

TP 6.2: A Low Power 128-Tap Digital Adaptive Equalizer for Broadband Modems

C.J. Nicol, P. Larsson, K. Azadet, J. H. O'Neill

Bell Laboratories, Lucent Technologies, Holmdel, NJ

This chip provides programmable fractional spacings and slicers making it suitable for 51Mb/s and 155Mb/s ATM over CAT3, as well as for the emerging 100Mb/s base-T2 fast Ethernet standard [1]. The primary design goal is to minimize the power consumption so that the equalizer may be integrated into low-cost single-chip transceivers. Two 64-tap adaptive FIR filters are configured in parallel as in-phase and quadrature filters (Figure 1). Each has a span of $16T$, where T is the symbol period, and is programmable to operate with $T/2$, $T/3$ or $T/4$ fractional spacing. On-chip programmable slicers enable slicing of up to 8×8 constellations. They use a reduced constellation for blind training and switch to the full constellation to obtain final convergence. The filters feature a zero latency cascaded systolic FIR structure that has the low power advantages of the direct form due to the reduced number of flip-flops in the output path, as well as the reduced critical path advantages of the transposed form (Figure 2). A programmable delay synchronizes the input data with the coefficients and the error for correct least mean squares (LMS) coefficient adaptation with different fractional spacings.

The architecture of an $8/12/16$ tap cascaded adaptive FIR filter block is shown in Figure 3. Eight of these modules are used. The filter contains four $10b \times 12b$ Booth encoded multipliers in the FIR block and four power-of-two multipliers in the update section. Programmable fractional spacing is made possible with the use of two-port $4 \times 10b$ register files to provide variable delays. The multipliers are combined into groups of four to further reduce the number of flip-flops in the carry-save output path. The output delays are relocated to the input data path where the word size is 60% smaller. The power of these delays is removed because the power of the register files is independent of their latency. The partial products of each group of multipliers are summed in a single cycle using a tree of 24b carry-save adders. This is added to a 27b cascaded input. A carry-propagate adder at the filter output performs sample rate accumulation with overflow detection and forwards 26b clipped results to the slicer at symbol rate. 26b NEXT/DFE inputs facilitate crosstalk cancellation and post-cursor ISI removal.

Booth encoding the 10b data input to the multiplier would produce fewer partial products and minimal delay, however the switching activity in the multiplier can be reduced by encoding the coefficient inputs because the taps change slowly with time. Further power savings are achieved in time-multiplexed multipliers using the scheme in Figure 4. A typical filter response in Figure 4(a) shows that neighboring outer taps are small in magnitude but changing in sign. A commonly-used Booth encoding circuit would add $-0X$ and $+0X$ in the most significant partial products as the coefficient inputs change sign resulting in unnecessary switching. The lowpow input to the Booth encoding circuit in Figure 4(c) selects either $+0X$ or $-0X$ in the bottom of the truth table so that power comparisons can be made. The graph in Figure 4(d) shows a 30% power reduction when lowpow=1 using the coefficients in Figure 4(a).

The power-of-two LMS tap update algorithm combines fast convergence with symbol rate tap updating and reasonable power consumption because it replaces the update multipliers with barrel shifters. Updating can be enabled on a per-tap basis. By setting a tap to zero and prohibiting its updating, the effective

filter length can be reduced as suggested in [2]. Each tap provides 23b precision, overflow detection and clipping — between the master and slave stages of the coefficient register file to minimize the critical path. Taps are randomly addressable through a reduced swing coefficient bus without halting the FIR.

An error monitor with each filter accumulates the error and disables tap updating whenever the average error falls below a programmable threshold (Figure 5). Hysteresis is provided by keeping the taps frozen until the error rises above a second threshold. When the taps are frozen, only the 12MSBs of the coefficient register files are active; the LSBs and all updating circuits are disabled to reduce power. A gain control is included on the output of the filter for adaptive bit-precision adjustment. When the error is small, the error monitor increases the gain of the filter which in turn reduces the amplitude of the taps. This reduces power in the FIR filter according to the graph in Figure 4(d).

