Soft Digital Signal Processing

Rajamohana Hegde, *Student Member, IEEE* and Naresh R. Shanbhag, *Member, IEEE*

***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 imposed by the requirement 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 subcritical voltage and the error control scheme is referred to as *soft DSP*. The effectiveness of the proposed scheme is 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 the 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 architecture and voltage scaling, with a maximum of 0.5-dB degradation in the output signal-to-noise ratio (SNR_o). It is also shown that the proposed algorithmic noise-tolerance schemes can also be used to improve the performance of DSP algorithms in presence of bit-error rates of up to 10^{-3} due to deep submicron (DSM) noise.**

***Index Terms*—Due to deep submicron (DSM) noise, low power, noise-tolerant design, voltage oversealing.**

I. Introduction

ENERGY-efficient very large scale integrated (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, beginning at the algorithmic level [1], architectural level [2], logic level [3], and finally at the circuit level [4]. Schemes at the lower levels of the design process such as the logic and circuit levels are usually application independent [3], [4]. At the algorithmic and architectural levels, features that are specific to a class of applications are exploited to develop application specific energy reduction techniques [5], [6]. Voltage scaling [5] 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 [7]. The static components of energy dissipation are also reduced when supply voltage is scaled without altering the device threshold voltage. However, the extent of voltage scaling [7] is limited by the critical path delay of the architecture and the throughput requirements of the application.

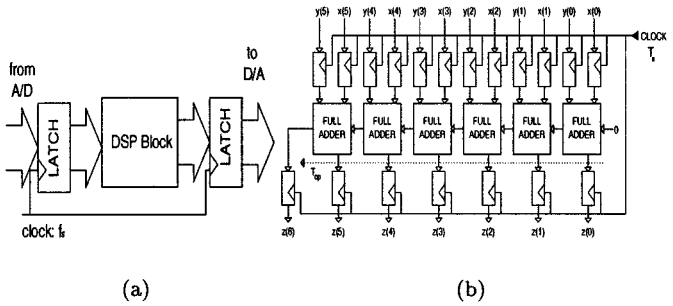

Fig. 1. A typical DSP implementation. (a) Block diagram of a typical DSP system. (b) Block diagram of a 6-bit ripple carry adder.

Consider the block diagram of a typical digital signal processing (DSP) system shown in Fig. 1(a), where the input/output (I/O) registers are clocked at the sample period $T_s = 1/f_s$, where f_s is the sample rate. The critical path delay [18] T_{cp} of the DSP block (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., $T_{\mathrm{cp}} \leq T_s$. Fig. 1(b), shows the block diagram of a 6-bit ripple carry adder. If the time period required for a single full-adder is represented by T_{FA} time units, then the critical path delay T_{cp} of an N-bit full adder is given by $T_{\mathrm{cp}} = NT_{\mathrm{FA}}$. Given T_s, the gates forming the adder are designed such that at the rated supply voltage V_{dd}, the *delay condition*, $T_{\mathrm{cp}} = NT_{\mathrm{FA}} \leq T_s$ is satisfied. The relationship between V_{dd} and circuit delay τ_d is given by [8]

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

where

- C_L load capacitance;
- α is the velocity saturation index,
- β gate transconductance;
- V_t device threshold voltage.

We refer to the voltage at which $T_{\mathrm{cp}} = T_s$ as the *critical supply voltage* $V_{\mathrm{dd\text{-}crit}}$ of a given architecture. Note that violating the delay condition by reducing V_{dd} beyond $V_{\mathrm{dd\text{-}crit}}$, i.e.,

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

where $0 < K_v \leq 1$ leads to erroneous output when the critical path is excited. Hence, $V_{\mathrm{dd\text{-}crit}}$ is seen as a lower bound on the supply voltage for a given architecture and throughput. In order to reduce energy dissipation, reduction in V_{dd} (without violating

Manuscript received June 3, 1999; revised April 9, 2001. This work was supported by NSF CAREER Award MIP 9623737.

The authors are with the Electrical and Computer Engineering Department, Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA (e-mail: rhegde@uivlsi.csl.uiuc.edu; shanbhag@uivlsi.csl.uiuc.edu).

Publisher Item Identifier S 1063-8210(01)07421-2.

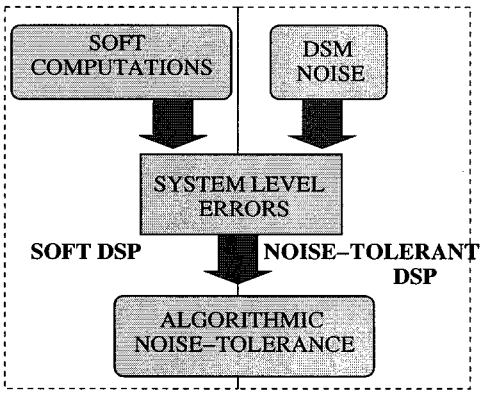

Fig. 2. The proposed soft and noise tolerant DSP framework.

the delay condition) can achieved by reducing the critical path delay of the VLSI implementation via architectural transformations such as pipelining [1] and parallel processing [7]. In this paper, we propose operating the DSP architecture at voltages lower than $V_{\text{dd-crit}}$. 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*.

Errors in system output can also be induced due to deep submicron (DSM) noise [9] due to phenomena such as ground bounce, cross-talk, process parameter variations [10], charge sharing, charge leakage, and slow and unpredictable interconnects. Current approaches to address the issue of DSM noise range from interconnect centric design methodologies [11] to systematic static noise analysis methodology [12].

In this paper, we propose *algorithmic noise-tolerance* (ANT) to compensate for degradation in the system output due to errors from either soft computations or DSM noise. 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 Fig. 2. The setup that comprises of the DSP architecture operating at a subcritical supply voltage (lower than $V_{\text{dd-crit}}$) 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 μm) favors low-power operation via soft DSP. In contrast to the existing solutions to the DSM noise problem [11], [12], we propose the use of ANT schemes to restore degradation in algorithmic performance due to DSM noise and refer to the resulting setup as *noise-tolerant DSP* as shown in Fig. 2. The motivation for the framework in Fig. 2 is derived from the recently proposed [13], [14] *information-theoretic approach* to jointly address the energy efficiency and reliability issues for DSM technology.

Several researchers have proposed algorithmic level error-control schemes. In algorithm based fault-tolerance (ABFT) [15], redundant computations are employed to detect and locate errors. The correct output is then recomputed. The fault-tolerant FFT processors [16] involves detection and isolation of permanent faults. In [17], redundant taps are provided

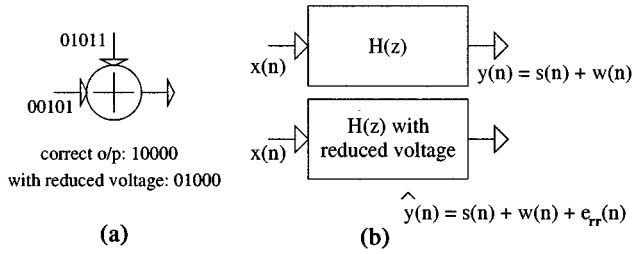

Fig. 3. Effect of errors on performance of DSP algorithms.

to restore performance degradation due to stuck-at faults at one or more bits in tap outputs. The ANT scheme proposed in this paper aims at restoring the degradation in SNR due to transient errors and not exact error correction. Hence, the error-control overhead is substantially smaller as compared to the existing fault-tolerant schemes.

