



École Polytechnique de Louvain
Department of Electrical Engineering

Design of a Serial Transceiver for Ultra-Low-Power Vision Sensor Nodes

Submitted by Charles Hovine

In partial fulfillment of the requirements for
the degree of Master of Science in Electrical
Engineering

June 2015

Advisors:

Pr. David Bol
Pr. Denis Flandre

Reader:

Guerric De Streel



ABSTRACT

Many multi-Gbps sub-pJ/b transceivers have been proposed as serial interfaces to ultra-low-power systems. If those solution are well adapted to high-speed systems, they come as excessive when targeting intermediate rate systems such as low-power vision sensor nodes working with several tens of Mbps. On the other hand, current low-rates solutions such as SPI or I²C present a very poor power efficiency. The literature describes many techniques to achieve high-speed and energy-efficient communication by the use of equalization, timing recovery and custom line drivers, but all those solutions targets very high rates and are thus much less efficient at lower rates. In this thesis, we introduce a 0.45pJ/b 150Mbps low-power transceiver enabling the transmission of data from chip to chip. Through the use of simple 1-tap discrete equalization, reduced swing and slew-rate controlled drivers, single clock regeneration and low-power decision circuit, we are able to propose a circuit working at rates higher than standard low-rates serial protocols while achieving the power efficiency of high-speed solutions. We expect this solution to be a viable alternative to currently existing systems and hence be suited for ultra-low-power vision sensor nodes.

ACKNOWLEDGMENTS

I am deeply grateful for the valuable advices and guidance I received from my thesis advisors, Denis Flandre and David Bol, without whom I could not have accomplished this work. I would also like to thank Gueric De Streeel for his availability and help with the tools I used and Julien De Vos for helping me understand the purpose of antenna rules. Furthermore I am thankful to my family for their support as well as to my friends who shared the same duties as I. Finally, I want to express my gratitude towards all the professors, researchers and assistants with whom I interacted over the past five years, as this thesis is the fruit of their teachings.

CONTENTS

Introduction	1
1 Background	5
1.1 Signal Propagation	5
1.1.1 Transmission Lines	5
1.1.2 FR-4 Microstrips	6
1.1.3 Low-Pass Response and ISI	8
1.1.4 Summary	10
1.2 State-Of-The-Art	11
1.2.1 Transmitter	12
1.2.2 Receiver	13
1.2.3 Limiting Factors	16
1.2.4 State-of-the-art of Low-power High-speed Transceivers	18
1.2.5 A Word About Classical Serial Buses	20
1.3 Problem Statement	21
2 Baseline Solution	23
2.1 Digital Power	23
2.1.1 Single Inverter	23
2.1.2 Tapered Buffer	26
2.2 Baseline Solution	28
2.2.1 Test Setup	28
2.2.2 Initial Architecture	29
2.2.3 Impact of Jitter	31
2.2.4 About Inverters Used as Input Stages	31
2.3 Summary	32
3 Design Of The Proposed Solution	33
3.1 Solution Architecture	33
3.1.1 Working with ISI	33
3.1.2 Global Architecture of the Solution	35
3.2 Tic-Toc Sampling	37
3.2.1 Sample & Hold	38
3.2.2 Source Follower	41
3.3 Regenerative Latch	41

CONTENTS

3.3.1	Performances	44
3.4	Drivers	46
3.4.1	Slew-rate limited driver	46
3.4.2	Bootstrap Circuit	47
3.4.3	Comparator	47
3.5	Clock Regeneration	49
3.6	Overall Performances	51
3.6.1	Data Eye Diagrams	51
3.6.2	Power Consumption	54
3.6.3	Measured BER and PVT Robustness	55
Conclusion and Future Prospects		59

LIST OF FIGURES

1.1	Infinitesimal section of a transmission line [Wikimedia]	6
1.2	Cross section of a microstrip [Wikimedia]	7
1.3	Microstrip transfer functions for various source resistances ($L = 10\text{cm}$)	8
1.4	Microstrip transfer functions for various lengths ($R_s = 10\text{k}\Omega$)	9
1.5	Decay time versus period	10
1.6	A transmission link [Yang, 1998]	11
1.7	Output drivers architecture [Poulton, 1998]	13
1.8	Input multiplexer [Yang, 1998]	13
1.9	Continuous-time equalization [Fu, 2012]	14
1.10	Decision feedback equalization [Kim et al., 2009]	14
1.11	Regenerative latch [Palermo, 2010b]	15
1.12	Receiver clocking architecture [Hu et al., 2010]	16
1.13	Data eye diagram [OnSemiconductors, 2014]	17
2.1	Energy per bit of a single $\times 2$ inverter charging a 40fF load	24
2.2	Output PAC of a single $\times 2$ inverter charging a 40fF load	25
2.3	Energy per bit of a single $\times 2$ inverter charging a 40fF load	26
2.4	Tapped chain of inverters [Frustaci et al., 2011]	27
2.5	Tapped chain switching consumption for an alternating input of 1's and 0's with $C_{in} = 1\text{fF}$, $V_{DD} = 0.8\text{ V}$ and $f = 100\text{MHz}$	28
2.6	Energy per bit of a minimum-delay tapered chain charging a 5.5pF load	28
2.7	Testbench	29
2.8	Transceiver architecture	30
2.9	MaTLAB simulation of the expected BER as a function of $f_{cut-off}/f_{signal}$	30
2.10	Power breakdown of the solution with optimally sized buffers	31
3.1	BER behavioral simulation for various peak-to-peak jitter amplitudes	35
3.2	Proposed receiver block-diagram	36
3.3	Proposed transmitter block-diagram	37
3.4	Tic-toc sequential sampling	38
3.5	Transient response of the Sample-and-Hold	39
3.6	NMOS switch with parasitic capacitances	39
3.7	Simulated scatter plot of the settling time and sampling error obtained by sweeping over the load and size parameters	40
3.8	Source-follower used to buffer the output	41

