LOW-POWER EQUALIZERS FOR 51.84 Mb/s VERY HIGH-SPEED DIGITAL SUBSCRIBER LOOP (VDSL) MODEMS

Manish Goel and Naresh R. Shanbhag
Coordinated Science Laboratory/ECE Department
University of Illinois at Urbana-Champaign
1308 West Main Street, Urbana, IL 61801.
E-mail: [mgoel,shanbhag]@uivlsi.csl.uiuc.edu.

Abstract - In this paper, we present low-power equalizers derived via dynamic algorithm transformations (DAT). These transformations achieve low-energy operation by reconfiguring the architecture and the supply voltage in response to channel non-stationarities. Practical reconfiguration strategies are derived as a solution to an optimization problem with energy as the objective function and a constraint on the algorithm performance (specifically the SNR). Simple energy models for multipliers are presented. The DAT-based adaptive filter is employed as an equalizer for 51.84 Mb/s very high-speed digital subscriber loop (VDSL) over 24-pair BKMA cable. On the average, 88% energy savings are achieved due to variations in cable length and number of far-end crosstalk (FEXT) interferers.

1 INTRODUCTION

Algorithm transformation techniques [1] have been employed for low-power digital signal processing system design. Existing techniques include *strength reduction* [1, 2, 3] and *variable-length vector quantizer (VQ)* [4]. We refer to these algorithm transformations as *static algorithm transformations* (SAT), because these are applied during the algorithm design phase assuming a worst-case scenario and their implementation is time-invariant.

In contrast, *dynamic algorithm transforms* (DAT) presented in this paper optimize energy dissipation via real-time energy-optimum *reconfiguration* in the presence of input non-stationarities. Traditionally, a signal processing system is designed for the *worst* case, so that it can meet the performance requirement for all cases including the *nominal* and the *best* case. For example, a broadband modem is typically designed for the longest cable length, the maximum cable temperature and the worst case near-end/far-end crosstalk interferers. This worst-case design is an "over-kill" in terms of energy consumption for nominal and best cases. If the worst-case occurs rarely, then much energy savings can be achieved by reconfiguration of the system for the best and the nominal cases. In the past, we have shown [5] that a DAT-based near-end crosstalk (NEXT) canceller can achieve $21\% - 62\%$ energy savings via *optimal choice of filter taps.*

In general, one needs to define a hardware platform for getting energy

0-7803-4997-0/98/$10.00

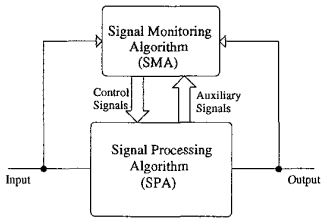

Figure 1: DAT based implementation.

savings via reconfiguration, such as application specific integrated circuits (ASIC) [6], multiprocessors [7, 8] and programmable gate arrays. In contrast, the proposed DAT-based approach is not tied to any particular hardware platform in the sense that platform-specific reconfiguration and energy parameters are made a part of the reconfiguration strategy. In this paper, we choose an ASIC-based hardware platform and present reconfiguration strategies for energy-optimum reconfiguration of the reconfigurable datapath. Related works include a 128 tap adaptive equalizer with varying filter order and precision [9], approximate signal processing [10], and a variable supply voltage methodology [11].

A DAT-based VLSI signal processing system shown in Fig. 1 has two subsystems: 1.) a signal processing algorithm (**SPA**) block that is a reconfigurable datapath (section 2); and 2.) a signal monitoring algorithm (**SMA**) block (section 3) that implements a reconfiguration strategy/control, based upon temporal/spatial variabilities in the input. In this paper, we propose a DAT-based adaptive filter, which is then employed as an equalizer for 51.84 Mb/s very high-speed digital subscriber loop (VDSL) [12] over BKMA cable in section 4.

