

Design of an integrated optical receiver in a standard CMOS process

Abstract—This paper presents an integrated optical receiver that allows for the “last-mile” access needed in today’s broadband networks. The circuit consists of an integrated photodetector, an amplification chain and a phase- and frequency-locked clock recovery system. It operates at a data rate of 155.52 Mbps with a sensitivity -20.45 dBm at a BER of 10^{-10} and was implemented in a 1.2 μm CMOS process. The circuit complies with relevant networking standards and synchronization of the data signal occurs within 1 μs . The entire receiver consumes approximately 26.5 mW from a 3.3 V supply and occupies an area of 2.9 mm^2 .

Index Terms—Clock and data recovery, CMOS analog integrated circuits, optical receivers, photodetectors.

I. INTRODUCTION

TODAY’S broadband networks support ever-increasing data rates in order to cope with the global demand for bandwidth. Although modern network backbones can usually service this traffic, bottlenecks occur at the nodes where end-users access these bandwidth resources. Thus, it is clear that the need for high-speed ‘last mile’ connectivity is rapidly becoming an important consideration in the implementation of telecommunication infrastructures.

Even though fibre-optic systems form the backbone of our telecommunication networks, optical communication circuits have been slow to follow the integration trend, exhibited by, for example, CMOS technologies. This has led to prohibitively expensive optical solutions for short-range communication links and has prevented fibre optics from becoming the ubiquitous means of data transport.

Seen in this context, the capability to produce a high-performance monolithic CMOS optical receiver with an integrated photo-detector would enable greater use of optics in short distance communication systems.

This paper describes the design of a fully integrated optical receiver that operates at a data rate of 155.52 Mbps and is compatible with the relevant optical telecommunication standards. The monolithic implementation of the receiver in a standard 1.2 μm CMOS process allows for the mass production of a mixed-signal circuit that can process analogue, digital and optical signals within the same Silicon structures.

A myriad of applications including wireless or fibre-based telecommunication links and high-speed inter-IC optical interconnects could be implemented with such a low-cost receiver.

The next section of the paper focuses on CMOS integrated photodetectors. Section III describes the design of the amplification chain and Section IV details the clock and data recovery system. Some simulation results are presented in Section V and Section VI concludes the paper.

II. INTEGRATED PHOTODETECTORS

When a burst of photons arrives at the surface of the photodetector, a fraction of the optical power is absorbed by the material and decays exponentially as a function of the penetration depth. As the field propagates into the material, it is attenuated and some of its energy is used to excite electron-hole pairs into the conduction and valence bands respectively. To harness these photo-generated charge carriers they have to be separated by an electric field before they recombine. These accelerated electron-hole pairs then give rise to a photocurrent that can be processed by the remaining optical receiver circuitry.

A. Photodetector Structures in CMOS

If a photodetector is to be implemented in a standard CMOS process the photo-generated charge carriers have to be accelerated by electric fields that exist within the material. Significant electric fields are present in the depletion regions formed at all pn-junctions and can be used to harvest electron-hole pairs. Figure 1 illustrates the basic structures present in a standard CMOS process.

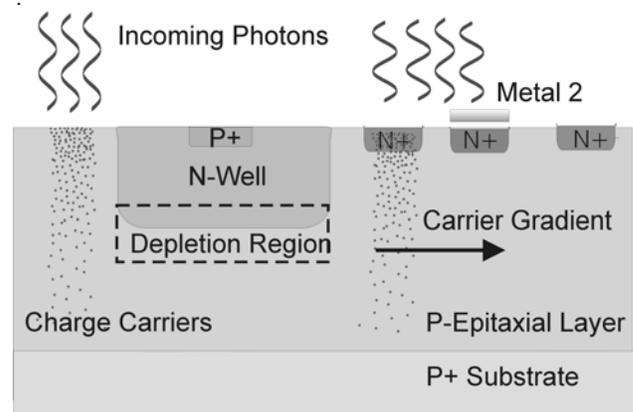


Fig. 1. Photo-detection structures available in a standard CMOS process. Incoming photons excite charge carriers within the substrate.

