

A 50 MHz Eight-Tap Adaptive Equalizer for Partial-Response Channels

Caesar S. H. Wong, Jacques C. Rudell, Gregory T. Uehara, and Paul R. Gray, *Fellow, IEEE*

Abstract—A new architecture for digital implementation of the adaptive equalizer in Class IV partial-response maximum likelihood (PRML) channels employing parallelism and pipelining is described. The architecture was used in a prototype integrated circuit in a 1.2 μm CMOS technology to implement an eight-tap adaptive equalizer and Viterbi sequence detector which consumes a total of 70 mW from a 3.3 V supply operating at an input sampling rate of 50 MHz.




Fig. 1. Block diagram of a typical PRML magnetic disk read channel.

I. INTRODUCTION

SAMPLED-DATA techniques such as Class IV partial response with maximum likelihood detection (PRML) are being applied to magnetic disk storage channels in order to increase transfer rates and recording densities [1]–[2]. In order to provide a robust implementation of the functions required in the read path of these channels, implementation of key blocks such as timing recovery, adaptive equalization, and sequence detection is often in the digital domain. The power consumed by these blocks can be appreciable due to the high speeds of operation required with data rates currently on the order of 50–100 Mb/s and exceeding this in the future. For example, the power dissipation of the CMOS logic alone in a recently reported BiCMOS 65 Mb/s read channel IC implementing digital equalization, sequence detection, and other associated functions was 1 W [3]. Powerful battery-operated portable personal computers are becoming increasingly prevalent, and the performance requirements of the storage devices in these systems continually increases. Power consumption is a critical parameter in such systems for increased battery life and to minimize heat dissipation effects due to the proximity of the electronics to the magnetic media in drives with decreasing form factors. The net result is the need for increased speed performance of the electronics at a reduced overall power consumption.

Manuscript received August 5, 1994; revised November 14, 1994. This work was supported by NSF Grant MIP 9109525, the California Micro Program, and National Semiconductor.

C. S. H. Wong was with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720 USA. He is now with Rockwell International, Newport Beach, CA 92658 USA.

J. C. Rudell and P. R. Gray are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720 USA.

G. T. Uehara was with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720 USA. He is now with the Department of Electrical Engineering, University of Hawaii at Manoa, Honolulu, HI 96822 USA.

IEEE Log Number 9408733.

A block diagram of a typical PRML read channel is shown in Fig. 1. The output of the magnetic disk is first amplified by the read amplifier before being passed on to the analog front end which includes a variable gain amplifier (controlled in an automatic gain control loop not shown), low-pass filter, sampler, and analog-to-digital converter followed by the functions of adaptive equalization, sequence detection, and timing recovery in the digital domain. The adaptive equalizer operates on the 6 b samples coming from the A/D converter, equalizing these samples for subsequent detection by the sequence detector and use by the timing recovery block.

This paper describes a low-power architecture for digital implementation of an adaptive equalizer suitable for use in PRML magnetic disk drive channels. The new architecture was used in a prototype IC containing an eight-tap adaptive equalizer and Viterbi detector which, at a 50 MHz input sampling rate, consumes a total power of 70 mW operating from a 3.3 V power supply in a 1.2 μm CMOS technology [14]. The reduction in operating power relative to the current practice was made possible through the use of an optimum mixture of parallelism and pipelining, which in turn allowed the use of a 3.3 V supply while meeting the throughput requirements.

II. LOW-POWER ADAPTIVE EQUALIZER ARCHITECTURE AND DESIGN

As has been recently demonstrated in other work [4], major reductions in power dissipation in such digital functions are possible through the use of lower power supply voltages since the power consumed by a CMOS digital circuit is proportional to CV^2f . The corresponding increase in gate delay is only linear with supply voltage, all else being equal, so that the effective power–delay product is improved at lower voltages. The increased gate delay can be accommodated through the use of a combination of parallelism and pipelining, requiring more gates for the implementation, but giving major reductions in overall power dissipation. The optimum combination of




Fig. 2. Four multipliers operating staggered in phase by the output period T to realize an increased multiplying rate.

parallelism and pipelining is highly dependent on the particular function being performed, and the exploration of these trade-offs for the implementation of an adaptive equalizer suitable for use in PRML magnetic disk drive read channels is the main topic of this paper.

