LOW-POWER SIGNAL PROCESSING VIA ERROR-CANCELLATION

Lei Wang and Naresh R. Shanbhag
Coordinated Science Laboratory,
Department of Electrical and Computer Engineering,
University of Illinois at Urbana-Champaign,
1308 West Main Street, Urbana, IL 61801.
Email: {leiwang, shanbhag}@uiuc.edu

Abstract – We present an *algorithmic noise-tolerance* (ANT) technique for designing low-power DSP systems. The proposed technique achieves substantial energy savings via *voltage overscaling*, whereby the supply voltage is scaled beyond the minimum supply voltage $V_{dd-crit}$ at which the architecture operates correctly for a given throughput specification. The resulting input-dependent soft errors are corrected via a low-complexity error canceller and hence is referred to as *adaptive error-cancellation*. The trade-off between energy savings and algorithmic performance is illustrated by employing a reduced-order least mean square (LMS) algorithm to compensate for the design overhead. Simulation results in a $0.35\mu m$ CMOS technology demonstrate that the proposed technique achieves up to 73% energy savings in a multiuser communication scenario over present-day voltage-scaling, with a $3dB$ algorithmic performance loss. Moreover, a 40% energy reduction is obtained over conventional DSP systems without algorithmic performance degradation.

1 INTRODUCTION

The recent growth in demand for portable and wireless computing platforms has made power dissipation a great concern for general VLSI as well as digital signal processing (DSP) systems [1]–[3]. For a given silicon technology, supply voltage scaling is a widely employed low-power technique, as it enables a linear reduction in static power dissipation and a quadratic reduction in capacitive power dissipation [4]. However, scaling the supply voltage increases the propagation delay of a logic gate. Therefore, the achievable energy reduction via voltage scaling of a conventional DSP system is limited by a critical supply voltage $V_{dd-crit}$, which is determined by the system throughput requirements and the architecture. Scaling supply voltage beyond $V_{dd-crit}$ (*voltage overscaling* (VOS)) leads to soft errors at system output and a degradation in algorithmic performance. Recently, we have shown [5] that this degradation can be compensated for via *algorithmic noise-tolerance* (ANT) techniques. The resulting low-power DSP system is referred to as a *soft DSP* system.

This research was supported by the NSF grant CCR-0000987.

To guarantee substantial energy savings, an ANT technique should be computationally efficient yet effective in mitigating algorithmic degradation. In order to derive the conditions under which ANT can lead to energy reductions (over present-day voltage scaling), consider the dynamic power dissipation of a DSP system operating at $V_{dd-crit}$, i.e., conventional voltage scaling,

$$P_{orig} = C_{eff} {V_{dd-crit}}^2 f, \tag{1}$$

where C_{eff} denotes the effective (lumped) switching capacitance that accounts for the transition activities, voltage swing ranges and load/parasitic capacitances at each circuit node, and f is the clock frequency.

The power dissipation of the corresponding soft DSP system is given by

$$P_{soft} = (C_{eff} + C_{ANT}) \left(\frac{V_{dd-crit}}{k_v} \right)^2 f, \tag{2}$$

where C_{ANT} is the effective (lumped) switching capacitance of the overhead circuitry due to ANT, and $k_v > 1$ is the VOS factor. Thus, $P_{soft} < P_{orig}$ provided

$$C_{ANT} < C_{eff}({k_v}^2 - 1). \tag{3}$$

In practice, the condition above is easily satisfied by making C_{ANT} as small as possible (reducing the overhead of the ANT scheme) and/or by making k_v as large as possible. Past work [5] has reported a *prediction-based error-control scheme* which achieves 60% – 80% energy savings (in a frequency-selective filtering application) over present-day voltage scaling while being subject to $0.5dB - 1dB$ output *signal-to-noise ratio* (*SNR*) loss. In this paper, we will present a new ANT technique, referred to as *adaptive error-cancellation*, by exploiting the inherent input-dependent nature of soft errors caused by voltage overscaling. The prediction-based ANT scheme [5] works well if error frequency is small (up to 5%). The proposed ANT scheme overcomes this limitation.

In section 2, we briefly review the past work on the use of soft DSP for low-power digital filtering. In section 3, we present a new ANT scheme – the *adaptive error-cancellation* scheme and a practical architecture for its implementation. Simulation results for frequency selective filters are presented and evaluated in section 4.

