

Analog Timing Recovery for a Noise-Predictive DFE

John P. Keane, Michael Q. Le, and Paul J. Hurst

John P. Keane and Paul J. Hurst are with Department of Electrical and Computer Engineering,
University of California, Davis, CA 95616. (e-mail jpkeane@ece.ucdavis.edu).

Michael Q. Le was with Department of ECE, University of California, Davis, CA 95616. He is
now with Broadcom Corporation, Irvine, CA 92619.

Abstract

A timing recovery architecture and its CMOS implementation are described for a noise-predictive decision-feedback equalizer (NPDFE). The $0.5\mu\text{m}$ CMOS prototype includes timing recovery and the NPDFE and operates at 160Mbps. The timing recovery blocks dissipate 27mW from 3.3V, occupy 0.2 mm^2 , and achieve a rms jitter of 50 ps, which is 0.8% of a bit period.

I. INTRODUCTION

Timing recovery is an important block in any digital-communication system that uses a sampling detector. At high data rates, timing recovery is often based on baud-spaced samples of the signal [1]. This paper describes a CMOS implementation of baud-spaced timing recovery for the noise-predictive decision-feedback equalizer (NPDFE) in [2]. The prototype includes timing recovery circuits and the NPDFE. It is targeted at the magnetic recording channel [3] but could be used in many communication applications.

The NPDFE consists of the unshaded blocks in Fig. 1: an analog discrete-time forward equalizer $B(z)$, a decision-feedback equalizer (DFE) $C(z)$, and an analog discrete-time noise predictive equalizer $A(z)$ [2]. (In the figures, thin lines represent analog signals, and thick lines represent digital signals.) A , B and C are adaptive finite impulse response (FIR) filters of order 3, 3 and 5, respectively. The equalizer outputs are summed to generate the analog voltage $y[k]$ that is sliced to produce the binary decision $\hat{a}[k] = \pm 1$.

In a conventional DFE-based read channel, only the FIR equalizer $B(z)$ and the DFE $C(z)$ are used and the FIR equalizer operates on samples of the input before the cursor. In this case, $B(z)$ will provide high-frequency boost to the input signal to eliminate precursor ISI, but that boost also affects the input noise. This noise enhancement reduces the signal-to-noise ratio (SNR) at the slicer input, which increases the bit-error rate. With the addition of the noise predictive equalizer $A(z)$, the transfer function from the input to the slicer for the noise is all-pass. Therefore, the noise boost is eliminated, which can improve the SNR at the slicer input by about 2dB for the magnetic recording channel [2].

Each coefficient in $A(z)$ is also a coefficient in $B(z)$. Therefore, only three adaptive loops are needed to determine the coefficients in $A(z)$ and $B(z)$. When the FIR equalizer transfer function $B(z)$ has maximum phase, as is the case when $B(z)$ is adapted to remove precursor ISI only, the feedback loop that contains the noise predictive equalizer $A(z)$ is stable [2].

II. TIMING RECOVERY ARCHITECTURE

The shaded blocks in Fig. 1 form the timing recovery subsystem. A recovered bit-rate clock controls the input sampler. An estimate $z[k]$ of the timing error in the recovered clock is fed through a first-order loop filter. The filter output adjusts the clock phase by varying the current in the current controlled oscillator (ICO).

The timing error estimate $z[k]$ is based on the product of a sample error e and the signal slope [4]. This estimate is not a direct measure of the sample time error, but it is a gradient estimate used to minimize the mean-square sampling error. Intuitively, the slope is used because the error in the sampled value due to a sample-time error will depend on the slope of the signal at that point. In order to simplify the implementation, a quantized version of the slope, \hat{slope} , is often used, so that the timing error estimate z is

$$z = e \cdot \hat{slope}. \quad (1)$$

The calculation of the slope of the signal and the error in the sampled value depends on whether the loop is acquiring initial phase and frequency lock or is performing steady-state tracking.

A. Acquisition mode

Initial acquisition of the correct timing phase and frequency is made during a training sequence. In the magnetic recording channel, this sequence is generated by repeated recording of $\{+1, +1, -1, -1\}$, which produces a sinewave with period $4T$ (where T is 1 bit period) when read. Acquisition is achieved by adjusting the ICO phase so that every other sample of $x(t)$ is at a zero crossing. The timing error estimate is calculated from the samples $x[k]$ of $x(t)$. Near a zero crossing, the sample error $e[k]$ is the sample $x[k]$ since the ideal sample value is 0. The sign of the slope of $x(t)$ near the zero crossing is given by the sign of the difference between the following and previous samples of $x(t)$, i.e. by the sign of $x[k+1] - x[k-1]$. In this case, since $x(t)$ is a sinusoid with period $4T$, $x[k+1] \approx -x[k-1]$, so the quantized slope estimate is

