A High Speed, Low Power PRML Read Channel Device

Jeff Sonntag, Oscar Agazzi, Pervez Aziz, Harley Burger, Vittorio Comino, Martin Heimann, Terence Karanink, John Khoury, Gary Madine, Krishnaswami Nagaraj, Glen Offord, Robert Peruzzi, Joe Plany, Narendra Rao, Necip Sayiner, Palaksha Setty, and Kevin Threadgill

AT&T Bell Laboratories, Allentown, PA & Murray Hill, NJ

Abstract-**A complete read channel device using PR-IV signalling and maximum likelihood detection is described. Descriptions, simulated performance and measured results of the digital adaptive feedback loops (AGC, FIR tap weights, DFE tap weights, frequency and phase) in the read channel are presented. Analog FIR filter and flash A/D converter performance is presented. The weighted averaging servo demodulation technique used in the device is described and compared to integration. This monolithic CMOS device consumes less than 0.75 W when reading at 85 Mbps, and occupies a step-and-repeat area of 24.6 mm^2.**

I. Introduction

While PRML (Partial Response, Maximum Likelihood) read channels have existed [1] for many years, monolithic implementations of the technique [2-5] have appeared only recently, and have suffered from high read mode power consumption, low maximum data rates, and high chip area. A large portion of these problems can be traced to the use of digital filtering, in particular, to the use of digital FIR equalization, which, at current linewidths, and at the required speeds, is a very power and area intensive technique; in order to avoid this expense, an inadequate number of FIR taps has sometimes been used [2], but this is not an appealing solution. In addition to the area and power benefits, it will be shown that the required resolution of the A/D converter can be reduced when the FIR filter is analog.

As track spacing in rotating media reduces each year, it should be obvious that a high-quality read channel is useless if it is coupled to a lower quality servo demodulator. While integrating servo demodulation is becoming more common, it has a complexity disadvantage when compared to peak detection. Weighted averaging demodulation, which will be presented here, is only slightly more complicated than peak detection, and can produce performance superior to integrating demodulation.

As Viterbi detection has been well described elsewhere [1,6-7], it will not be discussed here. This paper focuses on the read channel and the servo demodulation; the ENDEC, write precompensation, clock synthesizer, pulse position detector, serial interface, power control, and testability circuitry will also not be discussed here.

Manuscript received July 29, 1994.

J. Sonntag, P. Aziz, H. Burger, V. Comino, M. Heimann, T. Karanink, G. Madine, K. Nagaraj, G. Offord, R. Peruzzi, J. Plany, N. Sayiner, P. Setty, and K. Threadgill are with AT&T Bell Laboratories, Allentown, PA 18103.

O. Agazzi, J. Khoury, and K. Nagaraj are with AT&T Bell Laboratories, Murray Hill, NJ 07974.

N. Rao was with AT&T Bell Laboratories, Allentown, PA 18103. He is now with Silicon Systems Inc., Tustin, CA.

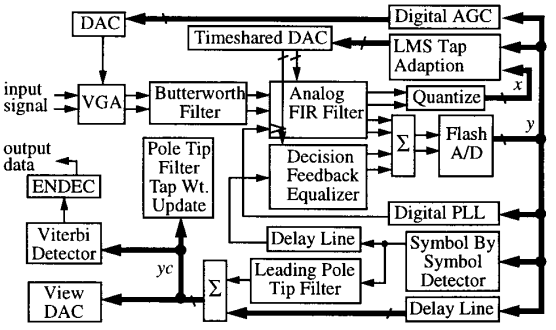

Fig. 1. A simplified block diagram of the read channel. Heavy lines denote digital busses.

In section II, the read channel of the device is described and the decision directed adaptation algorithms of the various feedback loops are described. Section III provides analysis and simulated performance of these systems; in section IV, experimental results are presented. Simulated and measured performance of the servo demodulation are examined in section V. Naturally, a summary is provided as section VI.

II. Read Channel Description

Referring to fig. 1, the signal path through the read channel starts with the VGA where the digitally closed AGC loop directs the addition of sufficient gain to force the signal level coming out of the A/D converter to a particular level. AC coupling networks at the input and output of the VGA remove both applied offset and offset introduced by the VGA.

The signal then passes into the Butterworth filter where high frequency noise components are attenuated, and boost is introduced for pulse slimming. The purpose of the boost is to shorten the non-zero pulse duration, making a 10 tap FIR filter sufficient for the rest of the equalization. If there were no FIR filter, or if the FIR filter's tap weights were constrained to be symmetrical, the phase delay dispersion introduced by the Butterworth response would be a problem; as this is not the case, this phase delay dispersion is irrelevant.

At the input of the FIR filter, the continuous time filter's output is sampled with the recovered clock. This sampled signal is made available to a two level quantizer; the result is the signal labeled x in fig. 1. This signal is necessary for the least mean squares (LMS) adaption technique used for the FIR tap weights. Ten FIR taps and a DC tap (used to remove all offsets referred to the A/D converter output) are summed. The tap weights are refreshed and updated periodically by the time-shared DAC.

At the output of the FIR filter, the equalized signal is summed with the output of a five tap decision feedback equalizer (DFE) which uses a delayed version of a symbol-by-symbol detector's output as its input. The same timeshared DAC mentioned earlier provides the tap weights for the DFE. The resulting summed signal, having the trailing undershoot characteristic of thin film heads cancelled by the DFE output is presented to the flash A/D converter.

The output of the flash A/D converter, referred to in fig. 1 as y, is used by the AGC loop, the tap weight adaptation circuitry, and the DPLL in updating these decision directed feedback loops. y is also used by the pole tip filters, both the analog DFE (for the trailing undershoot) and the digital DFE (for the leading undershoot)

The leading pole tip filter is a DFE structure which produces an inverted version of the leading pole tip undershoot. When this is added to an appropriately delayed version of y, the result (labelled yc) is a signal from which the leading pole tip response has been cancelled.

yc is used by the Viterbi detector for recovering the data, by the view DAC, and by the leading pole tip filter tap weight update mechanism.