The rest of this paper is organized as follows. In Section II, the proposed notion of soft DSP is introduced with a motivational example. It is also dependent on the path-delay distribution of the architecture employed. A new multiply-accumulate (MAC) architecture that augments the effectiveness of the proposed approach in digital filter implementations is presented. In Section III, 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 IV, we study the energy savings due to the proposed approach in the context of frequency selective filtering. In particular, the effectiveness of the proposed scheme for existing and proposed filter architecture is studied. 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 Section V conclusions and scope for future work on this topic are presented.

II. Energy Savings via Soft DSP

The efficacy of the soft-DSP approach is a function of: 1) supply voltage scaling beyond $V_{\text{dd-crit}}$; 2) error frequency; and 3) the overhead due to error-control. 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 MAC architecture that improves the effectiveness of the proposed scheme for the filtering algorithm is presented.

A. Motivational Example

Consider the 5-bit adder shown in Fig. 3(a), where the input operands are 00 101 and 01 011. Assuming that $T_{\text{FA}} = 3$ ns, the critical path delay of this adder is 15 ns. Note that the time taken to compute the output corresponding to the two operands is also 15 ns as the carry generated in adding the least significant bit (LSB) propagates all the way to the most significant bit (MSB). Let $T_s = 15$ ns. If the supply voltage is now reduced such that $T_{\text{FA}} = 5$ ns, the adder output at the end of the sample period will be 01 000 as shown in Fig. 3(a). Hence, the numerical value of the adder output will be 8 instead of 16. These errors in turn result in a wrong system output. If the inputs do not excite the

longer paths (e.g., 00 001 and 00 010), then the adder provides correct outputs. We refer to such an adder as a *soft adder* and, in general, we refer to arithmetic units operating at subcritical voltages as soft computational blocks.

In the absence of errors, the algorithmic performance of a filter transfer function $H(z)$ [as shown in Fig. 3(b)] is measured in terms of the output SNR given by

$$\text{SNR}_o = 20\log\left(\frac{\sigma_s}{\sigma_w}\right) \tag{3}$$

where σ_s^2 and σ_w^2 are the signal and noise powers, 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

$$\hat{y}(n) = y(n) + e_{rr}(n) \tag{5}$$

where $e_{rr}(n)$ is the error introduced in the output sample at the nth instant. Note that $e_{rr}(n)$ will be nonzero only when errors occur in the filter output. The output SNR is given by

$$\hat{SNR}_o = 20\log\left(\frac{\sigma_s}{\sigma_{w+c}}\right) \tag{6}$$

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 extent of the degradation depends on the magnitude of the error $e_{rr}(n)$ when it is nonzero and the frequency with which it occurs. As shown in the adder example above, errors from to soft computations occur in the MSBs due to longer path delays. These errors can cause substantial degradation in algorithmic performance and are easily detectable. 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. Note that intermittent errors could also be due to DSM noise. In this paper, we propose employing low-complexity algorithmic error-control schemes to detect and mitigate the effect of errors on algorithmic performance, where the errors could be due to either soft computations or DSM noise.

B. Path-Delay Distribution of Adders

A soft DSP-friendly architecture is one where the number of paths that fail, increases gradually as the supply voltage is scaled beyond $V_{\text{dd-crit}}$. 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 $x = 01\,010\,101$ and $y = 10\,101\,010$ (N is assumed to be 8) are evaluated in just T_{FA} time units. Other combinations such as $x = 11\,111\,111$, $y = 00\,000\,001$ excite the critical path requiring $8T_{\text{FA}}$ time units. In Fig. 4(a), the path delay distribution (histogram of the evaluation time) for all possible input combinations for an 8-bit ripple carry adder is shown. It can be seen that more than 95% of the possible input combinations are evaluated in $5T_{\text{FA}}$ time units. Assuming

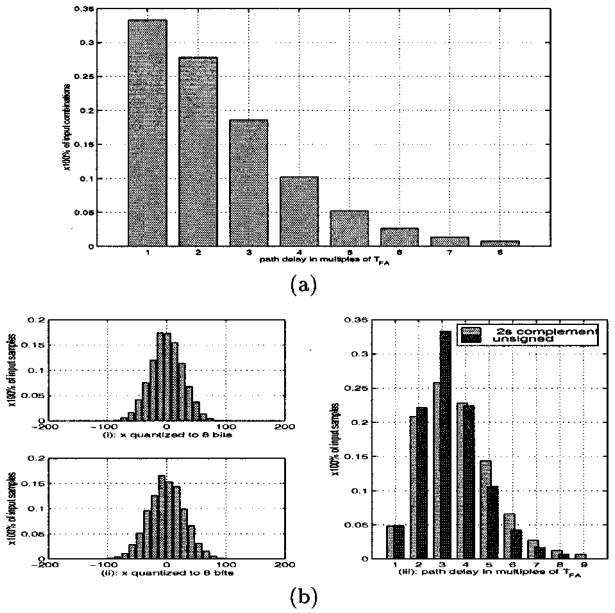

Fig. 4. Path delay histograms of 8-bit ripple carry adder: (a) Path delay statistics over all input combinations for an 8-bit ripple carry adder. (b) Input and output histograms of an 8-bit 2's complement ripple carry adder.

that the input distribution is uniform, reducing V_{dd} such that the new full-adder delay $T'_{\text{FA}} = (5/8)T_{\text{FA}}$, leads to an error probability of 0.05. This can be reduced even further provided the input combinations that excite the critical paths occur infrequently. The path delay histogram of an 8-bit two's complement ripple carry adder where the input operands are generated randomly, is shown in Fig. 4(b), (iii). Fig. 4(b), (i) and 4(b), (ii) show the probability distribution of the adder operands x and y, respectively, with the X axis being the integer value of the corresponding binary combination. Also shown in Fig. 4(b), (iii) is the histogram of path delays for an unsigned 8-bit adder. The operands are the same as in Fig. 4(b) (i) and 4(b) (ii). However, 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 the two's complement number representation, small negetive numbers have a higher number of ones which require larger path delays. Therefore, it is better to use unsigned number representation in order to further improve the effectiveness of soft DSP. In this paper, we propose a MAC architecture, presented in the following subsection, that employs an unsigned array multiplier for digital filter implementations. We show in Section IV, that this MAC architecture gives an additional 20% energy savings as compared to the 2's complement architecture when used to implement an ANT-based system.