$$\hat{slope}[k] = -\hat{x}[k-1] \quad (2)$$

where $\hat{x}[k-1]$ is a quantized version of $x[k-1]$. Thus, from (1), the timing error estimate is [5]

$$z[k] = -x[k]\hat{x}[k-1] \quad (3)$$

However, rather than use $\hat{x}[k] = \pm 1$ as in [5], we use $\hat{x}[k] \in \{-1, 0, +1\}$ by quantizing $x[k]$ to 3 levels using comparators with thresholds set to approximately $\pm 50\%$ of the full scale voltage. This signal and

slope quantization is illustrated in Fig. 2, where $s\hat{l}o\hat{p}e[1] = -\hat{x}[0] = -1$ and $s\hat{l}o\hat{p}e[2] = -\hat{x}[1] = 0$. Quantizing to 3 rather than 2 levels avoids adjusting the clock phase when $x[k-1]$ is small (when $x[k]$ is far away from a zero crossing); this results in a maximum initial phase offset which is less than the scheme in [5], which reduces the maximum acquisition time. At a phase offset close to $0.5T$, a repeating sample sequence $x \approx \left\{ +\frac{1}{\sqrt{2}}, +\frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}} \right\}$ is possible. Oscillation (hangup) of the timing loop at this timing offset can occur as $z \approx \left\{ +\frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}}, +\frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}} \right\}$ in this case. Such oscillation is avoided by preventing the timing loop from updating in two consecutive bit periods. This is implemented by forcing $s\hat{l}o\hat{p}e[k] = 0$ when $s\hat{l}o\hat{p}e[k-1]$ is non-zero.

B. Steady-state mode

In steady state, the input $x(t)$ consists of recorded data with large amounts of ISI; the ISI makes extraction of timing information from $x(t)$ difficult. Therefore, decision-directed timing recovery is used in steady state. Signals at the slicer, where equalization by the NPDFE has removed the ISI, are used to estimate the timing error. A popular timing recovery scheme [4] uses the slicer input $y[k]$ and decisions \hat{a} to calculate the timing error estimate. The difference between decisions is used to estimate the signal slope:

$$s\hat{l}o\hat{p}e[k-1] = (\hat{a}[k] - \hat{a}[k-2])/2 \quad (4)$$

Hence, the steady-state timing error estimate, from (1), is given by:

$$z[k] = e[k-1](\hat{a}[k] - \hat{a}[k-2])/2 \quad (5)$$

where $e[k-1] = y[k-1] - \hat{a}[k-1]$ is the decision-slicer error. In general, the nonlinear effect of the DFE affects the accuracy of the slope estimate using (4). However, due to the coarse quantization of the slope estimate, the effect was found to be negligible for a channel with $PW_{50} = 2.5T$. For higher channel densities, filtering of the decision levels to compensate for the effect of the DFE may be necessary [6].

Interaction between the adaptive equalizer and the timing recovery loop may give rise to stability concerns. However, this is not a problem in practice as the equalizer coefficients are only adapted for a short period after initial timing acquisition, and are then frozen while reading the remainder of the sector [2].

III. IMPLEMENTATION

A. Timing Error Estimator and Loop Filter

A simplified block diagram of the timing-error estimator and loop filter is shown in Fig. 3. While signals are shown as single-ended here for simplicity, all analog circuits are fully differential. The FIR

forward equalizer $B(z)$ uses sample-and-hold amplifiers (SHAs) that hold present and past samples of the input $x[k]$, and the noise predictor $A(z)$ holds past values of the slicer input $y[k]$. These held samples are used to generate the timing error estimate $z[k]$, so no additional SHAs are needed.

In steady-state mode, the error current $e[k - 1]$ is generated by applying $y[k - 1]$ and $\hat{a}[k - 1]$ to a differencing transconductance (G_m) cell. In acquisition mode, the voltage $x[k]$ and zero are input to the G_m cell to generate the error current. This current is multiplied by the slope estimate according to (3) or (5) using switches to generate $z[k]$. This slope estimate is generated from quantized signals (\hat{a} or \hat{x}) and takes one of 3 values $\{+1, 0, -1\}$.

Current $z[k]$ is integrated onto a capacitor. The voltage output of the integrator v_c is converted to a current by the transconductance cell β . The resulting current is added to $\alpha z[k]$ to form the loop filter output I_{ICO} . The integration allows for frequency offsets between the local ICO and the read data while maintaining zero phase offset in steady state.

In steady state, the loop filter coefficients were $\alpha = 0.2$ and $\beta = 40\mu\text{A/V}$. The integration capacitor value C_{int} was 40pF, and $G_m = 0.1\text{mA/V}$. The ICO gain was $K_{ICO} = 140\text{kHz}/\mu\text{A}$. This gives the timing recovery an approximate loop bandwidth of 100kHz.