A. The AGC Loop

The digital AGC loop operates by integrating a gain error signal. The resulting integral is used to control the gain, and reaches steady state only when the mean gain error is reduced to zero. As the goal of the AGC loop is to force y values to become $\in$ { +14 LSBs, 0, -14 LSBs}, the gain error is defined as:

$$gainerror = \hat{y}e, \tag{1}$$

where:

$$\hat{y} = (y > 7) - (y < -7) \quad \text{and} \tag{2}$$

$$e = y - 14\hat{y}\,. \tag{3}$$

With *gainerror* defined, the equation for updating the gain control value, *vga*, is written as:

$$vga_{n+1} = vga_n - gug \times gainerror_n \quad . \tag{4}$$

Where *gug* (for gain update gain) is a programmable AGC loop bandwidth control.

B. The Analog FIR Filter

The analog FIR filter, shown in fig.2, samples input signals into sequentially addressed, offset cancelled sample-and-hold

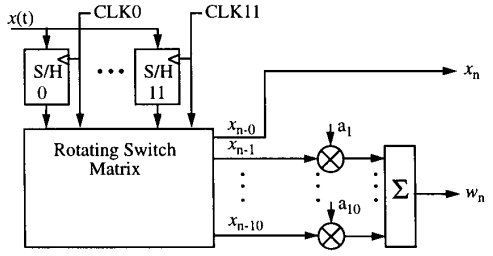

Fig. 2. A block diagram of the analog FIR filter.

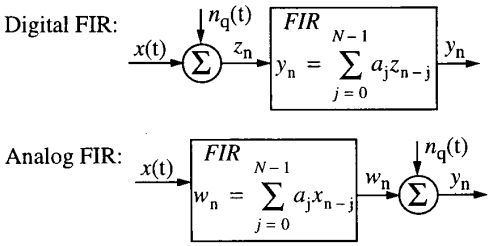

Fig. 3. Quantization noise sources before the FIR (digital FIR) vs after the FIR (analog FIR).

(S/H) cells. The outputs of these S/H cells are passed through a rotating switch matrix. The outputs of the rotating switch matrix appear as though they were generated by passing a signal through an analog delay line, but without the cumulative errors such a technique generates. These outputs are multiplied by the tap weights, using analog multipliers; the products are summed together.

One important difference between an analog and a digital FIR filter, is that quantization noise is added before a digital FIR filter, rather than afterwards. This is illustrated in fig.3.

The gain which the white quantization noise undergoes in the digital FIR filter is:

$$G_{noise} = \sqrt{\sum_{j=0}^{N-1} a_j^2} \tag{5}$$

In order to make a fair comparison between the two cases, it will be necessary to normalize the tap weights so that the peak signal level in and out of the FIR filter at the sampling times is the same. This would imply that the two A/D converters shown in fig. 3 could use the same range, and with the same number of quantization levels, produce the same amount of quantization noise. In fig. 4, G_{noise} is plotted versus user bit density (UBD, defined in (16)). The tap weight vectors used in this plot were generated by adapting a 10 tap FIR filter to a signal generated by convolving an 8/9 density coded random data stream with a 50% Lorentzian, 50% Gaussian isolated pulse shape, and filtering the resulting signal with a fourth order Butterworth filter having a corner frequency at 37% of the symbol rate, and a single right half plane zero providing 7 db of boost. Note that for UBD=2, a noise gain of 4.35 db is

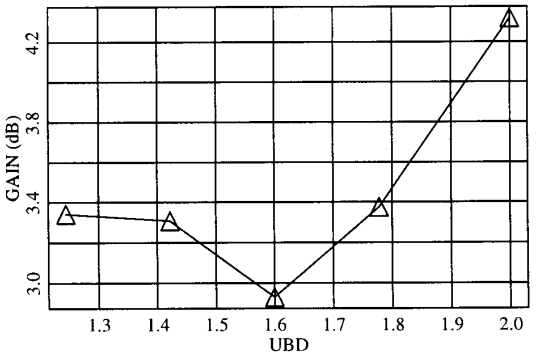

Fig. 4. Quantization noise gain produced by a 10 tap digital FIR vs user bit density.

produced, which means that an additional 0.725 bits of resolution are required in the A/D conversion process.

The update equation for the LMS algorithm used to adapt the FIR tap weights is:

$$a_{j,n+1} = a_{j,n} - twug \times x_{n-j} \times e_n \quad , \tag{6}$$

where: *twug* (for tap weight update gain) is a programmable adaptation bandwidth control, and x (shown in fig. 1) is a two level quantized version of the signal present at the input to the FIR filter.

C. The Decision Feedback Equalizer (DFE)

The analog DFE consists of five taps into a shift register, with the five single bit signals controlling the switching of a positive or negative reference voltage into five analog multipliers. The five products are summed to produce the output.

The update equation for the algorithm used to adapt the DFE tap weights is:

$$b_{j,n+1} = b_{j,n} - twug \times d_{n-j} \times e_n \quad , \tag{7}$$

where d is the delayed output of the symbol-by-symbol detector, and is also the input to the DFE.

D. The Flash A/D Converter

The flash A/D converter contains 37 offset cancelled comparators, slicing at levels of {-18, -17, ... ,18}. The comparator outputs are demetastabilized and encoded into two's complement binary.

The number of quantization levels was chosen to minimize power and area while impacting system performance very little. The total noise power coming out of the A/D converter is

$$N_{total} = N_{ADC} + N_{other} \quad . \tag{8}$$

The noise power produced by the ADC (in square LSBs) is 1/12, while the total noise power, if the bit error rate is *BER*, and a Viterbi detector is to be used can be determined by solving (9) for N_{total}.

$$BER = 2Q\left(\sqrt{2}\frac{L/2}{\sqrt{N_{total}}}\right) \tag{9}$$

where L is the number of quantization levels between 0 and

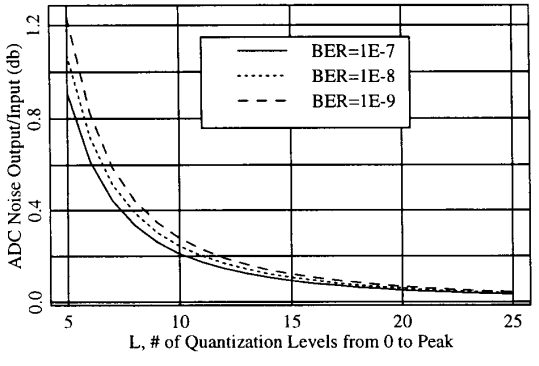

Fig. 5. Impact of quantization on the SNR present at the A/D converter output.

