Low-Power AEC-Based MIMO Signal Processing for Gigabit Ethernet 1000Base-T Transceivers*

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
[leiwang, shanbhag]@uiuc.edu

ABSTRACT

Presented in this paper is a low-power technique, denoted as MIMO-AEC, to reduce energy dissipation in multi-input-multi-output (MIMO) signal processing systems. The proposed technique extends a previously proposed adaptive error-cancellation (AEC) technique to MIMO systems by employing an algorithm transformation denoted as MIMO-DECOR. The purpose of MIMO-DECOR is to reduce complexity by exploiting correlations inherent in MIMO systems, thereby improving the effectiveness of AEC. We employ the MIMO-AEC in the design of a low-power Gigabit Ethernet 1000Base-T device. Simulation results demonstrate 44.3% – 25.2% overhead reduction due to MIMO-DECOR and 69.1% – 64.2% energy savings over conventional implementations with no loss in algorithmic performance.

1. INTRODUCTION

Power reduction is essential for high performance signal processing and communication systems such as Gigabit Ethernet, next generation digital subscriber loop (DSL) and future 3G wireless. With feature sizes being reduced towards the deep submicron (DSM) regime [1], the emergence of DSM noise [2] as well as increasingly stringent requirements on performance have raised concerns about our ability to maintain the reliability in an affordable manner and hence to ensure the performance/energy-efficiency trends in future CMOS technologies. Our past research in *algorithmic noise-tolerance* (ANT) [3], [4] enables reliability and energy-efficiency to be jointly addressed in order to push the limits of energy-efficiency.

ANT techniques have the potential of approaching the lower bounds on energy dissipation where aggressive design techniques create DSM noise-like behavior at the algorithmic level. One such example is the *voltage overscaling* (VOS) scheme [4] for dedicated DSP implementations. For the purpose of illustration, we assume that the critical path delay of the filter in Fig. 1 is equal to $22T_{a,V_{dd}}$, where $T_{a,V_{dd}}$ is the full adder delay at a voltage of V_{dd}. Scaling supply voltage reduces energy dissipation but increases circuit delay. The minimal voltage $V_{dd-crit}$ necessary for correct operation is thus determined by $T_s = 22T_{a,V_{dd-crit}}$, where T_s is the sample period of the input. If the supply voltage is overscaled to $V_{dd-sub} = V_{dd-crit}/k_{vos}$ ($V_{dd-sub} \geq 2 \sim 3V_t$, where V_t is the threshold voltage of MOS transistors and $k_{vos} > 1$) such that $T_{a,V_{dd-sub}} = 1.2T_{a,V_{dd-crit}}$, then $T_s \geq 18T_{a,V_{dd-sub}}$. This indicates that, while the filter still functions correctly at the lower LSBs, the top four MSBs of the output will be in error provided input patterns exciting the critical paths and other longer paths are applied. Therefore, the energy benefits of VOS can be reaped by developing low-complexity ANT techniques that correct output errors. This approach to low-power design is referred to as *soft DSP*, as shown in Fig. 1. In comparison with conventional low-power techniques, soft DSP avoids increase in signal latency and complex clocking. Furthermore, it works best in the context of DSP and communication systems, where system performance metrics are measured in terms of *signal-to-noise ratio* (SNR) and/or *bit-error rate* (BER).

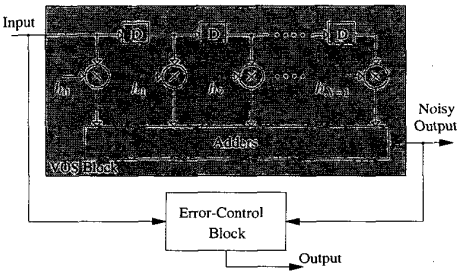

Figure 1: ANT-based low-power filtering scheme.

Past work [3] has reported an *adaptive error-cancellation* (AEC) scheme as a practical ANT technique. In this paper, we propose the MIMO-AEC technique for the design of low-power multi-input-multi-output (MIMO) systems. MIMO signal processing is employed in many modern-day DSP and communication systems, such as multi-user detection

*This research was supported by the NSF grant CCR-0000987 and CCR-9979381.

ISLPED'01, August 6-7, 2001, Huntington Beach, California, USA.

in wireless applications, interference suppression in Gigabit Ethernet 1000Base-T transceivers, etc.. Power reduction is important to MIMO systems as they usually require intensive filtering computations. The proposed MIMO-AEC technique exploits the inherent correlations in MIMO systems via MIMO-decorrelating (MIMO-DECOR) transform to improve the energy-efficiency of AEC. We employ the MIMO-AEC technique in the design of a low-power Gigabit Ethernet 1000Base-T device. Simulation results demonstrate $44.3\% - 25.2\%$ overhead reduction due to MIMO-DECOR and $69.1\% - 64.2\%$ energy savings over conventional implementations at the same algorithmic performance.

In section 2, we review our past work on energy-optimum AEC. In section 3, we present the MIMO-AEC technique. Simulation results of 1000Base-T transceivers are provided in section 4.

