A Low-Power, Reconfigurable Adaptive Equalizer Architecture *

Jim Tschanz and Naresh R. Shanbhag
VLSI Information Processing Systems (VIPS) Group
Coordinated Science Laboratory / ECE Department
University of Illinois at Urbana-Champaign
1308 West Main Street, Urbana, IL 61801.
Phone: (217) 244-0041, Fax: (217) 244-1946
Email: {jtschanz,shanbhag}@uivlsi.csl.uiuc.edu

Abstract

This paper presents an architecture for an adaptive equalizer that is dynamically reconfigurable for low-power operation. The equalizer is composed of a signal processing block which accomplishes the filtering operations and a signal monitoring block which controls reconfiguration by monitoring the equalizer performance and dynamically powering up or down filter taps in order to conserve energy. This reconfigurable equalizer is used in the design of a 51.84 Mb/s VDSL receiver core, and simulation results are shown which demonstrate the power savings accomplished through reconfiguration.

1. Introduction

Excessive power consumption has always been a concern for integrated circuit designers due to the problems of operating temperature and device reliability it causes. It is becoming even more important, however, as a larger number of communications systems are made wireless. Because the battery life of such a system is directly related to the amount of power it consumes, low-power designs are necessary in order to create a successful product. A major component in most communications receivers is the adaptive equalizer which filters the incoming data samples from the communications channel in order to recover the transmitted information. This paper presents an architecture for a low-power adaptive equalizer which achieves energy-efficient operation through the use of dynamic reconfiguration. This equalizer architecture is used in the design of a 51.84-Mb/s very high speed digital subscriber loop [1] (VDSL) receiver core and simulation results demonstrate the power savings.

*This work was supported in part by NSF CAREER award MIP 96-23737 and Analog Devices, Inc.

In the following sections, an overview is given of dynamic algorithm transforms and the VDSL environment. Then in Section 2 the architecture of the reconfigurable equalizer is presented, while simulation results are shown in Section 3.

1.1 Dynamic algorithm transformations (DAT)

Traditionally, signal processing systems have been designed for low-power operation by applying certain algorithm transforms [2, 3] in order to optimize the architecture. For example, pipelining [4] may be used to reduce the critical path of a design, thereby allowing the supply voltage to be reduced. Once the algorithm is sufficiently optimized, custom circuits are designed which provide the necessary balance between power consumption and speed. These optimizations are made during the design stage, and therefore are known as *static algorithm transformations*. Because they are static, no provision is made for the algorithm to change at run-time in response to the environment in which the system is used. To allow this dynamic reconfiguration, the concept of *dynamic algorithm transforms* (DAT) [5] was introduced. By applying DAT to a system, the system can adapt to its environment and therefore operate efficiently whereas a worst-case system always consumes the same amount of power no matter what the environment. In order to design reconfigurable features into a system, however, it is necessary to know the variabilities in the environment and how that affects the complexity of the algorithm. For example, in many communications systems the amount of complexity and power consumption required is directly related to the distance between the transmitter and receiver (the channel length) and the amount of interference present. This is especially true of VDSL, which will be explained next.

0-7803-5700-0/99/$10.00©1999 IEEE

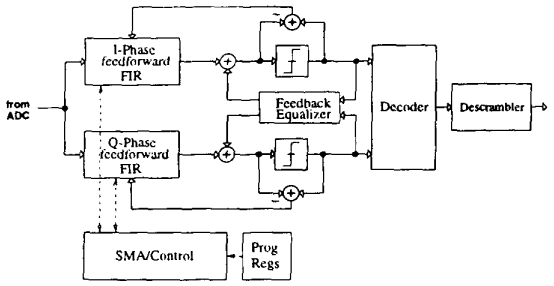

Figure 1. VDSL receiver block diagram.