C. Energy Savings Versus Probability of Soft Errors

The relationship between the supply voltage V_{dd} and the gate delay is given by (1). A plot of V_{dd} versus the increase in the normalized delay for several values of the velocity saturation index α is shown in Fig. 5(a). Note that reduction in α (due to short channel effects) leads to a reduction in the delay penalty suffered due to voltage reduction. The rate at which these errors would occur is input dependent. A plot of the factor by which

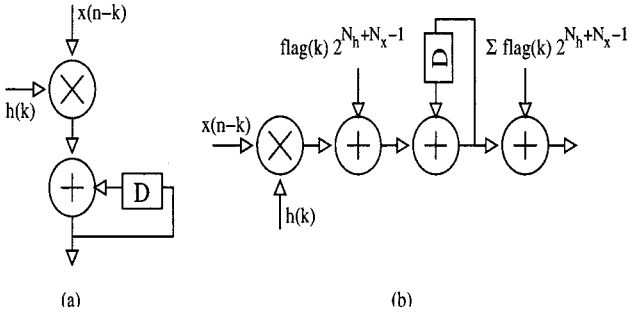

Fig. 5. MAC architectures. (a) Traditional. (b) Proposed.

voltage is reduced versus the probability of error for an 8-bit adder is shown in Fig. 5(b). We have set $|V_t| = 0.62$ V (typical for 0.35 μm CMOS technology) , and the value of C_L/β was chosen such that T_d is normalized to 1 ns at $V_{\rm dd} = 3.3$ V. Once again, note that reduction in α enables higher voltage scaling for a given probability of error. Finally, the relationship between energy dissipation E and the probability of error [shown in Fig. 5(c)] was obtained by assuming a quadratic relationship between $V_{\rm dd}$ and E, given by [18]

$$E = tC_L V_{\rm dd}^2 \tag{7}$$

where t is the probability of a $0 \rightarrow 1$ transition. Note that for $\alpha = 2.0$, about 50% reduction in energy dissipation can be achieved with a probability of error of 0.1. The energy savings possible are about 80% for velocity saturated devices with $\alpha = 1.2$. Hence, the proposed approach leads to higher energy savings with technology scaling. The impact of the probability of error can be reduced via the ability to detect and correct errors in the output of DSP systems with a low-complexity overhead. This will allow for substantial reduction in energy dissipation with marginal degradation in performance.

D. MAC Architecture for Soft DSP

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

$$y(n) = \sum_{k=0}^{N-1} h(k)x(n-k) \tag{8}$$

where

$h(k)$	kth coefficient of the filter;
$x(n-k)$	input at $n-k$th instant;
N	order of the filter.

A typical MAC structure to compute the filter output is shown in Fig. 6(a), where the multiplier computes the product $h(k)x(n-k)$, which are accumulated by the adder. 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 [5], [23], to reduce transition activity in correlators for wireless applications. The proposed MAC structure [refered to as the sign-magnitude architecture (SMA)] shown in Fig. 6(b) 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

$$\begin{aligned} y'(n) &= \sum_{k=0}^{N-1} h(k)x(n-k) \\ &\quad + \text{flag}(k)2^{(N_h+N_x-1)} \\ \text{where, flag}(k) &= 0 \text{ if } h(k)x(n-k) \text{ is } +\text{ve} \\ &= 1 \text{ if } h(k)x(n-k) \text{ is } -\text{ve} \end{aligned} \tag{9}$$

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 $\sum_{k=0}^{N-1} \text{flag}(k)2^{(N_h+N_x-1)}$ 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 two adders.

III. 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 Fig. 7(a), 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 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 [9] 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.

A. 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 III-B. Let $y(n)$ denote the filter output when the filter error free and is given by (8). The difference in consecutive samples of the filter output is given by

$$y_d(n) = y(n) - y(n-1). \tag{10}$$

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

$$\hat{y}(n) = y(n) + y_{\rm err}(n) \tag{11}$$

where $y_{\rm err}(n)$ denotes the the error in the filter output due to soft computations. Note that $y_{\rm err}(n)$ is nonzero only when the input

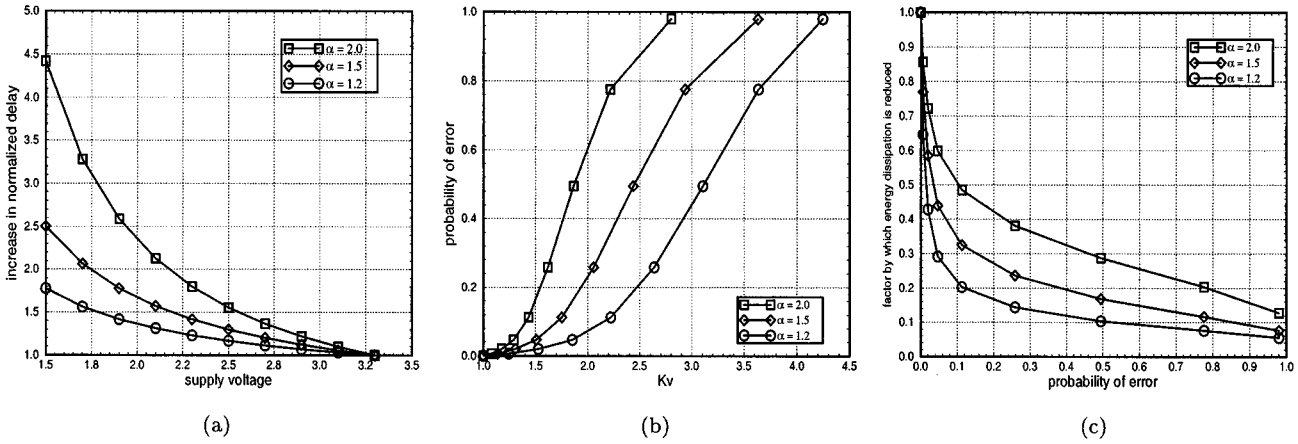

Fig. 6. Energy dissipation-error probability of 8-bit ripple carry adder. (a) $V_{\rm dd}$ versus normalized delay. (b) K_v versus probability of error, and (c) probability of error versus% reduction in energy dissipation.

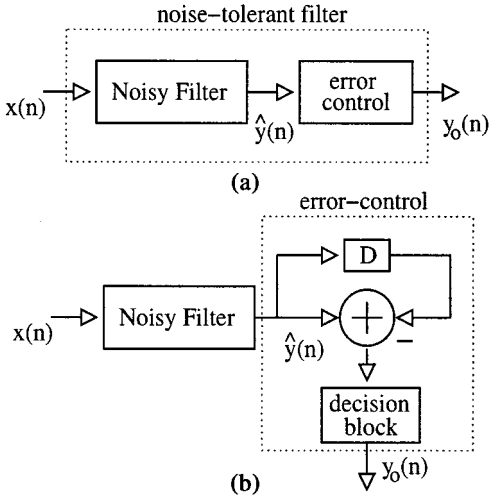

Fig. 7. Algorithmic noise-tolerant digital filtering. (a) The proposed scheme. (b) Difference-based ANT scheme for LPF.

pattern is such that longer paths in the filter implementation are excited. Assuming that the past input is noiseless, i.e, $y_{\rm err}(n-1)=0$, we have

$$\hat{y}_d(n) = y_d(n) + y_{\rm err}(n) \tag{12}$$

where $\hat{y}_d(n)$ is the difference in the filter output in presence of errors. From Schwartz inequality and (12), it can be easily shown that

$$|\hat{y}_d(n)| \geq |y_{\rm err}(n)| - |y_d(n)|. \tag{13}$$

Assuming that $|y_d(n)| < E_{\rm th}$ for all n, where $E_{\rm th}$ is a suitably chosen difference threshold as described later, the following difference-based error-control scheme (shown in Fig. 7(b)) is derived from (13):