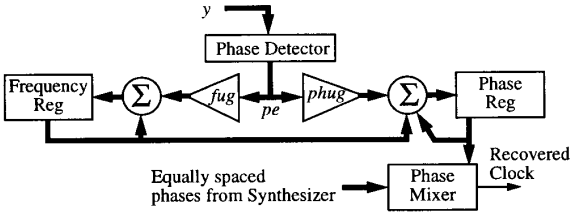

Fig. 6. A simplified block diagram of the DPLL.

equalized peak signal levels. Substituting N_{total} back into (8), and solving for N_{total}/N_{other}, we get:

$$N_{total}/N_{other} = 6L^2/(6L^2 - [Q^{-1}(BER/2)]^2) \quad , \tag{10}$$

which is the actual increase in noise power caused by the ADC quantization noise. This is plotted in figure 5 for three bit error rates.

As L=14 for this device, fig. 5 shows that 0.135 db of system performance degradation is caused by the ADC quantization noise when the total noise is sufficient to cause a BER of 1E-9.

E. The Digital Phase Locked Loop (DPLL)

The DPLL used for clock recovery (fig. 6) consists of a digital phase detector, a digital loop filter, an additional integrator acting as a digital "VCO", and a phase mixer. The phase detector's operation can be described by:

$$pe_n = e_{n-1} \times \mathrm{sgn}(\hat{y}_n - \hat{y}_{n-2}) \quad . \tag{11}$$

Where *pe* is the phase error, which is integrated in the frequency register with a gain of *fug* (for frequency update gain) applied. This integration models the capacitor in an analog loop filter. The other path, in which pe is multiplied by phug (for phase update gain), models the resistor in an analog loop filter.

The phase mixer produces one of 72 possible equally spaced clock phases under the control of the phase register. The quantization noise in the time domain produced by the discrete nature of this process is translated into an amplitude error when the recovered clock is used for sampling at the FIR input. The noise power produced can be quantified as:

$$N_{sampling} = N_{jitter} \times E[slope^2] \quad . \tag{12}$$

N_{jitter} (in clock-cycles squared) arising from the discrete-time nature of the phase mixer is:

$$N_{jitter} = 1/(12 \times 72^2) \tag{13}$$

If the signal slope (in units of LSBs/clock-cycle) through a sampling point is approximated as:

$$slope_{n-1} = L(\hat{y}_n - \hat{y}_{n-2}) \quad , \tag{14}$$

then, assuming random data,

$$E[slope^2] = L^2((0)(\tfrac{2}{8}) + 2^2(\tfrac{2}{8}) + 1^2(\tfrac{4}{8})) \quad . \tag{15}$$

Substituting (13) and (15) back into (12), the rms noise

caused by the quantized time effects is 0.069 LSBs, which is less than one fourth of the noise caused by ADC quantization.

F. The Leading Edge Undershoot Filter (LEUF)

The LEUF is a four tap transversal digital filter which uses the single bit output of the symbol-by-symbol detector as its input. Adaptivity is accomplished using a time-shared processor which estimates the correlation between the error in the *yc* signal and the signal at the tap on which it is working. If the correlation is small, the tap remains unchanged; if sufficiently large, the tap is incremented or decremented.

G. Layout

Layout quality is critical in making it possible to operate at high speeds and low power. The FIR filter and the flash A/D utilize elements which are carefully hand laid out and then repeated periodically. Critical digital speed path circuits, such as the AGC control, the tap weight processors, the DPLL, LEUF, and the ENDEC, were autorouted individually, and made use of customized adders. Only non-speed-critical elements were autorouted in the final assembly job. Fig. 7 is a microphotograph of the device.

III. Read Channel Simulations

The simulator which was used for these simulations links and schedules the use of C-language modules. All of the blocks in fig. 1 except for the ENDEC, the Viterbi detector, and the view DAC were included in the simulations. The width of digital busses, rounding, the number of bits in each DAC and in the ADC, the timeshared nature of the adaptation processors, latency in each element, and the discrete-time nature of the recovered clock are all modeled faithfully.

The input signals used for the simulations were generated by adding white gaussian noise to the convolution of a pulse shape with a data stream generated by 8/9 (0,4,4) coding and

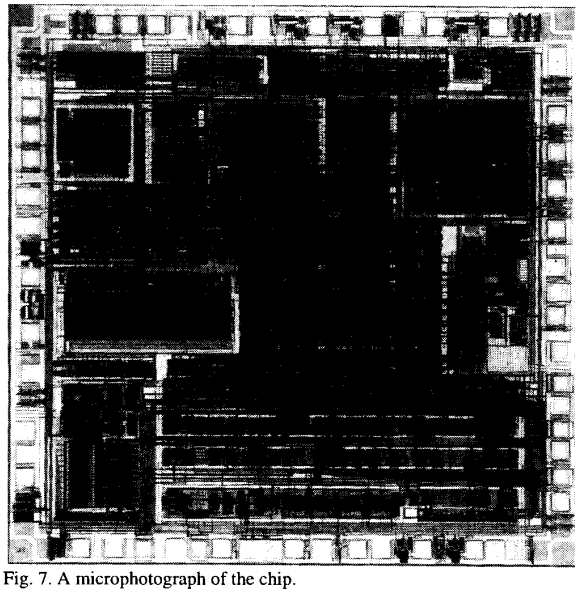

Fig. 7. A microphotograph of the chip.

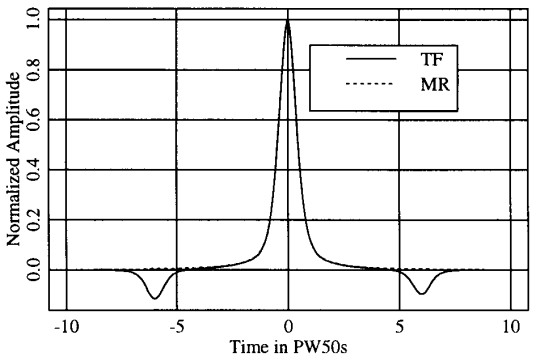

Fig. 8. Idealized pulse shapes used for thin film and jmagnetoresistive head simulations.

precoding of a scrambled data stream. A 50%-50% mix of Lorentzian and Gaussian pulse shape with no amplitude asymmetry was used to model an MR head response. The thin film head response is modeled as the same shape with leading and trailing undershoots having the same shape and with amplitudes of 12 and 10% respectively. These pulse shapes are shown in fig. 8.