2. LOW-POWER DIGITAL FILTERING VIA AEC

The previously proposed AEC technique [3] employs the fact that the soft error $e_s[n]$ due to VOS at the output of an N-tap filter $H(z)$ is induced by the input samples $x[n], x[n-1], \cdots, x[n-N+1]$. Thus, we can generate a statistical replica of soft errors from these input samples and then subtract it from the output to mitigate performance degradation. This can be done using the popular *least mean square* (LMS) algorithm [5]. In addition, we proposed an energy-optimum AEC scheme as given below

$$\widehat{e}_s[n] = \sum_{k=0}^{N-1} b_k h_{c,k}[n-1]x[n-k], \tag{1}$$

$$e_c[n] = e_s[n] - \widehat{e}_s[n], \tag{2}$$

$$h_{c,k}[n] = h_{c,k}[n-1] + \mu_c e[n]^* x[n-k], \tag{3}$$

where $h_{c,k}$'s are the coefficients of the error canceller $H_c(z)$, $\widehat{e}_s[n]$ is the estimate of soft error $e_s[n]$, $e_c[n]$ is the residual soft error after AEC, and μ_c is the step size. The vector $\{b_0, b_1, \cdots, b_{N-1}\}$, $b_k \in \{0,1\}$, determines the trade-off between system performance and achievable energy savings. The energy-optimum solution that minimizes the overall energy dissipation while being subject to a performance constraint can be derived via the *Lagrange multiplier* method [6] as

$$b_j^* = \begin{cases} 1 & \text{if } \frac{\mathcal{E}_{F,j}}{{h_{c,j}}^2 {\sigma_x}^2} < \lambda^*, \\ 0 & \text{if } \frac{\mathcal{E}_{F,j}}{{h_{c,j}}^2 {\sigma_x}^2} \geq \lambda^*, \end{cases} \tag{4}$$

where λ^* is the optimum *sensitivity vector* of the Lagrange multiplier, ${\sigma_x}^2$ is the variance of input $x[n]$, $\mathcal{E}_{F,j}$ is the energy dissipation due to the j^{th} tap of the error canceller $H_c(z)$ and can be estimated as a function of the coefficient $h_{c,j}$. From (4), $b_j^* = 1$ indicates that the j^{th} tap of $H_c(z)$ is powered up, otherwise $b_j^* = 0$.

It was shown [3] that the tap-length of the energy-optimum AEC, given by $N_c^{\mathbf{opt}} = \sum_{j=0}^{N-1} b_j^*$, is much smaller than N. Hence, the energy overhead incurred in AEC can be easily compensated by the energy savings due to VOS. In addition, the AEC block has a shorter critical path, thus being error-free at V_{dd-sub}. The proposed AEC technique can be employed in the design of low-power frequency-selective filters and adaptive filters. For AEC-based adaptive filters, we assume that 1.) the original adaptive filter $H_a(z)$ and the error canceller $H_c(z)$ are calibrated separately (in fact, this is necessary as otherwise the estimation errors due to $H_a(z)$ and $H_c(z)$ become indistinguishable, thereby preventing the convergence of the two filters) and 2.) the **WUD** block is powered-down during the steady-state filtering operation. Hence, a common weight-update (**CWUD**) block can be shared by both filters in order to reduce the hardware overhead of AEC. The operation of AEC-based adaptive filters includes three phases, as described below with reference to Fig. 2:

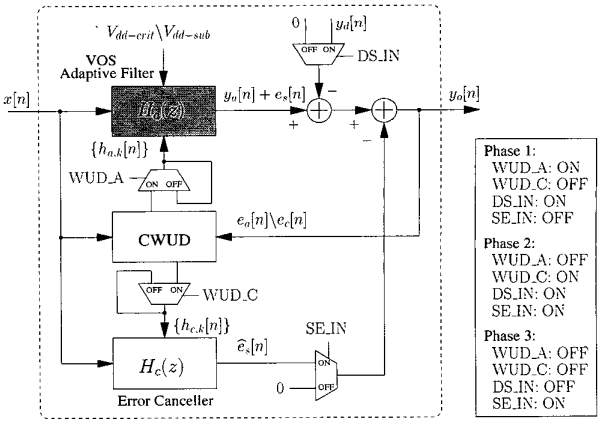

Figure 2: AEC-based adaptive filter.

1). Filter Calibration Phase: During this phase, the supply voltage is set to $V_{dd-crit}$ and a predefined training signal is fed into $H_a(z)$. The coefficients $h_{a,k}$'s of $H_a(z)$ get updated by the error signal $e_a[n]$ between the filter output $y_a[n]$ and a precomputed desired signal $y_d[n]$, as given below

$$y_a[n] = \sum_{k=0}^{N-1} h_{a,k}[n-1]x[n-k], \tag{5}$$

$$e_a[n] = y_a[n] - y_d[n], \tag{6}$$

$$h_{a,k}[n] = h_{a,k}[n-1] + \mu_a e_a[n]^* x[n-k]. \tag{7}$$

2). AEC Calibration Phase: During this phase, the supply voltage is overscaled to V_{dd-sub}. Thus, soft errors start to appear at the output of $H_a(z)$. The coefficients $h_{c,k}$'s of the energy-optimum error canceller $H_c(z)$ are computed according to (1)−(3). Note that $e_s[n]$ in (2) also contains residual error $e_a[n]$ (6) from the first phase. However, as $e_a[n]$ is much smaller than $e_s[n]$ (as $H_a(z)$ has already converged), it has a minor effect on the optimum configuration of $H_c(z)$.