2 LOW-POWER SIGNAL PROCESSING VIA ALGORITHMIC NOISE-TOLERANCE

In this section, we introduce the soft DSP framework and the prediction-based error-control scheme for low-power digital filtering. We also provide the motivation for the proposed error-cancellation scheme.

2.1 Soft DSP

In general, a DSP system is designed subject to an application-specific throughput requirement. Specifically, the critical path delay T_{cp} of the DSP architecture should be less than or equal to the sample period T_s of the application, i.e.,

$$T_{cp} \leq T_s. \tag{4}$$

On the other hand, the propagation delay τ_d of a logic gate can be expressed as a function of supply voltage, as [6]

$$\tau_d = \frac{C_L V_{dd}}{K(V_{dd} - V_t)^\alpha}, \tag{5}$$

where C_L is the load capacitance, V_{dd} is the supply voltage, V_t is the device threshold voltage, K is the transistor transconductance, and α denotes the velocity saturation index.

Thus, scaling the supply voltage V_{dd} for energy reduction increases the gate delay τ_d and consequently the critical path delay T_{cp} of the system. The supply voltage at which $T_{cp} = T_s$ is referred to as the critical supply voltage and is denoted as $V_{dd-crit}$. Note that present-day voltage scaling stops at the point where $V_{dd} = V_{dd-crit}$. Overscaling supply voltage beyond $V_{dd-crit}$ results in output errors if the critical paths are excited, i.e., soft errors occur. This degrades algorithmic performance measures such as the output *signal-to-noise ratio* (*SNR*), i.e.,

$$SNR = 10\log_{10}\left(\frac{{\sigma_s}^2}{{\sigma_n}^2 + {\sigma_e}^2}\right) < 10\log_{10}\left(\frac{{\sigma_s}^2}{{\sigma_n}^2}\right), \tag{6}$$

where ${\sigma_s}^2$, ${\sigma_n}^2$ and ${\sigma_e}^2$ are the variances of the desired signal, signal noise and the soft output error (due to VOS), respectively.

2.2 Energy Savings

Algorithmic noise-tolerance (ANT) techniques are able to compensate for the soft errors, thereby achieving substantial energy savings with a marginal performance loss. The prediction-based error-control scheme in [5] (see Fig. 1(a)) is suitable for DSP architectures with a path delay distribution and input statistics such that the soft errors due to VOS are of large magnitude and occur infrequently (see Fig. 1(b)). This condition is easily met for narrowband filters implemented using delay-imbalanced arithmetic units such as ripple-carry adders and array multipliers. Another factor that affects the frequency of soft errors is the input signal statistics. For example, a wideband input signal is likely to excite longer paths more frequently as compared with a narrowband input signal. Thus, it is expected that the error frequency increases rapidly with reduction in V_{dd} for wideband inputs, thereby requiring more complex ANT schemes which in turn can defeat the original goal of energy reduction.

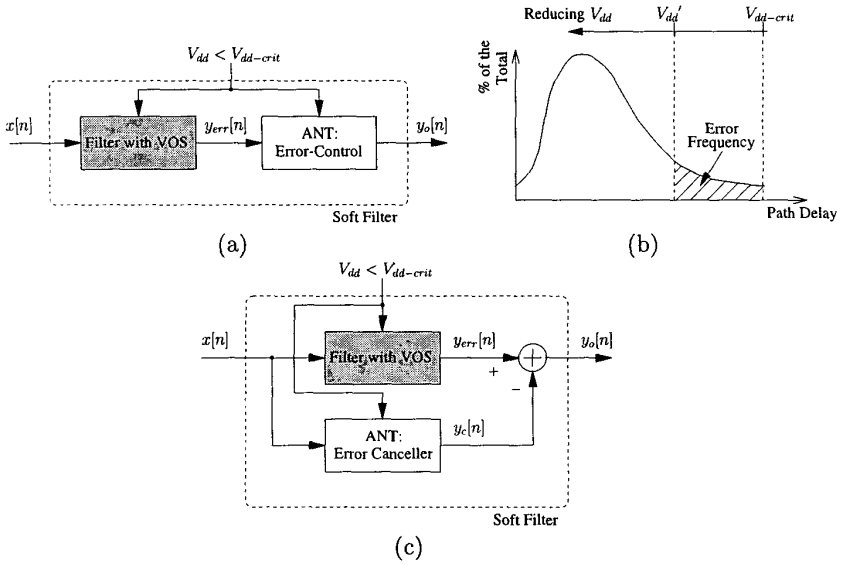

Figure 1: Soft DSP framework: (a) the error-control scheme in [5], (b) typical path delay distribution of a ripple-carry adder and (c) the proposed error-cancellation scheme.