3.9	Regenerative latch	42
3.10	Operation of the regenerative latch	43
3.11	Digital latch	43
3.12	Instantaneous power of the latch (<i>clk</i> is in red)	44
3.13	High-speed operation of the regenerative latch	45
3.14	Offset distribution of the regenerative latch	45
3.15	Output driver	46
3.16	Bootstrap circuits implementations	48
3.17	Output driver equivalent resistance	48
3.18	High-speed comparator	49
3.19	Clock regeneration circuit	49
3.20	Filter attenuation versus settling-time	50
3.21	TX clock generation and phase-shift	51
3.22	Receiver input eye diagram for $L = 1\text{cm}$ and $R = 50\text{Mbps}$ (maximum slew rate), time is normalized to the symbol period	52
3.23	Receiver input eye diagram for $L = 5\text{cm}$ and $R = 150\text{Mbps}$ (maximum slew rate), time is normalized to the symbol period	52
3.24	Upper graph: Red – swing reference, Blue – receiver input. Lower graph: Red – Comparator output, Blue – Input data. (A) receiver input above both high and low reference (B) receiver input between high and low references (C) comparator’s output does not change (D) comparator’s output changes	53
3.25	Comparator’s input eye diagram (ΔV) for $L = 1\text{cm}$ and $R = 50\text{Mbps}$ (maximum slew rate), time is normalized to the symbol period	54
3.26	Comparator’s input eye diagram (ΔV) for $L = 5\text{cm}$ and $R = 150\text{Mbps}$ (maximum slew rate), time is normalized to the symbol period	54
3.27	Receiver allowade swing and clock delay for $L = 1\text{cm}$ and $R = 50\text{Mbps}$ (maximum slew rate), time is normalized to the symbol period	55
3.28	Global consumption of the system. At 50Mbps, the slew-rate control voltage V_c was set to 0.35V, 0.4V at 100MHz and 0.8V at 150MHz	56
3.29	Transceiver power breakdown	56
3.30	Power consumption of our solution against the consumption of the optimized baseline solution both at 100MHz	57

LIST OF TABLES

1.1	Microstrip parameters	7
1.2	Summary of the performances of state-of-the-art systems	19
2.1	Channel RLC parameters	29
3.1	MUX/DEMUX control logic	37
3.2	Sizing and performances	40
3.3	Sizing and performances	41
3.4	Transistors' sizes of the reg. latch	44
3.5	Sizing and performances	48
3.6	Pass/fail test results for $V_{DD} : 0.7 \rightarrow 0.9$	57

LIST OF USED ABBREVIATIONS

IoT Internet-of-things

WSN Wireless sensor node

M2M Machine-to-machine

CDR Clock and data recovery

ISI Inter-symbols interferences

DAC Digital to analog converter

CTLE Continious-time linear equalization

DFE Decision-feedback equalization

RX Receiver

VMD Voltage-mode drivers

PVT Process, voltage, temperature

PLL Phase-locked loop

DLL Delay-locked loop

ILRO Injection-locked ring oscillator

PI Phase interpolator

BER Bit-error rate

SNR Signal-to-noise ratio

PCB Printed circuit board

FPS Frames-per-second

RF Radio frequency

ESD Electrostatic discharge

PAC Pulse amplitude closure

NRZ Non-return-to-zero

S&H Sample and Hold

RMS Root Mean Square

RL Regenerative latch.

INTRODUCTION

The Internet Of Things (IoT) will enable the connection of nearly every *things* through the deployment of up to trillions of Wireless Sensor Nodes (WSN) [Bryzek, 2012]. IoT opens the doors to many exciting applications such as health monitoring [Swan, 2012], smart cities [Sehgal et al., 2015], production line optimization [Severi et al., 2014], environment monitoring [Li et al., 2012] and more. With the data collected, we hope to optimize our usage of human and natural resources. For example, building on the anticipated success, 37 millions smart electricity meters have already been deployed in the United States and used to optimize the electrical consumption [EIA, 2015].

The phrase "Internet of things" was coined by Kevin Ashton in 1999 [Ashton, 2009] but the concept existed long before that. The first ever "thing" being connected to the Internet was a Coke machine modified by Carnegie Mellon University engineers in 1982. So why is it that we only hear about it now? Firstly, the cost of hardware is decreasing. That is as true for Systems On Chips (SoC) following Moore's Law [Bauer et al., 2013] than for the Micro Electromechanical Systems (MEMS) [Yole, 2013] used as actuators in IoT enabled devices. Secondly, cellular technologies are now suited for Machine-to-Machine (M2M) communications and extensive researches are conducted so that the next standard for cellular communications (namely 5G) will be able to handle the expected massive data transfers generated by M2M communications [Ratasuk et al.,]. Next, we now have the software and hardware capabilities to analyze the collected data and process them around the world through delocalized cloud computing [Edson, 2015]. Lastly, the envisioned economic impact is huge: major companies like Cisco estimates that the Internet of things could generate up to \$14.4 trillion revenue in the next decade [Cisco, 2015].

As the volume of wireless sensors envisioned is gigantic, circuits must be designed so to reduce both their production and operation carbon footprint. [Bol et al., 2013] estimates that the carbon footprint related to the production of 1000 WSN's amounts to approximately 8% of the average annual footprint of a European citizen. According to the same study, replacing batteries with energy harvesters could reduce the production footprint of about 12%. The success of WSNs thus mainly depends on our ability to integrate sensing, processing and communication capabilities in small and Energy Autonomous Systems (EAS). Due to the severe power constraints set by the energy harvesters used in EAS, there is a growing need for highly efficient and very low-power systems.

When reducing power and area, chip-to-chip interconnects are a severe limiting factor which is partly due to the important amounts of energy required to charge and discharge the large off-chip

capacitances. Current link-optimized implementations of low data rates serial protocols such as SPI or I2C consume up-to 50pJ/bit with rates limited to a few Mbps. Implementations of state-of-art protocols such as MBus [Pannuto et al., 2015] stand above a 22pJ/bit consumption¹. On the other hand, recent advances in the field of low-power high-speed transceivers opened new opportunities for multi-Gbps sub-pJ/b serial communications. But what is particular about high-speed systems? We talk about "high-speed" when "one or more digital abstractions fails, as a direct consequence of the circuit speed" [Diorio, 2004]. In the case of chip-to-chip interconnects, high-speed operation breaks our assumption about interconnects carrying signals undistorted. Unlike previous systems mostly relying on synthesizable cells, high-speed systems integrate many custom building blocks including drivers, timing recovery circuits, equalizers and more.

The techniques used in high-speed low-power serial communications can be applied to medium-rates (50-500Mbps) low-power systems such as low-power vision sensor nodes. On one hand, when working at medium rates, timing errors and channel attenuations are less severe and we can thus save power by disposing of the elements used to mitigate their effects. On the other hand, one of the benefits of high-rates is that, as the static power does not vary with increasing rates, its cost is distributed over more bits, thus reducing the mean energy required to send one bit (E_b). Current implementations of low-power wireless vision sensor nodes present resolutions no higher than 0.1 Mpixels [Choi et al., 2012] [Kim et al., 2014]. At 32 frames per second and 8-bit color depth, it translates to a continuous data-rate of about 25Mbps. Similarly, VGA cameras can produce a continuous rate of 80Mbps. In the case of the SunPixer SoC, another state-of-the-art ultra-low-power CMOS imager [Bol et al., 2014], an efficiency of 17pJ/pixel for 128x128 pictures taken at 32fps is achieved. The cost of one image bit is thus 2.1pJ. We cannot afford to transmit that information using the standard protocols mentioned above as their consumption would defeat the point of a low-power sensor. Consequently, we must design highly efficient sub-pJ/bit links to transmit the data off-chip.