An important design parameter that has to be considered when selecting a structure is the responsivity of the photodetector. Responsivity is a measure of the current that is generated for a given optical input power and is proportional to the pn-junction’s quantum efficiency. Figure 2 gives an indication of the responsivities obtainable through the use of various photodetector structures and is based on the 1.2 μm CMOS process that was used for this project.

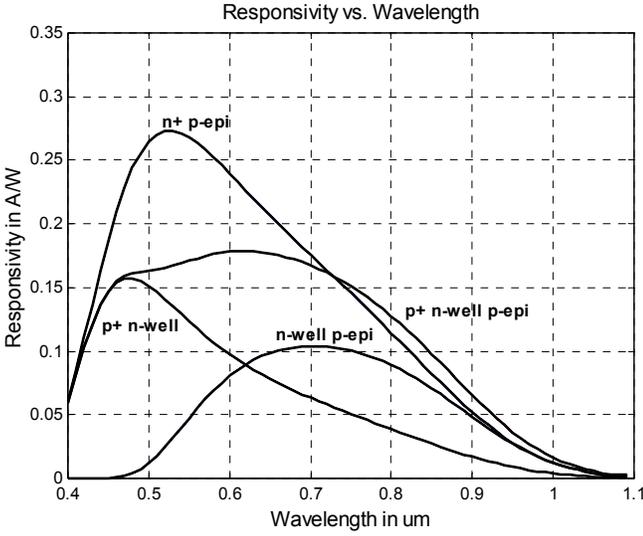


Fig. 2. Photo-detector responsivity vs. optical wavelength for various junction structures found in the 1.2µm CMOS process.

Even though detector responsivity is an important design parameter, it is the slow detector response times that have been found to prevent the implementation of fully integrated CMOS-based receivers. Electron-hole pairs that are generated deep within the substrate and have to diffuse a considerable distance to reach the depletion region slow down the photo-detector impulse response and result in a low frequency current gain. The effect of this diffusion tail is to increase the amount of intersymbol interference between consecutive light pulses, resulting in an increased bit error rate.

B. Spatially Modulated Detector Structures

To minimize the effect of this diffusion current on the detector frequency response, a spatially modulated light (SML) detector was implemented by placing a grid of identical pn junctions onto the Silicon surface and then masking alternate ones with a floating metal layer [1]. When this detector is illuminated, the absorbed light creates a spatial gradient in the carrier concentration that will relax by diffusion. The difference between illuminated and masked current responses will then result in an equivalent drift current response component that minimizes the diffusive tail. A spatially modulated light detector using the pn-junction formed between the n+ and p- epitaxial layer substrate was found to be the only structure in the 1.2 µm CMOS process that could accommodate the high data-rate required.

As the frequency response of a CMOS photo-detector is mainly determined by the transport of minority carriers [1], the time domain solution of the minority carrier distribution as described by the diffusion equation below, should give an indication of the speed obtainable with a CMOS photo-detector.

$$\frac{\partial n_p}{\partial t} - Dn \left(\frac{\partial^2 n_p}{\partial x^2} + \frac{\partial^2 n_p}{\partial y^2} \right) - \frac{n_p}{\tau_p} = g(t, x) e^{-\alpha y} \quad (1)$$

where Dn is the diffusivity of the minority carrier; n_p is the minority carrier concentration; τ_n is the minority carrier lifetime and $g(t, x)$ is the electron generation rate at the lower boundary of the depletion region.

The following figures show Matlab simulations of the time-domain solution of the minority carrier distribution for a simple SML structure with only one exposed junction.

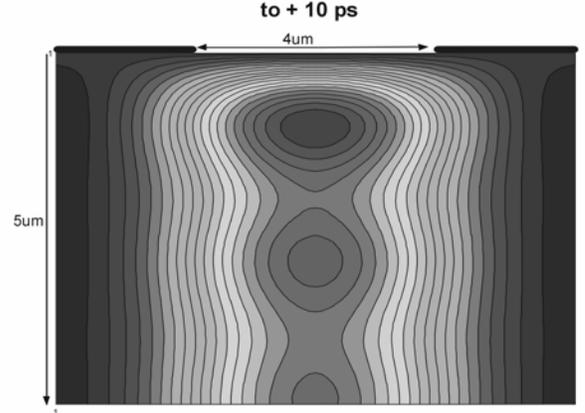


Fig. 3. Minority carrier concentration 10 ps after photon impact.