2 RECONFIGURABLE DATAPATH : THE SPA BLOCK

In this section, we describe a reconfigurable adaptive filter architecture that comprises the **SPA** block and a variable supply voltage scheme [11]. The *least mean square (LMS) algorithm* is employed to implement the adaptive filter architecture shown in Fig. 2. The filter (**F**) block implements a finite-impulse response (FIR) filter, whose coefficients (w_k) are updated by the weight-update block (**WUD**)-block. Assume that $x(n)$ is the input signal to the adaptive filter. The multiplication of $x(n-k)$ by w_k and updating of w_k constitute the k^{th} tap as shown in Fig. 2. The signals α_k and β_k in Fig. 2 are employed to power up/down the k^{th} tap of the adaptive filter. For example, setting $\alpha_k = 0$ forces a zero at the input to the **F**-block multiplier of the k^{th} tap and bypasses the **F**-block adder. Thus, the k^{th} tap in the **F**-block is powered down. Similarly, $\beta_k = 0$ powers down the k^{th} tap in the **WUD**-block. Additional energy savings can be achieved by changing the precisions of the input signal and the coefficients. This can be done by forcing least

significant bits of the input signal and the coefficients to zero. If we assume that input signal $x(n)$ is uncorrelated, the mean squared error (MSE) $\mathcal{J}$ of the reconfigurable adaptive filter in Fig. 2 can be computed as,

$$\mathcal{J} = \sigma_d^2 - \sum_{k=1}^{N} \alpha_k w_k^2 \sigma_x^2, \tag{2.1}$$

where σ_d^2 is the desired signal power, α_k is the control signal for the k^{th} tap, σ_x^2 is the input signal power and w_k are the filter coefficients. In many

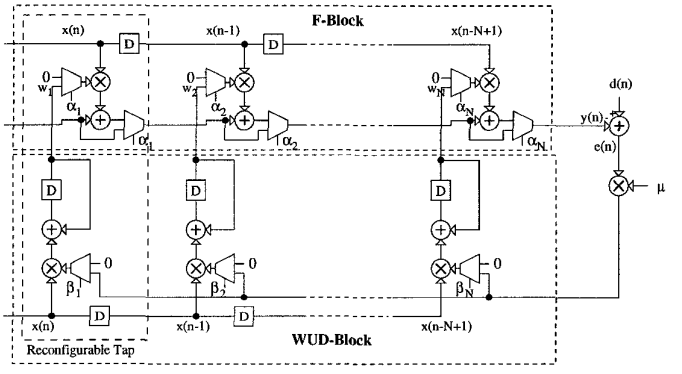

Figure 2: Reconfigurable Adaptive Filter Architecture.

applications, the **WUD**-block is powered down (to conserve energy) after the optimum coefficients have been obtained, i.e., after convergence. In that case, the critical path delay T_{cp} is due to the adders and multipliers in the **F**-block and is given by,

$$T_{cp} = T_m + \sum_{k=1}^{N} \alpha_k T_a, \tag{2.2}$$

where T_m and T_a are the computation times for the multiplier and the adder blocks, respectively. Here, the computation time of the multiplexers has been assumed to be much smaller than an adder or a multiplier computation time. The critical path delay in (2.2) will be maximum when all the taps in the adaptive filter are powered up. Thus, the maximum critical path delay $T_{cp,max}$ is obtained by substituting $\alpha_k = 1$ $(k = 1, 2, \cdots, N)$ in (2.2). If T_s is the sample period, then we choose $T_{cp,max} = T_s$, so that the sample rate requirements are met for the worst case. Now consider the scenario when some of the taps in the **F**-block are powered down. In this case, T_{cp} in (2.2) will be smaller than $T_{cp,max}$. The reduced critical path delay can be traded-off with the supply voltage to save energy. It can be shown [11] that the supply voltage can be lowered to $V_{dd}(r) \leq V_{dd,max}$, where $V_{dd}(r)$ is given by,

$$V_{dd}(r) = V_t + \frac{r}{2} V_o + \sqrt{\frac{r^2}{4} V_o^2 + r V_t V_o}, \tag{2.3}$$

and $V_o = (V_{dd,max} - V_t)^2/V_{dd,max}$, V_t is the threshold voltage, $V_{dd,max} = V_{dd}(1)$ and $r = T_{cp}/T_s$ is the normalized processing rate. Lowering the supply voltage V_{dd} enables significant energy savings as energy dissipation is a quadratic function of the supply voltage [11]. Thus, in our work, energy savings are achieved by reducing switched capacitance (powering down taps and varying precisions), and by reducing supply voltage V_{dd}. In the next section, we present a reconfiguration strategy for achieving these savings.