1.2 VDSL overview

VDSL is designed for a fiber-to-the-curb (FTTC) network topology where fiber-optic cable is run from the central office to an optical network unit (ONU) located in a neighborhood or basement of a high-rise building. Individual users are then connected to the ONU by unshielded twisted pair (UTP) wire carrying data at a downstream rate of 51.84 Mb/s. Because of the poor frequency response of UTP cable, the distance from the user to the ONU must be small: around 1000 ft or less. Therefore a VDSL receiver designed for the worst-case environment must be able to receive data from a 1000-ft channel while meeting the bit-error rate (BER) requirement of $< 10^{-7}$ [6]. Shorter cable lengths suffer reduced propagation loss and require less equalization than the worst-case. In addition to propagation loss, VDSL receivers are affected by far-end crosstalk (FEXT) caused by other transceivers operating over the same cable bundle. A standard BKMA cable consists of 12 wire pairs and therefore a worst-case receiver design must be able to operate in the presence of 11 other interfering transmitters. This worst-case design, however, operates inefficiently when used in an environment with less interference or a shorter cable length. The goal of designing a dynamically reconfigurable equalizer is to eliminate this inefficiency.

A block diagram of the VDSL receiver core is shown in Fig. 1. The reconfigurable equalizer consists of the two parallel feedforward filter blocks which make up the feedforward equalizer (FFE), and the complex-valued feedback equalizer (FBE). The feedforward equalizer output is added to the output from the feedback filter, and then sliced to get the transmitted symbol. This symbol is then decoded and descrambled to retrieve the information bits. The block labeled "SMA" in Fig. 1 refers to the signal monitoring algorithm block which controls reconfiguration of the equalizer. The architecture of the equalizer and the SMA block will be explained in the next section.

2 Reconfigurable equalizer architecture

Before dynamic algorithm transforms can be applied to an architecture, the target architecture must operate as efficiently as possible in the worst-case environment. System simulations of the VDSL receiver core determined the data precisions and the size and complexity of the equalizer required. These requirements are shown in Table 1.

Table 1. Receiver core features.

Feedforward equalizer taps	48 each I and Q-phase
Feedback equalizer taps	10 complex
Input data precision	8 b
FFE output precision	18 b
FFE coefficient precision	20 b
FBE coefficient precision	18 b

2.1 Static algorithm transforms

Next, static algorithm transforms are applied to the system to optimize the architecture for speed, area, and power consumption. Extensive use is made of folding which reduces the area requirement of an algorithm by multiplexing several operations onto one computation node. The feedforward equalizer which totals 96 taps is folded by a factor of four resulting in 24 folded tap structures. One method of folding this equalizer involves folding the I-phase and Q-phase filters separately and is explained in [7]. This method however does not take advantage of the fact that the data input to both filters is identical. By folding the two filters together, the data input to each multiplier can be held constant for two clock cycles rather than changing every clock cycle as in [7]. This results in an average power savings of 15% with no additional hardware required. Folding itself generates an architecture which requires only one-fourth of the area of the un-folded architecture. The folded FFE structure consists of two blocks "ffe0" and "ffe1" and is shown in Fig. 2. Each of the 12 taps in these filter blocks is identical except for the initial coefficient which is loaded into the coefficient registers upon chip reset.

The feedback filter is also folded by a factor of four to conserve area. This is a complex-data, complex-coefficient filter so that four real multiplications are required per tap in order to compute the complex output. Strength reduction [8] is used to reduce the number of multiplications required, resulting in an architecture with three real filters rather than four. This provides both area and energy savings.

Finally, as an additional power-saving strategy, burst-mode coefficient updating is employed. It is assumed that

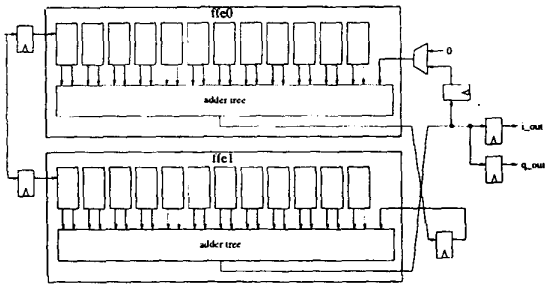

Figure 2. Folded FFE structure.

the VDSL channel characteristics change slowly with time and therefore once the equalizer has converged, the filter coefficients can be updated less frequently. When the coefficient update is disabled, the weight-update section of the adaptive filters can be powered-down, providing energy savings. The length of the burst cycle is programmable, and using the default values the coefficients update for $2k$ symbols and then are frozen for $16k$ symbols.