An alternative to using held analog samples to generate the timing error estimate would be to use 1-bit quantized values of $x[k]$ and $e[k] = y[k] - a[k]$, which are used in the adaptation logic of the NPDFE [2]. To investigate the performance when using these quantized values, a mode of operation was implemented in the prototype where the analog inputs to the G_m cell in Fig. 3 are changed. The analog inputs for $x[k]$ and $y[k - 1]$ are replaced by voltages representing 1-bit quantized values of $x[k]$ and $e[k - 1]$, respectively, and the input $\hat{a}[k - 1]$ is set to zero. With these changes, e in Fig. 3 equals $\text{sign}(x[k])$ during acquisition and equals $\text{sign}(e[k - 1])$ in steady state.

The multiplication and integration in Fig. 3 are realized together by the fully differential structure shown in Fig. 4(a). The error e is represented by the differential current

$$e = I_{ep} - I_{en} \quad (6)$$

The switches in Fig. 4(a) are controlled by the slope estimate \hat{slope} . For a positive slope estimate, $up = 1$ and $dn = 0$ so the error current is integrated onto the capacitors. For a negative slope estimate, $up = 0$ and $dn = 1$ and the negative of this current is integrated. For a slope estimate of zero, all switches remain open and no integration is performed. When all switches in Fig. 4(a) are open, all current sources are switched to a replica circuit (not shown) so that the current-source transistors remain in saturation. The signals up and dn have a duty cycle of 50% to avoid integrating glitches and allow time for the error e

to settle before being multiplied by the slope. These control signals are illustrated in Fig. 4(b).

The four cascode devices with gates connected to V_{CM} buffer the outputs from clock feedthrough due to the switches. The integration capacitors ($C_{int} = 40\text{pF}$) are implemented by PMOS devices whose drain, source and body are connected to the positive supply.

The common-mode (CM) voltage at the output of the integrator is set to V_{CM} by a CM feedback circuit (not shown) that controls the current I_{CM} [7]. However, there is no CM feedback when the slope is zero because the current sources are not connected to the outputs in this case. In practice, this is acceptable since the amount of CM voltage drift due to leakage currents is small, even during a long period of successive zero slopes.

Any dc offset that can be referred to the input of the loop filter will result in a steady-state phase error in the recovered clock. One potential source of such offset is due to offset in the Gm cell generating the error current e . Another potential source is mismatch between the two current sources generating the common-mode feedback current I_{CM} . When $\hat{slope} = 1$, these offsets will be integrated onto the capacitors. However, when $\hat{slope} = -1$, the complement of such offsets will be integrated. Since, for binary data, a positive slope estimate must be followed by a negative slope estimate (with an arbitrary number of zero slope estimates in between), the net offset integrated is zero.

The slope multiplication and scaling by α for the direct path in the loop filter is accomplished using the circuit in Fig. 5. The slope multiplication is performed by switching currents controlled by the same up and dn pulses that were used in the integrator. Matched current sources are used to generate the I_{ep} and I_{en} currents. When $\hat{slope} = 1$, $up = 1$ and $dn = 0$, and both I_{ep} current sources are steered to the left branch, while both I_{en} current sources are steered to the right branch. When $\hat{slope} = -1$, $up = 0$ and $dn = 1$, and the current flow is reversed. Therefore, the effect of offsets in this path is reduced in a similar manner to Fig. 4(a). When $\hat{slope} = 0$, $up = dn = 0$, and the I_{ep} and I_{en} currents are steered to both the left and right branches, so the differential output current is zero. Note that the sum of the currents flowing to both branches is $2I_{ep} + 2I_{en}$ in all cases, so the common-mode current is always constant. This avoids nonidealities due to common-mode to differential-mode conversion in the generation of I_{ICO} .

Multiplication by the factor α is achieved using the two coupled differential pairs in Fig. 5 controlled by an external differential voltage $V_{\alpha p} - V_{\alpha n}$. This allowed the value of α to be adjusted during testing. In practice, a fixed factor α could be implemented by a simple scaling of the I_{ep} and I_{en} current source transistors.

B. Current Controlled Oscillator

The current-controlled oscillator is shown in Fig. 6(a). It is a 3-stage ring oscillator; its frequency is controlled by tail current I_{ICO} in each stage. A replica bias circuit in Fig. 6(b) generates bias voltage V_B that makes the drain currents of both conducting PMOS load devices in each inverter equal when all the tail current is steered to one side. This biasing ensures a large output swing over a wide range of operating frequencies. When $v_{op} = v_{on}$, the PMOS devices controlled by V_B are triode. The other PMOS load devices are diode-connected. The conductances of the triode and diode-connected transistors have opposite voltage coefficients. Therefore, the two load impedances in each inverter are fairly well matched over the output swing [9], which helps reduce supply noise coupling that can cause jitter.