Unless otherwise denoted, the simulations were performed with a PW_{50}/T of 2.0 and with the Butterworth filter corner frequency set to 37% of the symbol rate with a single right half plane zero providing 7 db of boost. The term "input SNR" is used to denote the SNR of the input signal with both the signal and the noise bandlimited to one half of the symbol rate, while the term "detection SNR" is used to denote the ratio of signal to noise at the input to the Viterbi detector times the detection gain of the Viterbi detector, calculated as a function of the autocorrelation of the error signal present there [6]. This improvement due to the Viterbi detector is typically between 2.2 and 2.8 db. The signal component of the detection SNR is calculated based on an assumed ones density of 50%, rather than the slightly higher density caused by the block code.

A. Read Channel

The simulation results presented in fig. 9 display the performance of the entire read channel for both MR and thin film heads at a variety of user bit densities. User bit density for an 8/9 density code is:

$$UBD = \frac{PW_{50}}{T} \times \frac{8}{9} \quad . \tag{16}$$

B. AGC Loop Dynamics

The dynamics of the AGC loop were studied by inserting +2db and -2db steps into the input signal. Fig. 10 shows the results for each of the three programmable AGC loop bandwidths.

C. FIR Adaptation

Speed for the FIR adaptation was studied by forcing the adaptive tap weights to zero and watching the tap weight and

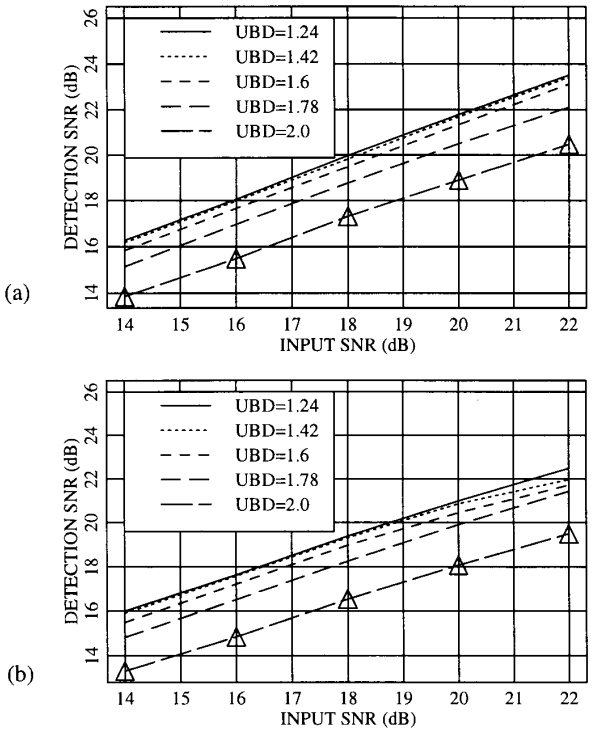

Fig. 9. Detection SNR vs Input SNR at a variety of user bit densities. (a) MR heads (b) thin film heads.

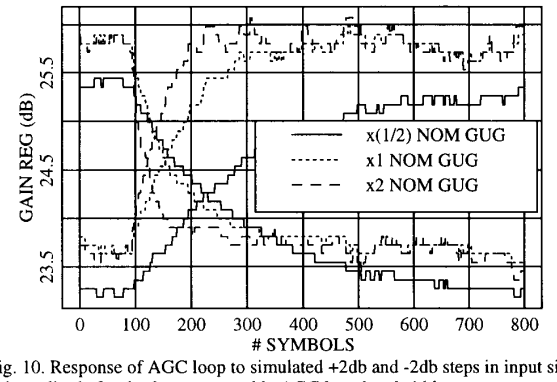

Fig. 10. Response of AGC loop to simulated +2db and -2db steps in input signal amplitude for the 3 programmable AGC loop bandwidths.

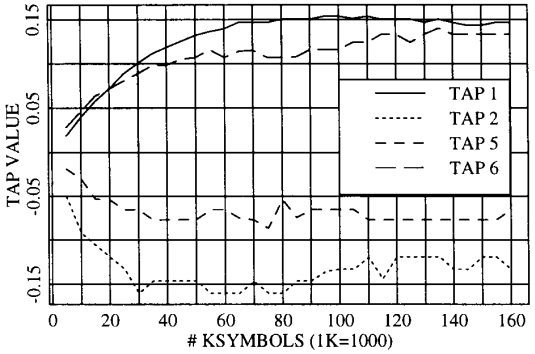

Fig. :11. Adaptation of taps number: 1, 2, 5 and 6 vs number of symbols.

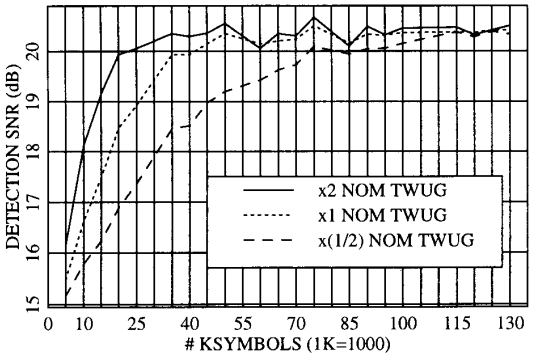

Fig. :12. Detection SNR vs number of symbols.

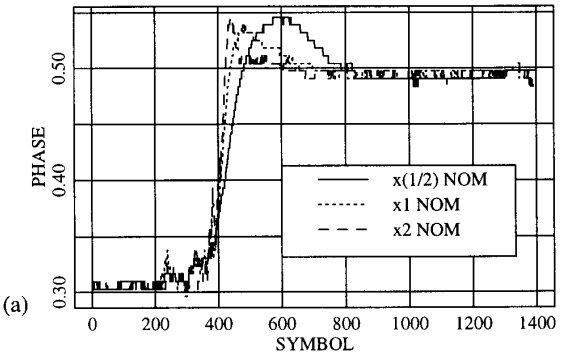

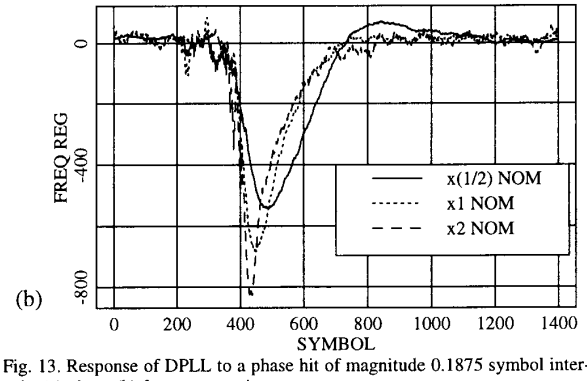

Fig. 13. Response of DPLL to a phase hit of magnitude 0.1875 symbol intervals. (a) phase (b) frequency register.