2.2 Reconfigurable datapath

It has been noted that the size of the equalizer required in a VDSL receiver depends on the length of the channel and the amount of interference present. Therefore, in order to adapt to different environments, the equalizer length must be made dynamically reconfigurable. This is accomplished by defining a control signal α for each feedforward filter tap which indicates whether the tap is enabled ($\alpha = 1$) or powered down ($\alpha = 0$). The feedback filter is not designed to be reconfigurable because the data precision is much smaller (3 b) and the filter consumes less power. Because there are a total of 96 feedforward filter taps, there are 96 alpha signals which must be controlled by the SMA block.

In a non-folded equalizer, a tap could be powered-down by simply setting the inputs to the multipliers and adders in the tap to zero, ensuring that the tap output is zero and therefore contributes nothing to the rest of the filter. However, using this strategy for disabling a filter tap in the folded structure presented here could actually result in increased power consumption. This occurs because each multiplier is used by four different coefficients in four consecutive clock cycles. Setting one coefficient to zero while the others remain at (possibly) large values would cause a large number of transitions to occur even though the tap is disabled. The solution used here is to "freeze" the inputs to the multiplier when the tap is powered down, and then ignore the tap output for one clock cycle. The filter tap structure which employs this reconfiguration mode is shown in Fig. 3.

The α signal that controls each tap must be carefully timed to ensure that there are no spurious transitions when

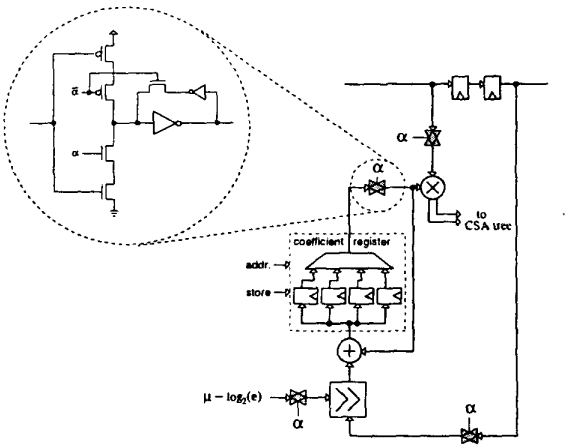

Figure 3. FFE filter tap with DAT features.

a coefficient is powered down. If the current coefficient is enabled ($\alpha_k = 1$) and the next coefficient is powered down ($\alpha_{k+1} = 0$), the alpha signal needs to transition to logic '0' *before* the rising clock edge so that the previous coefficient and data are latched before changing. On the other hand, if the current coefficient is disabled ($\alpha_k = 0$) and the next coefficient is enabled ($\alpha_{k+1} = 1$), the alpha signal must transition to logic '1' *after* the rising clock edge. Therefore, each filter tap also contains a block to produce the correctly-timed alpha signal.

2.3 Reconfiguration strategy

The reconfiguration of the feedforward equalizer is controlled by the SMA block which monitors the environment and determines the optimum equalizer configuration. Because it is difficult on-chip to measure the cable length and number of FEXT interferers, the error across the slicers is taken as a measure of the input signal-to-noise ratio (SNR). Long cable lengths or many interfering transmitters results in reduced SNR at the input and this is manifest in large slicer error. The goal of the SMA block is to keep the slicer error within a specified window by powering up or down specific filter taps. To accomplish this, the SMA block first accumulates the error across the slicers for several thousand symbols and finds the average error E_{avg}. Next this average error is compared to two programmable thresholds T_l and T_h which represent the lower and upper error thresholds, respectively. Three cases result:

1. $T_l \leq E_{avg} \leq T_h$: the slicer error falls inside the window. Therefore the receiver is operating efficiently and no reconfiguration is necessary. The SMA block sits idle for a period, and then begins accumulating the error once again.