3). Soft Filtering Phase: After $H_c(z)$ has converged, the supply voltage is kept at V_{dd-sub}. The filter output $y_o[n]$ can be expressed as

$$y_o[n] = y_a[n] + e_s[n] - \widehat{e}_s[n], \tag{8}$$

where $y_a[n]$ and $\widehat{e}_s[n]$ are given by (5) and (1), respectively. This starts the steady-state filtering operation where significant energy reduction is achieved via VOS while the required algorithmic performance is guaranteed by the AEC. It was shown [3] that the energy-optimum AEC is well-suited for wideband signal processing. In particular, we obtained $43\% - 71\%$ energy savings in the context of single-input-single-output (SISO) systems.

3. MIMO-AEC FOR LOW-POWER MIMO SIGNAL PROCESSING

In this section, we present the MIMO-AEC technique for low-power MIMO systems. A matrix representation of MIMO systems is given in section 3.1 and the MIMO-AEC is proposed in section 3.2. In section 3.3, we derive the MIMO-DECOR transform for practical MIMO systems.

3.1 MIMO Model

Consider a generic p-input, q-output MIMO system composed of N-tap filters and expressed in matrix form as

$$\begin{bmatrix} y_1[n] \\ y_2[n] \\ \cdots \\ y_q[n] \end{bmatrix} = \begin{bmatrix} \underline{\mathbf{h}}_{11} & \underline{\mathbf{h}}_{12} & \cdots & \underline{\mathbf{h}}_{1p} \\ \underline{\mathbf{h}}_{21} & \underline{\mathbf{h}}_{22} & \cdots & \underline{\mathbf{h}}_{2p} \\ \cdots & \cdots & \cdots & \cdots \\ \underline{\mathbf{h}}_{q1} & \underline{\mathbf{h}}_{q2} & \cdots & \underline{\mathbf{h}}_{qp} \end{bmatrix} \otimes \begin{bmatrix} \underline{\mathbf{x}}_1[n] \\ \underline{\mathbf{x}}_2[n] \\ \cdots \\ \underline{\mathbf{x}}_p[n] \end{bmatrix}, \quad (9)$$

where $\underline{\mathbf{x}}_j[n] = [x_j[n], x_j[n-1], \cdots, x_j[n-N+1]]^T$ is the j^{th} input sequence, $\underline{\mathbf{h}}_{ij} = [h_{ij}[0], h_{ij}[1], \cdots, h_{ij}[N-1]]$ is the impulse response of the filter having the j^{th} input and the i^{th} output, $y_i[n] = \sum_{j=1}^{p} \underline{\mathbf{h}}_{ij} \otimes \underline{\mathbf{x}}_j[n]$ is the i^{th} output, and "$\otimes$" denotes the element-by-element convolution operation, i.e., $\underline{\mathbf{h}}_{ij} \otimes \underline{\mathbf{x}}_j[n] = \sum_{k=0}^{N-1} h_{ij}[k] x_j[n-k]$. A special case of (9) is a system with a diagonal transfer matrix (i.e., $\underline{\mathbf{h}}_{ij} = \underline{\mathbf{0}}$ for $i \neq j$) representing independent SISO filtering operations, such as the channel equalizer in Gigabit Ethernet 1000Base-T transceivers (see section 4).

Henceforth, we refer to the filter with coefficient vector $\underline{\mathbf{h}}_{ij}$ as filter $\underline{\mathbf{h}}_{ij}$. Note that for practical MIMO systems the filters in (9) usually have correlated time-frequency characteristics. For example, the interference suppression scheme in a 1000Base-T device (see section 4) contains twelve NEXT cancellers, among which every three cancellers have the same input. These three cancellers have similar impulse responses, as they are designed to cancel three similar NEXT interferences which are induced by the same input signal on three spatially correlated crosstalk paths.

3.2 MIMO-AEC

We assume that all the filters in (9) operate in parallel and have matched critical path delays. Energy reduction via VOS induces soft-error degradation at all the filter outputs. This necessitates a bank of error cancellers, of which the Weiner-Hopf solution [5] is given by

$$\begin{bmatrix} \underline{\mathbf{h}}_{c,11} & \underline{\mathbf{h}}_{c,12} & \cdots & \underline{\mathbf{h}}_{c,1p} \\ \underline{\mathbf{h}}_{c,21} & \underline{\mathbf{h}}_{c,22} & \cdots & \underline{\mathbf{h}}_{c,2p} \\ \cdots & \cdots & \cdots & \cdots \\ \underline{\mathbf{h}}_{c,q1} & \underline{\mathbf{h}}_{c,q2} & \cdots & \underline{\mathbf{h}}_{c,qp} \end{bmatrix} = E\left(\begin{bmatrix} e_{s1}[n] \\ e_{s2}[n] \\ \cdots \\ e_{sq}[n] \end{bmatrix} \begin{bmatrix} \underline{\mathbf{x}}_1[n] \\ \underline{\mathbf{x}}_2[n] \\ \cdots \\ \underline{\mathbf{x}}_p[n] \end{bmatrix}^T \right) \cdot E\left(\begin{bmatrix} \underline{\mathbf{x}}_1[n] \\ \underline{\mathbf{x}}_2[n] \\ \cdots \\ \underline{\mathbf{x}}_p[n] \end{bmatrix} \begin{bmatrix} \underline{\mathbf{x}}_1[n] \\ \underline{\mathbf{x}}_2[n] \\ \cdots \\ \underline{\mathbf{x}}_p[n] \end{bmatrix}^T \right)^{-1}, \quad (10)$$