3 RECONFIGURATION STRATEGY : THE SMA BLOCK

In the last section, we presented a reconfigurable datapath. In this section, we present a reconfiguration strategy for computing the energy-optimum configuration for this datapath.

3.1 Input Non-Stationarity Model

We propose to distinguish between the worst, nominal and the best case scenarios via the definition of input state $\mathbf{s}(n)$ as follows:

Definition 1 : *The* **input state** $\mathbf{s}(n) \in \mathcal{S} = \{\mathbf{s}_1, \mathbf{s}_2, \cdots, \mathbf{s}_{N_s}\}$ *(where $\mathcal{S}$ is the state-space), at time instant n is a vector of input-dependent parameters where $\mathbf{s}(n) = \mathbf{s}_i$ with a probability $p(\mathbf{s}_i)$.*

Assume that a digital subscriber loop provides connections to 100 homes, out of which 80% are approximately 0.6 kft from the transmitter and the remaining are distributed equally between 0.1 kft and 1 kft. In this case, the state set $\mathcal{S} = \{\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3\}$, $p(\mathbf{s}_1) = 0.1$, $p(\mathbf{s}_2) = 0.8$ and $p(\mathbf{s}_3) = 0.1$, where $\mathbf{s}_1$, $\mathbf{s}_2$ and $\mathbf{s}_3$ are the input states corresponding to the cable length of 0.1 kft, 0.6 kft and 1 kft, respectively.

Definition 2 : *The* **configuration** $\mathbf{c}(n) \in \mathcal{C} = \{\mathbf{c}_1, \mathbf{c}_2, \cdots, \mathbf{c}_{N_c}\}$ *(where $\mathcal{C}$ is the configuration space) at time instant n is defined as a vector of reconfiguration control signals. Each configuration $\mathbf{c}_i$ corresponds to a particular value of the control signals.*

For example, for an N-tap reconfigurable adaptive filter (see Fig. 2), the configuration can be defined as:

$$\mathbf{c}_i = [\underline{\alpha}_i, \underline{\beta}_i, \mathbf{B}_{w,i}, \mathbf{B}_{x,i}, \mathbf{B}_{V,i}]$$

where $\underline{\alpha}_i$ and $\underline{\beta}_i$ are N-bit control words containing *specific values* of the control signals α_k and β_k, respectively, in Fig. 2. Similarly, $\mathbf{B}_{w,i}$ and $\mathbf{B}_{x,i}$ are control words indicating the coefficient precision and data precision, respectively. Finally, $\mathbf{B}_{V,i}$ is the control word indicating the value of the supply voltage V_{dd}. If the supply voltage V_{dd} can have values from $1.5V$ to $3.3V$ with a step of $0.6V$, then $\mathbf{B}_{V,i}$ is a 2-bit control word. Thus, with $N = 8$, $B_w = 10$ bits, $B_x = 4$ bits and four supply voltage V_{dd} levels, we need a 24-bit configuration vector $\mathbf{c}_i$ implying that the number of configurations $N_c = 2^{24}$. The above example indicates how easily the number of configurations explodes with increase in reconfigurability options.

For every input state $s_i \in \mathcal{S}$, there exists an energy-optimum configuration $\mathbf{c}^*(\mathbf{s}_i)$ defined as follows:

Definition 3 : *The* **energy-optimum configuration** $\mathbf{c}^*(\mathbf{s}_i) \in \mathcal{C}$ *for a given input state* $\mathbf{s}_i \in \mathcal{S}$ *is defined as:*

$$\begin{aligned} \mathbf{c}^*(s_i) &= \arg\min_{\mathbf{c}\in\mathcal{C}} \mathcal{E}_{SPA}(\mathbf{c}), \\ &s.t.\ \mathcal{J}_{SPA}(\mathbf{s}_i, \mathbf{c}) \leq \mathcal{J}_o, \end{aligned} \tag{3.1}$$

where $\mathcal{E}_{SPA}(\mathbf{c})$ *is energy dissipated by the* **SPA** *block in configuration* $\mathbf{c}$, $\mathcal{J}_o$ *is the specified MSE and* $\mathcal{J}_{SPA}(s_i, \mathbf{c})$ *is the MSE achieved by the* **SPA** *block when the input is in state* $\mathbf{s}_i$ *and the* **SPA** *block is in configuration* $\mathbf{c}$.