Considering the above, the guiding question of this work will be: What is the most efficient way to send and recover datas at an intermediate rate on a highly capacitive serial link? We will see that injecting more power in the channel by increasing the output driver's size is not the answer for power and area-efficient communications. Current state-of-the-art solutions have moved their focus from *how to send* to *how to recover* by increasing the receiver complexity. In this context, we propose a novel 2-wire sub-pJ/bit 150Mbps serial transceiver implemented in 65nm LP CMOS. It borrows techniques from the field of high-speed I/O's and aims at being integrated to low-power WSN's for the IoT. It achieves a power consumption as low as 0.52pJ/bit at a rate of 150Mbps on 5cm FR-4 microstrips. Low-power operation on the transmit side is accomplished by limiting the swing of the output drivers through a continuous sensing of their outputs. Furthermore, the input signals of the drivers are bootstrapped in order to limit the driver's size and, as a consequence, dynamic power. On the receive side, a simple 1 tap discrete-time linear equalizer was implemented and a regenerative-latch comparator was used to resolve the received datas to valid logic levels with no static currents. The clock is regenerated from

¹Note that MBus is a point-to-multipoint interconnect. The consumption of 22pJ/bit was reported when only 2 nodes were communicating.

the highly attenuated received clock by a simple comparison with the received signal's mean. Finally, the consumption was further reduced by decreasing the supply voltage down to 0.8V.

This thesis is organized as follows. The following chapter introduces the properties of FR-4 microstrips. Starting by a brief introduction to transmission lines theory, we define useful rules of thumb relative to the design of transceiver without impedance matching. We then go over the past and present architectures of high-speed serial transceivers, detailing the motivations between each blocks, before presenting three state-of-the-art solutions. We conclude this chapter by stating the problem and the specifications of its solution.

In chapter 2, the properties in terms of speed and power of basic CMOS inverter drivers are thoroughly studied. We extract the parameters influencing consumption and their impact on the system's speed. We then present our initial and naive solution and show its limitations in terms of power, maximum rate and reliability. We conclude the chapter by discussing the requirements of an improved and more efficient solution.

In the third chapter, we present our solution. After a global explanation and justification of the chosen architecture, each block is explained, meticulously studied and characterized. We conclude the chapter by an evaluation of the overall performances of our design and compare them to the performances of the baseline solution.

The fourth and last chapter suggests other prospects and perspectives of improvements of the solution and finishes by a brief conclusion.

BACKGROUND

In the first part of this chapter, we shortly present the ABC's of transmission-lines theory, required to understand the behavior of FR-4 microstrips. We then show the origins of inter-symbol interferences and how to avert it. In a second part, we describe the basic elements of a high-speed serial transmission system. We first describe the main architectural choices defining the link before detailing the sub-systems encountered. We then expose the figures of merit of a transmission link before enumerating the factors limiting the performances of such systems. We conclude this chapter by presenting the current state-of-the-art of low-power high-speed serial transceivers and defining the constraints on the solution.

SECTION 1.1.

SIGNAL PROPAGATION

When a high-frequency signal travels along a medium whose length is of the same order of magnitude than the considered signal's wavelength, the wave nature of the signal must be taken into account in order to properly model the effects of the medium. The most prominent consequence of the *transmission line* nature of the link will be the apparition of reflections, potentially introducing interferences between, the successive symbols.

1.1.1. TRANSMISSION LINES

Neglecting reflexion (considering a matched line, that is a line whose terminations have the same impedance than the line characteristic impedance), the transfer function of a transmission line of length l can be expressed as

$$H(j\omega) = e^{-\gamma l} \quad (1.1)$$

with γ being the complex propagation constant defined as

$$\gamma = \sqrt{(R + j\omega L)(G + j\omega C)} \quad (1.2)$$

R , L , C and G corresponding to the resistance, inductance, capacitance and conductance per unit length as depicted on figure 1.1 on the following page [Schrader et al., 2006]. We see that the magnitude of (1.1) decreases with increasing length, frequency and parasitic elements.




Figure 1.1: Infinitesimal section of a transmission line [Wikimedia]

The main sources of imperfection of a conductor are the skin effect and the dielectric loss. The skin effect consists in the electromagnetic waves not penetrating all the way to the center of the conductor, thus only propagating in the border region of the material [Schelkunoff, 1934]. As a consequence, due to the reduced area through which the wave is propagating, the effective resistance of the conductor increases as well as its internal inductance. Dielectric loss is linked to the imperfection of the isolating dielectric. When considering an ideal capacitor, its impedance is purely complex but in practice, a real term modeled as an equivalent series resistance appears. This term corresponds to the loss due to the dielectric conduction as well as the bipolar moments present in the dielectric [Schrader et al., 2006]. Another source of impairment comes from improperly terminated lines. Let us first link the characteristic impedance of a line to its RLGC parameters:

$$Z_c = \sqrt{\frac{R + j\omega L}{G + j\omega C}} \quad (1.3)$$

When a wave traveling in medium 0 of characteristic impedance Z_0 reaches the interface to medium 1 of characteristic impedance Z_1 , part of the wave power will be transmitted to medium 1 while the rest will be reflected back to medium 0. The impact of reflections can be quantified by the use of the reflection coefficient:

$$\Gamma = \frac{V_{reflected}}{V_{transmitted}} = \frac{Z_1 - Z_0}{Z_1 + Z_0} \quad (1.4)$$

If the characteristic impedances of the mediums are matched, no reflections will occur. Note that the reflection coefficient Γ can be complex and frequency dependent and will thus introduce different attenuations and phase-shifts for different frequencies. In other words, when matching the impedances of two lines, the matching will only be effective for a specific range of frequencies.

1.1.2. FR-4 MICROSTRIPS

The anatomy of a microstrip can be seen on figure 1.2 on the next page. It is composed of two conductive layers with the lower one serving as a reference ground plane. Between those layers, a dielectric material ensures proper insulation. On most printed circuit boards, FR-4 glass epoxy is used as an insulating material. Its RLGC parameters, and thus its characteristic impedance, are entirely defined by the microstrip cross-section thus allowing the line to be properly matched independently of its length. Rather than computing the analytical expressions of the transfer functions, we simulated microstrips of varying lengths through ADS (*Advanced Design System*) from *Agilent technologies*. We can see on figure 1.4 on page 9 and 1.3 on page 8 the impact of the line length and source resistance on the system's response. The simulations were performed

with the parameters presented in table 1.1 which correspond to a 50Ω microstrip laid on a FR-4 substrate.