The chip operates at an 80MHz sample rate to satisfy ATM 155Mb/s UTP and 100Mb/s Ethernet T2. The adaptive FIR blocks occupy $18mm^2$. The total active area is $21mm^2$ in three-level-metal $0.5\mu m$ CMOS. The 440k-transistor design is completed using full-custom symbolic layout. Simulations show an expected power dissipation of 975mW at 80MHz 3.3V with symbol rate updating. This drops to 390mW with tap updates every 10 symbols and reduced precision taps (error permitting). The power in the FIR filter is 415mW with maximum precision converged coefficients. This corresponds to 13.1mW per multiplier (which could be reduced depending on the application). Table 1 compares these estimates with previous work (Numbers are scaled to reflect differences in power supply, technology, multiplier size and speed. Hard-coded numbers correspond to this implementation).

$$P(mult) = \frac{Total\ Power}{\#mults} \times \frac{12}{\#bits\ coeff} \times \frac{10}{\#bits\ sample} \times \left(\frac{33}{Vdd}\right)^2 \times \frac{05}{Tech} \times \frac{80}{clk\ freq}$$

This assumes linear scaling which favors older technologies due to the wire fringing field capacitance. The filters in [4, 5, 6, 7] all have one multiplier per tap, but the filters in this work and in Reference 3 have time-multiplexed multipliers where the coefficient inputs are switching. This work demonstrates techniques to reduce the power of adaptive filters employing time-multiplexed multipliers to a level below that of previous filters with fixed coefficients. Power savings are shown in Figure 6.

References:

- [1] Im, G. H., J.J. Werner, "Bandwidth-efficient digital transmission up to 155Mb/s over unshielded twisted-pair wiring," Proc. IEEE ICC, pp. 1797-1803, 1993.
- [2] Ludwig, J., H. Nawab, A. Chandrakasan, "Low-Power Digital Filtering Using Approximate Processing," IEEE JSSC, pp. 395-400, 1996.
- [3] Choi, J. R., et al., "Structured design of a 288-tap FIR Filter by Optimized Partial Product Tree Compression," Proc. IEEE CICC, pp. 79-82, 1996.
- [4] Abbot, W., et al., "A Digital Chip with Adaptive Equalizer for PRML Detection in Hard-Disk Drives," IEEE ISSCC Digest of Technical Papers, pp. 284-285, 1994.
- [5] Pearson, D., et al., "250MHz Digital FIR Filters for PRML Disk Read Channels," IEEE ISSCC Digest of Technical Papers, pp. 80-81, Feb., 1995.
- [6] Thon, L., et al., "A 240MHz 8-Tap FIR Filter for Disk-Drive Read Channels," IEEE ISSCC Digest of Technical Papers, pp. 82-83, Feb., 1995.
- [7] Hatamian, M., S. K. Rao, "A 100MHz 40-Tap Programmable FIR Filter Chip," Proc. IEEE ISCAS, pp. 3053-3056, 1990.

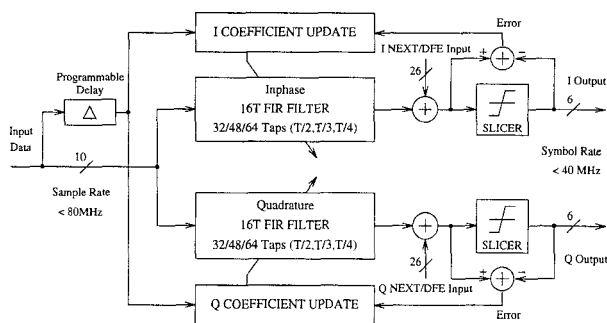


Figure 1: System architecture.

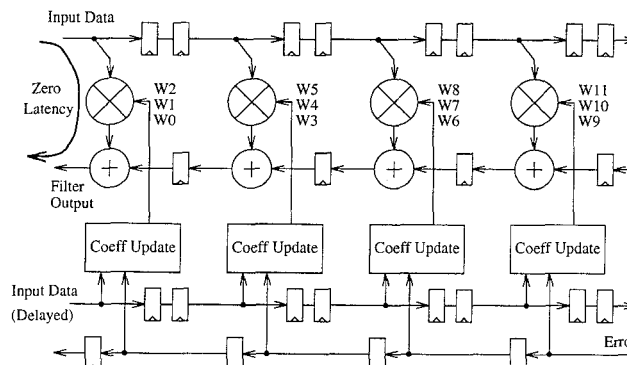


Figure 2: Zero latency adaptive filter structure (T/3).