Capacitors formed from the three metal layers form a linear differential capacitive load and set the oscillation frequency range to include the expected operating rate of the NPDFE. Common-mode (CM) oscillation is avoided by using a cascoded tail current source for high CM rejection. The measured operating range of the ICO is 40MHz-240MHz.

IV. MEASURED DATA

The $0.5\mu\text{m}$ CMOS prototype contains the blocks in Fig. 1. A die photograph is shown in Fig. 7. Measurements were taken using a signal that models a Lorentzian read channel with $PW_{50} = 2.5T$. The signal was loaded into an arbitrary waveform generator and band-limited white noise was added to form the test input signal. The measured bit-error rate versus input SNR with and without timing recovery enabled is plotted in Fig. 8 at 160Mbps. The SNR loss with timing recovery is 0.2-0.4dB compared to without timing recovery (i.e., operating the input sampler from an external clock synchronized to the test data). The measured jitter at 160Mbps is 50ps(rms), or 0.8% of a bit period, and 180ps(pk). A jitter histogram is shown in Fig. 9. Performance is summarized in Table I. The timing recovery circuits consume only 13% of the area and power while providing jitter low enough to allow near optimum performance of the NPDFE.

As mentioned in Section III-A, an alternative mode of operation was implemented where 1-bit quantized values of $x[k]$ and $e[k]$ are used to generate the timing error estimate. The resulting acquisition times were much longer, although they could be reduced by using gear-shifting to gradually reduce α and β . In steady state, a lower loop bandwidth was needed due to the noisier timing error estimate. The measured jitter increased to approximately 90ps(rms) at 160Mb/s for an input SNR of 16dB. This jitter resulted in a SNR loss of 0.4-0.6dB when compared to the case without timing recovery. The performance using

the 1-b values of $x[k]$ and $e[k]$ could be improved by using a hybrid analog-digital implementation [8] optimized for this mode; our design was not optimized for this mode of operation.

In acquisition mode, testing confirmed that frequency and phase lock was achieved within 100 bit periods for frequency offsets of up to 5%. A simulated plot is shown in Fig. 10 showing convergence of I_{ICO} to its final value I_o from a 5% frequency and $T/2$ phase offset. This simulation was for a channel with PW_{50} of $2.5T$ and input SNR of 18dB. The timing recovery switches from acquisition mode to steady-state mode after 100 bit periods.

V. CONCLUSION

A timing recovery scheme for a NPDFE has been demonstrated. The addition of on-chip timing recovery and clock generation increased the maximum data rate to 160 Mbps, compared to 100 Mbps for the prototype in [2], where all clock generation was performed off-chip. The maximum speed is limited by the NPDFE circuits and not by the timing recovery loop, which is capable of operating about 50% faster, based on simulations and the measured ICO oscillation range.

ACKNOWLEDGEMENT

The authors would like to thank R. Stephens and K. Dyer of Intel for performing a focused ion beam (FIB) repair on the prototype. We would also like to thank C. Cazana of Texas Instruments for allowing us to borrow test equipment to make jitter measurements,

REFERENCES

- [1] K.H. Mueller and M. Muller, "Timing Recovery in Digital Synchronous Data Receivers," *IEEE Trans. Comm.*, vol. COM-24, pp. 516-531, May 1976.
- [2] M.Q. Le, P.J. Hurst, and J.P. Keane, "An adaptive analog noise-predictive decision-feedback equalizer," *JSSC*, pp. 105-113, Feb. 2002.
- [3] A.S. Hoagland and J.E. Monson, *Digital Magnetic Recording, 2nd Edition*, Wiley, 1991.
- [4] S.U.H. Qureshi, "Timing Recovery for Equalized Partial-Response Systems," *IEEE Trans. Comm.*, vol. COM-24, pp. 1326-1331, Dec. 1976.
- [5] W.L. Abbott and J. M. Cioffi, "Timing recovery for adaptive decision feedback equalization of the magnetic storage channel," *GLOBECOM*, pp. 1794-1799, 1990.
- [6] J.W.M. Bergmans, *Digital Baseband Transmission and Recording*, Kluwer Academic Publishers, pp. 555-557, 1996.
- [7] M. Le, P. Hurst, and K. Dyer, "An analog DFE for disk drives using a mixed-signal integrator," *JSSC*, pp. 592-598, May 1999.
- [8] B. Kim, D.N. Helman, and P.R. Gray, "A 30-MHz hybrid analog/digital clock recovery circuit in 2- μ m CMOS". *JSSC*, pp.1385-94, Dec. 1990.

- [9] J.G. Maneatis, "Low-jitter process-independent DLL and PLL based on self-biased techniques," *JSSC*, pp. 1723-32. Nov. 1996.