Magnetic disk drive read channels are unique both because of the high sampling and bit rates involved, and because of the relatively modest SNR of the signals detected off the disk. Most implementations of PRML channels use a 6 b signal representation into the adaptive equalizer dictated by the off-channel signal-to-noise ratio. This small word width has major implications for the minimization of power through the use of parallelism since multipliers for this width are relatively compact. Extensive simulations were performed to explore various permutations of parallelism and pipelining in the implementation of a conventional multiplier [5] for filter sampling rates above 50 MHz and a power supply of 3.3 V. It was found that for implementation of a 6×6 b multiplier, the use of four multipliers operating in parallel and staggered in phase by the output period T as shown in Fig. 2 resulted in the solution dissipating the lowest overall power. This advantage of low power comes at the cost of increased silicon area which, however, will scale with technology. For the required resolution of 6 b, the overhead associated with pipeline latches in a pipelined implementation as opposed to a parallel implementation of the multiplier increases the power while decreasing the attainable speed (due to latch setup times).

While the parallel architecture helps to increase overall throughput for a given supply voltage, system latency is also increased. This latency will have implications from a system standpoint, particularly in applications in which the block employing parallelism is nested in one or more feedback loops. The resulting delay will affect the stability of the feedback loops and must be considered in the system design.

A. Parallel Filter Architecture

A block diagram of a *filter stage* used to implement the equalizer is shown in Fig. 3(a), and is comprised of a delay line, a set of multipliers, and an accumulator. The morphology of the filter stage resembles the block diagram of an FIR filter. In order to use multipliers which take four output periods to perform multiplication, four filter stages are used in parallel. The multiplier and accumulator sections are each clocked at one-fourth the output rate of the equalizer and operate staggered in phase by one output period. The parallel architecture and timing diagrams for this approach are shown




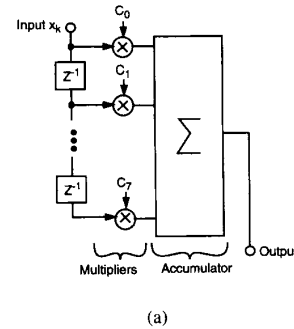


Fig. 3. Conceptual illustration of parallel filter structure and the use of pipelining within each filter stage to achieve high throughput. (a) FIR filter stage with delay line, multipliers, and accumulator. (b) High-rate filter realized by four parallel interleaved quarter-rate filter stages. (c) Timing diagram of parallel filter stages shown in (b). (d) Conceptual illustration of pipelining of the multiplication and accumulation required for one filter stage.

in Fig. 3(b) and (c), respectively. The use of input latches in both the multipliers and accumulator enable the functions of multiplication and accumulation to be pipelined, as shown for one filter stage in Fig. 3(d). During each filter stage clock cycle, the current eight inputs into a filter stage are multiplied by their respective tap weights. The resulting products are summed by the accumulator during the next clock cycle of that filter stage.

B. Multiplier Circuit Design

Much of the discussion thus far has focused on block level considerations to reduce power consumption of the multiplier function. Further optimization can occur with appropriate design choices at the circuit level. The multipliers were constructed using the Baugh–Wooley two’s complement parallel array algorithm [5]. The benefits of two logic styles,