Figure 1.2: Cross section of a microstrip [Wikimedia]

Substrate height	1 mm
Relative permittivity	4.6
Conductor conductivity	$5.96 \cdot 10^7$ S/m
Loss tangent	0.014
Conductor thickness	$60 \mu\text{m}$
Conductor width	1.7 mm
Load capacitance	10 fF

Table 1.1: Microstrip parameters

From figure 1.4 on page 9 and 1.3 on the following page, we can distinguish two parts in the channel response: for low frequencies, the channel acts as a smooth second order low-pass filter decaying at -40dB per decade, while for higher frequencies, the response becomes spurious. Those spurious peaks are the consequence of improperly terminated lines leading to signal reflections. From those figures we can derive the following rules of thumb:

$$f_{reflection} \approx \frac{c}{10l} \quad (1.5)$$

$f_{reflection}$ being the transition frequency between the low-pass and reflections part of the frequency response, l being the microstrip length and c , the speed of light. We see that R_s does not affect $f_{reflection}$ but impacts the 3-dB cut-off frequency f_c ¹. We see that

$$f_c \propto 1/(lR_s) \quad (1.6)$$

Lastly, eq. (1.4) can be used to determine the amplitude of the reflections at the source and load. If the reflections are located in the signal's bandwidth, inter-symbol interference due to reflected symbols will occur. That is why, when designing a communication system relying on

¹On the other hand, R_s definitely impacts the amplitude of reflections through eq. (1.4). In the worst tested case, $R_s = 10\Omega$, the channel attenuation and reflection coefficient are such that the energy injected into the channel is greater than the energy transmitted and dissipated, resulting in an amplification of the signal at higher frequencies




Figure 1.3: Microstrip transfer functions for various source resistances ($L = 10\text{cm}$)

unmatched transmission lines, an upper-bound on the line length and the transmission rate must be defined.

1.1.3. LOW-PASS RESPONSE AND ISI


A transient impulse response $h(t)$ is associated with the transfer function (1.1). If transmitting NRZ (Non-Return-to-Zero)² data, the resulting received signal will be (if sampling is performed at the end of the bit-time T , and assuming $h(t)$ to be causal³)

$$y[m] = \sum_{n=-\infty}^m b_n g((1+m-n)T) \quad (1.7)$$

with b_n , the m^{th} sent bit and $g(t)$ the total system response $g(t) = h(t) \otimes p(t)$, $\otimes p(t)$ denoting the convolution with the pulse shaping filter (a rectangular window of width T in the case of NRZ modulation). Equation (1.7) can also be written as

²NRZ data encode a one with a positive voltage (or positive current) and a zero with a negative or null voltage (or negative current)

³ $h(t) = 0$ for $t < 0$

Figure 1.4: Microstrip transfer functions for various lengths ($R_s = 10\text{k}\Omega$)

$$y[m] = b_m g(0) + \underbrace{\sum_{n=0}^{m-1} b_n g((1+m-n)T)}_{ISI} \quad (1.8)$$

where the second term corresponds to ISI. In order to avoid ISI, the pulse response associated to any bit m must be zero when sampling any other bit $n \neq m$, this condition can be expressed as

$$g(nT) = \delta[n] \quad (1.9)$$

which is known as the Nyquist criterion. Is this condition fulfilled if $H(j\omega)$ behaves as a low-pass filter in the signal bandwidth? We can express a first-order approximation of the pulse response of a low-pass system as [Buckwalter et al., 2004]

$$g(t) = \begin{cases} 0 & \text{if } t \leq 0 \\ 1 - e^{(T-t)/\tau} & \text{if } 0 < t \leq T \\ e^{-t/\tau} (1 - e^{-T/\tau}) & \text{if } t > T \end{cases} \quad (1.10)$$

where τ the channel time-constant (which can be approximated by $R_s l C_{MS}$ with C_{MS} the capacitance per unit length of the microstrip). For correct reception, there is of course a constraint on the ratio T/τ . Figure 1.5 depicts in blue the time it takes for the response to decay to 0.05 and in dashed red, the time between the start of the decay and the sampling of the next bit (i.e. $T/2$). ISI will be inferior to 0.05 as long as $t_{0.05}$ is inferior to $T/2$, that is approximately, $T/\tau > 6$ which is nearly equivalent to $1/T < 2\pi/\tau$ or $f < f_c$. This constraint does not exactly satisfy equation (1.9) but is close enough for ISI to be neglected.




Figure 1.5: Decay time versus period

1.1.4. SUMMARY

In this first section, we showed that inter-symbol-interferences stem from two phenomenon: the low-pass response of the channel and high-frequency reflections. In order to completely suppress

reflections, impedance matching should be done both at the source and load of the line. This is done in practice by the addition of 50Ω resistors. Unfortunately, those resistors will dissipate significant static power and are thus avoided in low-power applications. In order to minimize the impact of ISI, the signal bandwidth should not exceed significantly the cut-off frequency. Furthermore, reflections should not occur in the signal bandwidth and the signal maximum frequency should consequently not exceed $c/10l$.

SECTION 1.2.

STATE-OF-THE-ART

Many architectures are available when designing serial transmission links. The choice will be driven by the application's constraints: low-power, high-speed, area-efficient, cheap, etc. Figure 1.6 shows the three basic elements constituting a transmission link (serial or otherwise): the transmitter which shapes the data and drives them onto the transmission medium, the channel, and the receiver which regenerates the distorted signals to their proper digital levels.




Figure 1.6: A transmission link [Yang, 1998]

From a broader system point of view, multiple topologies can be encountered. Differential systems, where the data are sent with a known reference or their complementary, demonstrate a very good common-mode noise rejection, allowing for smaller amplitude signals to be sent and thus, lower power systems. Still, matching the length of the two paths can be a challenging problem. Although it requires more area, it is widely believed that transmission can be as much as two times faster [Poulton, 1998] [Palermo, 2010a]. On the matter of clocking, the three basic structures are: forwarded clock, common clock and embedded clock. Common clock presents the benefit of having an unique oscillator for both the receiver and transmitter. If the paths' delays are equal the system is termed "synchronous". In the case where the delays are not equal, the phase must be recovered from the data with the help of additional circuitry (typically a DLL or PLL), those systems are called "mesochronous". In forwarded clock systems, the clock is transmitted along the data. That architecture dispenses for the need of a timing recovery circuit as long as the delays are matched. Besides, the channel will attenuate the clock signal and amplify already present jitter. Its main advantages are its simplicity and low cost in terms of required components. This leads us to the embedded clock architecture, the heaviest in terms of required hardware. Transmitter and receiver both receive a separate clock with possibly different phases and frequency and CDR (Clock and Data Recovery) is performed at the receiver in order

to recover the clock matching the data. [Palermo, 2010b].