Process	0.5 μ m SPTM CMOS
Active Area (NPDFE + Timing Rec.)	1.5mm ²
Active Area (Timing Rec. only.)	0.2mm ²
Data Rate = $1/T$	160Mbps
Power (NPDFE + Timing Rec.)	213mW
Power (Timing Rec. only)	27mW
Frequency Lock Range	$\pm 5\%$
Maximum Acquisition Time	$100T$
Steady-State Jitter	50ps(rms), 180ps(pk)

TABLE I

MEASURED PERFORMANCE AT 3.3V AND 25°C.

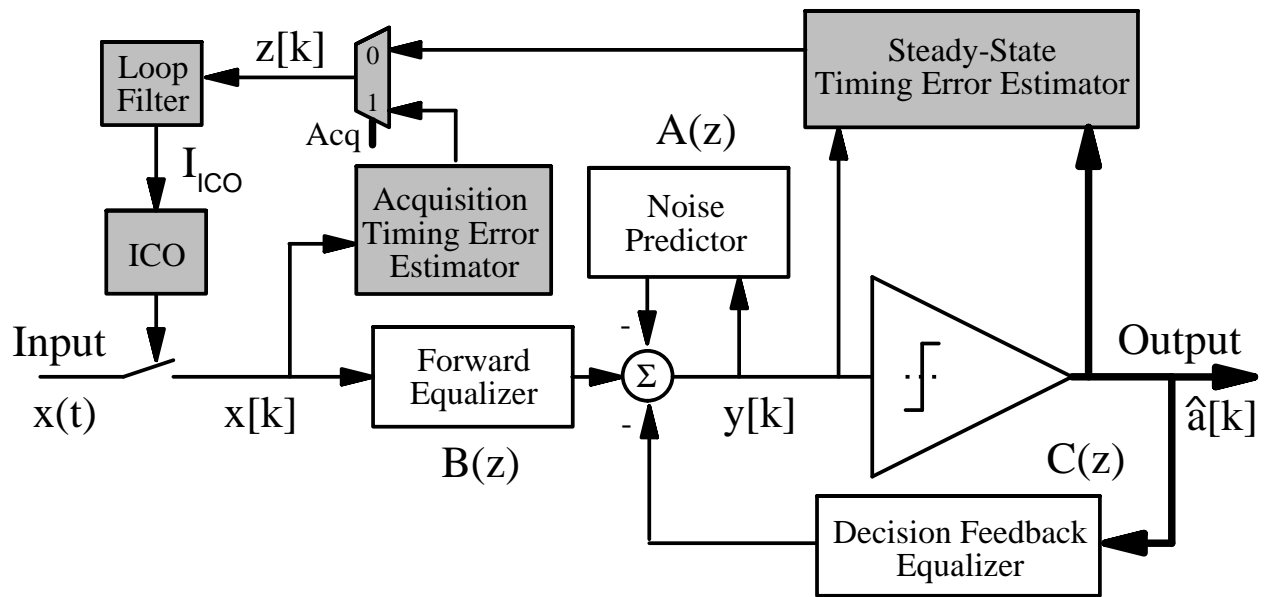


Fig. 1. Block diagram of the complete NPDFE-based read channel.