where $\underline{\mathbf{h}}_{c,ij}$ is the tap-weight vector of the error canceller for the filter $\underline{\mathbf{h}}_{ij}$ and $e_{si}[n]$ is the soft error at the i^{th} output $y_i[n]$. Note that $e_{si}[n]$ contains p soft-error components, i.e.,

$$e_{si}[n] = \sum_{j=1}^{p} e_{si,j}[n], \quad (11)$$

where $e_{si,j}[n]$ denotes the soft error induced by $\underline{\mathbf{x}}_j[n]$ in the VOS filtering operation $\underline{\mathbf{h}}_{ij} \otimes \underline{\mathbf{x}}_j[n]$.

Assume that the input sequences $\underline{\mathbf{x}}_j[n]$'s are zero-mean and from independent data sources. Thus, $e_{si,j}[n]$ in (11) is uncorrelated to $\underline{\mathbf{x}}_k[n]$ for $k \neq j$, resulting in

$$\underline{\mathbf{h}}_{c,ij} = \frac{E\left(e_{si,j}[n]\underline{\mathbf{x}}_j[n]^T\right)}{\sigma_{x_j}^2}, \quad (12)$$

where $\sigma_{x_j}^2$ is the variance of $\underline{\mathbf{x}}_j[n]$. Note that this result is the same as that obtained for SISO filters. Hence, we can decouple the $p \times q$ error cancellers in (10) and implement each of them independently via the energy-optimum AEC given in section 2. We denote this approach as *direct-AEC*.

While the energy-optimum AEC guarantees the minimum energy overhead for an individual error canceller, the overall error-control scheme consisting of $p \times q$ independent error cancellers, one for each VOS filter in (9), may not be energy-efficient. This is due to the fact that the possible correlations among the original filters $\underline{\mathbf{h}}_{ij}$'s may introduce computational redundances. In order to improve the energy-efficiency, we propose the MIMO-AEC technique where an algorithm transformation, denoted as MIMO-DECOR, is employed prior to VOS phase to reduce the correlation-induced complexity in the original MIMO systems.

3.3 MIMO-DECOR

The goal of MIMO-DECOR is to shorten the critical path for some filters in (9) so that they become error-free during VOS. This reduces the number of error cancellers needed for error-control. Note that the MIMO-DECOR derived below differs from the previous work [7], [8] in that it is employed to reduce the correlations among the filters in a MIMO system. These filters can be wideband or narrowband, whereas the previous DECOR can only be employed to narrowband filters.

In its most general form, the MIMO-DECOR can be expressed as

$$\left[\tilde{\underline{\mathbf{y}}}[n]\right]_{q\times 1} = \left[\mathcal{T}\left(\left[\underline{\mathbf{h}}\right]_{q\times p}\right)\right]_{q\times p} \otimes \left[\underline{\mathbf{x}}[n]\right]_{p\times 1}, \quad (13)$$

$$\left[\underline{\mathbf{y}}[n]\right]_{q\times 1} = \left[\mathcal{T}^{-1}\left(\left[\tilde{\underline{\mathbf{y}}}[n]\right]_{q\times 1}\right)\right]_{q\times 1}, \quad (14)$$

where $\mathcal{T}(\cdot)$ denotes the MIMO-DECOR transform, $\left[\tilde{\underline{\mathbf{y}}}[n]\right]_{q\times 1}$ is the output of the transformed system, $\left[\underline{\mathbf{h}}\right]_{q\times p}$, $\left[\underline{\mathbf{x}}[n]\right]_{p\times 1}$ and $\left[\underline{\mathbf{y}}[n]\right]_{q\times 1}$ are the short forms for the transfer matrix, inputs and outputs, respectively, of the original MIMO system (9). An inverse MIMO-DECOR transform, denoted by $\mathcal{T}^{-1}(\cdot)$, is employed to convert the output $\left[\tilde{\underline{\mathbf{y}}}[n]\right]_{q\times 1}$ back to the desired $\left[\underline{\mathbf{y}}[n]\right]_{q\times 1}$.

Consider a general class of MIMO systems where the filters with the same input (say, filters $\underline{\mathbf{h}}_{ij}$ and $\underline{\mathbf{h}}_{kj}$ in (9)) have correlated time-frequency characteristics (i.e., bandwidths, impulse responses, etc.). This is typical for many MIMO systems, e.g., Gigabit Ethernet 1000Base-T transceivers. We note that the similar impulse responses imply a smaller precision for the difference between the coefficients $\underline{\mathbf{h}}_{ij}$ and $\underline{\mathbf{h}}_{kj}$ than that for $\underline{\mathbf{h}}_{ij}$ and $\underline{\mathbf{h}}_{kj}$ themselves. In addition, fewer taps might be sufficient for representing $(\underline{\mathbf{h}}_{ij} - \underline{\mathbf{h}}_{kj})$. Hence, we can employ the following scheme to compute the filter

outputs as

$$y_{ij}[n] = \underline{\mathbf{h}}_{ij} \otimes \underline{\mathbf{x}}_j[n], \quad (15)$$