1.2.1. TRANSMITTER

A serial transmitter is usually made of an output driver and, optionally, a forward equalizer, input serializer and output impedance regulator used to dynamically match the line impedance.

FORWARD EQUALIZER

In most cases, the transfer function of the transmission medium is non-flat and the sent symbols are thus distorted. In order to overcome the imperfect characteristic of the channel, equalization can be used, either at transmit or receive side. Forward equalization, also called transmit equalization, can be implemented as continuous or discrete-time (we then talk of feed-forward equalization). Feed-forward equalization pre-emphasizes the symbols according to the previously sent symbols in order to compensate for ISI before sending the symbols onto the channel. Continuous transmit equalizers pre-amplify the high frequency content of the signals using an integrated active filter while discrete equalizers use DAC to compute the transmitted symbol as the weighted sum of consecutive bits [Peffer, 2003].

OUTPUT DRIVER

The output driver is the most essential component of the transmitter, it drives the signal as a current or voltage on the channel. Its drive strength will determine the cut-off frequency of the channel and thus the output distortion. Figure 1.7 on the next page shows the two basic output driver architectures, voltage and current-mode. In voltage-mode, the line is charged to V_{high} or V_{low} and the symbols are encoded onto the voltage levels. Current-mode, on the other hand, signals through different current levels then converted to voltage at the receiver termination [Poulton, 1998]. Current-mode drivers were historically known for their better power efficiency over voltage-mode drivers, but the recent attempts at aggressively decreasing the supply voltage showed that voltage-mode drivers performed better at low voltages [Green and Singh, 2003]. Nonetheless, VMD suffer from switching noise and can induce crosstalk when multiple parallel drivers share the same supply [Poulton, 1998]. Some have suggested to reduce the voltage swing of drivers in order to minimize output switching power. Voltage swings as low as 100mV_{pp} are used but a separate supply is thus required for those drivers. It is common when signaling at a lower swing to also reduce the common-mode output voltage of the driver. In this configuration, two NMOS devices are used as pull-up/pull-down and working in the triode region. By working in the triode region, impedance matching can easily be implemented by adjusting the value of the gate voltage [Song et al., 2009] [Song et al., 2013] [Wong et al., 2004]. Impedance matching can also be obtained by adjusting the resistor used in the implementation of current-mode drivers (see fig. 1.7b on the facing page).

SERIALIZER

In the case where multiple data streams are to be transmitted on the same link, they are multiplexed through a high bandwidth multiplexer [Palermo, 2010a]. A typical implementation is shown on figure 1.8 on the next page. The capacitance of the multiplexing node, A , will grow

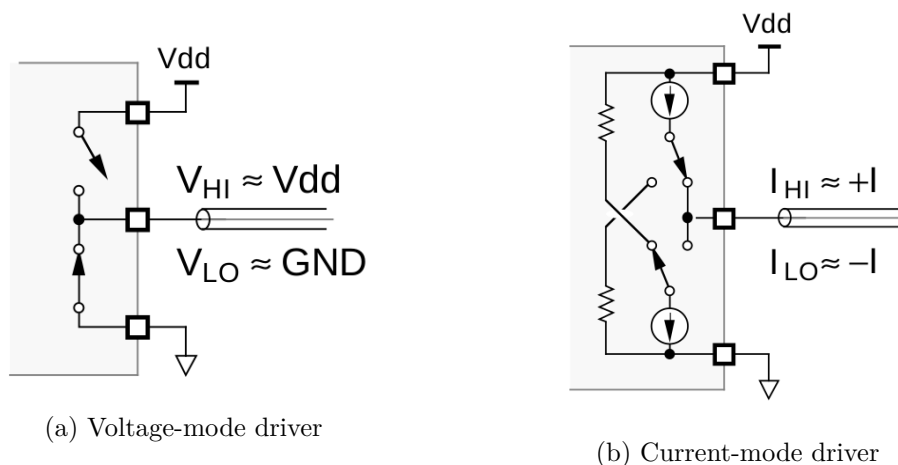


Figure 1.7: Output drivers architecture [Poulton, 1998]

linearly with the fan-in of the multiplexer, potentially introducing ISI before the application of the channel response. One of the challenges for those designing high-bandwidth multiplexer is the generation of the select signals. Those signals must be non-overlapping short pulses of width T_{bit}^{in}/N , T_{ck}^{in} being the input bit time and N the number of parallel streams. Due to the process intrinsic speed limit, achieving sharp pulses with widths inferior to t_{FO4} ⁴ is really challenging and special techniques must often be deployed [Yang, 1998] [Kyeongho et al., 1995].




Figure 1.8: Input multiplexer [Yang, 1998]

1.2.2. RECEIVER

A serial receiver is usually comprised of an input buffer, an equalizer, a timing recovery circuit and possibly a deserializer. State-of-the-art implementations of serial link often only consist of a very efficient receiver, independent of any particular transmitter. Rather than trying to lower the channel time-constant by using power-hungry output stages, an efficient receiver capable of recovering very small and distorted input is often preferred.

⁴The FO-4 delay, or Fan-out of 4 delay, is the delay through an inverter loaded by a capacitance four times larger than its input capacitance, and driven by an inverter four times smaller.

EQUALIZER

Two major architectures are usually considered when doing receive equalization. Figure 1.9 illustrates the concept behind continuous-time linear equalization. CTLE is usually implemented through an integrated high-pass filter [Hu et al., 2012] and is more efficient at equalizing channels with a smooth frequency response [Fu, 2012]. It presents the benefit of being easily tunable and thus easily adaptable to varying channels [Hu et al., 2012]. Decision feedback equalization, however, takes previous decisions into account in order to remove ISI⁵. That is, the currently received value is estimated as a linear combination of the the received signal and the previously made decisions (discrete values, 0 or 1 in the case of NRZ coding). More complex structures exist where the equalizer coefficient are dynamically computed during a channel estimation phase by method such as the gradient descent [P. Sobieski, 2010]. A simple two-taps DFE is represented on fig. 1.10. Those equalizers obviously need as much DAC's to generate the feedback coefficients than the total number of coefficients but are more efficient at equalizing channels with a "bumpy" frequency response [Fu, 2012].