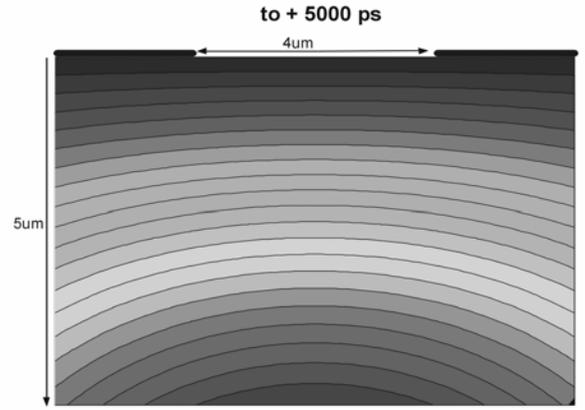


Fig. 4. Minority carrier concentration 5 ns after photon impact.

As can be seen in the figures above, the charge gradient relaxes very quickly near the surface of the material. As the SML structure uses the relaxation of the gradient to extract information from the incoming burst of photons, the data rate that can be supported is much higher than can be realized with other CMOS photo-diode structures. Table I summarizes some detector parameters.

TABLE I
PHOTODETECTOR PARAMETERS

Responsivity	0.041 A/W
Cutoff Frequency	442.5 MHz
Detector Capacitance	149 fF
Photodetector Area	3906 µm ²

III. AMPLIFIER DESIGN

Once a photo-current has been extracted by the detector structure, it has to be amplified for further processing. In order to amplify the equivalent drift current response, the receiver has to output the difference between the immediate and deferred current responses. This implies the need for two identical transimpedance-type preamplifier structures that feed into a high-speed difference amplifier, subtracting the immediate from the deferred current responses. As input signal strengths may vary considerably, a postamplifier that incorporates a limiting function for automatic gain control was included in the design.

Figure 5 illustrates the structure of the complete amplifier.

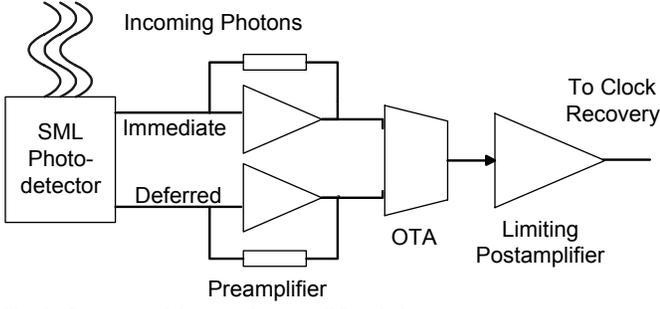


Fig. 5. Structure of the complete amplifier chain.

A. Transimpedance Preamplifier

Having established that a single gain stage would be used for the transimpedance amplifier due to a limited gain-bandwidth product, equation 2 was solved for the open loop gain and bandwidth required for a given feedback resistance and transfer function Q .

$$\frac{V_{out}}{I_{ph}} = \frac{-Z_{fb}(j\omega)}{1 + \frac{1}{A_v(j\omega)} \left(1 + \frac{Z_{fb}(j\omega)}{Z_d(j\omega)} \right)} \quad (2)$$

where Z_{fb} is the feedback impedance, Z_d is the input impedance of the amplifier and A_v represents the open loop gain. A feedback resistor was selected and the required open loop gains and cutoff frequencies were plotted against transimpedance bandwidth values. Figure 6 shows these design graphs plotted for a Q of 0.707 and 0.9 respectively.

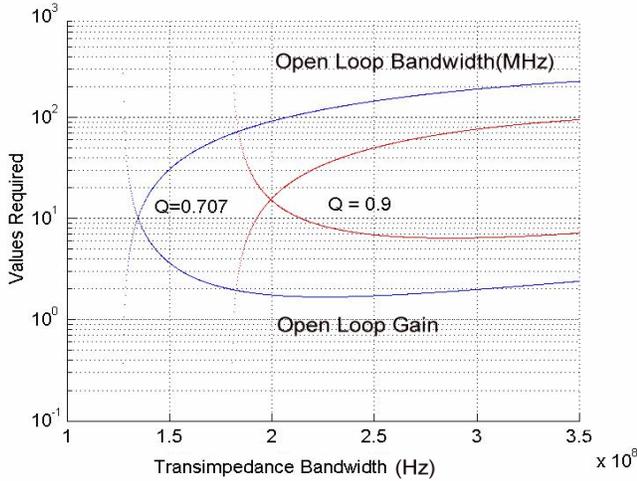


Fig. 6. Transimpedance amplifier design graphs.