$$\Delta y[n] = (\underline{\mathbf{h}}_{kj} - \underline{\mathbf{h}}_{ij}) \otimes \underline{\mathbf{x}}_j[n], \quad (16)$$

$$y_{kj}[n] = y_{ij}[n] + \Delta y[n], \quad (17)$$

where $y_{ij}[n]$ and $y_{kj}[n]$ are the outputs of the filters $\underline{\mathbf{h}}_{ij}$ and $\underline{\mathbf{h}}_{kj}$, respectively. Obviously, the critical path delay of (16) is much less than that of (15) due to its reduced complexity. This also leads to additional energy savings that easily compensate for the overhead of extra computations in (17). Moreover, when applying VOS for further energy reduction, only (15) will induce soft output errors and thus require an error canceller, whereas (16)−(17) will be error-free (if their critical paths are sufficiently short, which is the case for NEXT cancellers in 1000Base-T transceivers). This reduces the AEC overhead to just one error canceller as compared to two in a direct-AEC implementation, as illustrated in Fig. 3.

The effectiveness of the above filtering scheme is determined by the relative configuration of $\underline{\mathbf{h}}_{ij}$ and $\underline{\mathbf{h}}_{kj}$. If $\underline{\mathbf{h}}_{ij}$ and $\underline{\mathbf{h}}_{kj}$ are identical (which implies a maximum correlation scenario), then $(\underline{\mathbf{h}}_{kj} - \underline{\mathbf{h}}_{ij}) = \underline{\mathbf{0}}$, resulting in the maximum energy savings as the computations in (16)−(17) can be avoided. On the other hand, if $\underline{\mathbf{h}}_{ij}$ and $\underline{\mathbf{h}}_{kj}$ are so different (uncorrelated) that the complexity of (16) is comparable to that of (15), then no energy savings can be obtained over direct-AEC implementations. As will be shown in section 4, for practical MIMO systems such as 1000Base-T transceivers, the computations in (16) are typically simple enough to guarantee substantial energy savings via MIMO-AEC.

The filtering scheme given by (15)−(17) can be equally applied to other filters in (9) as well. Let $i = 1$, $k = 2, \cdots, q$ and $j = 1, \cdots, p$, we obtain the MIMO-DECOR transform as

$$\left[\mathcal{T}\left([\underline{\mathbf{h}}]_{q\times p}\right)\right]_{q\times p} = \begin{bmatrix} 1 & 0 & \cdots & 0 \\ -1 & 1 & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots \\ -1 & 0 & \cdots & 1 \end{bmatrix}_{q\times q} \cdot [\underline{\mathbf{h}}]_{q\times p}$$

$$= \begin{bmatrix} \underline{\mathbf{h}}_{11} & \underline{\mathbf{h}}_{12} & \cdots & \underline{\mathbf{h}}_{1p} \\ (\underline{\mathbf{h}}_{21} - \underline{\mathbf{h}}_{11}) & (\underline{\mathbf{h}}_{22} - \underline{\mathbf{h}}_{12}) & \cdots & (\underline{\mathbf{h}}_{2p} - \underline{\mathbf{h}}_{1p}) \\ \cdots & \cdots & \cdots & \cdots \\ (\underline{\mathbf{h}}_{q1} - \underline{\mathbf{h}}_{11}) & (\underline{\mathbf{h}}_{q2} - \underline{\mathbf{h}}_{12}) & \cdots & (\underline{\mathbf{h}}_{qp} - \underline{\mathbf{h}}_{1p}) \end{bmatrix}, \quad (18)$$

and the inverse MIMO-DECOR transform is given by

$$\mathcal{T}^{-1} = \begin{bmatrix} 1 & 0 & \cdots & 0 \\ 1 & 1 & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots \\ 1 & 0 & \cdots & 1 \end{bmatrix}_{q\times q}. \quad (19)$$

From (18)−(19), the MIMO-DECOR involves coefficient precomputation and the inverse MIMO-DECOR adds up the differential outputs of the transformed system to give the desired ones. It is easy to see that both transforms incur very small overhead and are easy to implement. On the other hand, there are $(q-1) \times p$ out of $q \times p$ filters in the transformed system (18) performing low-complexity filtering. In addition, the number of error cancellers needed for error-control during VOS can be reduced from $q \times p$ for direct-AEC to only p for MIMO-AEC. In summary, the proposed MIMO-AEC technique achieves substantial energy reduction via: 1.) MIMO-DECOR resulting in low-complexity filtering, 2.) VOS and 3.) energy-optimum AEC for restoring the algorithmic performance.

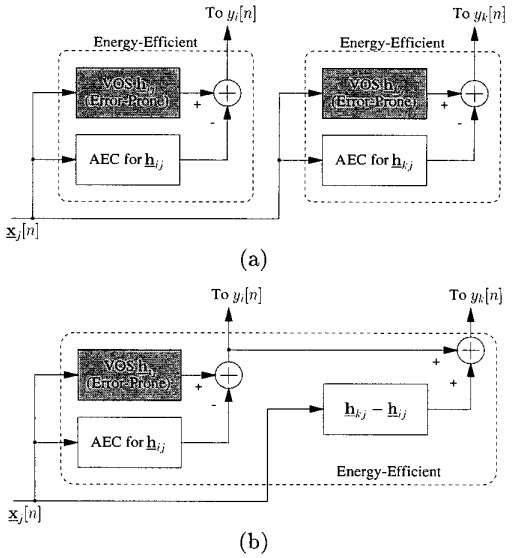

Figure 3: The proposed MIMO-AEC technique: (a) direct-AEC and (b) MIMO-AEC with MIMO-DECOR.