Figure 1.9: Continuous-time equalization [Fu, 2012]




Figure 1.10: Decision feedback equalization [Kim et al., 2009]

INPUT STAGE

The input-stage is usually composed of a pre-amplifying stage⁶ and a sampler/comparator. The most basic input buffer encountered uses a simple inverter for pre-amplification and the result is sent to a simple flip-flop. This simple solution unfortunately have a major drawback: amplitude changes of the input signal due to noise and channel loss might lead to timing error due to the

⁵As for feed-forward equalization, it computes a weighted sum of the previously received symbols and decision to compute the expected sent symbol.

⁶The CTLE is often implemented in the pre-amp stage.

constant DC reference used by the inverter [Baker, 2010] [Sidiropoulos and Horowitz, 1997]. Even more, it suffers from a high supply voltage sensitivity as well as high PVT variation translating in high input offset [Palermo, 2010b]. In practice, differential structures are used in order to account for those losses. More elaborate solutions use a differential pair for pre-amplification (which can be coupled with a linear equalizer [Hu et al., 2012]). The most commonly seen clocked comparator in high-speed low-power applications is the regenerative latch (see fig. 1.11). A regenerative latch presents the benefit of providing a clocked output without drawing any DC current. Yet, it must be precharged prior to resolving an input. Furthermore, the output are already at fully valid logic levels.




Figure 1.11: Regenerative latch [Palermo, 2010b]

TIMING RECOVERY

In a practical system, there will always be a unknown time-varying offset between the considered clock and the reference clock known as jitter. Low pass channels introduce data-dependant deterministic jitter which can fortunately be equalized [Buckwalter and Hajimiri, 2006] [Analui et al., 2005]. Besides, random jitter will appear due to noise and crosstalk in the circuit (including the oscillator circuit). Furthermore, mismatch between PCB traces and various circuit delays introduce a systematic timing offset. Typically, three structures are used to perform timing recovery: phase-locked loops, delay-locked loops and injection-locked ring oscillators. As seen on figure 1.12 on the next page, a typical DLL/PLL architecture generates multiple phases from the received clock which are then interpolated (or simply rotated amongst themselves) to provide the receiver with a clock in phase with the data. ILRO's work by injecting a signal of known frequency in an oscillating system. The frequency difference between the injected signal and the system will induce a phase-shift in the system. ILRO's require less area, power and can work at smaller voltages. Besides, no PI (phase interpolator) is required [Hu et al., 2010]. Unfortunately, ILRO are hard to model and a considerable number of steps is required to properly simulate them through SPICE [Bhansali and Roychowdhury, 2009].

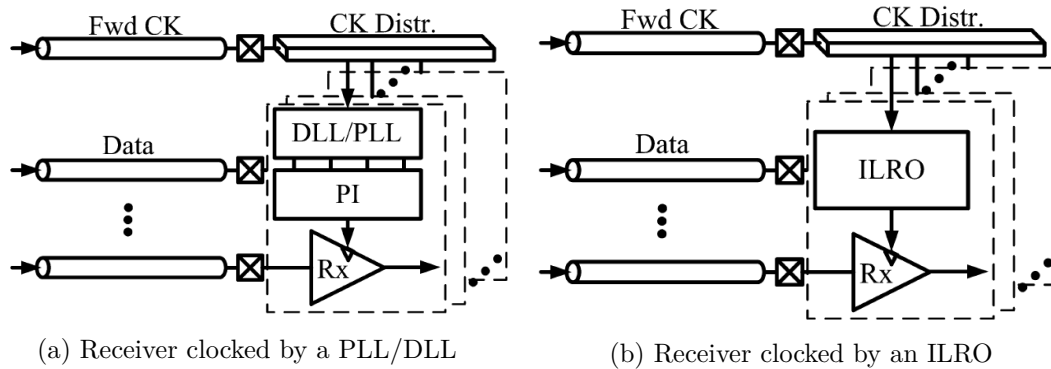


Figure 1.12: Receiver clocking architecture [Hu et al., 2010]

1.2.3. LIMITING FACTORS

In this section, we first cite the various performance metrics of a serial communication system before enumerating the elements restricting those performances.

PERFORMANCE METRICS

The most common figure of merit of any communication system is undoubtedly the bit-error rate (BER). It is simply defined as the ratio of correctly decoded bit over the total number of processed bits. Generally, communication is often considered error-free for BER's inferior to 10^{-12} . It is common to compute the BER as a function of the signal-to-noise ratio. Second-hand to BER, the maximum achievable rate will determine the speed of the communication. For system-on-chips, area and supply voltage will be critical figures but, more importantly when talking about low-power systems, so is the power consumption. It is often translated into an energetic efficiency defined as the required energy to transmit one bit and can be expressed as the ratio of dissipated power to the rate. State-of-the-art systems can reach bit energies as low as 0.1 pJ/bit. It is important to note that, even though the metric is normalized to the data rate, it is not independent of the data rate and, as will be shown in another chapter, usually decreases with the rate. Lastly, we should mention the use of the data eye diagram to evaluate the received signal integrity. Eye diagrams are generated by overlapping multiple periods of the signals in order to extract some other performance metrics. Figure 1.13 on the facing page shows those metrics. Eye width will give information about the timing margin, eye crossing percentage is directly linked to the data duty-cycle while the height gives good insight about the output SNR [Anritsu, 2010].

CHANNEL BANDWIDTH

The communication channel will most likely have a finite bandwidth, outside this bandwidth, the signal will suffer from attenuation. In 1928, Harry Nyquist stated that symbols sent at a rate R could be unambiguously decoded as long as the system bandwidth was at least $2R$ [Nyquist, 1928] [Black, 1953]. Of course, if the system's signaling rate does not meet this condition, the input signal can still be recovered. Indeed, as stated before, the deterministic transformation




Figure 1.13: Data eye diagram [OnSemiconductors, 2014]

made by the channel can be inverted (i.e equalized) and the signal recovered. The maximum rate allowable will thus rather be defined by the equalization strategy's performances rather than the channel cut-off frequency.