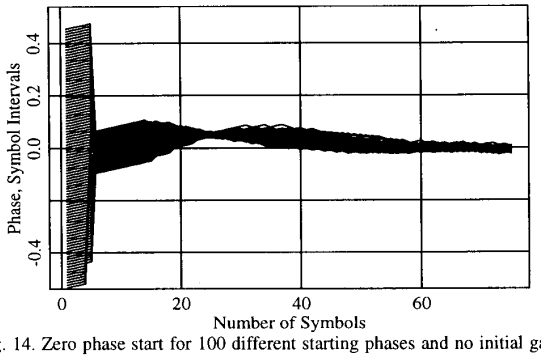

Fig. 14. Zero phase start for 100 different starting phases and no initial gain error and 0.5% initial frequency error.

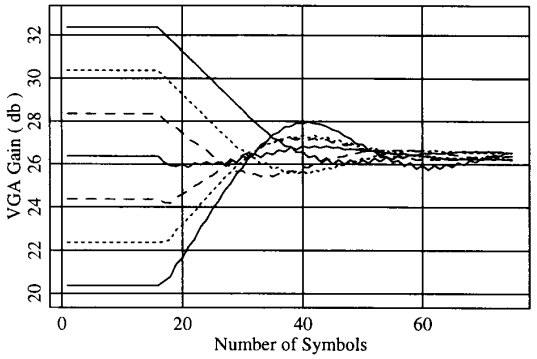

Fig. 15. Gain acquisition during zero phase start for initial gain errors of: -6, -4, -2, 0, 2, 4, and 6 db.

SNR transients which resulted. These are shown in fig. 11-12.

D. Digital PLL Dynamics

DPLL dynamics were excited by introducing a phase step of 0.1875 symbol intervals into the input signal. The resulting phase and frequency transients were plotted in fig. 13.

E. Zero Phase Start

In order to quickly acquire phase, frequency, and gain at the beginning of a new read event, both the AGC and DPLL use higher than normal adaptation bandwidths for a short time at the beginning of each read event. In addition, an estimate of the initial phase error is made and used to correct the phase error at the very beginning of each read event. In fig. 14, 100 simulation results using different initial phases were overlaid. These simulations included an initial frequency error of 0.5%.

In a another set of simulations, initial gain errors from -6 db to +6db caused gain transients shown in fig. 15.

IV. Read Channel Experimental Results

Those results shown in this section which are not oscilloscope traces were obtained by reading register contents back from the device with the help of a PC and software customized for driving the device's serial interface in a manner synchronous to the read events. Slow transient events (e.g. FIR adaptation) were studied by reading the relevant register values back after each successive read event. Faster transient events were studied by repeating the event periodically, reading the relevant registers after each read event and computing an average just to obtain a single data point; this process was repeated after increasing the width of the read event in order to obtain a second data point, and so on.

In all cases, the input signals used were generated in the same way as those used by the simulator, with the addition of an arbitrary waveform generator.

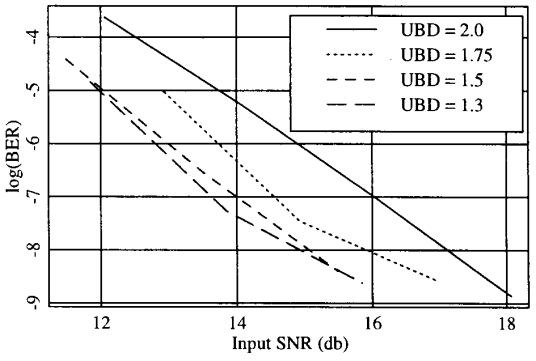

Fig. 16. Measured BER vs Input SNR at a variety of user bit densities and MR heads.

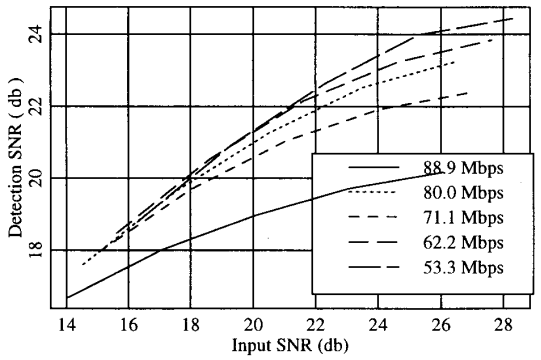

Fig. 17. Measured detection SNR vs input SNR at a variety of data rates for MR heads at a user bit density of 2.0.

A. Read Channel

In this section, the input signal was corrupted by the addition of white gaussian noise generated by a Noise Com model NC6107 noise generator. Fig. 16 shows the measured BER for MR heads at a variety of user bit densities.

In order to gauge how well performance is maintained at higher speeds, the detection SNR (using MR heads and a user bit density of 1.77) was measured as a function of input SNR for a variety of data rates. The results obtained were plotted as fig. 17.

B. AGC Loop Dynamics

Measurements of the AGC loop dynamics were made by using an input signal with an amplitude step of +2db and

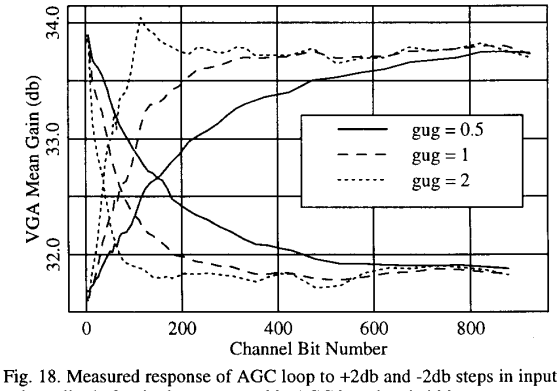

Fig. 18. Measured response of AGC loop to +2db and -2db steps in input signal amplitude for the 3 programmable AGC loop bandwidths.