4. APPLICATION TO GIGABIT ETHERNET

In this section, we study the performance of the proposed MIMO-AEC technique in a Gigabit Ethernet 1000Base-T system [10]. We first give an overview of 1000Base-T standard and then employ the MIMO-AEC technique to design a low-power 1000Base-T device.

4.1 1000Base-T Transceivers

Figure 4 illustrates the structure of 1000Base-T transmission scheme. The 1000Mb/s, full duplex data throughput is achieved by using four pairs of wire in Category 5 (CAT-5) cable, each pair transmitting a 250Mb/s data stream encoded into a 4-dimension 5-level Pulse Amplitude Modulation (4-D PAM-5) signal constellation. Hence, each 1000Base-T device contains four identical transceivers, one for each pair of physical wire. The bidirectional data transmission on the same wire is made possible by hybrid circuits.

On the receive side, each receiver confronts a physical channel of minimum 100m CAT-5 cable. As illustrated in Fig. 4(a), the three major causes of signal distortion are propagation loss (due to channel attenuation), echo (generated by a self-returned signal due to impedance mismatch in hybrid circuits) and NEXT noise (caused by near-end crosstalk between adjacent wires). The IEEE 802.3ab Standard [9] specifies the models for the worst-case noise environment as

$$L_P(f) = 2.1f^{0.529} + 0.4/f, \quad (20)$$

$$L_N(f) = 27.1 - 16.8\log_{10}(f/100), \quad (21)$$

$$L_E(f) = \begin{cases} 15, & 1 \le f < 20, \\ 15 - 10\log_{10}(f/20), & 20 \le f \le 100, \end{cases} \quad (22)$$

where $1 \le f \le 100$ in MHz, $L_P(f)$, $L_N(f)$ and $L_E(f)$, all

expressed in dB/100m, are the squared magnitude of the propagation loss, NEXT and echo transfer function, respectively.

The 1000Base-T data transmission requires a BER $\leq 10^{-10}$. Using 4-D PAM-5 coding scheme, the SNR at the slicer for achieving this BER is $19.3dB$. To overcome considerable signal distortion caused by cable attenuation, echo and NEXT, advanced digital signal processing and filtering techniques are needed for signal recovery. Fig. 4(b) depicts the block diagram of a 1000Base-T device which consists of four identical transceivers operating simultaneously. At each receiver, the incoming signal is first filtered by a feed-forward equalizer (FFE) to cancel the intersymbol interference (ISI) introduced by the channel. As each received signal is also corrupted by one echo and three NEXT interferences from the adjacent wires, one echo canceller and three NEXT cancellers are needed correspondingly to perform interference suppression. In total, each 1000Base-T device needs four FFEs, four echo cancellers and twelve NEXT cancellers, all of which are LMS adaptive filters. This involves intensive filtering operations which necessitate effective energy reduction techniques to alleviate power dissipation.

4.2 Simulation Results

From Fig. 4(b), the signal processing scheme in a 1000Base-T device can be expressed as a MIMO system, i.e.,

$$\begin{bmatrix} Y_1[n] \\ Y_2[n] \\ Y_3[n] \\ Y_4[n] \end{bmatrix} = \begin{bmatrix} \underline{\mathbf{h}}_{e1} & 0 & 0 & 0 \\ 0 & \underline{\mathbf{h}}_{e2} & 0 & 0 \\ 0 & 0 & \underline{\mathbf{h}}_{e3} & 0 \\ 0 & 0 & 0 & \underline{\mathbf{h}}_{e4} \end{bmatrix} \otimes \begin{bmatrix} \underline{\mathbf{R}}_1[n] \\ \underline{\mathbf{R}}_2[n] \\ \underline{\mathbf{R}}_3[n] \\ \underline{\mathbf{R}}_4[n] \end{bmatrix} - \begin{bmatrix} \underline{\mathbf{w}}_e^1 & \underline{\mathbf{w}}_x^{12} & \underline{\mathbf{w}}_x^{13} & \underline{\mathbf{w}}_x^{14} \\ \underline{\mathbf{w}}_x^{21} & \underline{\mathbf{w}}_e^2 & \underline{\mathbf{w}}_x^{23} & \underline{\mathbf{w}}_x^{24} \\ \underline{\mathbf{w}}_x^{31} & \underline{\mathbf{w}}_x^{32} & \underline{\mathbf{w}}_e^3 & \underline{\mathbf{w}}_x^{34} \\ \underline{\mathbf{w}}_x^{41} & \underline{\mathbf{w}}_x^{42} & \underline{\mathbf{w}}_x^{43} & \underline{\mathbf{w}}_e^4 \end{bmatrix} \otimes \begin{bmatrix} \underline{\mathbf{T}}_1[n] \\ \underline{\mathbf{T}}_2[n] \\ \underline{\mathbf{T}}_3[n] \\ \underline{\mathbf{T}}_4[n] \end{bmatrix}, \quad (23)$$