PROCESS INTRINSIC SPEED

Two metric are used to evaluate the process intrinsic speed: the Fan-Out of 4 time discussed earlier and the process characteristic time constant which is defined as

$$\tau_n = R_n C_{ox} \quad (1.11)$$

with R_n being the equivalent switching resistance of a NMOS device and C_{ox} its oxide capacitance [Baker, 2010]. Those two metrics are directly linked to the maximum switching frequency of a CMOS process and will thus limit the maximum working frequency of the transceiver. High-speed systems are more often limited by the process speed rather than by the channel's time constant. Again, as equalization is used to compensate for the channel response, very high-rates can be achieved. On the contrary, distorted on-chip signals cannot be as easily corrected. That is the *raison d'être* of serializers and deserializers (also known as SerDes), on-chip signals are kept at frequencies lower than the process maximum frequency by using parallelization but are transmitted serially on the channel at a much higher frequency [Yang, 1998]. More recent architecture such as [Song et al., 2013] overcome the limited process speed by implementing some of the critical on-chip elements such as the SerDes with current-mode logic (CML) as it presents much lower propagation delays than standard CMOS logic [Allam and Elmasry, 2001].

JITTER

Jitter, also known as timing noise, limits the actual working frequency of a digital system by adding uncertainty on the position of clock edges. The two main sources of jitter are the circuit noise (particularly in oscillator circuits) and the channel response. Random jitter will be caused by crosstalk between signals, shot noise, thermal noise and flicker noise [Hajimiri et al., 1999]. Deterministic jitter on the other hand, will be caused by ISI and low-pass channels [Analui et al.,

2005] [Buckwalter and Hajimiri, 2006]. Jitter and the systematic timing offset will determine the system timing margin. The timing margin can be expressed as [Palermo, 2010a]

$$t_{margin} = T_b - t_{so} - t_{jd} - t_{jc} \quad (1.12)$$

where T_b is the bit duration, t_{so} the systematic sampling offset and t_{jd} and t_{jc} are the uncertainties on the clock and data edges.

NOISE

Amplitude noise will directly affect the BER as it depends on the output SNR. As said above, for low-power systems, it can be profitable to drive data on the channel with a low-swing or allow the channel to attenuate the signal, thus leading in smaller amplitude signals and lower SNRs. Noise can have multiple sources. Random noise is intrinsic to the devices used to implement the drivers and input buffers but can also come from other circuits on the chip/PCB or even from electromagnetic radiations (EMI for electromagnetic interferences). Cross-talk noise is caused by capacitive coupling between the lines, introducing correlated noise between the channels. Switching noise (or power supply noise) on the other hand is, as cross-talk noise, deterministic and is caused by multiple driver sharing a common power supply [Song et al., 2009] [Baker, 2010]. With the receiver offset, noise will determine the amplitude margin of the system and thus the minimum allowable swing for near error-free communication.

1.2.4. STATE-OF-THE-ART OF LOW-POWER HIGH-SPEED TRANSCEIVERS

State-of-the-art solutions often presents most of the elements presented above. Those solutions are well-adapted for multi-Gbps communications over relatively long links (5-30cm) but are not as efficient when used for our target rates and link lengths (as will be shown in the next chapter). In this section we present four state-of-the-art transceivers.

[**Song et al., 2013**] is a typical example of high performance and high efficiency transceiver. At 0.8V, it achieves a transmission rate of 8Gbps while only dissipating 0.47pJ/b. The transmitter is based on a 8:1 CML multiplexer driving an all-NMOS reduced-swing voltage-mode driver, thus limiting the on-chip data rate to 1Gbps. Reduced swing is achieved by using a dedicated voltage regulator to supply the driver. In order to limit the output driver's size, a level shifter, implemented as a feed-forward capacitor is used to increase the gate voltage and thus the drive strength. Data are equalized by a CTLE driving 8 sense amplifiers (implemented as strong-arm latches). This design stands out by using an ILRO. The ILRO generates eight clock phases used to control the sampler at a minimal power cost.

[**Wong et al., 2004**] is a bit older and does not achieve the performances of more recent designs⁸ but present an interesting feed-forward equalizer paired with low-common mode drivers. The drivers have some noteworthy features. First, impedance matching circuitry has been implemented and can adjust the output by controlling the gate voltage of the drivers. Secondly, in order to limit the switching noise, the output slew-rate is controlled by limiting the pre-driver

⁸Though, It has been shown that this design's performances would benefit a lot from technology scaling

	[Song et al., 2013]	[Wong et al., 2004]	[Pamuto et al., 2015]	[Kim et al., 2009]	This Work
Technology	65nm	0.18 μ m	0.18 μ m	65nm	65nm
Supply Voltage	0.6-0.8V	1.8V	1.2V	1V	0.8V
Data Rate	6.4Gbps	3.6Gbps	50 Mbps	8.9Gbps	150Mbps
Clocking	Forwarded clock	CDR	Forwarded clock	Ideal common clock	Forwarded clock
Channel	8.9cm FR-4	8cm FR-4	2pF standard wire model ⁷	4cm silicon carrier	5cm FR-4
Area	0.057mm ²	0.158mm ²	0.0372mm ²	0.0228mm ²	-
Energy efficiency	0.47pJ/b	7.5pJ/b	50pJ/b	1.9pJ/b	0.45pJ/b
Receiver					
Equalization	CTLE	None	None	DFE-IIR	1-tap DTE ^a
Energy efficiency	0.17pJ/b	3.25pJ/b	22.5pJ/b	0.68pJ/b	0.12pJ/b
Transmitter					
Driver	Voltage-mode	Voltage-mode	Not reported	Voltage-mode	Voltage-mode
Output swing	100-200V _{pp}	250V _{pp}	1.2V _{pp}	0.5V _{pp}	100V _{pp}
Equalization	None	2-Tap feed-forward	None	None	None
Energy efficiency	0.3pJ/b	4.25 pJ/b	27.5 pJ/b	1.22pJ/b	0.33pJ/b

Table 1.2: Summary of the performances of state-of-the-art systems

^aLinear discrete-time equalizer^{a7} Equivalent to 1.8 cm FR-4

(controlling the impedance) slew-rate. The actual driver consists in 4 binary weighed drivers controlled by the sent bits $x[n]$ to $x[n-3]$, thus implementing digitally controllable feed-forward equalization right into the driver. On the receive side, no equalization is performed, the received signals are directly fed to the comparators. The comparator is implemented by a pre-amp featuring offset-compensation circuitry and a variant of a regenerative-latch. Contrarily to the previous design, this system is single-ended in order to save consumption.

[Pannuto et al., 2015] was designed as a bus transceiver and serial protocol adapted to die-stacked chip. As this work, it is aimed at being integrated to low-power vision sensor nodes. Contrarily to the other transceivers presented in this state-of-the-art, the Mbus is fully synthesizable with standard cells. It uses only 2 wires per chip (the clock and a bi-directional data-link) and features a sleep-mode. As the design is aimed at low rates and short links, simple oversized NAND gates were used to drive the datas, resulting in a higher consumption. The main benefit of this design is its reliability as it features a complete protocol relying on the acknowledgment of correct reception, ensuring error-free transmission. Being only implemented with standard cells, it present a relatively small area compared to other designs.

