error-control block has been designed

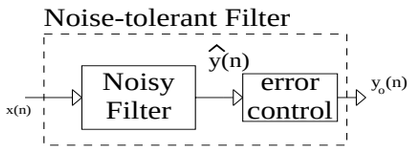

Figure 5: The proposed algorithmic noise-tolerant digital filtering scheme

to be error-free. For soft DSP, as it will be shown later, this assumption holds as the critical path delay of the error-control block will be small compared to that of the filter. Similarly, in case of DSM noise, a noise elimination design strategy [7] can be adopted to obtain an error-free error control block. As the complexity of the error-control block is much lesser than that of the filter, the design overhead will be significantly smaller.

3.1 A Difference-based Error-Control Scheme

In this subsection, we present a simple error-control scheme suitable for lowpass digital filters with a relatively narrow passband. This scheme can be shown to be a special case of the more general predictive error-control scheme presented in section 3.2. Let $y(n)$ denote the filter output when the filter error-free and is given by (7). The difference in consecutive samples of the filter output is given by,

$$y_d(n) = y(n) - y(n-1). \quad (9)$$

Let $\hat{y}(n)$ denote the filter output when the filter is operating under reduced voltage, with

$$\hat{y}(n) = y(n) + y_{err}(n), \quad (10)$$

where $y_{err}(n)$ denotes the the error in the filter output due to soft computations. Note that $y_{err}(n)$ is non-zero only when the input pattern is such that longer paths in the filter implementation are excited. Assuming that the past input is noiseless, i.e, $y_{err}(n-1) = 0$, we have,

$$\hat{y_d}(n) = y_d(n) + y_{err}(n), \quad (11)$$

where $\hat{y_d}(n)$ is the difference in the filter output in presence of errors. From Schwartz inequality and (11), it can be easily shown that,

$$|\hat{y_d}(n)| \geq |y_{err}(n)| - |y_d(n)|. \quad (12)$$

Assuming that $|y_d(n)| < E_{th}$ for all n, where E_{th} is a suitably chosen difference threshold as described later, the following difference-based error-control scheme (shown in Figure 6) is derived from (12):

- compute $\hat{y_d}(n) = \hat{y}(n) - \hat{y}(n-1)$ (from (9)).
- Error detection:
  if $|\hat{y_d}(n)| \geq E_{th}$, an error is declared.
- Error correction:
  if an error is declared, $y_o(n) = \hat{y}(n-1)$.
  else $y_o(n) = \hat{y}(n)$.

If an error is detected, the past output sample is taken to be the estimate for the current output sample. The performance of the above algorithm is based on the choice of E_{th}, the relative magnitudes of $y_d(n)$ and $y_{err}(n)$, and the frequency with which errors occur. The value of E_{th} is chosen such that $|\hat{y}(n)| < E_{th}$ when $y_{err}(n) = 0$

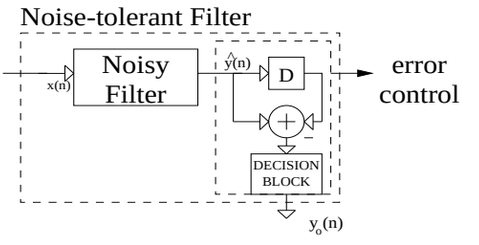

Figure 6: Difference-based ANT scheme for LPF.

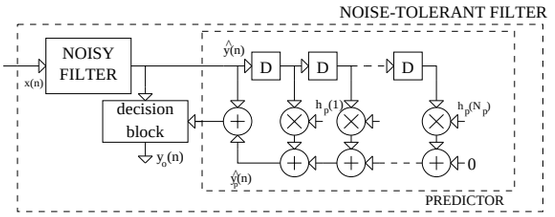

Figure 7: Prediction-based algorithmic noise-tolerance.

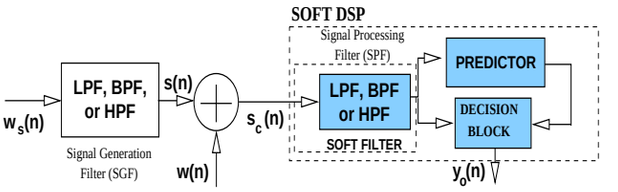

Figure 8: Simulation setup to evaluate the proposed scheme.

(in absence of error) and $|\hat{y}(n)| > E_{th}$ when $y_{err}(n) \neq 0$ (in presence of error). In this paper, we have chosen $E_{th} = 5\sigma_d$, where σ_d^2 is the variance of $y_d(n)$. As the variance in $y_d(n)$ increases with bandwidth, the effectiveness of the above approach in performing error detection deteriorates. Hence for larger bandwidths, a more sophisticated prediction-based scheme presented in section 3.2 is employed.