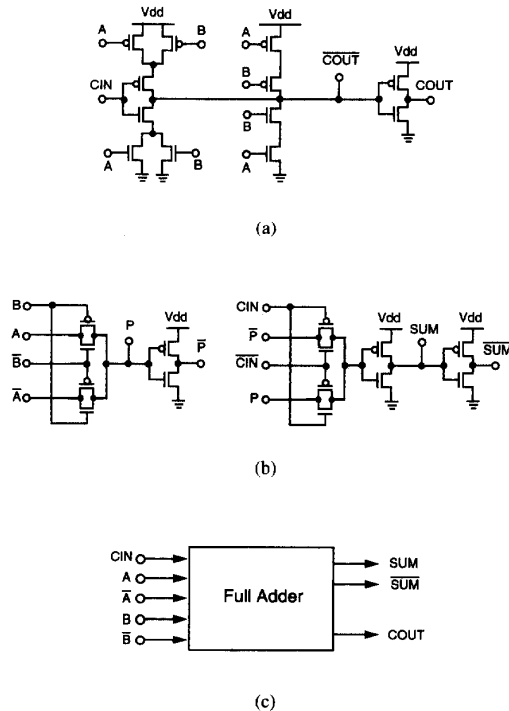


Fig. 4. Full adder design utilizing both combinatorial and complementary pass-transistor logic. (a) Carry generation. (b) Sum generation. (c) Block diagram.

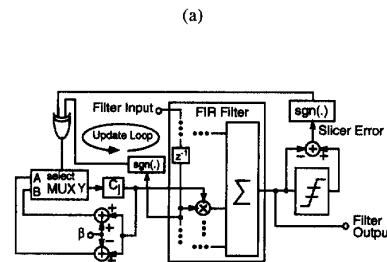

complementary pass logic and static CMOS, were combined to implement the full adders, as shown in Fig. 4. To reduce the number of transistors and subsequently the power consumption associated with extra parasitic drain and source capacitances, complementary pass logic was used in the sum generation portion of the full adder, while in order to increase the speed and drive capability, static CMOS was used at both the output of the sum generator and to implement the carry generation signal.

C. Coefficient Update

The least mean square (LMS) algorithm is often used to update the tap weights in the adaptive equalizer. In the LMS algorithm, each coefficient C is updated using the equation

$$C_{k+1} = C_k + \beta e_k x_k \quad (1)$$

where β is the step size, e_k the input error to the slicer at time k , and x_k the input to the particular tap weight at time k . Implementation of this algorithm requires two multiplies following generation of the slicer error in order to obtain the correction term which is added to the current value of the coefficient as shown in Fig. 5(a). The time required for these multiplies adds directly to the latency in the coefficient update. Simulation results indicated that the latency associated with the parallel multipliers degraded stability in the adaptation loop, in turn requiring very small values of β and long adaptation times to eliminate stability problems.



Full LMS		Sign-Sign LMS	
1) Multiply inputs and tap weights:	4T	1) Multiply inputs and tap weights:	4T
2) Sum products:	2T	2) Sum products:	2T
3) Generate error e_k :	1T	3) Generate $\text{sgn}(\cdot)$ of the error e_k and XOR with $\text{sgn}(\cdot)$ of the input and choose $\alpha_k + \beta$ or $\alpha_k - \beta$:	2T
4) Multiply input and e_k :	4T		
5) Multiply result of (4) by β :	4T		
6) Sum product with current value of the tap weight:	1T		
Total:	16T	Total:	8T

Fig. 5. Block diagram of the coefficient update circuitry. (a) Coefficient update loop using the full LMS algorithm. (b) Implementation of the coefficient update loop using the sign-sign LMS algorithm. (c) Latency in number of output periods T in the coefficient update loop using the full and sign-sign LMS algorithms.

In order to reduce this feedback latency in the update of the coefficients, the sign-sign LMS algorithm was employed where the coefficients are updated using the equation

$$C_{k+1} = C_k + \beta \text{sgn}(e_k) \text{sgn}(x_k) \quad (2)$$

where $\text{sgn}(\cdot)$ is the signum function which is $+1$ or -1 for a positive or negative argument, respectively. The coefficient update then reduces to

$$C_{k+1} = C_k \pm \beta \quad (3)$$

depending on the product of $\text{sgn}(e_k) * \text{sgn}(x_k)$. In the implementation, both $c_k + \beta$ and $c_k - \beta$ are computed in parallel, while $\text{sgn}(e_k) * \text{sgn}(x_k)$ is computed using an exclusive-OR operation on the sign bits of both the error and the input to the respective tap. Depending on the outcome of the exclusive-OR operation, $c_k + \beta$ or $c_k - \beta$ is selected. This is shown schematically in Fig. 5(b). This approach reduces the feedback latency from 16 periods down to 8, as listed in Fig. 5(c), and performs robustly. Another advantage of the sign-sign LMS algorithm is that the multiplications required in the full LMS algorithm reduce to a single exclusive-OR, which results in significant power and area savings. Information on the stability of the sign-sign LMS algorithm can be found in [7]–[10].