- compute $\hat{y_d}(n) = \hat{y}(n) - \hat{y}(n-1)$ (from (10));
- error detection: if $|\hat{y_d}(n)| \geq E_{\rm th}$, 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 $E_{\rm th}$, the relative magnitudes of $y_d(n)$ and $y_{\rm err}(n)$ and the frequency with which errors occur. The distributions of $y(n)$ and $y_d(n)$ obtained over 20 000 samples for several bandwidths (ω_b) of the frequency selective filter is shown in Fig. 8. Note that for lower bandwidths, the variance in $y_d(n)$ is an order of magnitude smaller than that in $y(n)$. Hence, the magnitude of the error in filter output $|y_{\rm err}(n)|$, which occurrs mainly in the MSBs, will be much larger than $|y_d(n)|$. Therefore, if $y_{\rm err}(n) \neq 0$, $|\hat{y}_d(n)|$ will be large. The value of $E_{\rm th}$ is chosen such that $|\hat{y}(n)| < E_{\rm th}$ when $y_{\rm err}(n) = 0$ (in absence of error) and $|\hat{y}(n)| > E_{\rm th}$ when $y_{\rm err}(n) \neq 0$ (in presence of error). In this paper, we have chosen $E_{\rm th} = 5\sigma_d$, 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 III-B is employed.

For the proposed error-control scheme for lowpass filters, note that the past sample of the filter output is used as the estimate of the current sample. Hence, as shown in Fig. 7(b) the overhead involved is just an adder, a delay element and a comparator for the decision block.

B. Prediction-Based Error-Control

A general prediction-based scheme can be proposed to handle different input correlation structures, as shown in Fig. 9. 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.,

$$y_p(n) = \sum_{k=1}^{N_p} h_p(k) y(n-k) \tag{14}$$

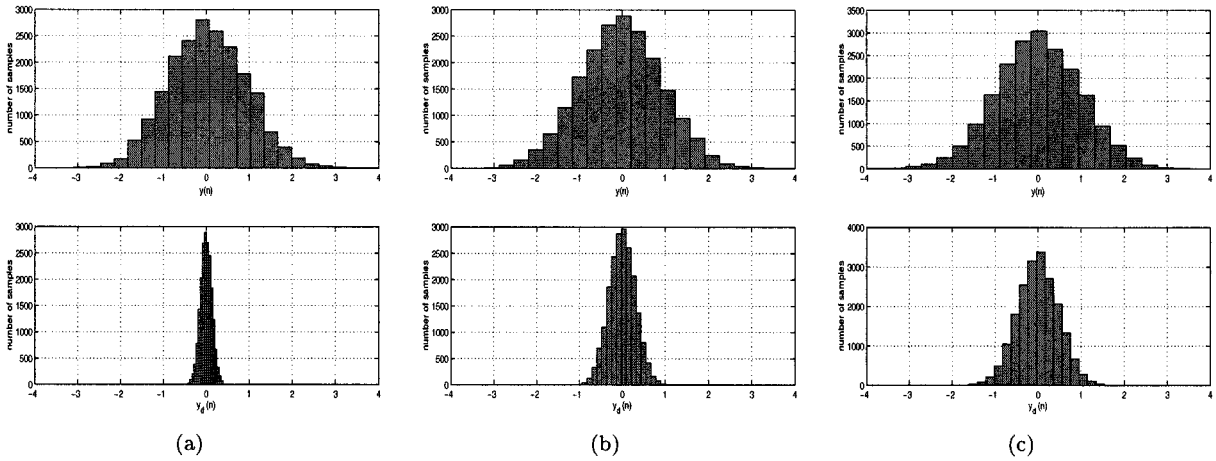

Fig. 8. Distribution of $y(n)$ and $y_d(n)$ for (a) $\omega_b = 0.1\pi$, (b) $\omega_b = 0.2\pi$, and (c) $\omega_b = 0.3\pi$.

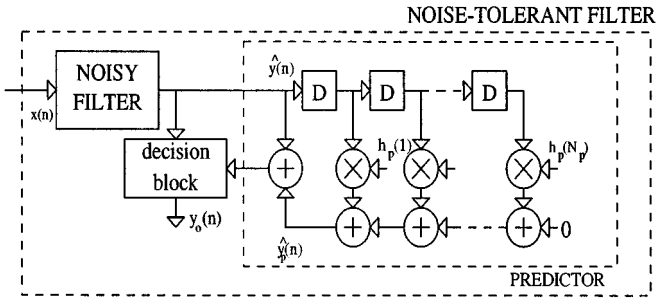

Fig. 9. Prediction-based algorithmic noise tolerance.

where $h_p(k)$ denotes the *optimum predictor coefficients* [24] that minimize the mean squared value (MSE) $\langle e_p^2(n)\rangle$ of the prediction error $e_p(n)$, given by

$$e_p(n) = y(n) - y_p(n). \tag{15}$$

The minimum MSE (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 $\hat{y}(n) = y(n) + y_{\rm err}(n)$, where $y_{\rm err}(n)$ denotes the error in $y(n)$ due to voltage reduction. From (11) and (15), we get

$$\hat{e}_p(n) = y_{\rm err}(n) + e_p(n). \tag{16}$$

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

$$\hat{e}_p(n+m) = -h_p(m)y_{\rm err}(n) + e_p(n+m) \tag{17}$$

for $m = 1, 2, \ldots, N_p$. Equations (16) and (17) can now be expressed in vector form as

$$\hat{\mathbf{e}}_p(n) = y_{\rm err}(n)\mathbf{h} + \mathbf{e}_p(n) \tag{18}$$

where

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

The following theorem derived from (18) is employed next to derive the error detection scheme.

Theorem 1: If $|\mathbf{h}^{\rm T}\mathbf{e}_p(n)| \le E_{\rm th}$, where $E_{\rm th}$ is positive, then

$$|\mathbf{h}^{\rm T}\hat{\mathbf{e}}_p(n)| \ge |y_{\rm err}(n)|\|\mathbf{h}\|^2 - E_{\rm th}. \tag{19}$$

Proof: Multiplying (18) on both sides by $\mathbf{h}^{\rm T}$, we get

$$\mathbf{h}^{\rm T}\hat{\mathbf{e}}_p(n) = y_{\rm err}(n)\|\mathbf{h}\|^2 + \mathbf{h}^{\rm T}\mathbf{e}_p(n). \tag{20}$$

Taking absolute values on both sides and applying Schwartz inequality, we obtain (19). ■

It can be seen from the above theorem that if $|y_{\rm err}(n)| > 2E_{\rm th}/\|\mathbf{h}\|^2$, then $|\mathbf{h}^{\rm T}\hat{\mathbf{e}}_p(n)| > E_{\rm th}$ and, hence, the error is detected.

C. Error-Control Algorithm

The following algorithm, derived from *Theorem 1* is employed for error control.

- Let $\sigma_{e_p}^2$ be the variance of the prediction error with noiseless digital filter.
- Error detection: If $|\mathbf{h}^{\rm T}\hat{\mathbf{e}}_p(n)| > \sigma_{e_p}$, then an error is declared.
- Error correction: If an error is declared, then $y_o(n) = y_p(n)$, else $y_o(n) = \hat{y}(n)$.