3.2 Prediction-based Error-Control

A general prediction-based scheme that can handle different input correlation structures is shown in Figure 7. In this scheme, a low complexity linear forward predictor is employed to get an estimate of the current sample of the filter output based on its past samples. In the absence of errors in the filter output, the prediction error is usually small. A large error in the filter output due to excessive voltage reduction leads to an increase in the magnitude of prediction error and this phenomenon is employed to detect errors in the filter output.

Let $y_p(n)$ denote the output of an N_p-tap predictor when the filter is noiseless, i.e.,

$$y_p(n) = \sum_{k=1}^{N_p} h_p(k) y(n-k), \tag{13}$$

where $h_p(k)$ denotes the *optimum predictor coefficients* [13] that minimize the mean squared value (MSE) $< e_p^2(n) >$ of the prediction error $e_p(n)$, given by,

$$e_p(n) = y(n) - y_p(n). \tag{14}$$

The minimum mean square error (MMSE) depends on the autocorrelation function of $y(n)$ and the order of the predictor. Let $\hat{y}(n)$, $\hat{y_p}(n)$, and $\hat{e_p}(n)$ denote the filter output, the predictor output, and the prediction error, respectively, in presence of errors due to soft computations.

Define $\hat{y}(n) = y(n) + y_{err}(n)$, where $y_{err}(n)$ denotes the error in $y(n)$ due to voltage reduction. From (10) and (14), we get

$$\hat{e_p}(n) = y_{err}(n) + e_p(n). \tag{15}$$

Assuming that no more errors occur in the next N_p output samples, we can show that,

$$\hat{e_p}(n+m) = -h_p(m) y_{err}(n) + e_p(n+m), \tag{16}$$

for $m = 1, 2, \cdots, N_p$. Equations (15) and (16) can now be expressed in vector form as

$$\hat{\mathbf{e}}_p(n) = y_{err}(n)\mathbf{h} + \mathbf{e}_p(n) \tag{17}$$

where $\hat{\mathbf{e}}_p(n) = [\hat{e_p}(n) \;\; \hat{e_p}(n+1) \;\; \cdots \;\; \hat{e_p}(n+N_p)]^T$,
$\mathbf{h} = [1 \;\; -h_p(1) \;\; -h_p(2) \;\; \cdots \;\; -h_p(N_p)]^T$,
and $\mathbf{e}_p(n) = [e_p(n) \;\; e_p(n+1) \;\; \cdots \;\; e_p(n+N_p)]^T$.

3.3 Error-Control Algorithm

The following algorithm, derived from (17) is employed for error control:

- Error detection:
  If $|\mathbf{h}^T \hat{\mathbf{e}}_p(n)| > E_{th}$, then an error is declared.
- Error correction:
  If an error is declared, then $y_o(n) = y_p(n)$,
  else $y_o(n) = \hat{y}(n)$.

In this case, we set $E_{th} = \sigma_{ep}$, where σ_{ep}^2 is the variance of the prediction error with a noiseless digital filter. Hence, if an error is detected, the predictor output based on the past correct samples is declared as the system output. The performance of the prediction-based error control algorithm depends upon the choice of E_{th} and the frequency with which errors occur. As, the magnitude of the error in filter output $|y_{err}(n)|$ will be several orders larger than $|\mathbf{h}^T \mathbf{e}_p(n)|$, if $y_{err}(n) \neq 0$, $|\mathbf{h}^T \hat{\mathbf{e}}_p(n)|$ will be large. It can be seen from (17) that, when $y_{err}(n) \neq 0$, the $||\mathbf{h}||^2$ term amplifies the effect of $y_{err}(n)$ on the product $\mathbf{h}^T \hat{\mathbf{e}}_p(n)$. This enables the prediction-based algorithm to detect errors of smaller magnitude and hence we choose a smaller decision threshold σ_{e_p}.

The effectiveness of the error detection and correction scheme described above depends on the following assumptions:

1. The magnitude of $y_{err}(n)$ is relatively large. Errors with higher magnitudes lead to a higher value of $|\mathbf{h}^T \hat{\mathbf{e}}_p(n)|$ and the error is easily detected.
2. The probability that $\hat{y}(n) \neq y(n)$ is small enough such that the frequency of the errors in the filter output is less that $1/2N_p$. The performance of the above scheme deteriorates when multiple errors occur at the filter output in the span of $2N_p$ samples.