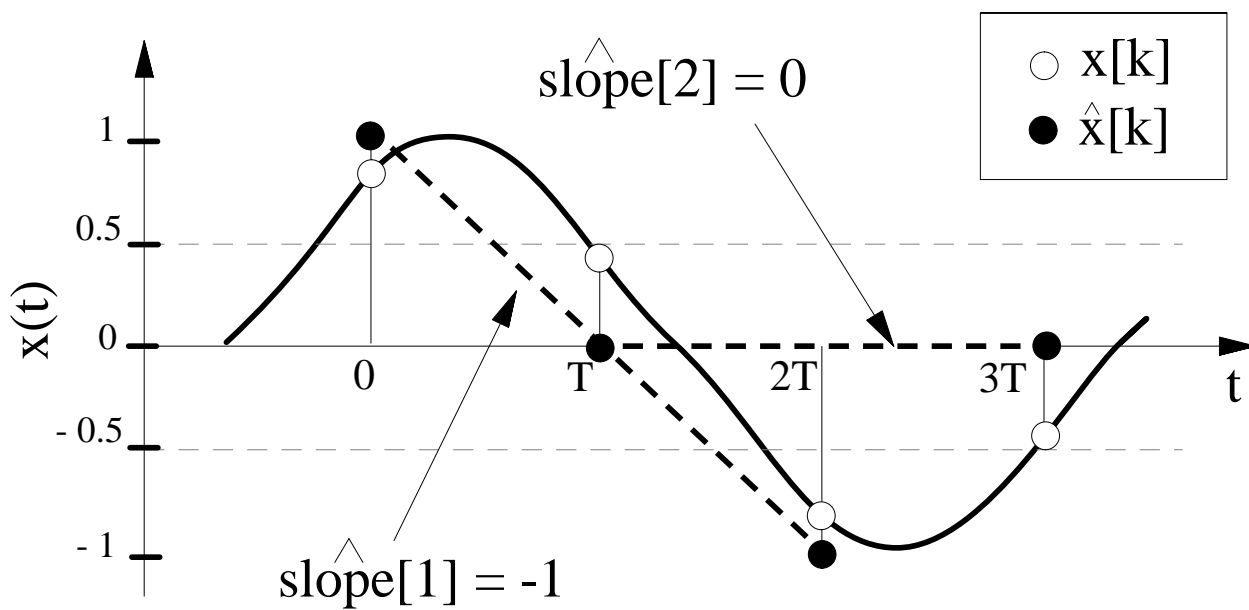


Fig. 2. Signal and slope quantization in acquisition mode.

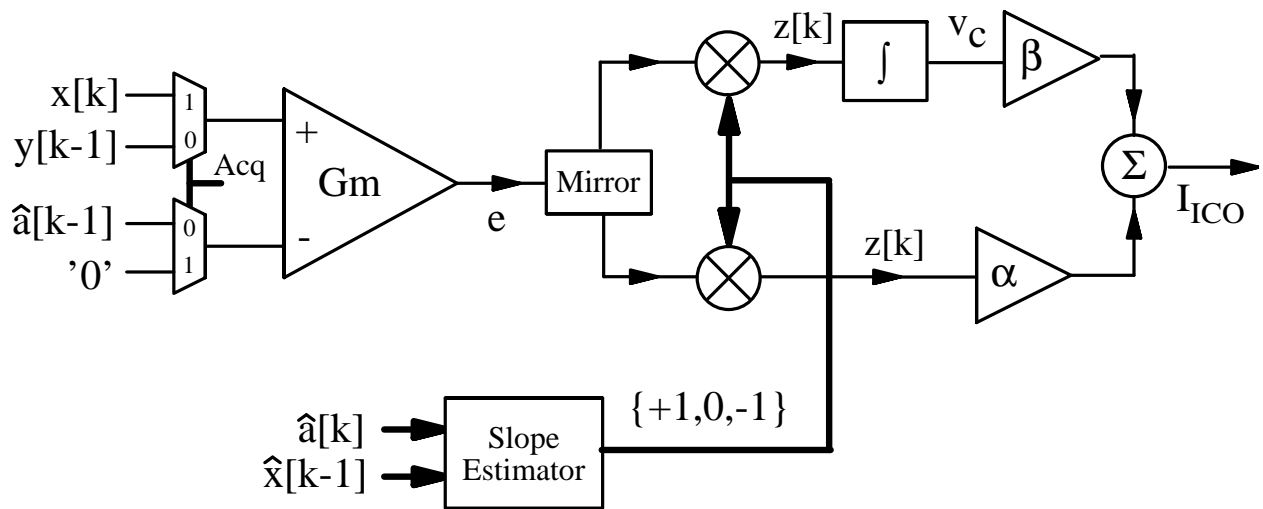


Fig. 3. Timing error estimation and loop filter. Signals with arrows in the middle are currents.