As can be seen from the above figure, choosing an amplifier with a higher Q relaxes the open loop bandwidth requirements. Unfortunately this choice comes at the expense of a flat passband and leads to increased distortion of the signal.

The transistor implementation of the transimpedance amplifier is based on a structure developed by [2] and consists of a PMOS transistor with a diode-connected load biased by an NMOS current mirror. The input is connected to the output by

a feedback resistance that acts as a transimpedance element and converts the photo-current to an output voltage.

B. Operational Transconductance Amplifier

The high-speed difference amplifier is based on a differential pair of modified push-pull inverters that is biased at its switching threshold by active resistors. The inverter input stage is loaded with NMOS transistors to decrease the gain and hence increase the bandwidth of the amplifier. The differential signal is then converted to a single-ended output voltage for further processing.

C. Replica-biased Postamplifier

The postamplifier, based on a design by [3], was implemented by cascading two replica-biased inverter strings.

Replica biasing is a feedback strategy that aims to minimize the effect of duty cycle degradation through the optimal biasing of the various gain stages in the signal path. It involves level-shifting the information-modulated signal through a gain stage, to a set-point determined by a replica of the subsequent amplification stage. Figure 7 illustrates the basic structure of the replica-biasing scheme.

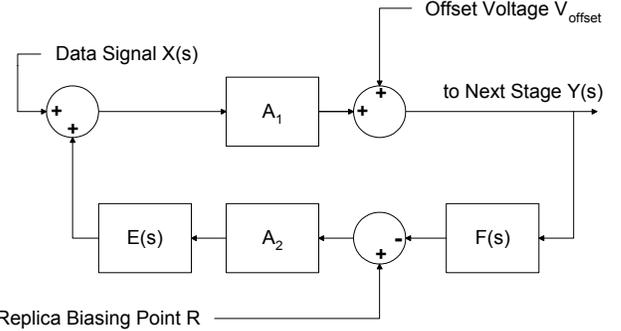


Fig. 7. Block diagram representing the replica-biasing scheme.

From the above figure, the loop dynamics can be described by equation 3 below.

$$Y(s) = A_1 [X(s) + A_2 E(s)(R - Y(s)F(s))] + V_{offset} \quad (3)$$

where $E(s)$ and $F(s)$ have low-pass characteristics.

Thus at low frequencies $\lim_{s \rightarrow 0} E(s) = \lim_{s \rightarrow 0} F(s) = 1$ and the output signal reduces to

$$Y(s) \approx X(s) \frac{A_1}{1 + A_1 A_2} + R \frac{A_1 A_2}{1 + A_1 A_2} + V_{offset} \frac{1}{1 + A_1 A_2} \approx R \quad (4)$$

Thus, the low frequency output of the loop, i.e. the biasing point of the next amplifier stage is driven to a value that is equal to the set-point determined by the replica circuit. As can be seen from equation 4, all low frequency components of the data signal $X(s)$ are attenuated.

Similarly for high frequencies $\lim_{s \rightarrow \infty} E(s) = \lim_{s \rightarrow \infty} F(s) = 0$ and

$Y(s)$ can be approximated as

$$Y(s) \approx A_1 X(s) + V_{offset} \quad (5)$$

Thus at high frequencies relative to the filter cutoff points, $Y(s)$ is equal to an amplified version of the input signal in addition to a small dc offset voltage. The effect of this offset can be minimized by increasing the open loop gain until $A_1 X(s) \gg V_{offset}$ and cascading identical replica biasing loops.