The errors due to soft computations occur in the MSBs and hence are of large magnitude. This validates assumption 1. Assumption 2, limits the factor by which voltage can be reduced as the set of error inducing input combinations grows with increase in delay due to voltage reduction. The experimental results presented in this paper demonstrate that substantial energy savings can be obtained before assumption 2 is violated. In case of soft computations, note that the errors occur in the MSBs and hence are of higher magnitude. In such a case, the proposed error-control algorithm can be relaxed further to reduce the complexity of the error-control circuitry.

4 Experimental Results

The setup used to measure the performance of the proposed scheme in which the filtering algorithm is employed in the frequency selective filtering configuration is shown in Figure 8. A lowpass, bandpass or a highpass filter (LPF, BPF or HPF), denoted as the signal generation filter (SGF), is used to generate a bandlimited signal $s(n)$ from a wideband input $w_s(n)$. The signal $s(n)$ is the corrupted by wideband noise $w(n)$, i.e., the signal $s_c(n)$ is obtained

as $s_c(n) = s(n) + w(n)$, where $s(n)$ is the output of the SGF for a wideband input $w_s(n)$. As $s(n)$ is bandlimited, the SNR can be improved by passing $s_c(n)$ though a frequency selective filter with bandwidth ω_b. This filter is denoted as the signal processing filter (SPF) in Figure 8, and it supresses the out-of-band components of the noise signal $w(n)$. We employ the proposed soft DSP implementation of the filtering algorithm to perform frequency selective filtering on $s_c(n)$ as shown in Figure 8. Note that this setup simulates several practical scenarios for signal processing where the task is to extract a bandlimited signal embedded in wideband noise. We employ a folded implementation for the signal processing filter containing N taps where all the taps are mapped on to a single MAC. The coefficient and the input data precisions are chosen to be 10 and 8 bits, respectively. We have chosen $N = 29$ for all the experimental results presented in this paper as it was sufficient in providing the required SNR improvement for several bandwidths considered in this paper.

4.1 Performance Measures

The performance of the proposed scheme is measured via two experiments. In the first experiment, we study the performance in restoring the SNR degradation due to soft computations. We also measure the resulting savings in energy dissipation, present the energy-performance relationship, and compare it to that of the conventional TCA. In the second experiment, we measure the performance of the proposed scheme in presence of DSM noise by introducing errors randomly at the SPF output. The SNR at the output of the filter in presence of errors is given by,

$$\hat{SNR}_o = 20 \log_{10} \left(\frac{\sigma_s}{\sigma_n + \sigma_c} \right), \tag{18}$$

where σ_s^2 is the variance of the signal component (due to $s(n)$), σ_n^2 is the variance of the noise component (due to $w(n)$),and σ_c^2 is the variance of error in the output due soft computations or DSM noise (i.e., $< y_{err}(n)^2 >$).

In order to estimate the energy savings obtained via voltage reduction as proposed, the energy dissipation values are obtained by using MED [14], a gate level energy estimator. Note that the simulator uses a real delay model and hence takes into account the glitching activity in the circuit. An extended simulation for 2000 input vectors is performed for the arithmatic blocks employed in both the traditional and the proposed schemes to obtain energy estimates. The gate library parameters comprised of delay and capacitance values that are typical of a $0.5\mu m$ CMOS technology. When the supply voltage is scaled, the V_{dd} values corresponding to a given path-delay are obtained by solving (1) with $\alpha = 2.0$ (no velocity saturation) and for $\alpha = 1.5$ and 1.2 (with velocity saturation). The reduction in energy dissipation is characterized by *energy savings* (ES), defined as

$$ES = \frac{E_{original} - E_{proposed}}{E_{original}} \times 100\%, \tag{19}$$

where $E_{original}$ is the energy dissipation with conventional voltage scaling (i.e., with $V_{dd} = V_{dd-crit}$), and $E_{proposed}$ is the energy dissipation with the proposed scheme.

4.2 Effect of velocity saturation on soft DSP

The plot of K_v vs. SNR_o for a lowpass filter employing the proposed SMA, with filter bandwidth $\omega_b = 0.2\pi$ for several values of α is shown in Figure 9(a). Note that smaller α enables a smaller value of K_v, the voltage scaling factor. With $\alpha = 1.2$, V_{dd} scaling by a factor of 0.72 is possible with a degradation of less than $0.5dB$ in SNR_o using ANT. The difference-based ANT scheme is employed here as the filter bandwidth is small. It can be seen that reduction in α due to velocity saturation enables higher reduction in K_v and hence higher energy savings.