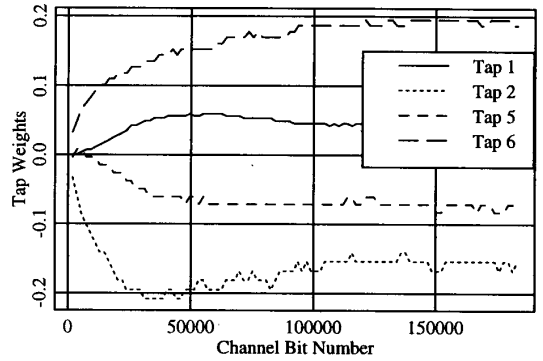

Fig. 19. Adaptation of taps number: 1, 2, 5 and 6 vs number of symbols.

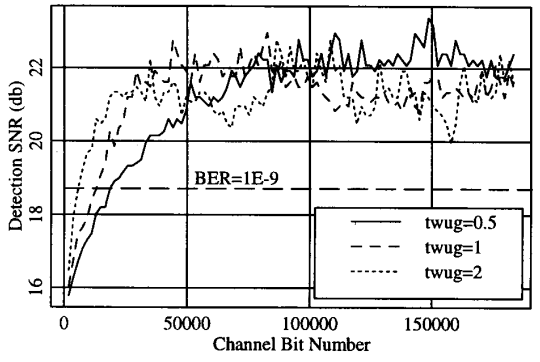

Fig. 20. SNR vs number of symbols after all adaptive tap weights were zeroed at start of simulation.

another amplitude step of -2db. The results are shown in fig.18.

C. FIR Adaptation

Measured results showing four of the FIR tap weights settling out after all of the tap weights were programmed to zero are shown in fig. 19.

The detection SNR measured during the same experiment is plotted in fig. 20. As the actual Viterbi detector contribution depends on the autocorrelation function of the error sequence, and this is not actually monitored on chip, an estimate of 2.4 db was added to the raw SNR in order to calculate the detection SNR.

D. Digital PLL Dynamics

The response of the DPLL to a phase step was measured by introducing a phase step of 0.1875 symbol intervals into the input signal. The frequency register contents as a function of time were read from the device via the serial interface, while the phase measurements were made between the recovered clock and the arbitrary waveform generator's master clock. These plots are shown in fig. 21.

E. Zero Phase Start

Fig. 22 is a persistence mode sampling oscilloscope capture

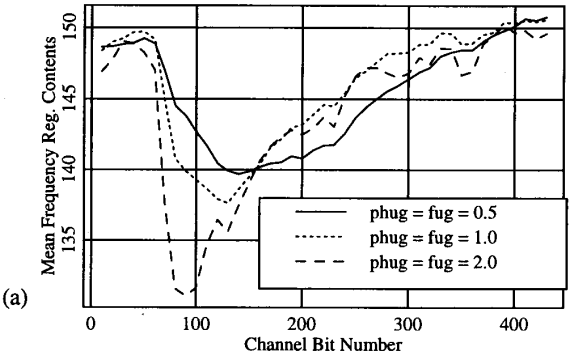

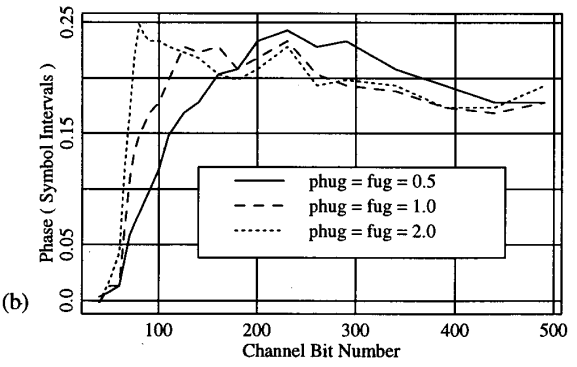

Fig. 21. Response of DPLL to a phase hit of magnitude 0.1875 symbols. Frequency register (a) and phase (b) are plotted versus number of symbols for three of the nine possible loop filter settings.

of the FIR filter output at the beginning of a read event. The oscilloscope was triggered synchronously to the input signal, while the reference clock was asynchronous to the input signal.

Fig. 23 was generated by forcing the VGA register to levels representing gain errors from -6 db to +6 db after each read event. The resulting gain transients were captured via the serial interface.

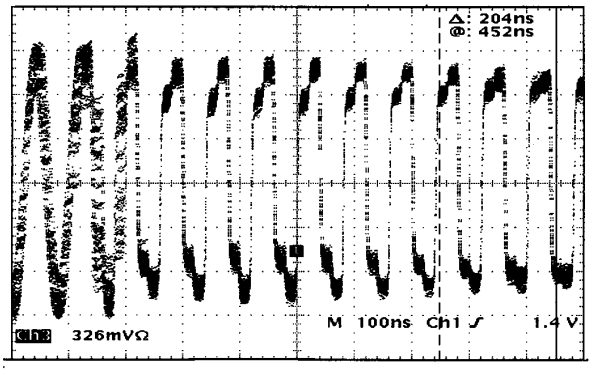

Fig. 22. persistence mode sampling oscilloscope capture of zero phase start with a 0.5% initial frequency error.

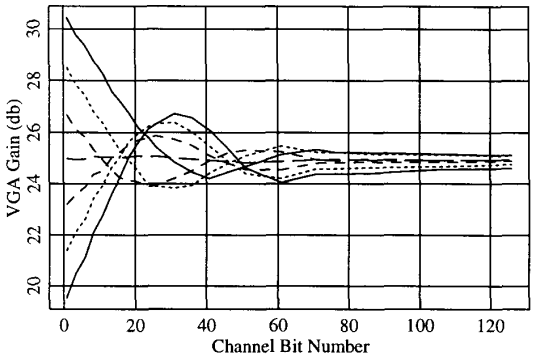

Fig. 23. Gain acquisition during zero phase start for initial gain errors of: -6, -4, -2, 0, 2, 4, and 6 db.

F. Flash A/D Converter

Flash A/D converter outputs were captured via the serial interface for a variety of input signal frequencies. An arbitrary waveform generator was used to generate full scale input signals at frequencies which were exact multiples of the Fourier frequency for the entire capture of 1024 samples so that Fourier analyses could be done without spectral smearing.

In fig. 24, the S/N, the S/D, and the S/(N+D) was plotted as a function of input signal frequency. In calculating THD, harmonics up to the fifth were considered. Note that in the many cases where harmonic frequencies were above the 511th FFT bin, the bin number to which that harmonic energy had been folded was calculated so that the THD could be calculated correctly.