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 $Y_{\rm th}$ and $E_{\rm th}$ and the frequency with which errors occur. The distributions of $y(n)$ and $|\mathbf{h}^{\rm T}\mathbf{e}_p(n)|$ obtained over 20 000 samples for several bandwidths of the frequency selective filter is shown in Fig. 10. Note that for all the bandwidths, the variance in $|\mathbf{h}^{\rm T}\mathbf{e}_p(n)|$ is several orders of magnitude smaller than that in $y(n)$, making it possible to choose a smaller value of $E_{\rm th}$ that satisfies the assumption $|\mathbf{h}^{\rm T}\mathbf{e}_p(n)| < E_{\rm th}$ for all n. As, the magnitude of the error in filter output $|y_{\rm err}(n)|$ will be several orders larger than $|\mathbf{h}^{\rm T}\mathbf{e}_p(n)|$, if $y_{\rm err}(n) \neq 0$, $|\mathbf{h}^{\rm T}\hat{\mathbf{e}}_p(n)|$ will be large. It can be seen from (20) that when $y_{\rm err}(n) \neq 0$, the $\|\mathbf{h}\|^2$ term amplifies the effect of $y_{\rm err}(n)$ on the product $\mathbf{h}^{\rm T}\hat{\mathbf{e}}_p(n)$. This enables the prediction-based algorithm to detect errors of smaller

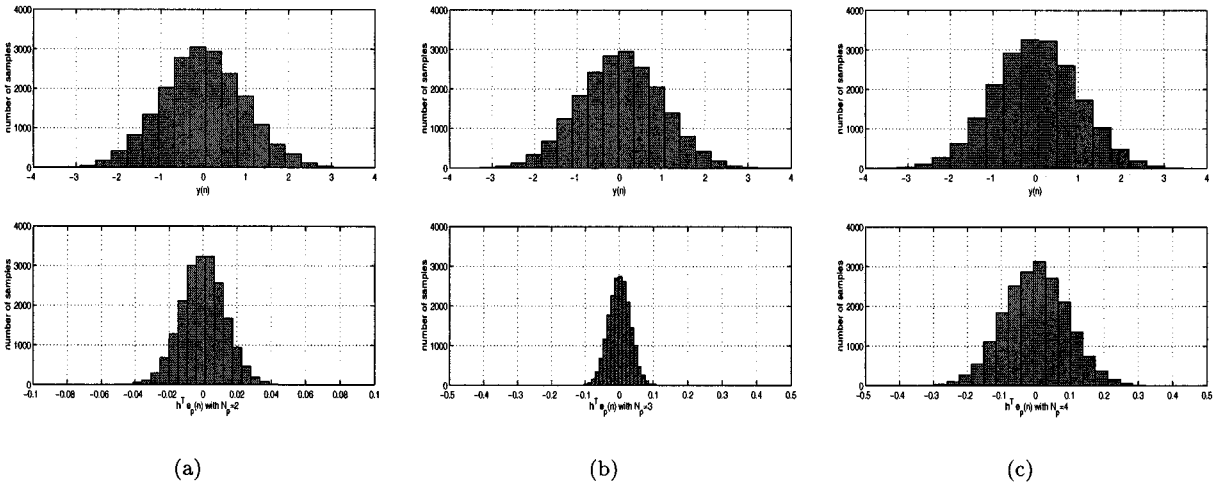

Fig. 10. Distribution of $y(n)$ and $y_d(n)$ for (a) $\omega_b = 0.2\pi$, (b) $\omega_b = 0.4\pi$, and (c) $\omega_b = 0.6\pi$.

magnitude and, hence, we choose a smaller decision threshold σ_{e_p}.

Hence, the effectiveness of the error detection and correction scheme described above depends on the following assumptions.

1) The magnitude of $y_{\rm err}(n)$ is relatively large. Errors with higher magnitudes lead to a higher value of $|\mathbf{h}^{\rm T}\hat{\mathbf{e}}_p(n)|$ 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 reductions. The experimental results presented in this paper demonstrate that for long channel devices ($\alpha = 2.0$), up to 25% reduction in supply voltage can be achieved before assumption 2 is violated. The corresponding value for short channel devices ($\alpha = 1.2$) is about 52%. Hence, in both the cases, substantial energy savings can be obtained before assumption 2 is violated.

D. Complexity of the Error-Control Algorithm

The algorithm presented above involves an N_p-tap filter for the linear predictor and N_p multipliers and N_p-1 adders (equivalent to an N_p-tap filter) to compute the term $|\mathbf{h}^{\rm T}\hat{\mathbf{e}}_p(n)|$. 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 detection circuitry.

1) Relaxation via Coefficient Reduction: Here, the coefficients of the predictor that are small in magnitude are reduced to zero. This allows for elimination of an entire tap from the prediction filter reducing its complexity and energy dissipation. As the coefficients with smaller magnitude have an insignificant impact on $||\mathbf{h}||$, the reduction in performance is usually marginal.

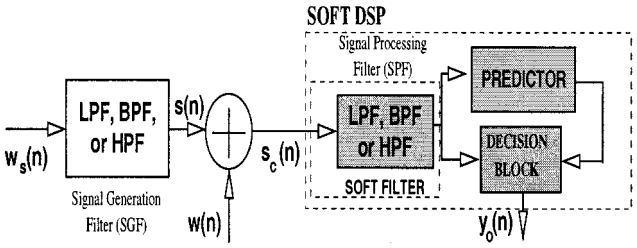

Fig. 11. Simulation setup to evaluate the proposed scheme.

2) Relaxation for Error Detection: It can be seen from (16) that when $|y_{\rm err}(n)|$ is large relative to $|e_p(n)|$, the magnitude of $e_p\hat{(n)}$ is large and, hence, an error can be detected by a jump in the magnitude of $e_p\hat{(n)}$. Hence, step 2 in the error-control algorithm presented above can be modified to the following.

- If $|e_p(n)| > E_{\rm th}$, then an error is declared.

In this case, we set $E_{\rm th} = 4\sigma_{e_p}$.

In the next section, the performance of the error detection and correction schemes described above are simulated in the scenario where a digital filter is employed to reduce out-of-band noise in a bandpass signal. It is also shown that the proposed approach provides for substantial reduction in energy dissipation with marginal performance degradation.

IV. 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 Fig. 11. A lowpass, bandpass, or a high-pass filter (LPF, BPF, or HPF), denoted as the signal generation filter (SGF), is used to generate a bandlimited signal $s(n)$ from a wide-band input $w_s(n)$. The signal $s(n)$ is the corrupted by wide-band noise $w(n)$, i.e., the signal $s_c(n)$ is obtained as

$$s_c(n) = s(n) + w(n) \tag{21}$$

where $s(n)$ is the output of the SGF for a wide-band input $w_s(n)$. The SNR of $s_c(n)$ is given by

$$\text{SNR}_{\rm in} = 20\log_{10}\left(\frac{\sigma_s}{\sigma_w}\right) \tag{22}$$

where σ_s^2 is the variance of $s(n)$ and σ_w^2 is the variance of $w(n)$. 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 Fig. 11 and it supresses the out-of-band components of the noise signal $w(n)$. The SNR at the filter output, denoted by SNR_o, is given by

$$\text{SNR}_o = 20\log_{10}\left(\frac{\sigma_s}{\sigma_n}\right) \tag{23}$$

where σ_n the component of $w(n)$ that occupies the same band as $s(n)$ and, hence, is not suppressed.