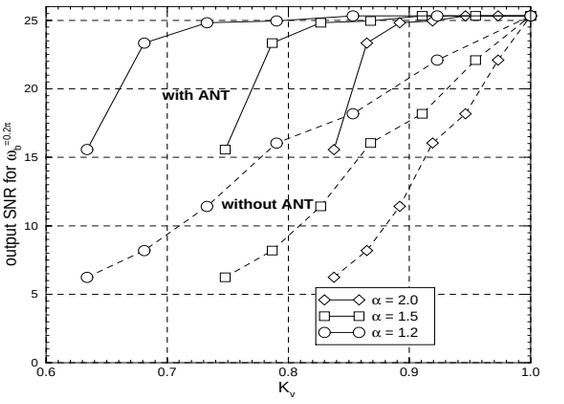

Figure 9: K_v vs. SNR_o for several values of α for the proposed SMA

4.3 Energy-Performance Characteristics

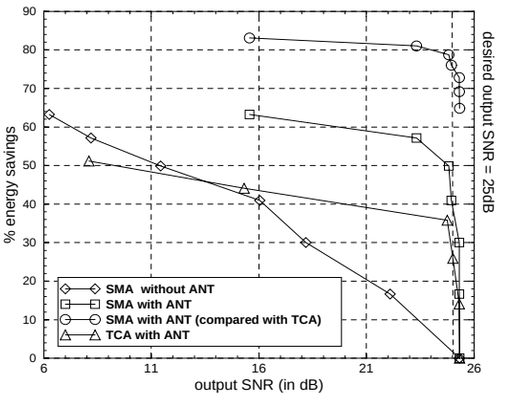

Figure 10: Performance vs. Energy Savings ($\alpha = 1.2$) with filter bandwidth $\omega_b = 0.2\pi$ for the difference-based ANT scheme.

The plot of SNR_o vs. energy savings of the proposed soft DSP scheme for $\omega_b = 0.2\pi$ employing the difference-based ANT scheme is shown in Figure 10. For comparison purposes, we choose a conventional TCA architecture operating at its $V_{dd-crit}$ as a reference. Note that this is the best that traditional voltage scaling achieves. The proposed SMA architecture with ANT leads to 80% energy savings over the conventional TCA. The energy-savings via soft DSP employing the TCA is 34%, whereas the same with the proposed SMA is 51% when the performance degradation allowed is less than $0.5dB$. The proposed architecture leads to savings over conventional TCA due to the following reasons: 1.) the critical path delay of the proposed architecture is reduced due to the unsigned multiplier, and 2.) the transition activity of the proposed architecture is reduced due to the employment of the signed magnitude representation in the multipliers. Note that the reduced transition activity in the MSBs allows for higher voltage scaling for a given error frequency.

Figure 11 shows the *energy-performance* relationship for $\omega_b = 0.5\pi$. In this case, the possible energy savings over the conven-

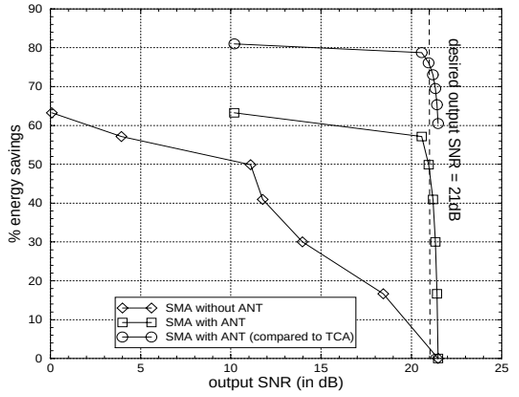

Figure 11: Performance-Energy relationship (with $\alpha = 1.2$) of the prediction-based ANT scheme with $N_p = 3$ for filter bandwidth $\omega_b = 0.5\pi$.

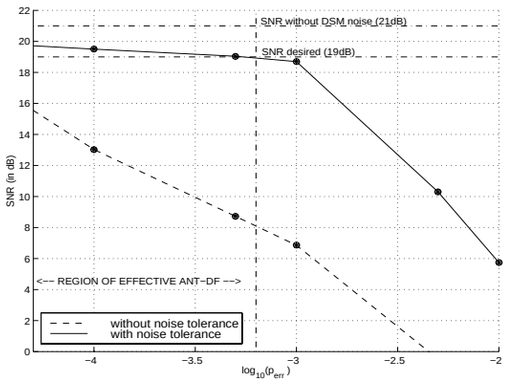

Figure 12: Performance of the proposed ANT scheme for filter bandwidth $\omega_b = 0.4\pi$ with predictor tap-length, $N_p = 4$.

tional TCA operating at $V_{dd-crit}$ is 78% with a performance loss of about $0.5dB$. Note that the energy savings compared to the proposed architecture operating at $V_{dd-crit}$ is about 58% and the corresponding savings for filter bandwidth = 0.3π is 64%. The drop in energy savings for higher bandwidths is due to the reduction in correlation of the filter output which requires a higher predictor length for a given error-control performance.