where $\{Y_1[n], Y_2[n], Y_3[n], Y_4[n]\}$ is the recovered 4-D PAM-5 signal, $\underline{\mathbf{T}}_i[n]$ and $\underline{\mathbf{R}}_i[n]$ are the transmitted and received data, respectively, at the i^{th} transceiver, $\underline{\mathbf{h}}_{ei}$, $\underline{\mathbf{w}}_e^i$ and $\underline{\mathbf{w}}_x^{ij}$ denote the FFE, echo canceller and NEXT canceller (to cancel the NEXT noise generated by the j^{th} transmitter), respectively, for the i^{th} receiver. We assume the SNR requirements for the FFE, echo cancellation and NEXT cancellation are $25dB$, $28dB$ and $30dB$, respectively. This results in a $21dB$ SNR at $Y_i[n]$, which is $1.7dB$ higher than the minimum of $19.3dB$ needed to achieve a BER of 10^{-10}.

From (23), the FFE and echo cancellation involve independent SISO filtering operations, thus enabling direct-AEC for energy reduction. As mentioned in section 3.1, every three NEXT cancellers having the same input (e.g., $\underline{\mathbf{w}}_x^{14}$, $\underline{\mathbf{w}}_x^{24}$ and $\underline{\mathbf{w}}_x^{34}$ with $\underline{\mathbf{R}}_4[n]$ as their input) are similar in impulse response, thus enabling MIMO-AEC to further improve the energy-efficiency. Note that in practical systems the frequency responses of NEXT interferences vary away from the bound given by (21) due to physical variations of CAT-5 cable. In this paper, we emulate these variations by introducing a disturbance $\Delta L_N(f)$ uniformly distributed between $[-L_N^{\min}, 0]$ onto the transfer function $L_N(f)$. Thus, an instance of NEXT interference with a frequency response of $L_N(f)+\Delta L_N(f)$ is generated and utilized to calibrate the associated NEXT canceller.

We use a full adder with $T_a = 0.3ns$ at $V_{dd-crit} = 2.5V$ to implement these filters. All the simulations employ the

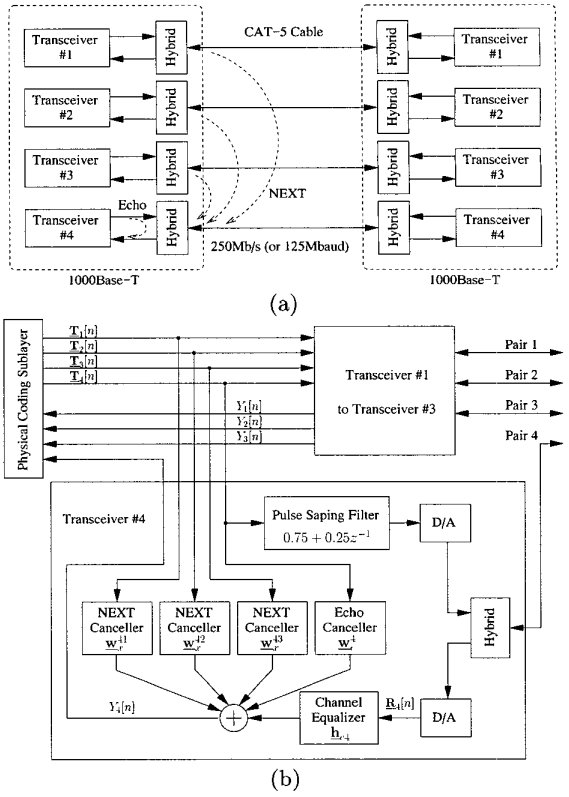

Figure 4: 1000Base-T transmission scheme: (a) signal impairments in one transceiver and (b) transceiver block diagram.

filter architecture shown in Fig. 1, where two's complement carry-save Baugh-Wooley multipliers and ripple-carry tree-style adders are being employed. It was found (see Table 1) that the critical path delays T_{cp} for the FFEs, echo cancellers and NEXT cancellers are no more than $26T_a$. Thus, these filters meet the sample period requirement which is $8ns$ for 1000Base-T devices. We employ a logic level simulation [3] to detect delay violations due to VOS on every path to the filter output given a sequence of inputs. Thus, all paths and not just the critical paths are included. 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 [11] for a $0.25\mu m$ CMOS technology.

Figure 5 plots the energy-performance trade-offs achieved via the direct-AEC for an individual filter as well as for the 1000Base-T device. It is shown that in comparison with conventional implementations, energy savings of 63.1%, 65.7% and 59.5% are achieved for the FFEs, echo cancellers and NEXT cancellers, respectively, at the desired SNR. The overall energy savings for the 1000Base-T device is found around 60.2% at $21dB$ SNR. These energy savings are obtained at $k_{vos} \approx 1.7$ and $T_{a,V_{dd-sub}} \approx 0.4ns$. Table 1 provides design specifications for these filters and the associated AECs. As indicated, the energy-optimum AECs have a critical path delay of less than $20T_a$, thus being error-free ($T_{cp} < 8ns$) at V_{dd-sub}. These results are consistent with those obtained for frequency-selective filters in [3].