Both feedback loops consist of five inverting stages: three modified inverters, a comparator, and a level shifter. The modified high-speed inverters consist of a standard push-pull structure that is gain-limited by a NMOS diode-coupled transistor in order to increase the bandwidth. Two RC filters and a comparator complete the replica feedback loop. The comparator is based on a simple low-frequency operational transconductance amplifier (OTA) structure. The circuit implementation of figure 7 is shown below.

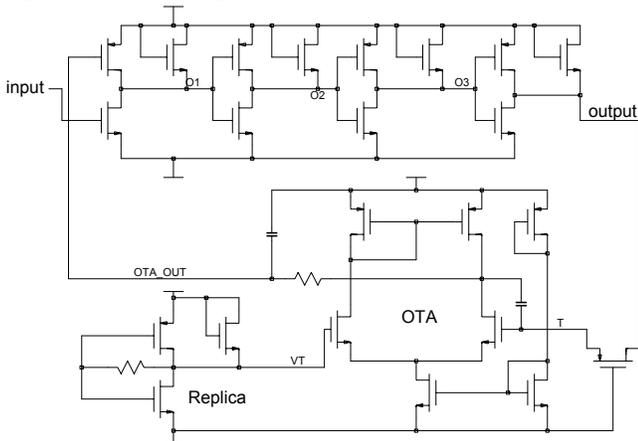


Fig. 8. CMOS implementation of one replica-biased gain stage

Table II presents some simulation results of the entire amplification chain.

TABLE II
AMPLIFIER SIMULATION RESULTS

Number of Stages	13
Worst-Case Bandwidth	106 MHz
Transimpedance Gain	204 dB Ω
Dynamic Range	106 dB

IV. CLOCK AND DATA RECOVERY SYSTEM

A. CDR Architecture

The amplified and possibly gain-limited signal that is obtained at the output of the receiver amplifier stages is an unsynchronized stream of symbols that may not be processed correctly by the subsequent digital circuitry due to input jitter and phase uncertainty. For it to be of any use, this signal has to be processed by a decision circuit and synchronized with a recovered clock signal to output an estimate of the original data sequence sent by the information source.

Clock and data recovery (CDR) circuits for the nonreturn-to-zero (NRZ) data streams used in optical communications have to include a non-linear processing element, such as the phase-locked loop used for this implementation, to create energy at the missing bit rate. Once energy has been introduced at the clock rate, this signal can be extracted.

Process variations that are present in the production of integrated circuits require PLL's to have a relatively large acquisition range. This requirement is necessary in order to ensure that the signal can still be locked on to despite the unpredictable local oscillator centre frequency, which may cause significant frequency errors between the incoming data rate, and the local clock. This error prevents phase

synchronization of the VCO with the incoming signal and leads to false clock frequencies being extracted from the data. Unfortunately a large acquisition range increases the output jitter of the recovered clock [4] to unacceptable levels.

In order to solve this acquisition problem, a frequency-locked loop (FLL) was implemented in parallel with the PLL in order to compare the frequencies of the input and output signals and to adjust a voltage-controlled oscillator (VCO) until the error reaches a sufficiently small value. Once this value is reached, the FLL is disabled and the low jitter PLL takes over, acquiring lock. The CDR architecture used for this circuit is shown in figure 9.

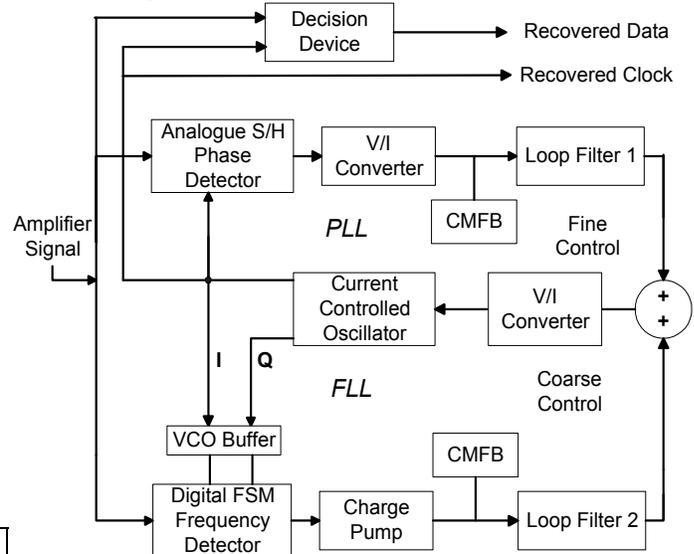


Fig. 9. Dual-loop clock and data recovery architecture.