2. $E_{avg} < T_l$: the slicer error is less than the lower error threshold. Therefore the receiver performance is higher than necessary and the equalizer size can be reduced, conserving power.

3. $E_{avg} > T_h$: the slicer error is greater than the upper error threshold. The receiver performance is lower than desired, and a powered-down tap must be once again enabled to increase the equalizer size.

If the average error falls outside the thresholds as described above, reconfiguration is necessary. The SMA block next needs to determine the optimum filter tap to power down (or power up). The strategy employed here is the simplified result of an optimization problem which is explained in [5]. In this strategy, the effect of a coefficient on the receiver performance is measured by the absolute value of the coefficient: $|w_k|$. Larger coefficients have a more dramatic effect on output SNR than coefficients small in magnitude. In addition, there is a function $\mathcal{E}_m(w_k)$ which gives an approximation to the amount of energy consumed by the filter tap containing coefficient w_k. Ideally, a detailed energy model for the Booth multiplier and adders in the filter tap would exist, and the energy consumed by each coefficient could be determined. A simplification is used for the SMA architecture described here: the energy model $\mathcal{E}_m$ is implemented as a small lookup table (LUT) containing four programmable constants. The lookup table contains four entries because it was noted through power simulations of the Booth multiplier that the power consumed falls roughly into four levels corresponding to which rows of the multiplier were active. The purpose of the lookup table is to map the coefficient into one of these four energy levels. When powering down a filter tap, it is desirable to choose the coefficient that has the least effect on receiver performance while consuming a large amount of power. This is quantified by picking the coefficient which minimizes the expression

$$\frac{|w_k|}{\mathcal{E}_m(w_k)}$$

Equivalently, the above expression should be maximized when powering up a coefficient.

3 Simulation results

The equalizer architecture described above, along with the other components in the VDSL receiver core, have been implemented in a 0.35 μm CMOS process which will be fabricated through the MOSIS service. The design methodology used is a combination of full-custom design (for the filter taps) and automatic place-and-route (for control logic, SMA block, and full-chip routing). Individual blocks are verified at the circuit level using *Avant! HSPICE* and

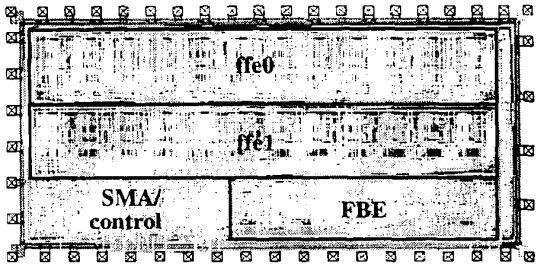

Figure 4. Receiver core layout

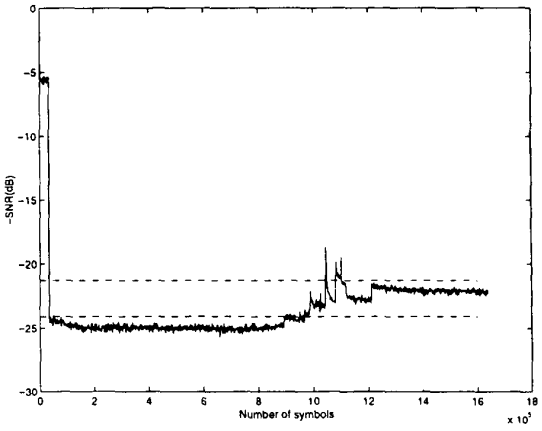

Figure 5. SNR plot for 500-ft cable, 7 FEXT interferers.

Cadence Spectre, while algorithm simulations are accomplished by matching gate-level VHDL simulations with a bit-precise C-language model. The final chip requires a core area of 4.76 mm × 2.14 mm and is shown in Fig. 4. Of this area, approximately 0.6 mm^2 is devoted to the SMA block which represents an area overhead due to DAT of 6%.