If **SPA** block consists of the adaptive filter in Fig. 2, then $\mathcal{J}_{SPA}(\mathbf{s}_i, \mathbf{c})$ is given by (2.1). This leads to the following optimization problem:

$$\begin{aligned} \min_{\alpha_k\in\{0,1\}} \quad & \sum_{k=1}^{N} \alpha_k \mathcal{E}_m(w_k), \\ s.t. \quad & \sigma_d^2 - \sum_{k=1}^{N} \alpha_k |w_k|^2 \sigma_x^2 \leq \mathcal{J}_o, \end{aligned} \tag{3.2}$$

where $\mathcal{E}_m(w_k)$ is the energy consumption of the **F**-block for the k^{th} tap (see Fig. 2). Note that we do not include β_k's in the optimization problem because $\beta_k = 0$ after the adaptive filter has converged, i.e., the **WUD**-block is powered down. The optimization problem in (3.2) is solved via the *Lagrange Multiplier Method* to get a reconfiguration strategy, which is described next.

3.2 Reconfiguration Strategy

It can be shown that the solution to (3.2) is given by:

$$\alpha_{k,opt} = \begin{cases} 1, & \frac{|w_k|^2}{\mathcal{E}_m(w_k)} \geq \tilde{\lambda}^* \\ 0, & \frac{|w_k|^2}{\mathcal{E}_m(w_k)} < \tilde{\lambda}^*, \end{cases} \tag{3.3}$$

where $\tilde{\lambda}^*$ is a constant. The solution (3.3) indicates that it is better to power down taps with small values of $|w_k|^2/\mathcal{E}_m(w_k)$. Intuitively, this makes sense as small values of $|w_k|^2/\mathcal{E}_m(w_k)$ imply that the k^{th} tap contributes less to the performance measure (as w_k is small) but consumes more energy ($\mathcal{E}_m(w_k)$ is large). In practice, we don't need to compute the constant $\tilde{\lambda}^*$ in (3.3) if we employ the strategy that $\alpha_{k,opt}$ can be obtained by powering down the taps starting with the smallest value of $|w_k|^2/\mathcal{E}_m(w_k)$ until the MSE constraint (see (3.2)) is violated.

The β_k can similarly be obtained by choosing $\beta_k = 0$ whenever $\alpha_k = 0$ or the filter has converged. The optimum supply voltage V_{dd} is obtained by substituting $\alpha_k = \alpha_{k,opt}$ in (2.2) and then employing (2.3) to obtain a coarse V_{dd} adjustment. Fine adjustments [11] can be done subsequently. The

optimum input precision $B_{x,opt}$ can be obtained by computing the input peak-to-average ratio (PAR) (defined as the ratio of the maximum to the root mean squared value of the input signal) and then determining the minimum precision required for the desired signal-to-quantization ratio (SQNR) at the input. In general, the input precision can be reduced by one bit for each $6dB$ reduction in PAR of the input signal. It was seen that for 51.84 Mb/s VDSL application, the input precision requirements do not change. The optimum coefficient precision $B_{w,opt}$ is chosen by deriving expression for the round-off error in terms of the coefficient precision and then determining the precision to achieve the desired SQNR at the output. It can be shown that SQNR at the output is a function of number of powered-up taps, and the coefficient precision can be reduced by 1 bit for each four-fold reduction in number of powered-up taps.

The reconfiguration strategy in (3.3) requires computation of energy values $\mathcal{E}_m(w_k)$. The energy $\mathcal{E}_m(w_k)$ is computed by evaluating the energy consumed by the multiplier in the k^{th} tap of the **F**-block. Next, we present models for $\mathcal{E}_m(w_k)$.