We employ the proposed soft DSP implementation of the filtering algorithm to perform frequency selective filtering on $s_c(n)$ as shown in Fig. 11. Note that this setup simulates several practical scenarios for signal processing where the task is to extract a bandlimited signal embedded in wide-band 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. The conventional MAC structure that employs 2's complement architecture (TCA) requires a 8×11 multiplier due to sign extension. The proposed structure however requires a 6×9 multiplier as there is no need for sign extension. The length of the ripple-carry adder in the MAC is 19 bits for the proposed architecture and 22 bits for the conventional TCA. Hence, the critical path delay T_{cp} of the conventional MAC structure is $30T_{\text{FA}}$. The critical path delay of the proposed MAC structure is $25T_{\text{FA}}$. The sample period is given by $T_s = NT_{\text{cp}}$, where N is the number of taps in the filter. 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.

A. Performance Measures

The complexity of the predictor and the error-control block depends on the bandwidth of the signal, the statistics of the errors introduced due to soft computations or DSM noise and the desired SNR_o. 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

$$\hat{SNR}_o = 20\log_{10}\left(\frac{\sigma_s}{\sigma_n + \sigma_c}\right) \tag{24}$$

where

- σ_s^2 variance of the signal component (due to $s(n)$);
- σ_n^2 variance of the noise component (due to $w(n)$);
- σ_c^2 variance of error in the output due soft computations or DSM noise (i.e., $\langle y_{\text{err}}(n)^2 \rangle$).

In order to estimate the energy savings obtained via voltage reduction as proposed, the energy dissipation values are obtained by using MED [25], 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 μm 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

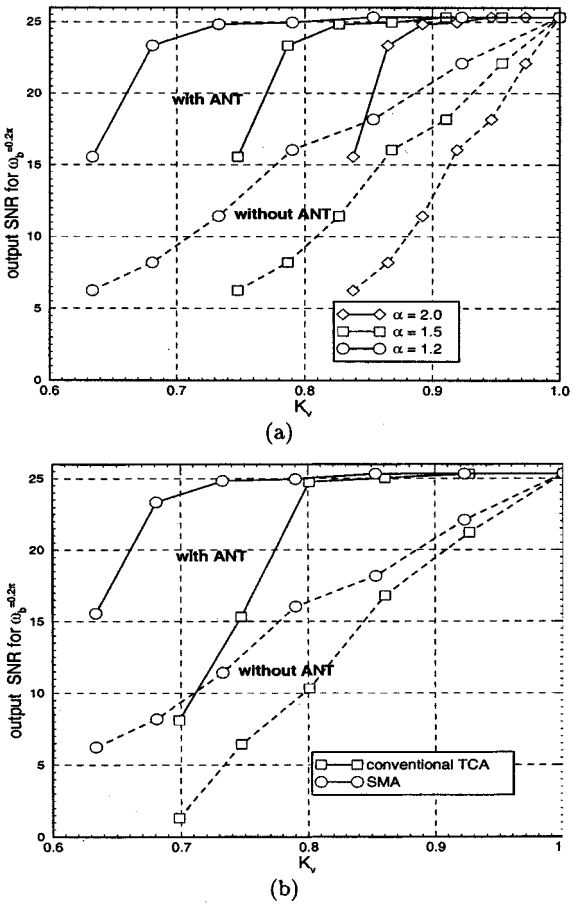

Fig. 12. Performance with filter bandwidth $\omega_b = 0.2\pi$ of the difference-based ANT scheme. (a) K_v versus SNR_o for several values of α for the porposed SMA. (b) K_v versus SNR_o for $\alpha = 1.2$ for the conventional TCA and the proposed SMA.

$$ES = \frac{E_{\text{original}} - E_{\text{proposed}}}{E_{\text{original}}} \times 100\% \tag{25}$$

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

B. Effect of Velocity Saturation on Soft DSP

The plot of K_v versus 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 Fig. 12(a). Note that smaller α en-

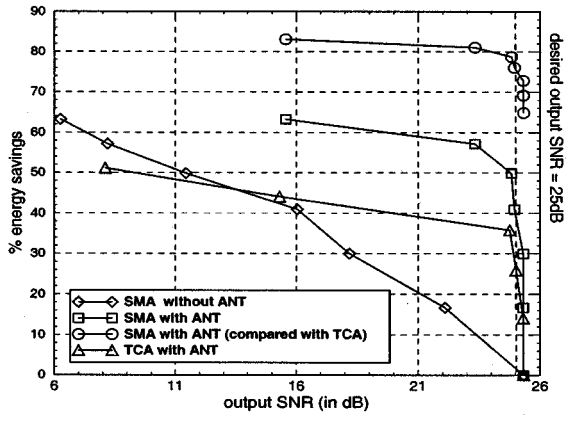

Fig. 13. Performance versus energy savings ($\alpha = 1.2$) with filter bandwidth $\omega_b = 0.2\pi$ for the difference-based ANT scheme.

ables a smaller value of K_v, the voltage scaling factor (see (2)). With $\alpha = 1.2$, $V_{\rm dd}$ scaling by a factor of 0.72 is possible with a degradation of less than 0.5 dB 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.

For the same bandwidth and the ANT scheme, the plot of K_v versus SNR$_o$ characteristics of the proposed signed-magnitude architecture (SMA) and the conventional TCA for $\alpha = 1.2$ are compared in Fig. 12(b). Note that the proposed architecture allows for higher $V_{\rm dd}$ scaling as compared to the conventional TCA. For a performance degradation of less than 0.5 dB, the conventional TCA allows for $K_v = 0.81$ as compared to the proposed scheme that allows for $K_v = 0.72$.

C. Energy-Performance Characteristics

The plot of SNR$_o$ versus energy savings of the proposed soft DSP scheme for $\omega_b = 0.2\pi$ employing the difference-based ANT scheme is shown in Fig. 13. For comparison purposes, we choose a conventional TCA architecture operating at its $V_{\rm dd\text{-}crit}$ 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 0.5 dB. 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. The difference-based scheme can only be employed for low bandwidth filters as its performance degrades rapidly with higher filter bandwidths. For larger filter bandwidths, the prediction-based ANT scheme is more effective in obtaining satisfactory energy-performance relationship despite the higher overhead.