B. VCO Design

Since the clock recovery circuit was implemented with a dual-loop architecture, the oscillator was designed to output quadrature clock signals as a provision for the digital frequency detector used in the FLL. The current-controlled oscillator (CCO) in figure 9 is based on a 2-stage ring oscillator with composite PMOS-loads [5] to ensure a -180° phase shift at the required unity-gain frequency.

In order to vary the oscillation frequency, the VCO incorporates delay interpolation, providing a tuning range wide enough to encompass process and temperature variations. As illustrated in figure 10, each oscillator stage consists of a slow and a fast path whose outputs are summed and whose gains are adjusted by a differential control voltage.

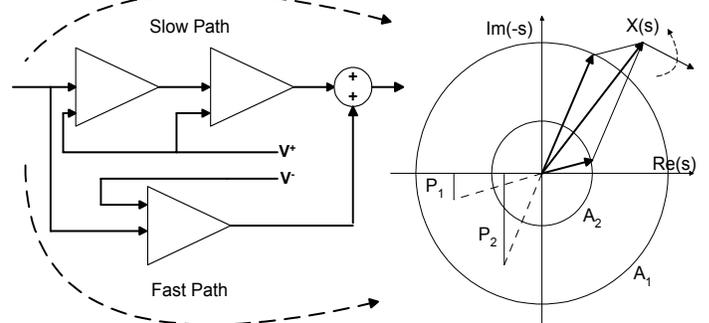


Fig. 10. Frequency variation by delay interpolation.

Let $X(s)$ be a signal applied to the input and let $e^{-j\omega T_d}$ be the delay experienced by $X(s)$ as it travels through one stage. Summing the contributions from the fast path with a gain A_1 and the slow path with a gain of A_2 yields an equivalent output signal that is described by

$$Y(s) = X(s) \left(A_1 e^{-j\omega T_d} + A_2 e^{-j2\omega T_d} \right) \quad (5)$$

Plotting this output signal in the complex plane seen in figure 10, indicates that the equivalent ‘delay vector’ can be modulated by varying the gains of the component ‘delay vectors’ having fixed phase angles P_1 and P_2 respectively.

The transistor implementation of figure 10 is shown below. The fast path consists of one differential pair, M5-M6, whereas the slow path consists of two differential pairs, M1-M2 and M3-M4. Interpolation is achieved by varying the tail currents of M5-M6 and M3-M4 in opposite directions and hence, modulating the respective differential pair’s transconductance. The output currents from the gain-modulated delay stages are then added to yield a signal, equal to the sum of the slow and fast path’s outputs.

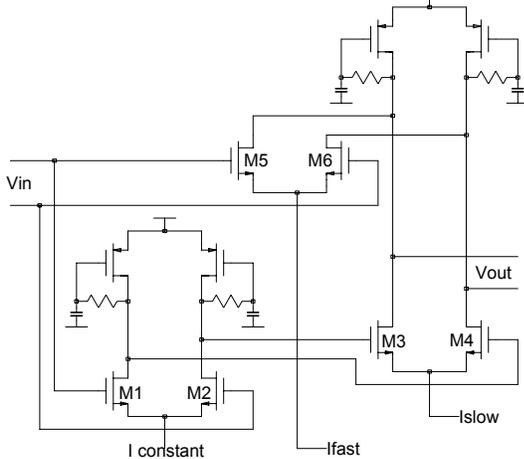


Fig. 11. Transistor implementation of figure 10.

A current folding circuit was designed to steer the tail currents and implement a coarse and fine control function with VCO gains for the FLL and PLL being approximately equal to 26 MHz/V and 7 MHz/V respectively. The CCO transfer function shown below validates the need for dual-loop architecture as the dramatic effects of process variations on the oscillator center frequency can clearly be discerned.