4.4 Performance of ANT in the Presence of DSM Noise

In order to experimentally verify the effectiveness of the proposed approach in presence of DSM noise, errors are introduced at the system level by flipping the output bits of the digital filter independently, with a fixed probability denoted by p_{err}. Note that more accurate performance results require detailed DSM noise models for the arithmetic units employed in the digital filter which are currently not available. The performance of the proposed algorithm for a digital filter with bandwidth $\omega_b = 0.4\pi$ and 48 taps is shown in Figure 12(a). As expected, without noise-tolerance, the degradation in performance increases with increase in p_{err} as expected. Also, the proposed scheme provides up to $10dB$ improvement in performance. The SNR with ANT stays almost constant above $19dB$ till $p_{err} = 10^{-3}$ and then reduces sharply. In this range, the probability of error is low enough that the assumption of infrequent errors (assumption 2 in section 3.3) is satisfied.

5 Conclusions & future work

In this paper, we have proposed soft DSP for reduction in energy dissipation. It was shown that the effectiveness of the proposed approach depends on two key features: 1) the path delay distribution of the architecture employed and 2) the effectiveness of the error-control schemes in restoring performance degradation. Past work in low-power arithmetic unit design focuses on delay balancing to reduce glitch power. However, this work suggests that employing delay imbalanced arithmetic units in a soft DSP framework leads to higher energy savings with marginal degradation in performance. Future work on this topic will involve exploring various arithmetic units and DSP architectures for their effectiveness in soft DSP scenario and developing error-control schemes for widely used DSP algorithms.

References

[1] S. Iman and M. Pedram, "An approach for multi-level logic optimization targeting low power," *IEEE Trans. on Computer Aided Design*, Vol. 15, No. 8 (1996), pages 889-901.

[2] R. Gonzalez, et. al., "Supply and threshold voltage scaling for low power CMOS," *IEEE Journal of Solid State Circuits*, vol. 32, No. 8, Aug. 1997.

[3] V. Gutnik and A. Chandrakasan, "Embedded power supply for low power DSP," *IEEE Trans. on VLSI Syst.*, Vol. 5, pp. 425-435, Dec. 1997.

[4] R. Hegde and N. R. Shanbhag, "A Low-power phase splitting passband equalizer," *IEEE Trans. on Signal Processing*, vol. 47, no. 3, March 1999.

[5] A. Chandrakasan and R. W. Brodersen, "Minimizing power consumption in digital CMOS circuits," *Proceedings of the IEEE*, vol. 83, no. 4, pp. 498-523, April 1995.

[6] P. J. Restle et. al. "Interconnect in high speed designs: problems, methodologies and tools," *ICCAD 98*, San Jose, CA.

[7] K. L. Shepard and V. Narayanan, "Noise in deep submicron digital design," pp. 524–531, *DAC 96*, San Francisco, CA.

[8] N. R. Shanbhag, "A mathematical basis for power-reduction in digital VLSI systems", *IEEE Trans. on Circuits and Systems, Part II*, vol. 44, no. 11, pp. 935-951, Nov. 1997.

[9] R. Hegde and N. R. Shanbhag,"Energy efficiency in presence of deep submicron noise," *ICCAD 98*, San Jose, CA, Nov. 1998.

[10] J. M. Rabaey, *Digital Integrated Circuits: A Design Perspective*. Prentice-Hall, New Jersey, 1996.

[11] D. Sylvester and K. Keutzer, "Getting to the bottom of deep submicron," *ICCAD 98*, San Jose, CA.

[12] M. D. Ercegovac and T. Lang, "Low-Power Accumulator (Correlator)," in proc. *International symposium on Low-Power Electronic Design (ISLPED)*,san Francisco, CA, 1995, pp. 30-31.

[13] J. G. Proakis and D. G. Manolakis, *Digital Signal Processing: Principles, Algorithms and Applications* Prentice-Hall, New Jersey, 1996.

[14] M. Xakellis and F. Najm, "Statistical estimation of the switching activity in digital circuits," in 31^{st} *ACM/IEEE Design Automation Conference*, San Diego, CA, June 6-10, 1994, pp. 33–36.