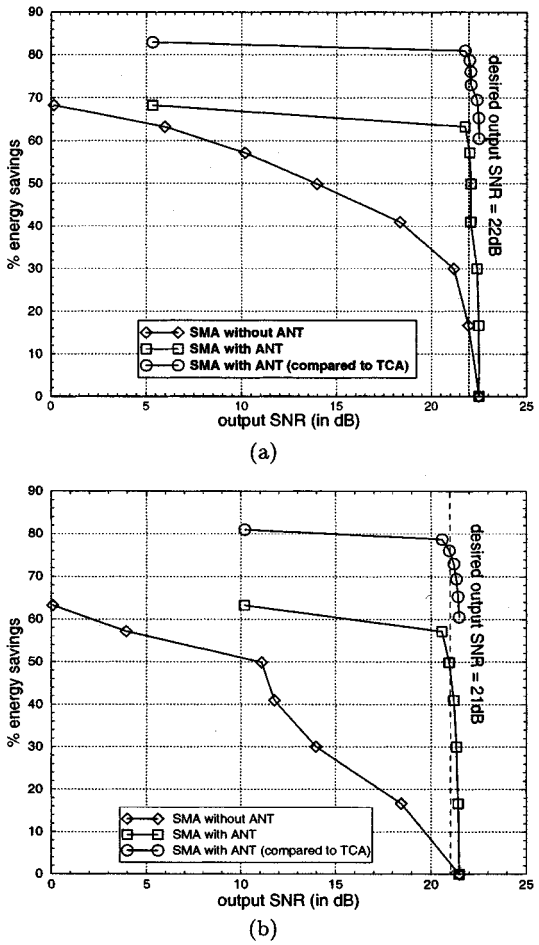

Fig. 14. Performance–energy relationship (with $\alpha = 1.2$) of the prediction-based ANT scheme with $N_p = 3$ for filter bandwidths. (a) $\omega_b = 0.3\pi$. (b) $\omega_b = 0.5\pi$.

Fig. 14(a) shows the plot of *energy-performance* curves (with $\alpha = 1.2$) for a lowpass filter with $\omega_b = 0.3\pi$ employing the prediction-based ANT scheme. The degradation in performance due to voltage scaling is nearly linear in the absence of ANT. It was observed that, in presence of ANT, the performance degradation is restored for scaling factors upto 0.49 (as compared to conventional TCA) resulting in 81% energy savings over conventional TCA. The performance loss suffered is less than 0.5 dB. Beyond this point multiple errors occur and, hence, the assumption 2 in Section III-C is violated. Fig. 14(b) shows the *energy-performance* relationship for $\omega_b = 0.5\pi$. In this case, the possible energy savings over the conventional TCA operating at $V_{\rm dd\text{-}crit}$ is 78% with a performance loss of about 0.5 dB. Note that the energy savings compared to the proposed architecture operating at $V_{\rm dd\text{-}crit}$ is about 58% and the corresponding savings for filter bandwidth $= 0.3\pi$ 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 increases effectiveness.

D. 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

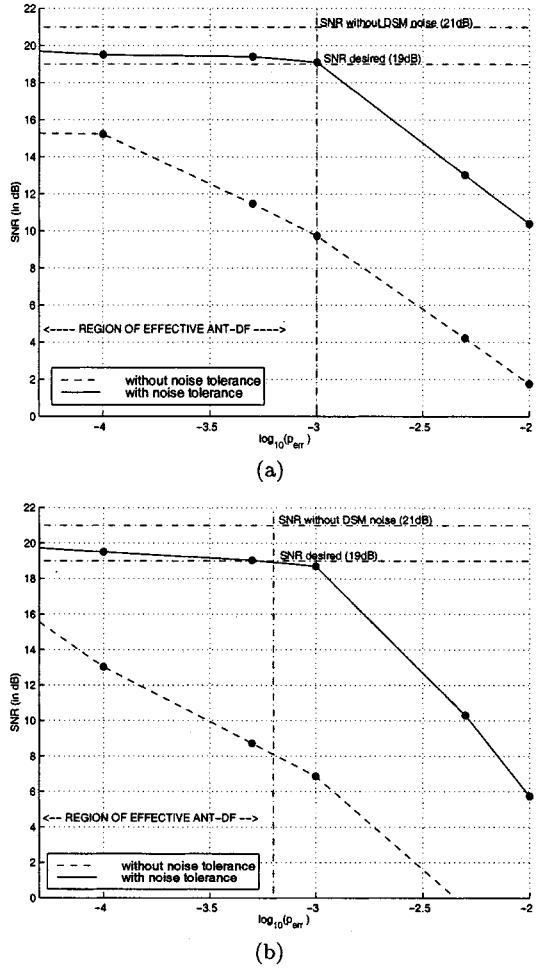

Fig. 15. ANT performance under DSM noise. (a) Performance of the proposed ANT scheme for filter bandwidth $\omega_b = 0.2\pi$ with predictor tap-length $N_p = 4$. (b) Performance of the proposed ANT scheme for filter bandwidth $\omega_b = 0.4\pi$ with predictor tap-length $N_p = 4$.

at the system level by flipping the output bits of the digital filter independently, with a fixed probability denoted by $p_{\rm err}$. 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 $\omega_b = 0.2\pi$ and 48 taps is shown in Fig. 15(a). As expected, without noise-tolerance, the degradation in performance increases with increase in $p_{\rm err}$ as expected. Also, the proposed scheme provides up to 10 dB improvement in performance. The SNR with ANT stays almost constant above 19 dB till $p_{\rm err} = 10^{-3}$ and then reduces sharply. In this range the probability of error is low enough that the assumption of infrequent errors (assumption 2 in Section III-C) is satisfied. Hence, error detection is performed effectively. As $p_{\rm err}$ is increased, this assumption is not satisfied any more and, hence, there is a rapid degradation in performance of the proposed algorithm. Similar results for a bandwidth of $\omega_b = 0.4\pi$ are shown in Fig. 15(b). Note that in both the cases, the proposed scheme is quite effective in combating DSM noise that is large enough to cause bit errors at a rate of 1 per 1000 samples.

V. Conclusion and Future Work

In this paper, we have proposed soft DSP for reduction in energy dissipation, where the supply voltage is reduced beyond that limited by the critical path delay of a given DSP architecture. The degradation in performance of the DSP algorithms is restored via *algorithmic noise tolerance*, where the signal statistics are exploited to develop low complexity error-control schemes. Note that the proposed approach is also a viable low-power technique in presence of DSM noise in future technologies, particularly for DSP and communications applications.

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. The ANT scheme presented in this paper is applicable to frequency selective finite impulse response (FIR) digital filters. However, the potential for energy reduction via soft DSP exists any DSP/ communications applications where the algorithmic performance is specified in terms of average metrics such as SNR or bit error rate (BER). Such applications include adaptive filtering for channel equalization, fast Fourier transform/inverse fast Fourier transform (FFT/IFFT) computations, image/video processing algorithms etc.. We also intend to study the impact of algorithmic noise-tolerance schemes in handling various DSM noise sources and study the effect of a noisy error-control block on the performance of ANT schemes for space-based applications that are prone to noise problems due to cosmic rays.

References

[1] N. R. Shanbhag and M. Goel, "Low-power adaptive filter architectures and their application to 51.84 mb/s ATM-LAN," *IEEE Trans. Signal Processing*, vol. 45, pp. 1276–1290, May 1997.

[2] P. E. Landman and J. M. Rabaey, "Architectural power analysis: The dual bit type method," *IEEE Trans. VLSI Syst.*, vol. 3, pp. 173–187, June 1995.

[3] S. Iman and M. Pedram, "An approach for multilevel logic optimization targeting low power," *IEEE Trans. Comput.-Aided Design*, vol. 15, pp. 889–901, 1996.