3.3 Multiplier Energy Models

The multiplier energy model is based on two estimates of the transition activity in an array multiplier. First estimate $\mathcal{N}_1(w)$ is given by number of ones in a binary representation of the coefficient w. The second estimate $\mathcal{N}_2(w)$ is given by the difference of B_w and number of zeros at the least significant bit (LSB) positions in the binary representation of w. It was found via a gate-level simulation that the estimates $\mathcal{N}_1(w)$ and $\mathcal{N}_2(w)$ respectively underestimates and overestimates the energy consumption of the multiplier. It was found that an energy model derived by taking average of $\mathcal{N}_1(w)$ and $\mathcal{N}_2(w)$, results in less than 8% error as compared to a zero-delay gate-level simulation. This energy model is also useful in computing the energy savings of the DAT-based system as compared to the conventional worst case design. The average energy dissipation of the DAT-based system, $\mathcal{E}_{DAT}$ is computed by averaging the energy dissipation corresponding to each of the states. To reduce the energy dissipation of the **SMA** block, it is activated at a lower frequency. If $\mathcal{E}_{WC}$ is the worst case energy dissipation, then the energy savings ($\mathcal{E}_{SAV}$) are obtained by the ratio of $\mathcal{E}_{WC} - \mathcal{E}_{DAT}$ to $\mathcal{E}_{WC}$. In the next section, we compute these energy savings for 51.84 Mb/s very high-speed digital subscriber loop (VDSL).

4 APPLICATION TO 51.84 Mb/s VDSL

In this experiment, we demonstrate the performance of the DAT-based adaptive filter as an equalizer for 51.84 Mb/s very high-speed digital subscriber loop (VDSL) [12] over BKMA cable. The channel carries asynchronous transfer mode (ATM) cells and uses synchronous optical network (SONET) framing. Therefore, the data rate for VDSL is 51.84 Mb/s, which is an integer divisor of SONET rate 622.08 Mb/s ($51.84 \times 12 = 622.08$).

4.1 Simulation Setup

First, we describe the transmitter and the receiver for 51.84 Mb/s VDSL. We will assume that the carrierless amplitude phase (CAP) modulation scheme is being employed. The scrambled 51.84 Mb/s data is passed through a 16-CAP (4 bits/symbol) encoder, and the encoded symbols are passed through in-phase and quadrature-phase shaping filters (defined by square-root raised cosine pulses with sampling frequency f_s=51.84 MHz, center frequency f_c=12.96 MHz, and excess bandwidth $\alpha = 38\%$). The shaping filter output is converted to an analog signal by a digital-to-analog (D/A) converter operating at 51.84 MHz.

In the receiver (see Fig. 3), the digital output of the analog-to-digital (A/D) converter (operating at 51.84 MHz) is processed by a decision feedback equalizer (DFE). The DFE consists of the two sections - a feedforward section and a feedback section. The feedforward section is a fractionally-spaced linear equalizer (FSLE), which is a pair of 48-tap adaptive filters. A complex adaptive filter with 10 taps is employed as the feedback filter. The strength-reduced architecture proposed in [2] is employed for a low-power implementation of the feedback filter. The energy dissipation of the DFE is dominated by the feedforward section because of the lower operating frequency (12.96 MHz), smaller tap-length,and lower precision requirements for the feedback filter. Therefore, we do not reconfigure the feedback section. The complex-

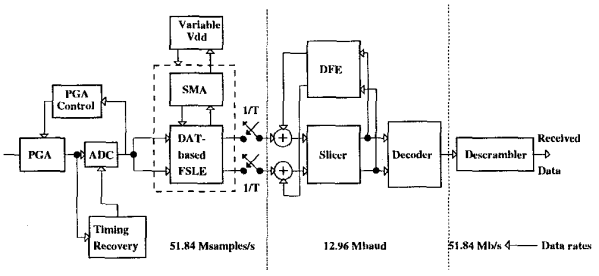

Figure 3: DAT-based Receiver for 51.84 Mb/s VDSL.

ity of the update algorithm is reduced by employing reduced constellation algorithm (RCA) for blind convergence and powers-of-two approximations for the slicer errors. The algorithmic performance measure in this case is the signal-to-noise ratio (SNR) at the slicer (SNR_{sl}). For 16-CAP, a performance specification of $SNR_{sl} = 21.5dB$ is sufficient to obtain a probability of error less than 10^{-7}.