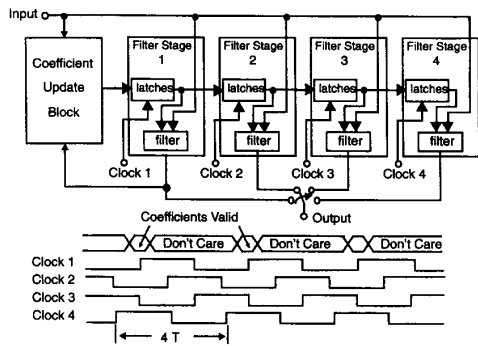


Fig. 6. Routing of the filter coefficients through each filter stage on to the next.

Because the disk drive channel is slowly time varying, the output of only one of the four parallel filter stages is used in the update of the coefficients. This same set of coefficients is used by all four filter stages. In order to reduce timing requirements on the availability of the coefficient set to each of the four filter stages, the approach shown in Fig. 6 is used. New coefficients from the coefficient update block need be available only during the latching instant of Filter Stage 1. The coefficients are held by the Filter Stage 1 latches for the multipliers there, as well as applied to the inputs of the latches in Filter Stage 2 where they are sampled one output period later. This process is repeated until the coefficients are latched into Filter Stage 4. Since the outputs from the coefficient update block need only be available during the latching instant of Filter Stage 1, their availability requirements are much less than if the coefficient update block were to feed each of the filter stages directly. Furthermore, the amount of parasitic capacitance which would otherwise have to be driven by one set of buffers if all four filter stages were to receive the coefficients directly is greatly reduced.

A three-tap raised-cosine equalizer response with tap weights $-K, 1.0, -K$ was assumed for the initial equalizer impulse response before convergence. An externally controlled initialization signal allows the fourth–sixth coefficients to be initialized to tap weight values provided off-chip. While these three taps are set to their respective values (fourth and sixth tap weights are the same), the other five tap weights are forced to zero. The ability to externally set each of the tap weights individually could easily have been implemented, but was avoided in the prototype for simplicity. The value of the step size β is also provided externally, and can be programmed to implement gear-shifting algorithms for faster convergence.

D. Input Data Flow

The chip receives two 6 b words at a one-half rate, relaxing the timing requirements for latching of the input data while having no effect on the equalizer performance. Input samples are supplied via a tapped delay line shown in Fig. 7. Samples are first latched into Filter Stage 1 on the rising edge of Clock 1. Seven of the samples residing in the Filter Stage 1 latches, while being held for the multipliers, are passed through to

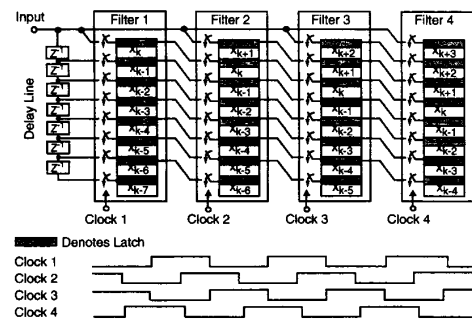


Fig. 7. Conceptual diagram of input data flow through the four filter sections.




Fig. 8. Viterbi detector implemented with two half-rate detectors.

Filter Stage 2 to be latched along with one new sample on the rising edge of Clock 2. This process is again repeated for Filter Stages 3 and 4 on the rising edge of Clocks 3 and 4, respectively. This approach reduces the amount of data needing to be passed around the chip at high speed and relaxes timing requirements on the movement of most of the input data throughout the chip. Note that the desired convolution occurs as a vector of input samples moves diagonally through the parallel stages.