In this paper, we will present a new ANT scheme – the *adaptive error-cancellation* (see Fig. 1(c)) which can tolerate high error frequencies that could be due to wideband inputs or due to excessive VOS. Note that in some systems such as multiuser wireless communications, a wideband input signal is composed of several bandlimited signals occupying adjacent frequency bands, e.g., frequency-division multiplexed (FDM) systems [7]. The task of digital signal processing is then to extract one of the embedded signals of interest. In the following sections, we will study the performance of the proposed ANT scheme in the context of FDM systems.

3 ADAPTIVE ERROR-CANCELLATION

Soft errors due to VOS are in fact a deterministic function of the input signal and the architecture of the underlying arithmetic circuits (e.g., adders and multipliers). Thus, one could conceive of a "logic-level" error-control scheme that monitors the input, compares it with a predetermined set of input sequences that are known to cause an error, and then sends the corresponding precomputed output through an output MUX that bypasses the normal output. However, the relationship between the input signal and the soft errors of a filter implementation is extremely complicated making us resort to ANT schemes. One possible approach is to model the input signal as a stochastic process and then exploit

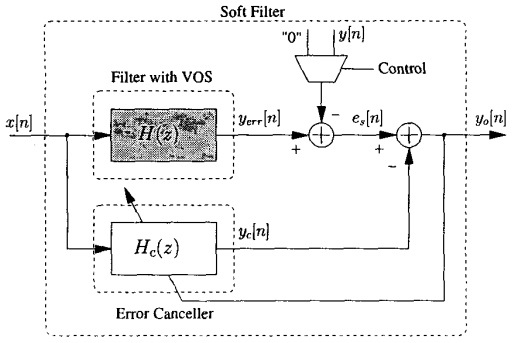

Figure 2: The proposed ANT scheme based on adaptive error-cancellation.

the cross-correlation between the input signal and the soft errors. This can be very efficiently done using adaptive filtering techniques [8]. As the soft errors are input-dependent, they can be regarded as an "echo" of the input signal and hence echo-cancellation schemes used in modern communication systems can be employed to restore the system performance. The proposed adaptive error-cancellation technique is one such scheme.

3.1 Error-Cancellation Algorithm

The proposed technique is shown in Fig. 2 where an adaptive error canceller $H_c(z)$ is employed to generate a (statistical) replica of the soft errors induced at the output of filter $H(z)$ and then subtract it from the output. In the presence of soft errors due to VOS, the noisy output $y_{err}[n]$ of the N-tap filter $H(z)$ can be expressed as

$$\begin{aligned} y_{err}[n] &= \sum_{k=0}^{N-1} h[k]x[n-k] + e_s[n] \\ &= y[n] + e_s[n] \end{aligned} \tag{7}$$

where $e_s[n]$ denotes the soft error in the output, $y[n]$ is the error-free output, $h[k]$ is the k^{th} coefficient of the filter, and $x[n-k]$ is the k^{th} delayed input sample.

For a given implementation of $H(z)$, the soft error $e_s[n]$ is mainly determined by the input samples $x[n], x[n-1], \cdots, x[n-N+1]$. The error canceller $H_c(z)$ adaptively estimates the value of $e_s[n]$, denoted by $y_c[n]$ and given by

$$y_c[n] = \sum_{k=0}^{N_c-1} w_k x[n-k], \tag{8}$$

where $\mathbf{w} = \{w_0, w_1, \cdots, w_{N_c-1}\}$ and N_c are the coefficient vector and the order, respectively, of the error canceller. At this time we assume $N_c = N$ to account

for the error contributions from all the taps of $H(z)$. Thus, the error-canceller $H_c(z)$ is as complex as $H(z)$. In the next subsection, however, we will derive a simplified algorithm which allows significant reduction in the complexity of the error canceller.