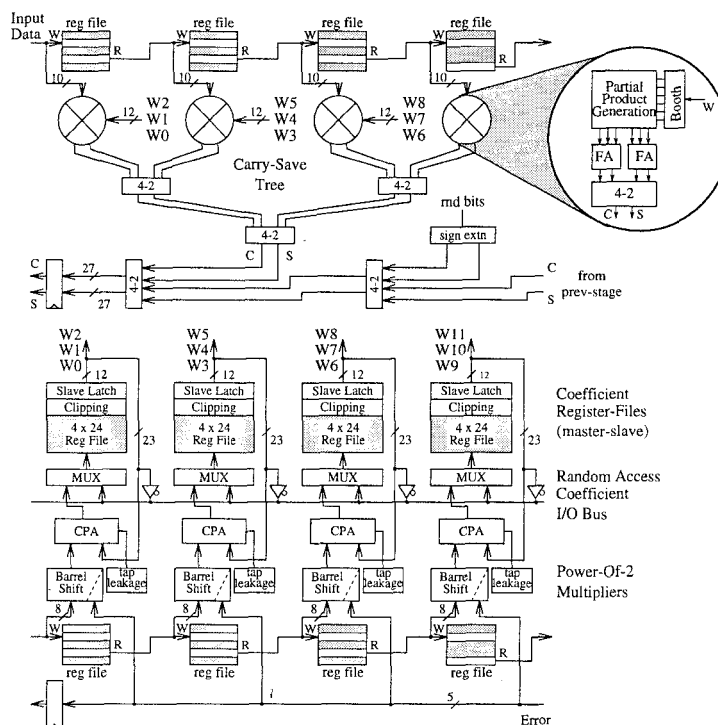




Figure 3: Architecture of 8/12/16 tap adaptive FIR filter.




(a) Raised Cosine Filter Response

Multiplier Bits			Operator	Output Bits		
Mi+1	Mi	Mi-1		X1	X2	NEG
0	0	0	0	0	0	0 (adds +0X)
0	0	1	+1X	1	0	0
0	1	0	+1X	1	0	0
0	1	1	+2X	0	1	0
1	0	0	-2X	0	1	1
1	0	1	-1X	1	0	1
1	1	0	-1X	1	0	1
1	1	1	0	0	0	1 (adds -0X)

(b) Booth Encoding Table



(c) Low-Switching NEG Circuit



(d) Power as a function of Coefficient Precision

Figure 4: Low-power Booth encoding.

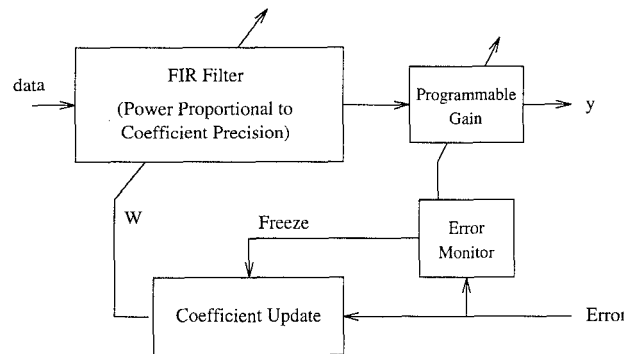


Figure 5: Adaptive coefficient precision.
 Figure 6, Table 1: See page 437.




Figure 6: Micrograph.

TP 6.2: A Low Power 128-Tap Digital Adaptive Equalizer for Broadband Modems
(Continued from page 95)

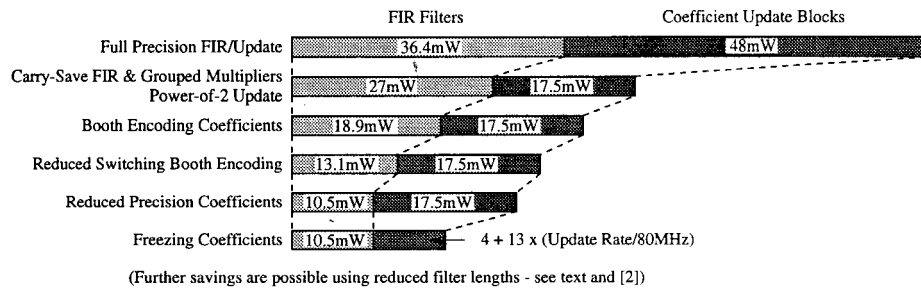


Figure 6: Power reduction in adaptive equalizer.

Ref	Description	Tech (μm)	Vdd (V)	Power (W)	P(Mult) (mW)
[3]	288 tap FIR, 72 10x10 mult@60MHz	0.6	3.3	3.5	72
[4]	9 tap FIR, 9 10x6 mult@72MHz	0.8	5.0	0.5	34
[5]	8 tap FIR, 6x6 mult@200MHz	0.5	3.3	0.25	42
[6]	8 tap FIR, 6x6 mult@240MHz	0.8	3.7	0.426	29
[7]	40 tap FIR, 40 10x12 mult@100MHz	0.9	5.0	3.1	15
*	128 tap FIR, 32 10x12 mult@80MHz	0.5	3.3	0.415	13.1

Table 1: Comparison with existing FIR filters.