We assume that the worst case design corresponds to 1 kft cable length and 11 far-end crosstalk (FEXT) interferers. This design corresponds to number of filter taps N=48, coefficient precision B_w=10 bits, data precision B_x=8 bits, and supply voltage V_{dd}=3.3 V. We will assume that the cable length varies from 0.1 kft to 1 kft. Also, number of FEXT interferers can be 4,

7 or 11. The exact probability distribution of the states requires a survey of the VDSL network installations. Since this information is not known at the present moment, we assume two extreme distributions. In first case, we assume that states have a Gaussian distribution with 0.55 kft and 4-FEXT as the mean values (nominal case) and standard deviation of 0.2 kft and 3-FEXT. In the second case, we assume a uniform distribution for the states. The channel models for BKMA cable are similar to those of a category 5 cable, which are specified in the TIA/EIA-568A Standard [13]. The number of supply voltage levels required depend upon the algorithm being implemented. Fewer the levels, simpler is the design for the control of the supply voltage V_{dd}. We assume V_{dd} can take value from $\{1.5V, 2.1V, 2.7V, 3.3V\}$. Next, we present the simulation results.

4.2 Simulation Results

The number of powered up taps for the in-phase and the quadrature-phase adaptive filters corresponding to each input state are given in Table 1. It can be seen all 48 taps in both filters are needed to be powered up in the worst case (1 kft cable length and 11-FEXT). However, only 4 taps are needed for

Table 1: Number of Powered-up Taps ($\sum \alpha_{k,opt}$)

Cable Length	11-FEXT		7-FEXT		4-FEXT	
	in-ph	q-ph	in-ph	q-ph	in-ph	q-ph
1.0 kft	48	48	14	20	11	11
0.9 kft	40	36	13	11	11	11
0.8 kft	25	25	9	11	8	11
0.7 kft	11	11	10	10	9	10
0.6 kft	8	11	7	11	7	10
0.5 kft	5	9	5	9	5	8
0.4 kft	5	8	5	7	5	7
0.3 kft	5	5	5	5	5	5
0.2 kft	4	4	4	4	4	4
0.1 kft	4	4	4	4	4	4

the best case, which corresponds to (0.1 kft cable length and 4-FEXT). The number of powered up taps vary from 48 to 4 as the cable length varies from 1 kft to 0.1 kft and for a variation in number of FEXT interferers from 11 to 4. The optimum precisions and supply voltage are determined from the knowledge of number of powered-up taps and critical path delay calculation.

In Fig. 4(a), we plot energy savings for a DAT-based equalizer when the cable length varies from 0.1 kft to 1 kft while the number of FEXT interferers is kept fixed at 11. These energy savings include the energy consumption in the **SMA** block. The energy savings range from -2% to 96% when cable

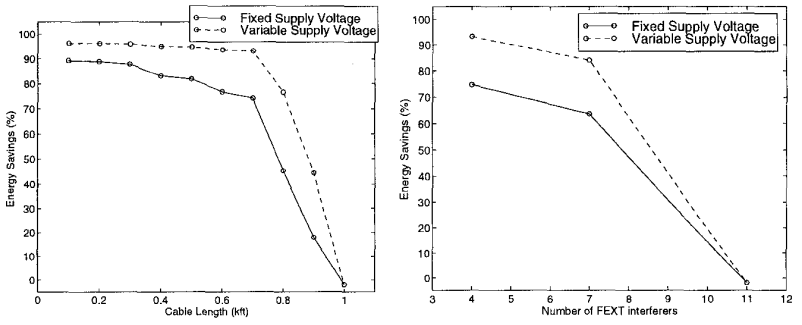

(a) (b)

Figure 4: Energy savings with variation in: (a) cable length (for 11-FEXT) and (b) number of FEXT interferers (for 1 kft cable length).

length varies from 1 kft to 0.1 kft, respectively. When the supply voltage is fixed, then these savings (see Fig. 4(a)) range from -2% to 89% for the corresponding cases. Also plotted in Fig. 4(b) are the energy savings when the number of FEXT interferers vary from 4 to 11 while the cable length is fixed at 1 kft. The energy savings for this case range from -2% to 75% when the supply voltage V_{dd} is fixed and from -2% to 93% when the supply voltage V_{dd} is also varied.