Let $e[n]$ be the estimation error, defined as

$$e[n] = e_s[n] - y_c[n]. \tag{9}$$

The commonly employed *least mean square* (LMS) algorithm that minimizes the *mean-square error* (MSE, defined as $\mathcal{J}[n] = E(\mid e[n] \mid^2)$, where $E(\cdot)$ denotes the expectation operator) is given by [8]

$$y_c[n] = \sum_{k=0}^{N_c-1} \widehat{w}_k[n-1]x[n-k], \tag{10}$$

$$\widehat{w}_k[n] = \widehat{w}_k[n-1] + \mu e[n]^* x[n-k], \tag{11}$$

where $\widehat{\mathbf{w}}[n] = \{\widehat{w}_0[n], \widehat{w}_1[n], \cdots, \widehat{w}_{N_c-1}[n]\}$ is the estimate of tap-weight vector of $H_c(z)$, $e[n]^*$ is the complex conjugate of $e[n]$, and μ is the step-size.

Note that the above LMS adaptive algorithm describes the computations during the training phase. The noisy output $y_{err}[n]$ is obtained by passing a sequence of predefined training data through the filter $H(z)$ subject to VOS, whereas the error-free output $y[n]$ is calculated off-line. After the tap-weight vector $\widehat{\mathbf{w}}[n]$ has converged, we can switch off the weight-update operation (11) and subtract $y_c[n]$ in (10) directly from the output $y_{err}[n]$ in (7) to cancel the soft errors. These operations can be employed to autocalibrate a soft DSP integrated circuit in the field so as to be able to account for variations in the process, temperature and statistical properties of the input signal.

3.2 Algorithm Simplification

The hardware complexity of the proposed ANT scheme can be significantly reduced by exploiting the fact that performance degradation is dominated by errors from a few of the filter taps. These taps have coefficients with large magnitudes and therefore contribute more to the output error. Assuming that soft errors due to different filter taps are statistically independent, the total output error energy $\mathcal{E}_s$ due to VOS can be expressed as

$$\mathcal{E}_s = \sum_{k=0}^{N-1} \mathcal{E}_k, \tag{12}$$

where $\mathcal{E}_k$ is the error energy due to the k^{th} tap of $H(z)$. Note that the above assumption about independence is not essential for correct operation. It only assists us in developing a low-complexity ANT scheme. Also, this assumption is reasonable for wideband input signals because such signals have a low correlation coefficient.

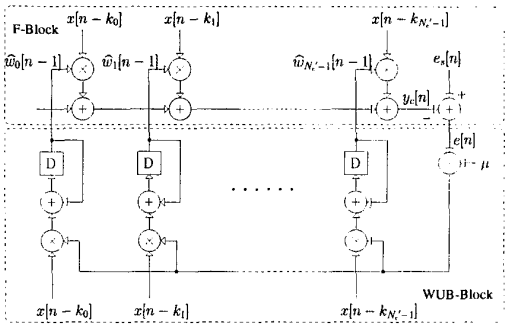

Figure 3: Reduced-order LMS error canceller architecture.

We select N_c' dominant error components from $\mathcal{E}_0, \mathcal{E}_1, \cdots, \mathcal{E}_{N-1}$, subject to an error-suppression requirement given by

$$\sum_{i=0}^{N_c'-1} \mathcal{E}_{k_i} = \mathcal{K}_e \mathcal{E}_s, \tag{13}$$

where $\mathcal{E}_{k_i}$ is the error energy due to the input $x[n-k_i]$ at the $k_i{}^{th}$ tap, and $0 < \mathcal{K}_e \leq 1$ is the *error-suppression factor* which trades off the achievable energy savings and performance enhancement. Note that the LMS algorithm shown in (10)–(11) can be regarded as a specific case with $\mathcal{K}_e = 1$ and $N_c' = N$, where we obtain the best algorithmic performance at the expense of a large design overhead. According to (13), the reduced-order LMS algorithm for error-cancellation can be expressed as

$$y_c[n] = \sum_{i=0}^{N_c'-1} \widehat{w}_i[n-1]x[n-k_i], \tag{14}$$

$$e[n] = y_{err}[n] - y[n] - y_c[n], \tag{15}$$

$$\widehat{w}_i[n] = \widehat{w}_i[n-1] + \mu e[n]^* x[n-k_i]. \tag{16}$$

From (14)–(16) we can see that the dominant error components (a fraction $\mathcal{K}_e$ of the total) from the $k_0{}^{th}, k_1{}^{th}, \cdots, k_{N_c'-1}{}^{th}$ tap have been accounted for. Thus, the order of the error canceller is reduced to N_c'. Simulation results in the next section demonstrate a $4X$ reduction in hardware complexity as compared to the conventional LMS algorithm (10)–(11), while the performance loss is negligible. Note that in most cases $k_0, k_1, \cdots, k_{N_c'-1}$ are consecutive tap indices. This simplifies the design of the error canceller and easily satisfies the constraint in (3).