III. SEQUENCE DETECTOR

The sequence (Viterbi) detector is designed for a $PR - IV$ target, and as such is realized by two interleaved independent half-rate Viterbi detectors [6] operating on the outputs of Filter Stages 1 and 3 and 2 and 4, as shown in Fig. 8. The input words into the Viterbi detector blocks are 6 b wide. The outputs of both detectors contribute to a two-wide one-half rate bit stream which is output off-chip (commutator shown in the figure). Descriptions of Viterbi detector implementations can be found in the literature [16]–[17].

IV. CHIP PLAN AND LAYOUT

Due to the regularity in the flow of both the input data and the filter taps throughout the chip (shown in Figs. 6 and 7), a close mapping of the block diagram to the layout




Fig. 9. Detailed block diagram of Filter Stage 1. Both signals CLO and ILO are output to Filter Stage 2.

is possible. A detailed schematic of a filter stage is shown in Fig. 9, indicating the flow of the input data and coefficient into each multiplier cell. The two words enter their respective latches where they are held for the duration of the multiply, as well as passed through the cell to the next filter stage. A high-level view illustrating the flow of the input data and coefficients through the filter is shown in Fig. 10. A die photo of the chip is shown in Fig. 11. Note that the layout closely resembles the block diagram shown in Fig. 10.

V. EXPERIMENTAL RESULTS

A prototype IC including the adaptive equalizer and sequence detector was fabricated in a $1.2 \mu\text{m}$ Mosis CMOS process.

To verify performance, the chip was supplied via a logic analyzer with pseudorandom channel samples generated using Ptolemy, a system simulation tool [11]. A linear channel was assumed and modeled using a Lorentz pulse [18] to represent transitions on the disk. PW_{50}/T is a parameter used in the Lorentz model as a relative measure of the intersymbol interference (ISI) in the channel, where higher numbers for PW_{50}/T correspond to more ISI. Typical disk drive channels today can be characterized with values of PW_{50}/T between 1.5 and 2.0, and will be approaching 2.5 in the future. Since it requires the most boost from the equalizer, a PW_{50}/T value of 2.5 was used to generate channel samples for all functional testing of the prototype. Experiments were performed by initializing the equalizer taps to intentionally introduce mis-equalization and feeding the generated channel samples via the logic analyzer to the prototype and observing the equalizer outputs and the tap weight values over time. The equalizer outputs approached the desired target response, while the tap weights converged to their desired values for pseudorandom data. Functionality

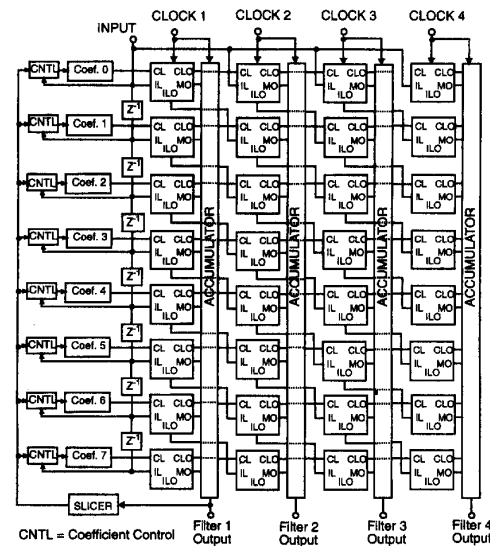


Fig. 10. Detailed schematic illustrating the flow of the input signals and coefficients through the filters. This block diagram maps into the layout, as can be seen in the die photo.




Fig. 11. Die photograph of prototype IC containing the eight-tap adaptive equalizer and Viterbi detector.

was further verified by comparing the equalizer outputs to previous system simulations.