V. Weighted Averaging Servo Demodulation

Weighted averaging servo demodulation, illustrated in fig.25 is a simple technique which uses a high speed peak detector (which would be needed in any event to close the servo AGC loop), resistor, a number of storage capacitors, and a timer. In the implementation described here, two peak detectors are used, responsive to positive and negative peaks, respectively. At the beginning of each servo strobe, during a short interval determined by the timer, a selected one of the storage capacitors is precharged to the output of the summer, which is the sum of a positive and a negative peak. Subsequently, during the rest of the servo strobe, the storage capacitor is connected to the summer output through the resistor. At the end of the servo strobe, when the capacitor is once again isolated from the summer, the voltage stored on the capacitor represents a weighted sum of the peak signal levels which occurred during the servo strobe time, as stated in (17), where *burst* is the demodulator output, $peak_j$ is the magnitude of the jth peak, N is the number of peaks within the strobe, and w is the weighting function.

$$burst = \sum_{j=0}^{N-1} w_j peak_j \tag{17}$$

The duration of the precharge interval must be sufficiently long to ensure the both a positive and a negative peak have been detected; beyond that, further precharging simply subtracts from the amount of time available for averaging during the (fixed) servo strobe time. It is therefore necessary to control the duration of the precharge timer closely.

The time constant during averaging must also be well controlled. It should be clear that for very long time constants, the result would depend almost entirely upon the peak levels present at the end of the precharge interval. Conversely, for very short time constants, the result depends almost entirely on the peak levels present at the end of the servo strobe. In both cases, little averaging is accomplished, and the demodulation noise is not improved over what is present at the summer output. The actual improvement over what is present at the peak detector outputs can be calculated if the assumption is made that the peaks are corrupted with uncorrelated noise samples with standard deviation σ_{peak}, resulting in (18). Note

$$\text{Noise Gain} = \frac{\sigma_{burst}}{2\sigma_{peak}} = \sqrt{\sum w_j^2} \tag{18}$$

that a factor of two is normalized out of the calculation because the process of summing the positive and negative peaks naturally yields a gain of two. The largest reduction in demodulation noise would occur if each element in w had the same value, $2/N$; in this case, the demodulation noise gain would be $1/\sqrt{N}$. While this cannot be achieved with the circuitry shown in fig. 25, the noise reduction for the optimum time constant is close to this bound.

Some simplifying assumptions allow the calculation of w: 1.) the burst (of length $N/2$ periods of the input sine wave)

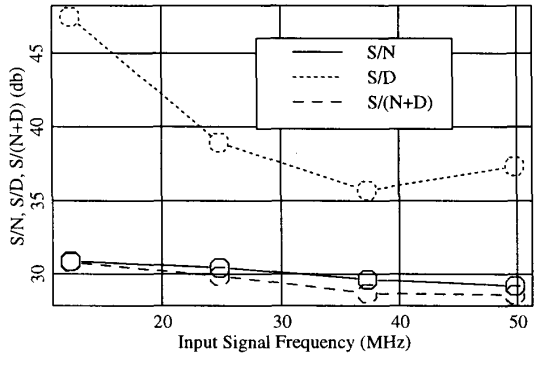

Fig. 24. S/N, S/D, and S/(N+D) vs input signal frequency, with the A/D running at 100MHz.

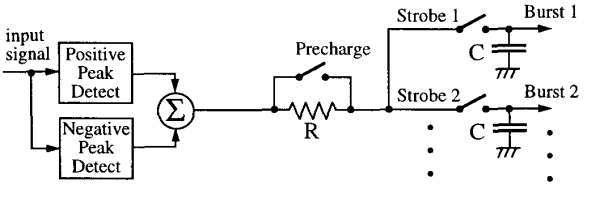

Fig. 25. Weighted averaging amplitude demodulation.

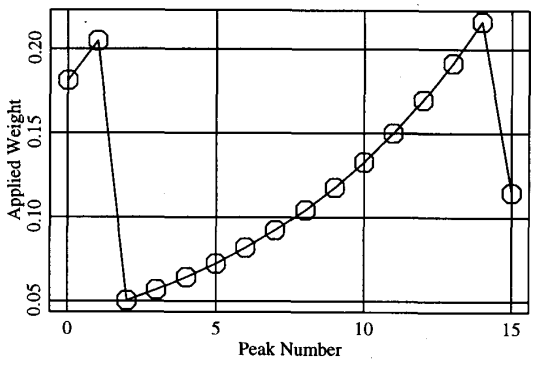

Fig. 26. Weighted averaging weights, *w*, for N=16, and RC/T=8.1.

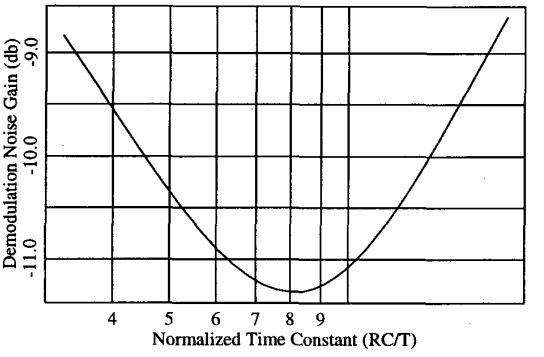

Fig. 27. Noise gain vs τ for N=16 (strobe asserted for 8 periods).

begins at $peak_0$ and ends at $peak_N$; 2.) each peak detector produces a sampled and held output which changes only at peak times; 3.) the precharge interval lasts exactly from $peak_0$ to just before $peak_2$. As the RC filter is a linear system, it is possible to consider the response to a single peak detector output pulse, starting at time j and ending at time j+2, having a level of $peak_j$ between these times and zero elsewhere, resulting in (20), where T is the time between peaks, and τ is RC. The first

$$w_j = (1 - e^{-2T/\tau})\, e^{-(N-(j+2))\,T/\tau} \quad , \qquad (20)$$

term in (20) corresponds to the amount the RC filter charged up towards $peak_j$ during the time between j and j+2, while the second term corresponds to the amount the RC filter decayed towards zero during the interval between j+2 and N.

Peaks 0 and 1 must be handled slightly differently because of the precharging; peak N-1 must be handled differently because the strobe ends before peak N-1 ends. The weights associated with these peaks are given in (21-23). Fig. 26

$$w_0 = e^{-(N-2)\,T/\tau} \qquad (21)$$

$$w_1 = e^{-(N-3)\,T/\tau} \qquad (22)$$

$$w_{N-1} = 1 - e^{-T/\tau} \qquad (23)$$

shows w for N=16 and τ=8.1T. If the noise gain for the weighted averaging process, as defined in (18) is plotted as a function of τ, the result is a broad optimum, as shown in fig. 27.