Fig. 4. (a) Multiplier and loop filter integrator. (b) Control signals

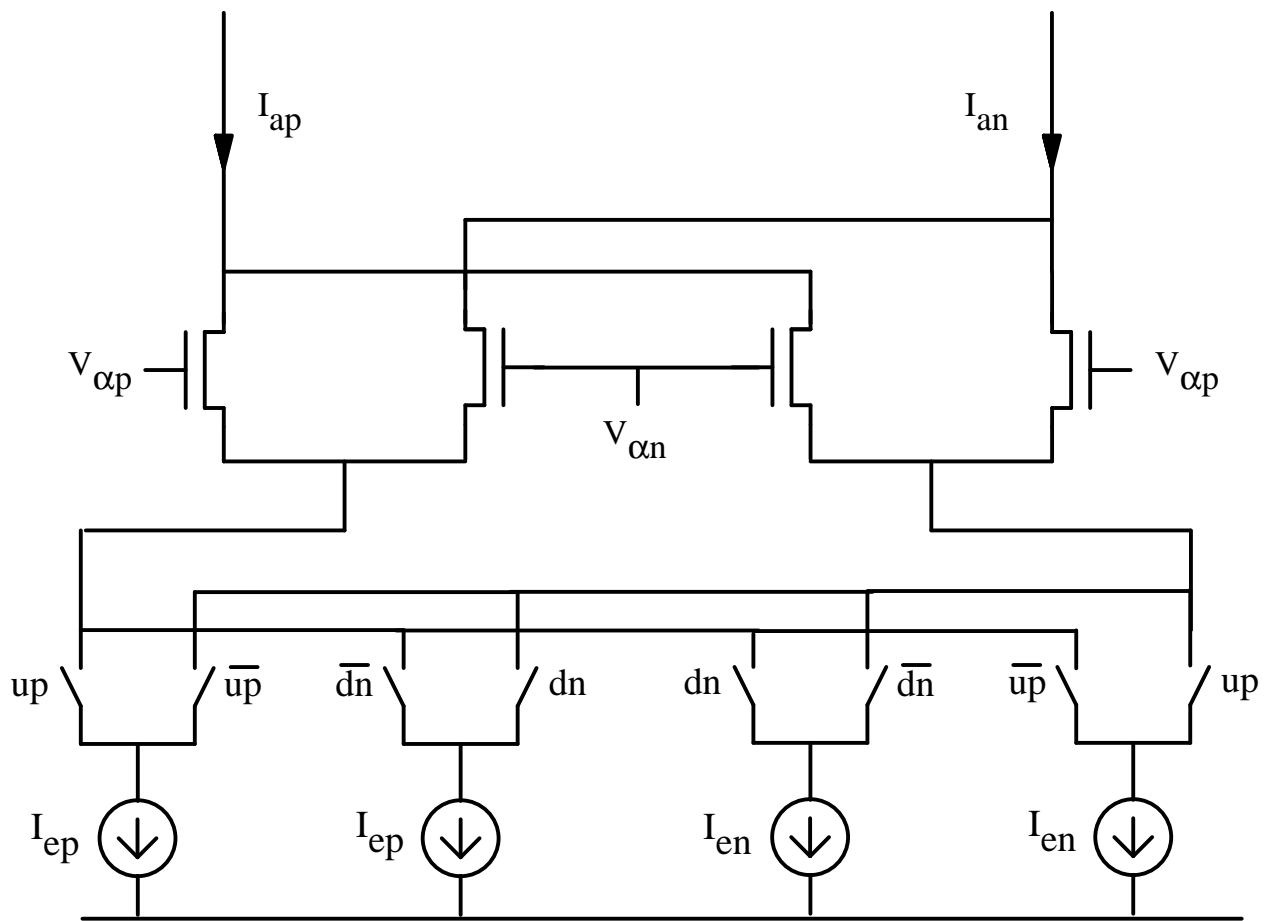




Fig. 5. Direct path slope multiplier and scaling by α .



(a)



(b)

Fig. 6. (a) Three stage ring oscillator. (b) Replica bias circuit.




Fig. 7. Die photograph.