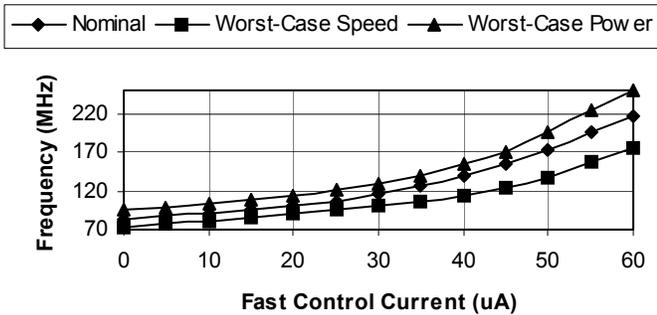


Fig. 12. VCO frequency as a function of a control current

C. Frequency Detector

The frequency detector is implemented using a digital 3-state machine to control the VCO frequency and is based on a system developed by [6]. The detector samples the in-phase and quadrature outputs of the VCO at the rising and falling edges of the unsynchronised data stream. By comparing these two sampled clock signals with values stored in its memory, the frequency detector can then go into either an UP, DOWN or RESET state. These states control a charge pump that can be used to pump the voltage on a floating loop filter up or down depending on what is required. The operation of the frequency detector can be explained by considering the phasor diagram in figure 13.

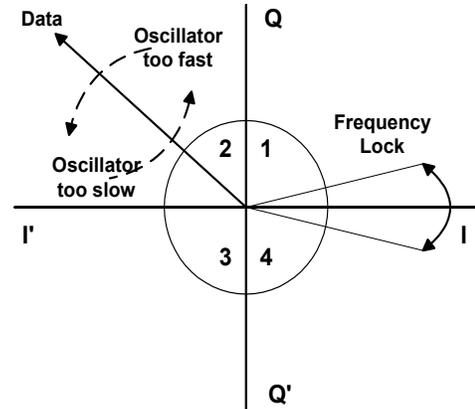


Fig. 13. Phasor diagram describing the operation of the frequency detector.

When the VCO is at a higher frequency than the data signal, the data phasor rotates counter-clockwise and the finite state machine (FSM) goes into the DOWN state. Similarly an oscillator that is too slow will cause the data phasor to rotate in the clock-wise direction and will initiate an UP output from the FSM.

D. Charge Pump

The charge pump is a switched current source that converts the logic levels of the 3-state frequency detector into an analog current signal that charges a loop filter.

A differential charge pump, based on [7], was designed. This charge pump provides two identical output paths for charging and discharging a floating loop filter to circumvent the problem of unequal currents flowing from the positive and negative current pumps respectively.

E. FLL Loop Filter and CMFB

In order for the FLL to lock onto the signal, despite process and temperature variations, the loop must have a capture range of about 40MHz. As the capture range is approximately equal to the loop bandwidth, the filter, based on a simple RC network, was designed with this cutoff frequency.

The output common-mode level of the floating loop filter is incompatible with the VCO control inputs and mandates the implementation of a common-mode feedback (CMFB) scheme based on a design by [7].

F. Phase Detector

A low jitter, analogue sample-and-hold phase detector based on [5] was implemented for this design. The phase detector is realized as a master-slave circuit that tracks the analog oscillator output continuously in response to a rising data transition. A falling data transition opens the first switch and the instantaneously sampled oscillator voltage is stored on parasitic capacitors at the input of the subsequent buffers. The second switch is also closed and the sampled oscillator voltage is transferred to the output and held. Thus the circuit generates an output that is linearly proportional to the input phase difference in the vicinity of lock.

The phase detector was implemented using two complementary differential sampling switches that are buffered by source-follower stages. To minimize the amount of charge that leaks out onto the parasitic capacitances of the buffers during data transitions, dummy switches are added on either side and are driven with an inverted clock to absorb a part of the injected charge.

G. V/I Converter and PLL Loop Filter

The output voltage of the phase detector needs to be converted to a current in order to charge the PLL loop-filter. A V/I converter, that also amplifies the phase detector signal in order to compensate for the attenuation resulting from the source-follower buffers, performs this function. Once the voltage that represents the phase error between the incoming data stream and the VCO signal is converted to a current, the control signal is channelled through a loop-filter to shape its spectral characteristics. The loop filter is based on a simple lead-lag network and was optimized for low-jitter operation.