A. Simulated Results

While the foregoing analysis shows that weighted averaging gives large improvements over simple peak detection, it was necessary to make an unhappily large number of assumptions in order to make the problem analytically tractable; also, the question of how weighted averaging stacks up against integration was not resolved.

A timestep simulator was used to investigate this issue. An input signal with addative white gaussian noise was passed through a low pass filter and through models of both weighted sum and integration based amplitude demodulators. The low pass filter was a 4th order non-canonical linear phase filter (one of the settings available on the device) set to a -3db frequency 120% of the frequency of the input sinusoid. The model for the peak detectors used in the weighted sum demodulation was substantially more accurate than the "sample and hold" model used in the previous analysis. The model for the integration based demodulator was an ideal integrator integrating throughout the strobe time. One hundred servo bursts were processed by both models with independent noise for each burst. The results were converted to a position error signal (PES), using the definition in (24). The rms component of

$$PES_j = 0.5 \times \left(\frac{burst_{2j} - burst_{2j+1}}{burst_{2j} + burst_{2j+1}} \right) \qquad (24)$$

the PES signal for both demodulators is plotted in fig. 29 as a function of input SNR for a near-optimum value of τ (8.1T). In fig. 28, the rms component of the PES signal is plotted as a function of τ for an input SNR of 14 db.

B. Experimental Results

Experimental results for the weighted averaging demodulation circuitry were obtained by adding white gaussian noise to a servo signal consisting of a 100% AGC field followed by two 50% servo bursts. Each burst consisted of 10 cycles of a 10MHz sinusoid; the filter was programmed for linear phase

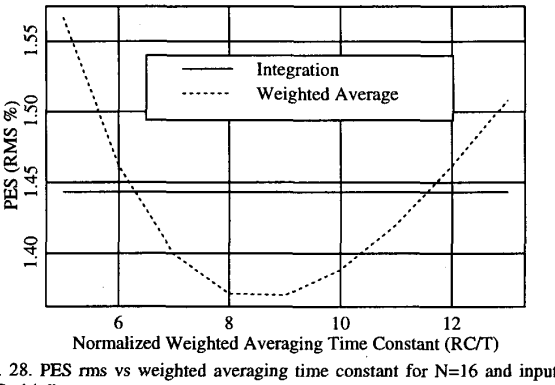

Fig. 28. PES rms vs weighted averaging time constant for N=16 and input SNR=14 db.

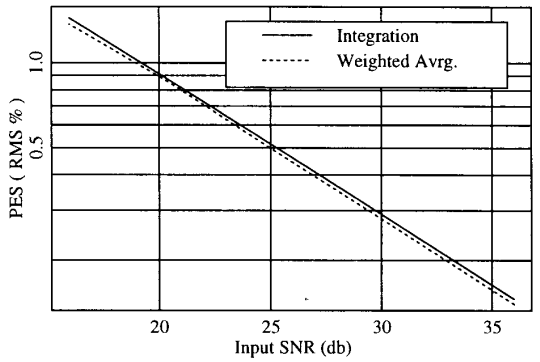

Fig. 29. Simulated performance comparison of weighted averaging and integrating servo amplitude demodulation.

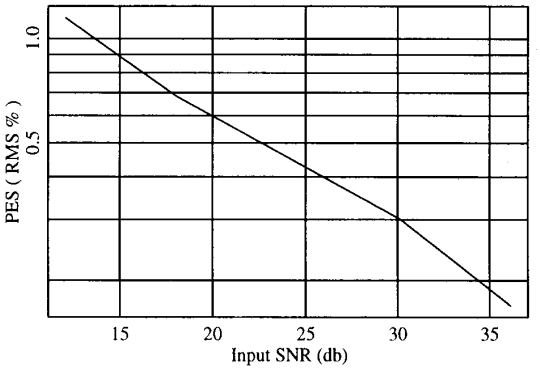

Fig. 30. Measured PES rms noise versus input SNR.

operation and a -3db frequency of 12MHz. The precharging timer was programmed to 170ns, and the weighted averaging time constant was programmed to 450ns. The results are plotted in fig. 30.

VI. SUMMARY

An area and power efficient 82 Mbps read channel device has been described. The decision directed algorithms used for adapting the characteristics of the device to the head and media characteristics were presented. Extensive simulation results were compared to experimentally obtained data.

An unusual method of servo demodulation, combining the best features of integration and peak detection was described. Both experimentally obtained results and simulator based results were presented.

ACKNOWLEDGMENT

The authors thank the design personnel who contributed to the layout of the device, the engineers developing the automated test software, and the writer of the PC interface software.

REFERENCES

[1] R. Wood and D. Petersen, "Viterbi detection of class IV partial response on a magnetic recording channel," *IEEE Transactions on Communications*, vol. COM-34, No. 5, pp. 454-461, May 1986.

[2] D. Welland et al., "A digital read/write channel with EEPR4 detection," *1994 IEEE International Solid-State Circuits Conference Digest of Technical Papers*, vol. 37, pp. 276-277, Feb. 1994.

[3] R. Yamasaki, T. Pan, M. Palmer and D. Browning, "A 72Mb/s PRML disk-drive channel chip with an analog sampled-data signal processor," *1994 IEEE International Solid-State Circuits Conference Digest of Technical Papers*, vol. 37, pp.278-279, Feb. 1994.

[4] R. Cideciyan, F. Dolivo, R. Hermann, W. Hirt and W. Schott, "A PRML system for digital magnetic recording," IEEE Journal on Selected Areas in Communications, vol. 10, No. 1, pp. 38-56, Jan. 1992.

[5] J. Coker, R. Galbraith, G. Kerwin, J. Rae, and P. Ziperovich, "Integrating a partial-response, maximum likelihood channel into the 0681 disk drive," in *Proc. 24th Asilomar Conf. Signals, Syst., Comput.*, Pacific Grove, CA, Nov. 1990, vol. 2, pp. 674-677.

[6] G. Forney, "The Viterbi algorithm," *Proc. IEEE*, vol. 61, no. 3, pp. 268-278, Mar. 1973.

[7] F. Dolivo, and G. Ungerboeck, "Viterbi decoders for partial-response class-IV signaling: theory and implementation," IBM Res. Rep. RZ 1177, Sept. 1982.