Gate-level simulations of the VDSL receiver for different environments verify that the dynamic reconfiguration achieves the desired BER requirement. Figure 5 shows the SNR plot for the case of a 500-ft cable with 7 FEXT interferers. Note that when the equalizer converges, the error is smaller than the desired window of operation (shown with dotted lines). As a result, the SMA block reconfigures the equalizer until the error rises to an acceptable value. The final equalizer for this case contains only 8 powered-up taps.

In order to quantify the power savings made possible through dynamic reconfiguration, the "ffe1" component (consisting of 12 folded taps) was simulated using the circuit simulator *Spectre*. This was performed for three different cable lengths and numbers of FEXT interferers, and the coefficients and alpha values used were obtained first

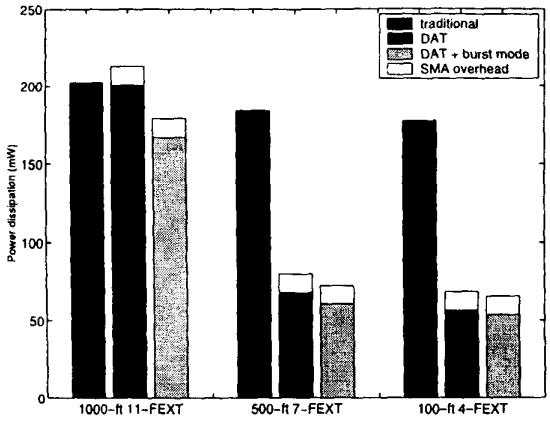

Figure 6. Power dissipation of feedforward equalizer

through C simulation of the desired environment. The total equalizer power consumption is found by taking twice the power of the "ffe1" block (since the two halves of the FFE consume approximately the same amount of power) and adding to that the SMA block power. The power consumption results are shown in Fig. 6. In the worst-case environment (1000-ft cable, 11 FEXT interferers), no reconfiguration is possible and therefore the DAT-based design does not result in power savings. However, for the other two cases the FFE size can be reduced and this results in significant power reduction. The best-case design with burst mode coefficient updating enabled consumes only 37% of the power of a traditional design.

4 Conclusion

Presented in this paper is an architecture for a low-power equalizer that employs dynamic reconfiguration in order to operate efficiently in any environment. This equalizer is used in the design of a 51.84-Mb/s VDSL receiver core and simulated for different combinations of cable length and amount of interference in the channel. These simulations reveal that the equalizer provides large power savings over a traditional design when used in environments that are not worst-case, and requires a small area and power overhead. Such an equalizer would be especially useful in wireless communications applications where battery life is of the utmost concern.

References

[1] 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.

[2] A. Chandrakasan, M. Potkonjak, R. Mehra, J. Rabaey, and R. W. Brodersen, "Minimizing power using transformations," *IEEE Trans. Computer-Aided Design*, vol. 14, no. 1, pp. 12–31, Jan. 1995.

[3] K. K. Parhi, "Algorithm transformation techniques for concurrent processors," *Proceedings of the IEEE*, vol. 77, no. 12, pp. 1879–1895, Dec. 1989.

[4] K. K. Parhi and D. G. Messerschmitt, "Pipeline interleaving and parallelism in recursive digital filters - part I & II," *IEEE Trans. Acoust., Speech, and Signal Processing*, vol. 37, pp. 1099–1134, July 1989.

[5] M. Goel and N. R. Shanbhag, "Dynamic algorithm transforms for low-power reconfigurable adaptive equalizers," *IEEE Trans. Signal Processing*, vol. 47, no. 10, pp. 2821–2832, Oct. 1999.

[6] Davic 1.0 Specification Part 8, *Lower Layer Protocols and Physical Interfaces*, Jan. 1996.

[7] 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.

[8] N. R. Shanbhag and K. K. Parhi, *Pipelined Adaptive Digital Filters*. Boston, MA: Kluwer Academic Publishers, 1994.