H. Decision Device

Once the clock has been recovered, it is used to clock out an estimate of the data from a decision device. This decision device is nothing more than a high-speed D-flip-flop biased at the optimal switching threshold.

V. SIMULATION RESULTS

Figure 14 shows the PLL's locking transient after frequency lock has been achieved.

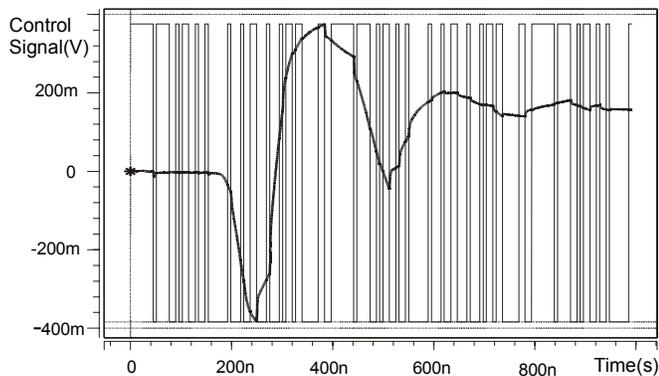


Fig. 14. PLL locking transient in response to a random bit-sequence.

TABLE III
OPTICAL RECEIVER SIMULATION RESULTS

Wavelength	770-880 nm
Data Rate	155.52 Mbps
Sensitivity	-20.45 dBm at a BER of 10^{-10}
Locking Time	1 μ s
Locking Range	100 MHz
Phase Noise	-76 dBc/Hz 130 kHz offset
Power Supply	3.3V
Power Dissipation	26.5 mW
CMOS Process	1.2 μ m (SAMES)
IC Area	2.9 mm ²

VI. CONCLUSION

This paper describes the implementation a fully integrated optical receiver that operates at a data rate of 155.52 Mbps in a standard 1.2 μ m CMOS process. A spatially modulated light detector feeds two identical transimpedance amplifiers whose outputs are processed by a difference amplifier to yield an equivalent drift current response. A post-amplifier that incorporates a limiting function for automatic gain control further processes the signal and feeds a parallel phase- and frequency-locked loop structure that recovers the clock from the NRZ data stream.

It was found that integrating optical functions in CMOS technology provides an inexpensive solution for short-range broadband applications. This project demonstrates that the design of an optical receiver that is fully integrated in a CMOS process and complies with a recognized telecommunications standard is indeed feasible and would allow for the so-called "last-mile" access needed in today's broadband networks.

VII. REFERENCES

- [1] D. Coppée, J. Genoe, J.H. Stiens, R.A. Vounckx, M. Kuijk, "Calculation of the Current Response of the Spatially Modulated Light CMOS Detector," *IEEE Transactions on Electron Devices*, vol.48, pp. 1892-1902, September 2001.
- [2] C. Rooman, D. Coppée, D. Kuijk, "Asynchronous 250-Mb/s Optical Receivers with Integrated Detector in Standard CMOS Technology for Optocoupler Applications," *IEEE Journal of Solid-state Circuits*, vol.5, pp. 953-958, July 2000.
- [3] M. Ingels, M. Steyaert, "A 1Gb/s, 0.7 μ m CMOS Optical Receiver with Full Rail-to-Rail Output Swing," *IEEE Journal of Solid-state Circuits*, vol.34, pp. 971-977, July 1999.
- [4] B. Razavi, *Design of Analog CMOS Integrated Circuits*, McGraw-Hill, p.550, 2001.
- [5] B.A. Anand and B. Razavi, "A CMOS Clock Recovery Circuit for 2.5-GB/s NRZ Data," *IEEE Journal of Solid-State Circuits*, vol. 36, pp. 432-439, March 2001.
- [6] H. Wang and R. Nottenburg, "A 1Gb/s CMOS Clock and Data Recovery Circuit," *IEEE International Solid-State Circuits Conference*, 2000.
- [7] H. Djahanshahi and C.A.T. Salama, "Differential CMOS Circuits for 622-MHz/933-MHz Clock and Data Recovery Applications," *IEEE Journal of Solid-state Circuits*, vol.35, pp. 847-855, June 2000.

Table III lists some results that characterize the receiver.