Fig. 3 shows the architecture of the proposed reduced-order error canceller. It is worth mentioning that the weight-update (**WUD**) block is switched off after convergence. Hence, the energy overhead involves the power dissipation due to the N_c' taps in the filter (**F**) block. We also need to mention that the

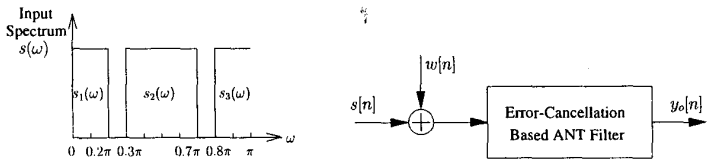

Figure 4: Simulation setup: (a) input signal spectrum and (b) lowpass filtering via the proposed ANT scheme.

reduced-order error canceller is error-free with VOS, as its critical path is much smaller than that of the filter $H(z)$.

4 SIMULATION RESULTS

In this section, we present the simulation results of the proposed ANT scheme. We first describe the simulation setup and then discuss the achievable energy savings and algorithmic performance.

4.1 Simulation Setup

Fig. 4(a) illustrates the spectrum of the input signal $s[n]$ which consists of a baseband signal $s_1[n]$ occupying the $[0, 0.2\pi]$ band and two bandpass signals $s_2[n]$ and $s_3[n]$ in the adjacent bands. This input signal emulates a frequency-division multiuser communication system [7]. The input signal $s[n]$ is also corrupted by a white noise signal $w[n]$. All the signals $s_1[n]$, $s_2[n]$, $s_3[n]$ and noise $w[n]$ are statistically independent of each other.

The goal of the receiver signal processing is to extract the primary signal $s_1[n]$. This is accomplished by passing the input signal through a lowpass filter to suppress the out-of-band signals and noise. We employ the proposed ANT scheme to perform this frequency selective filtering (see Fig. 4(b)). A 32-tap lowpass filter with passband edge frequency $\omega_p = 0.2\pi$ and stopband edge frequency $\omega_s = 0.3\pi$ is implemented. All the simulations employ an 8-bit sign-magnitude MAC structure [5], while the **WUD**-block is 16-bit precision.

The performance of the proposed ANT scheme is evaluated under various sub-critical supply voltages. A logic level simulation is used to calculate the number of full-adder delays on every path to the filter output given a sequence of inputs. Thus, all paths and not just the critical paths are included. The corresponding circuit delay in a $0.35\mu m$ CMOS technology is then obtained by computing the full-adder delay (5) when driving another full-adder as fan-out. The output delays are compared with the sample period T_s as in (4). If the constraint is violated, the corresponding output will not be able to settle to its new value but instead retain its previous value thereby resulting in an output error. The output *SNR* is calculated by averaging over the entire input data set. The energy dissipation is obtained via the gate-level simulation tool MED [9] for

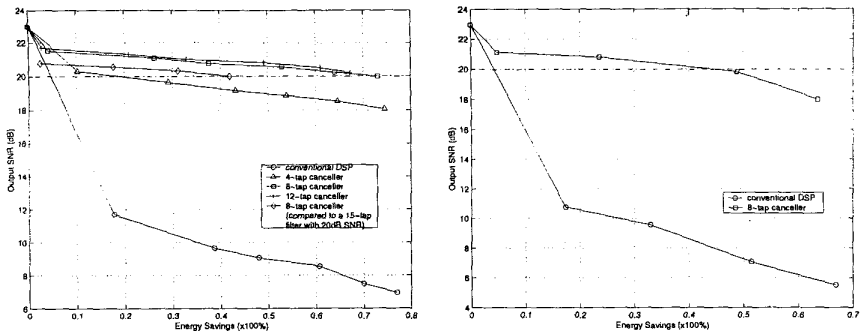

Figure 5: Performance of the proposed ANT scheme: (a) serial and (b) pipelined architecture.

a $0.35\mu m$ CMOS technology. This energy dissipation represents the operations when the **WUD**-block is switched off. The average energy savings E_{sav} is given by

$$E_{sav} = \frac{E_{orig} - E_{ANT}}{E_{orig}} \times 100\%, \tag{17}$$

where E_{orig} is the energy dissipation of a conventional implementation with $V_{dd} = V_{dd-crit}$, and E_{ANT} is the energy dissipation of the proposed ANT scheme.