In Fig. 5(a), we plot energy savings when the supply voltage V_{dd} is fixed. The cable length is varied from 0.1 kft to 1 kft and the number of FEXT interferers vary from 4 to 11. It can be seen that energy savings range from -2% to 91%. Average energy savings are obtained by a probability weighted

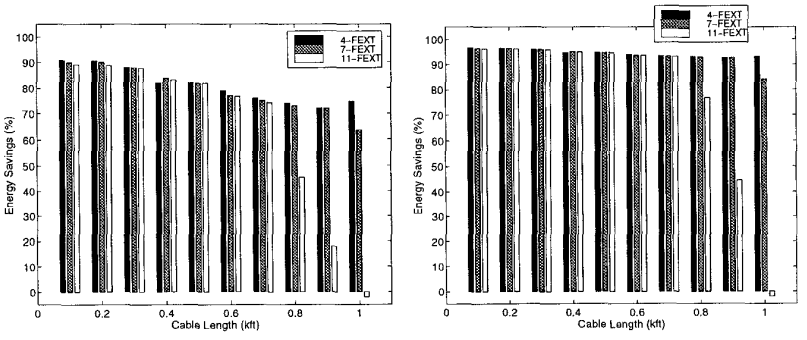

(a) (b)

Figure 5: Energy savings for: (a) fixed supply voltage and (b) variable supply voltage.

summation of the input state energy values. It was found that, the average energy savings with the fixed supply voltage is 75% for input states with a

Gaussian probability distribution. Similarly, for the variable supply voltage case (see Fig. 5(b)), energy savings range from -2% to 96%, and the average energy savings are 91%. Similarly, for the uniform state probability distribution, the average energy savings of 74% and 89% for the fixed and variable supply voltage designs are obtained.

Thus, it can be seen that the DAT-based approach is quite attractive from the viewpoint of energy savings for the VDSL application. Future work is being directed towards a silicon implementation of the DAT-based VDSL receiver mentioned in this paper and extension of the DAT-framework to more general systems.

5 ACKNOWLEDGMENTS

This work was supported via DARPA contract DABT63-97-C-0025 and NSF CAREER Award MIP-962377.

6 REFERENCES

[1] A. Chandrakasan *et al.*, "Minimizing power using transformations," *IEEE Trans. Computer-Aided Design*, vol. 14, no. 1, pp. 12–31, Jan. 1995.

[2] 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, no. 5, pp. 1276–1290, May 1997.

[3] D. A. Parker and K. K. Parhi, "Low-area/power parallel FIR digital filter implementations," *Journal of VLSI Signal Processing*, vol. 17, pp. 75–92, Sept. 1997.

[4] E. Tsern and T. H. Meng, "A low-power videorate pyramid VQ decoder," *IEEE J. of Solid State Circuits*, vol. 31, no. 11, pp. 1789–1794, Nov. 1996.

[5] M. Goel and N. R. Shanbhag, "Dynamic algorithm transformations (DAT) for low-power adaptive signal processing," *Proceedings of International Symposium on Low Power Electronics and Design*, pp. 161–166, Aug. 1997.

[6] M. S. B. Romdhane and V. K. Madisetti, "LMSGEN: A prototyping environment for programmable adaptive digital filters in VLSI," *VLSI Signal Processing*, pp. 33–42, 1994.

[7] "Pleiades: Ultra-low-power hybrid and configurable computing," *http://infopad.eecs.berkeley.edu/research/reconfigurable/*.

[8] "Programmable and configurable digit-serial digital signal processors," *http://www.ee.umn.edu/groups/ddp/dig_ser/*.

[9] C. J. Nicol, P. Larsson, K. Azadet, and J. H. O'Neill, "A low power 128-tap digital adaptive equalizer for broadband modems," *IEEE Journal of Solid-State Circuits*, vol. 32, no. 11, pp. 1777–1789, Nov. 1997.

[10] J. T. Ludwig, S. H. Nawab, and A. P. Chandrakasan, "Low-power digital filtering using approximate processing," *IEEE J. of Solid State Circuits*, vol. 31, no. 3, pp. 395–400, Mar. 1996.

[11] V. Gutnik and A. P. Chandrakasan, "Embedded power supply for low-power DSP," *IEEE Trans. on VLSI Systems*, vol. 12, pp. 425–435, Dec. 1997.

[12] N. R. Shanbhag and G.-H. Im, "VLSI system design of 51.84 Mb/s transceiver for ATM-LAN and broadband access," *IEEE Trans. Signal Processing*, vol. 46, pp. 1403–1416, May 1998.

[13] TIA/EIA-568-A Standard, *Commercial Building Telecommunications Cabling Standard*, 1994.