The power consumption for a supply voltage (V_{dd}) of 3.3 and 5.0 V is plotted as a function of output rate in Fig. 12. The prototype circuit was originally designed to operate at 100 MHz with a power supply of 3.3 V. Due to an error in the extraction procedure during the design process, the capacitance in the critical path through the accumulator and coefficient update circuitry was underestimated. As a result, in order to achieve proper operation at 100 MHz, the supply voltage had to be increased to 5.0 V. Extrapolating the power consumption at low frequencies suggests that a 3.3 V, it is feasible that 100 MHz operation could be achieved in




Fig. 12. Measured total power plotted as a function of operating frequency for V_{dd} equal to 3.3 and 5.0 V. Also indicated is the extrapolated power consumption for 100 MHz operation at 3.3 V.

TABLE I
PERFORMANCE RESULTS

Parameter	Value
Technology	1.2 μ m CMOS
Die Size	217 mil by 260 mil
Number of Taps in Adaptive Equalizer	8
Input Resolution	6 bits
Internal Equalizer Coefficient Length	10 bits
Number of Bits in the Multipliers	6 bits by 6 bits
Algorithm used for Updating Coefficients	Sign-Sign LMS
Total Power Dissipation @ 100 MHz $V_{dd} = 5$ V	338 mW
Total Power Dissipation @ 50 MHz $V_{dd} = 3.3$ V	70 mW
Parallel Filter Stages	40 mW
Coefficient Update Block	9 mW
Clock Generation	14 mW
Viterbi Detector	7 mW
Extrapolated Total Power @ 100 MHz $V_{dd} = 3.3$ V	140 mW

a redesign with a power consumption below 200 mW. A summary of the key performance characteristics is given in Table I. A comprehensive summary of test results can be found in [12] and [13].

VI. SUMMARY

In this paper, an architecture for implementation of adaptive equalizers as required in applications such as the PRML magnetic disk read channel was described. The architecture utilizes parallelism and pipelining which enables high-speed operation at a reduced power supply, resulting in a speed-power ratio of 1.4 mW/MHz (for eight taps), which compares favorably to conventional approaches. The design of a prototype digital detector including the functions of adaptive equalization and sequence detection was described. The prototype demonstrates the potential of the proposed architecture to implement these key functions at high speed in a relatively conservative technology at a low-power consumption.

The speed-power ratio advantage of the described approach comes at the cost of increased hardware, system complexity, and latency. The optimum combination of parallelism and pipelining depends on the requirements of a particular system and the available power supply values. In most applications, the power supply values are set by those already available in the particular system, although efficient dc-dc conversion

techniques are specifically being developed to optimize the combination of system performance and power consumption [15]. Latency affects adaptive algorithms such as those used to control the coefficients of the equalizer (in the case of the adaptive equalizer), as well as those which operate on the equalizer outputs. The allowable latency depends upon the convergence requirements of the different algorithms within a particular system, and simulations must be performed to explore the specific performance/power savings tradeoffs particular to that system.

ACKNOWLEDGMENT

The authors would like to thank A. Chandrakasan, S. Sheng, C. Conroy, members of the IC Group, and Prof. R. Brodersen and Prof. J. Rabaey and their research groups at UC Berkeley for their guidance, support, and encouragement.