4.2 Performance Comparison

In this example, the critical supply voltage $V_{dd-crit}$ was found to be $3.3V$. Thus, the conventional architecture operates at this voltage. Fig. 5(a) shows the energy-performance relationship of the proposed ANT scheme employing a 4-tap, 8-tap and 12-tap error canceller, respectively. Note that due to the presence of the adjacent-band signals, output errors occur frequently as soon as the V_{dd} is reduced. Thus, a sharp *SNR* drop is observed for the conventional implementation. Further V_{dd} scaling leads to erroneous outputs at the higher bits, thereby reducing the output *SNR* further. There is a trade-off between the achievable energy savings and algorithmic performance for the proposed adaptive error-cancellation scheme. Energy savings of 20% ($V_{dd} = 2.7V$), 73% ($V_{dd} = 1.7V$) and 68% ($V_{dd} = 1.7V$) are achieved at the expense of a $3dB$ loss in output SNR by using a 4-tap, 8-tap and 12-tap error canceller. Thus, an 8-tap error canceller gives the best energy-performance trade-off for this application. Moreover, the proposed soft filter (with an 8-tap canceller) achieves a 40% energy savings as compared with a 15-tap conventional filter that is designed with $20dB$ output *SNR* and operated at $V_{dd-crit}$.

Pipelined architectures [10] result in shorter critical paths and smaller delay-imbalance and thus do not favor soft DSP in general. Fig. 5(b) shows energy

savings vs. algorithmic performance of a three-level pipelined MAC architecture. In spite of this fact, we observe a 45% energy savings for an 8-tap error canceller at a $3dB$ performance degradation.

5 CONCLUSIONS

In this paper, we have proposed an adaptive error-cancellation algorithm for designing low-power soft DSP systems. In particular, we optimize the proposed error-cancellation scheme by employing a reduced-order LMS algorithm to compensate for the performance degradation due to voltage overscaling. Future work is being directed towards the application of the proposed noise-tolerance technique to other communication systems, an analytical study of this technique and an IC implementation. Soft DSP provides a new direction for research in the design of energy-efficient DSP algorithms and architectures, whereby DSP algorithms, architectures and circuit properties are jointly optimized to push the limits of energy reduction.

References

[1] A. P. Chandrakasan and R. W. Brodersen, "Minimizing power consumption in digital CMOS circuits," *Proceedings of the IEEE*, vol. 83, pp. 498-523, April 1995.

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

[3] J. T. Ludwig, S. H. Nawab and A. P. Chandrakasan, "Low-power digital filtering using approximate processing," *IEEE J. Solid-State Circuits*, vol. 31, pp. 395-400, March 1996.

[4] R. Gonzalez, B. M. Gordon, and M. A. Horowitz, "Supply and threshold voltage scaling for low power CMOS," *IEEE J. Solid-State Circuits*, vol. 32, pp. 1210-1216, August 1997.

[5] R. Hegde and N. R. Shanbhag, "Energy-efficient signal processing via algorithmic noise-tolerance," *Proc. of Intl. Symp. on Low-Power Electronics and Design*, San Diego, 1999.

[6] J. M. Rabaey, *Digital Integrated Circuits: A Design Perspective*, Prentice Hall, 1996.

[7] J. G. Proakis, *Digital Communications*, McGraw-Hill, 1995.

[8] S. Haykin, *Adaptive Filter Theory*, Prentice Hall, 1996.

[9] M. G. Xakellis and F. N. Najm, "Statistical estimation of the switching activity in digital circuits," *Design Automation Conf.*, pp. 728-733, June 1994.

[10] C. Asato, C. Ditzen and S. Dholakia, "A data-path multiplier with automatic insertion of pipeline stages," *IEEE J. Solid-State Circuits*, vol. 25, pp. 383-387, April 1990.