[4] R. K. Krishnamurthy and L. R. Carley, "Exploring the design space of mixed swing quadrail for low power digital circuits," *IEEE Trans. VLSI Syst.*, vol. 5, pp. 388–400, Dec. 1997.

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

[6] R. Hegde and N. R. Shanbhag, "A low-power phase splitting passband equalizer," *IEEE Trans. Signal Processing*, vol. 47, Mar. 1999.

[7] A. Chandrakasan and R. W. Brodersen, "Minimizing power consumption in digital CMOS circuits," *Proc. IEEE*, vol. 83, no. 4, pp. 498–523, Apr. 1995.

[8] S. M. Kang and Y. Leblebici, *CMOS Digital Integrated Circuits: Analysis and Design*. New York: McGraw-Hill, 1996.

[9] K. L. Shepard, "Conquering noise in deep-submicron digital IC's," *IEEE Design Test Comput.*, pp. 51–62, Jan./Mar 1998.

[10] R. Gonzalez *et al.*, "Supply and threshold voltage scaling for low-power CMOS," *IEEE J. Solid State Circuits*, vol. 32, pp. 1210–1216, Aug. 1997.

[11] P. J. Restle, J. Phillips, and I. Elfadel, "Interconnect in high speed designs: problems, methodologies and tools," in *ICCAD'98*, San Jose, CA, Nov. 1998, p. 4.

[12] K. L. Shepard and V. Narayanan, "Noise in deep submicron digital design," in *ICCAD'96*, San Francisco, CA, Nov. 1996, pp. 524–531.

[13] N. R. Shanbhag, "A mathematical basis for power-reduction in digital VLSI systems," *IEEE Trans. Circuits Syst.*, pt. II, vol. 44, pp. 935–951, Nov. 1997.

[14] R. Hegde and N. R. Shanbhag, "Energy efficiency in presence of deep submicron noise," in *ICCAD'98*, San Jose, CA, Nov. 1998.
[15] P. Banergee *et al.*, "Algorithm-based fault tolerance on a hypercube multiprocessor," *IEEE Trans. Comput.*, vol. 39, pp. 1132–1145, Sept. 1990.
[16] Y. Choi and M. Malek, "A fault-tolerant FFT processor," *IEEE Trans. Comput.*, vol. 37, pp. 617–621, May 1988.
[17] B. A. Schnaufer and W. K. Jenkins, "Adaptive fault tollerance for reliable LMS adaptive filtering," *IEEE Trans. Circuits Syst.*, pt. Part II, vol. 44, no. 12, pp. 1001–1014, Dec. 1997.
[18] J. M. Rabaey, *Digital Integrated Circuits: A Design Perspective*. Englewood Cliffs, NJ: Prentice-Hall, 1996.
[19] D. Sylvester and K. Keutzer, "Getting to the bottom of deep submicron," in ICCAD'98, San Jose, CA, Nov. 1998, pp. 203–211.
[20] J. D. Meindl, "Low power microelectronics: Retrospect and prospect," *Proc. IEEE*, vol. 83, no. 4, pp. 619–635, Apr. 1995.
[21] E. A. Vittoz, "Low-power design: Ways to approach the limits," in *ISSCC'94*, San Fransisco, CA, June 1994, pp. 14–18.
[22] M. Alidina *et al.*, "Precomputation-based sequential logic optimization for low power," *IEEE Trans. VLSI Syst.*, pp. 426–436, Dec. 1994.
[23] M. D. Ercegavoc and T. Lang, "Low-power accumulator (Correlator)," in *Proc. Int. Symp. Low-Power Electronic Design (ISLPED)*, San Francisco, CA, Aug. 1995, pp. 30–31.
[24] J. G. Proakis and D. G. Manolakis, *Digital Signal Processing: Principles, Algorithms and Applications*. Englewood Cliffs, NJ: Prentice-Hall, 1996.
[25] M. Xakellis and F. Najm, "Statistical estimation of the switching activity in digital circuits," in *Proc. 31st ACM/IEEE Design Automation Conf.*, San Diego, CA, June 1994, pp. 33–36.

Rajamohana Hegde (S'95) received the B.E. degree in electrical and electronics engineering from Manipal Institute of Technology, Manipal, India, and the M.S. degree in electrical engineering from Wright State University, Dayton, OH. He is currently working toward the Ph.D. degree in electrical engineering at the University of Illinois at Urbana-Champaign, Urbana.

From May 1999 to August 1999, he was a summer intern at the Circuits Research Laboratory of Microcomputer Research Laboratory, Intel Corporation, Hillsboro, OR, where his work invoved high-speed and noise-tolerant datapath design. His research interests are in the area of VLSI design of DSP and communications systems and design of noise-tolerance techniques for deep submicron VLSI systems.

Mr. Hegde has received the 2001 IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION SYSTEMS Best Paper Award from the IEEE Circuits and Systems Society.

Naresh R. Shanbhag (S'87–M'93) received the B.Tech. degree from the Indian Institute of Technology, New Delhi, India, in 1988, the M.S. degree from Wright State University, Dayton, OH, in 1990, and the Ph.D. degree from the University of Minnesota, in 1993, all in electrical engineering.

From July 1993 to August 1995, he was with AT&T Bell Laboratories, Murray Hill, NJ, in the Wide-Area Networks Group, where he was responsible of development of VLSI algorithms, architectures, and implementation of high-speed data communications transceivers. In particular, he was the Lead Chip Architect for AT&T's 51.84 Mb/s transceiver chips over twisted-pair wiring for asynchronous transfer mode (ATM)-LAN and broad-band access chip sets. In August 1995, he joined the Cordinated Science Laboratory and the Electrical and Computer Engineering Department at the University of Illinois at Urbana-Champaign, where he is now an Associate Professor. His research interests are in the design and integrated circuit implementation of low-power/high-performance signal processing and communications systems. He has published more than 70 journal/conference articles/book chapters and holds three U.S. patents on these topics. He is also a coauthor of the research monograph *Pipelined Adaptive Digital Filters* (Norwell, MA: Kluwer, 1994).

Dr. Shanbhag received the 2001 IEEE TRANSACTIONS ON VLSI SYSTEMS BEST PAPER AWARD 1999 IEEE Leon K. Kirchmayer Best Paper Award, the 1999 Xerox Faculty Award, the National Science Foundation CAREER Award in 1996, and the 1994 Darlington Best Paper Award from the IEEE Circuits and Systems Society. Since July 1997, he has been a Distinguished Lecturer for the IEEE Circuits and Systems Society. He is currently serving as an Associate Editor for the IEEE TRANSACTION ON CIRCUITS AND SYSTEMS: PART II. He is a member of the Design and Implementation of Signal Processing Systems (DISPS) Technical Committee of the IEEE Signal Processing Society, the VLSI Systems and Applications and VLSI in Communications Technical Committees of the IEEE Circuits and Systems Society. He has served on the technical program committees of various conferences including the 1998 and 1999 International Symposium on Low-Power Electronics and Design, the 1999 IEEE Workshop on Signal Processing Systems, and the 1999 IEEE International Symposium on Circuits and Systems.