REFERENCES

- [1] H. Kobayashi and D. T. Tang, "Application of partial-response channel coding to magnetic recording systems," *IBM J. Res. Develop.*, vol. 14, pp. 368-375, July 1970.
- [2] R. D. Cideciyan, F. Dolivo, R. Hermann, W. Hirt, and W. Schott, "A PRML system for digital magnetic recording," *IEEE J. Select. Areas Commun.*, vol. 10, pp. 38-56, Jan. 1992.
- [3] R. Philpott, R. Kertis, R. Richetta, T. Schmerbeck, and D. Schulte, "A 7 MB/sec (65 MHz) mixed-signal magnetic recording channel DSP using partial response signaling with maximum likelihood detection," in *Proc. 1993 CICC*, May 1993, pp. 10.4.1-10.4.4.
- [4] A. P. Chandrakasan, S. Sheng, and R. W. Brodersen, "Low power CMOS digital design," *IEEE J. Solid State Circuits*, vol. 27, pp. 473-484, Apr. 1992.
- [5] C. R. Baugh and B. A. Wooley, "A two's complement parallel array multiplication algorithm," *IEEE Trans. Comput.*, vol. C-22, pp. 1045-1047, Dec. 1973.
- [6] R. W. Wood and D. A. Peterson, "Viterbi detection of class IV partial response on a magnetic recording channel," *IEEE Trans. Commun.*, vol. COM-34, pp. 454-461, May 1986.
- [7] P. M. Clarkson, *Optimal and Adaptive Signal Processing*. Boca Raton, FL: CRC Press, 1993, pp. 281-287.
- [8] S. Dasgupta and C. R. Johnson, Jr., "Some comments on the behavior of sign-sign adaptive identifiers," *Syst. Contr. Lett.*, vol. 7, pp. 75-82, Apr. 1986.
- [9] S. Dasgupta, C. R. Johnson, Jr., and A. M. Baksho, "Sign-sign LMS convergence with independent stochastic inputs," *IEEE Trans. Inform. Theory*, vol. 36, pp. 197-201, Jan. 1990.
- [10] E. Eweda, "Convergence of the sign algorithm for adaptive filtering with correlated data," *IEEE Trans. Inform. Theory*, vol. 37, pp. 1450-1457, Sept. 1991.
- [11] E. A. Lee *et al.*, *Ptolemy Manual*. Berkeley, CA: UC-Berkeley, 1991.
- [12] C. S. H. Wong, "Low-power high-speed DSP architecture for magnetic disk PRML read channel," M.S. thesis, Univ. California, Berkeley, Memo. UCB/ERL M93/72, Oct. 1993.
- [13] J. C. Rudell, "A low-power high-speed digital adaptive equalizer for magnetic disk drive channels utilizing class IV partial response signalling," M.S. thesis, Univ. California, Berkeley, Memo. UCB/ERL M94/14, Mar. 1994.
- [14] G. T. Uehara, C. S. H. Wong, J. C. Rudell, and P. R. Gray, "A 50 MHz 70 mW 8-tap adaptive equalizer/Viterbi sequence detector in 1.2 μ m CMOS," in *Proc. 1994 CICC*, San Diego, CA, May 1994, pp. 4.2.1-4.2.4.
- [15] A. Stratakos, S. Sanders, and R. Brodersen, "A low-voltage CMOS DC-DC converter for a portable battery-operated system," in *Proc. IEEE Power Electron Specialists Conf.*, 1994, pp. 619-626.
- [16] P. J. Black and T. H. Meng, "A 140 Mb/s 32-state radix-4 Viterbi decoder," *IEEE J. Solid-State Circuits*, vol. 27, pp. 1877-1885, Dec. 1992.
- [17] P. A. Ziperovich and J. K. Wolf, "CMOS implementation of a Viterbi detector for hard disk drives," in *Proc. 1993 CICC*, San Diego, CA, May 1993, pp. 10.3.1-10.3.4.

- [18] J. M. Cioffi, W. L. Abbott, H. K. Thapar, C. M. Melas, and K. D. Fisher, "Adaptive equalization in magnetic-disk storage channels," *IEEE Commun. Mag.*, pp. 14-29, Feb. 1990.



Caesar S. Wong received the B.S. degree in electrical engineering from the University of Texas at Austin in 1991, and the M.S. degree in electrical engineering and computer sciences from the University of California, Berkeley, in 1993.

From 1990 to 1991, he was an Undergraduate Student Researcher at the Microelectronics Research Laboratory at the University of Texas. From 1992 to 1993, he worked as a Graduate Student Researcher at the Electronics Research Laboratory at the University of California, Berkeley, where he was involved in high-speed DSP design for disk drive application. In October 1993, he joined Rockwell International, Newport Beach, CA, where he is presently working on mixed-signal IC's for wireless communication applications.