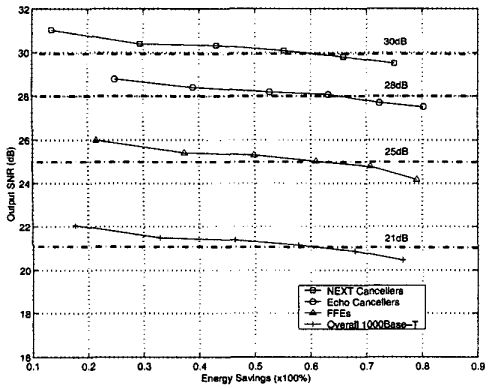

Figure 5: Energy savings via direct-AEC.

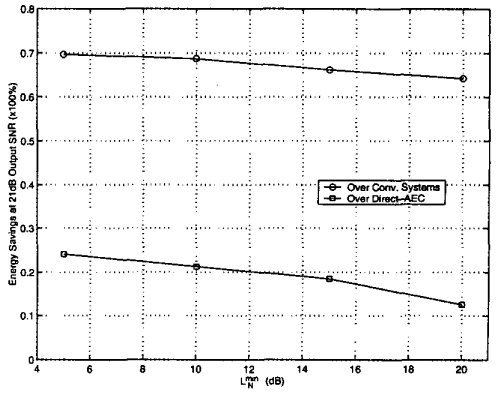

Figure 6: Energy savings via MIMO-AEC.

Table 1: Design specifications for the energy-optimum AEC-based filters.

	FFE	AEC for FFE	Echo	AEC for Echo	NEXT	AEC for NEXT
Tap #	16	5	64	15	70	18
B_{INPUT}	$8b$	$8b$	$3b$	$3b$	$3b$	$3b$
B_F	$10b$	$5b$	$12b$	$8b$	$10b$	$8b$
B_{WUD}	$12b$	$12b$	$14b$	$14b$	$14b$	$14b$
T_{cp}	$26T_a$	$19T_a$	$25T_a$	$17T_a$	$25T_a$	$19T_a$

Table 2: Design specifications and energy savings for the differential NEXT cancellers.

$L_N^{\min}$	$5dB$	$10dB$	$15dB$	$20dB$
Tap #	52	56	59	63
bit-width	$5b$	$5b$	$6b$	$7b$
T_{cp}	$18T_a$	$18T_a$	$19T_a$	$20T_a$
Overhead Reduction	44.3%	42.6%	37.1%	25.3%

Further energy reduction can be obtained for the NEXT cancellation via MIMO-AEC, resulting in four conventional NEXT cancellers and eight differential NEXT cancellers (see (16)). In addition, the number of AECs needed is reduced from twelve to only four. As shown in Table 2, a strong correlation (e.g., $L_N^{\min} = 5dB$) among the NEXT cancellers enables a large reduction in the AEC overhead (44.3%), otherwise the overhead reduction is small and the differential filters may become error-prone at VOS. This is consistent with our previous discussion. For practical 1000Base-T transceivers where $L_N^{\min}$ is typically around $10dB - 15dB$, the MIMO-AEC guarantees about 40% overhead reduction. Also shown in Fig. 6, a 1000Base-T device employing the MIMO-AEC can achieve energy savings of $22.1\% - 12.6\%$ over direct-AEC implementations and $69.1\% - 64.2\%$ over conventional implementations at the same output SNR.

5. CONCLUSIONS

In this paper, we propose a MIMO-AEC technique for the design of low-power MIMO signal processing systems. We employ the proposed technique in a Gigabit Ethernet 1000Base-T device and demonstrate substantial energy savings. MIMO-AEC provides a design paradigm for energy-efficient DSP algorithms and architectures, whereby DSP algorithms, architectures and circuit properties are jointly optimized to push the limits of energy reduction.

6. REFERENCES

[1] B. Davari, R. H. Dennard and G. G. Shahidi, "CMOS scaling for high-performance and low power - the next ten years," *Proceedings of the IEEE*, vol. 83, pp. 595-606, April 1995.

[2] K. L. Shepard and V. Narayanan, "Noise in deep submicron digital design," *ICCAD'96*, pp. 524-531, 1996.

[3] L. Wang and N. R. Shanbhag, "Low-power signal processing via error-cancellation," *Proc. of IEEE Workshop Signal Process. Syst. (SiPS)*, pp. 553-562, Oct. 2000.

[4] R. Hegde and N. R. Shanbhag, "Energy-efficient signal processing via algorithmic noise-tolerance," *Proc. of Intl. Symp. on Low-Power Electronics and Design*, pp. 30-35, Aug. 1999.

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

[6] D. P. Bertsekas, *Nonlinear Programming*, Boston, MA: Athena Scientific, 1995.

[7] S. Ramprasad, N. R. Shanbhag and I. N. Hajj, "Decorrelating (DECOR) transformations for low-power digital filters," *IEEE Trans. on Circuits and Systems II*, vol. 45, pp. 776-788, June 1999.

[8] N. Sankarayya, K. Roy and D. Bhattacharya, "Algorithms for low power and high speed FIR filter realization using differential coefficients," *IEEE Trans. Circuits and Systems II*, vol. 44, pp. 488-497, June 1997.

[9] IEEE Standard 802.3ab, 1999, *URL:http://www.manta.ieee.org/groups/802/3/ab/.*

[10] R. He, N. Nazari and S. Sutardja, "A DSP based receiver for 1000Base-T PHY," *ISSCC'2001*, pp. 308-309, San Francisco, CA, Feb. 2001.

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