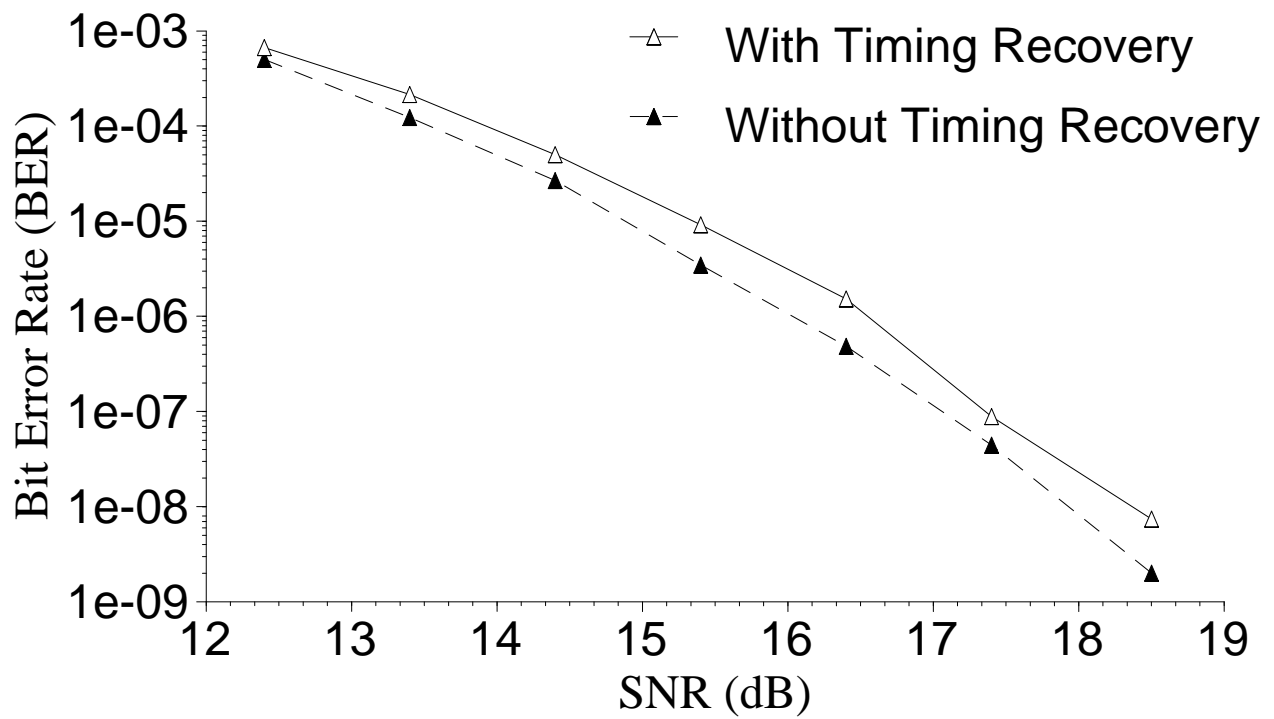


Fig. 8. Measured bit-error rate vs. input SNR with (solid) and without (dashed) timing recovery at 160Mbps.

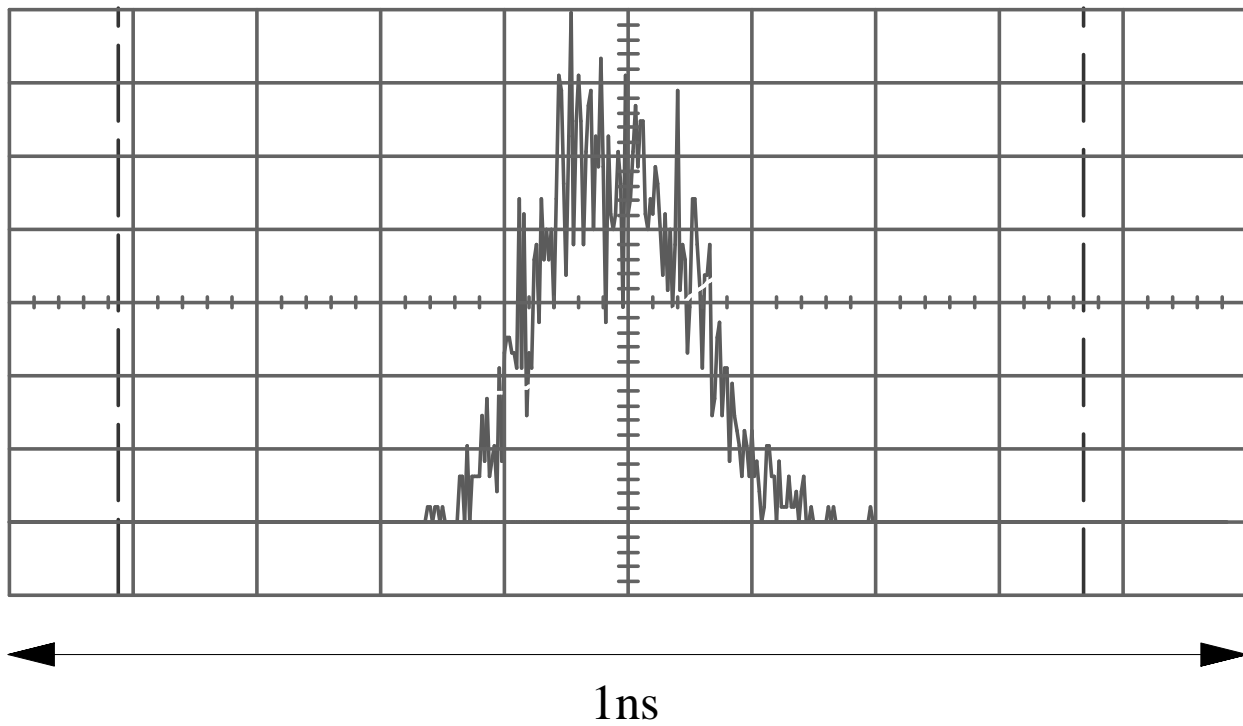


Fig. 9. Measured jitter histogram at 160Mbps and 16dB input SNR.




Fig. 10. Simulated convergence of I_{ICO} to I_o from 5% frequency offset in 100 bit periods.

LIST OF FIGURES

1	Block diagram of the complete NPDFE-based read channel.	10
2	Signal and slope quantization in acquisition mode.	11
3	Timing error estimation and loop filter. Signals with arrows in the middle are currents. . . .	12
4	(a) Multiplier and loop filter integrator. (b) Control signals	13
5	Direct path slope multiplier and scaling by α	14
6	(a) Three stage ring oscillator. (b) Replica bias circuit.	15
7	Die photograph.	16
8	Measured bit-error rate vs. input SNR with (solid) and without (dashed) timing recovery at 160Mbps.	17
9	Measured jitter histogram at 160Mbps and 16dB input SNR.	18
10	Simulated convergence of I_{ICO} to I_o from 5% frequency offset in 100 bit periods.	19

LIST OF TABLES

I	Measured performance at 3.3V and 25°C.	9
---	--	---