Mr. Wong was a recipient of a Regents Fellowship from the University of California.



Jacques C. Rudell was born in Grand Rapids, MI, on July 13, 1965. He received the B.S.E.E. degree from the University of Michigan, Ann Arbor, in 1989, and the M.S.E.E. degree from the University of California, Berkeley, in 1994, focusing on digital adaptive equalization for magnetic disk drive applications.

After graduation in 1989, he spent two years as an IC Designer and Project Manager at Delco Electronics, Kokomo, IN. While at Delco, his work focused mainly on bipolar analog circuits for automotive applications. He is currently working towards the Ph.D. degree at the University of California, Berkeley, concentrating on CMOS mixer circuits for RF applications.

Mr. Rudell is a member of Tau Beta Pi and Eta Kappa Nu. In his junior year of undergraduate studies, he was named a James B. Angell Scholar.



Gregory T. Uehara was born in Honolulu, HI, on July 1, 1960. He received the B.S. degree in pre-engineering from the College of Idaho, Caldwell, and the B.S. degree in electrical engineering from the University of Washington, Seattle, both in 1983, and the M.S. and Ph.D. degrees, both in electrical engineering, from the University of California, Berkeley, in 1989 and 1993, respectively.

He worked as a Design Engineer in mixed-signal CMOS integrated circuit design, first in the Telecommunications Group at Intel Corporation, Chandler, AZ, from 1983 to 1985, and then the Custom CMOS Circuit Design Group at MicroLinear Corporation, San Jose, CA, from 1985 to 1986. He was both a Teaching Assistant and Research Assistant during his studies at Berkeley. Since the fall of 1993, he has been an Assistant Professor in the Department of Electrical Engineering, University of Hawaii at Manoa, Honolulu. His research at Berkeley focused on CMOS mixed-signal integrated circuit implementation of high-speed magnetic disk read channels. His current research interests include the design of high-speed analog and digital integrated circuits for communication systems and magnetic storage systems.



Paul R. Gray (S'65-M'69-SM'76-F'81) was born in Jonesboro, AR, on December 8, 1942. He received the B.S., M.S., and Ph.D. degrees from the University of Arizona, Tucson, in 1963, 1965, and 1969, respectively.

In 1969 he joined the Research and Development Laboratory, Fairchild Semiconductor, Palo Alto, CA, where he was involved in the application of new technologies for analog integrated circuits, including power integrated circuits and data conversion circuits. In 1971 he joined the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, where he is now a Professor. His research interests during this period have included bipolar and MOS circuit design, electrothermal interactions in integrated circuits, device modeling, telecommunications circuits, and analog-digital interfaces in VLSI systems. He is the coauthor of a widely used college textbook on analog integrated circuits. During year-long industrial leaves of absence from Berkeley, he served as Project Manager for Telecommunications Filters at Intel Corporation, Santa Clara, CA, during 1977-1978, and as Director of CMOS Design Engineering at Microlinear Corporation, San Jose, CA, during 1984-1985. At Berkeley, he has held several administrative posts including Director of the Electronics Research Laboratory (1985-1986), Vice-Chairman of the EECS Department for Computer Resources (1988-1990), and Chairman of the Department of Electrical Engineering and Computer Sciences (1990-1993). He currently holds the Edgar R. and Harold H. Buttner Chair in Electrical Engineering at Berkeley.

Dr. Gray has been co-recipient of Best Paper Awards at the International Solid State Circuits Conference and the European Solid-State Circuits Conference, and was co-recipient of the IEEE R. W. G. Baker Prize in 1980, the IEEE Morris K. Liebman award in 1983, and the IEEE Circuits and Systems Society Achievement Award in 1987. In 1994 he received the IEEE Solid-State Circuits Award. He served as Editor of the *IEEE JOURNAL OF SOLID-STATE CIRCUITS* from 1977 through 1979, and as Program Chairman of the 1982 International Solid State Circuits Conference. He served as President of the IEEE Solid-State Circuits Council from 1988 to 1990. He is a member of the National Academy of Engineering.