[Kim et al., 2009] was designed for silicon carriers channels but it can easily be transposed to FR-4 substrates. In this work, impedance matching was done statically. It is particular in the sense that a perfect matching was not targeted but a compromise between static consumption and reflections amplitude. Similarly to [Song et al., 2013], they used reduced swing driver and CML logic for the critical elements. Their particularity comes from the receiver which features a decision-feedback equalizer where sampler and summers are integrated in a single circuit resulting in compact design. The results are regenerated by a double regenerative latch achieving greater speed and sensitivity than a standard latch.

The characteristics and performances of those systems are presented, along with the results obtained in this work, in table 1.2 on the previous page.

1.2.5. A WORD ABOUT CLASSICAL SERIAL BUSES

In the previous section, advanced solutions have been introduced but, when speed or power are not critical aspects for the target application, more classical implementations can be considered. The most commonly encountered are the Serial-Peripheral-Interface (SPI) by Motorola, Inter-Integrated Circuit (I²C) by Phillips Semiconductor or even 1-Wire by Dallas Semiconductor. When interfacing with a computer Universal-Serial-Bus (USB) is the most popular solution and, formerly, RS-232.

SPI is one of the most widely used communication protocol in embedded systems as it can be implemented using synthesizable cells only. The communication takes place on 4 wires (Select, Clock, Data-in, Data-out) at relatively low-rates (A few tens of Mbps) and short distances. We should also mention that SPI does not have any overhead bytes and that every bit transferred is a useful data bit.

I²C Contrarily to the common SPI implementation using push/pull drivers, I²C relies on an open drain design with pull-up resistors. Compared to SPI, significant power is dissipated in the pull-up resistors and additional circuitry required. Furthermore, speed is typically limited to less than 1Mbps (but some hosts support rates up to 3.4Mbps).

1-Wire Maxim-Dallas is a peculiar interface. It uses only 2 wires for data and ground and is limited to a few kbps. The particularity of this interface is that the slave devices do not need any external power, power from the data line is stored on a capacitor used to provide energy to the slave device.

RS-232 is the ancestor of USB, it is bulky (between 3 and 25 wires), slow and power hungry. Indeed, its large voltage swing (between 6 and 30V) is very costly in terms of energy. Besides, as opposite to the protocols above, it can be used to transmit data over several meters of cable.

USB The last iteration of the USB standard (USB 3.1) can reach rates up-to 10Gbps. As the 1-Wire interface, USB transmits power to the slave device by featuring a supplementary +5V wire and is thus not adapted at all to low-power applications. Furthermore, large amount of energy are required to preserve signal integrity at those speeds without equalization or other signal recovery techniques.

SECTION 1.3.

PROBLEM STATEMENT

The objective of this work is to design a transceiver compliant with the rates and power requirements of state-of-the-art low-power vision sensor nodes. As it displays cutting-edge performances, we chose the SunPixer developed at UCL, whose effectiveness has been demonstrated in [Bol et al., 2014], as our case study. In its 32 fps mode, it consumes 17pJ per 8-bit pixel. Considering a possible iteration of the sensor which could feature VGA resolution, the expected throughput would amount to 79 Mbps. In order not to shadow the outstanding power efficiency of the image sensor, each pixel should be transmitted off-chip at a cost much lower than 17pJ and thus energetic cost of each bit should be significantly below 2.1pJ. Furthermore, to avoid significant area overhead, we constraint the number of wires to 2.

BASELINE SOLUTION

In this chapter, we analyze the limitations of a naive serial transceiver. We first introduce some fundamental concepts through the case study of a single-inverter transmitter before analyzing a slightly more elaborated solution relying on tapered buffered chains.

SECTION 2.1.

DIGITAL POWER

In this section, we first study the power consumption and behavior of a single inverter driving a large capacitive node. We then extend our results to the case of a tapered inverter chain.

2.1.1. SINGLE INVERTER

Consider a simple inverter driving a capacitive load C . The dissipated power can be expressed as

$$\bar{P} = \underbrace{\alpha \frac{CV_{DD}^2}{2T}}_{P_{switch}} + \underbrace{\alpha \frac{\Delta t_{sc} \bar{I}_{sc} V_{DD}}{T}}_{P_{sc}} + \underbrace{\bar{I}_{leak} V_{DD}}_{P_{static}} \quad (2.1)$$

where V_{DD} is the supply voltage, Δt_{sc} and \bar{I}_{sc} are the switching time and mean switching current, \bar{I}_{leak} is the leakage current and α the activity factor (linked to the node switching probability at each clock cycle). The dynamic power P_{dyn} is related to the charge transferred to/from the load at each cycle. The short circuit power P_{sc} corresponds to the dissipated power when both the P and N-network are on, thus creating a low-impedance path between ground and V_{DD} . Lastly, the leakage or static power P_{static} emanates from sub-threshold currents allowing charges to flow in the channel even when V_{GS} is well under the threshold voltage. Another source of leakage must be considered when using technologies with very thin oxides: currents may appear due to tunneling effects in the gate oxide allowing some charges to flow through. Reducing the supply voltage has a significant effect on power consumption as switching power decreases quadratically with V_{DD} , however, when aggressively reducing voltage, static power may become dominant as \bar{I}_{leak} does not scale with the supply voltage. Frequency is also an interesting knob to control power, but it only effects dynamic power. To reduce leakage power, devices with thicker oxide

2.1. DIGITAL POWER

/ higher V_{th} can be used. Lets express eq. (2.1) as an "energy per bit" by multiplying it by the bit rate $T = 1/R$:

$$\overline{E}_b = \alpha \frac{CV_{DD}^2}{2} + \alpha \Delta t_{sc} \bar{I}_{sc} V_{DD} + \frac{I_{leak} V_{DD}}{R} \quad (2.2)$$

The efficiency is composed of a constant term, independent of frequency, and a leakage term, decreasing with frequency. Working at a higher rate is thus beneficial as it mitigates the impact of leakage current on the total power. We can obtain more insight about this equation by looking at figure 2.1.




Figure 2.1: Energy per bit of a single $\times 2$ inverter charging a 40fF load

Three regions can be distinguished: the energy first decreases before reaching a plateau and then decreases again. In the first region, switching occurs so rarely that the dynamic power is negligible in regard to the static power, we thus only observe the decrease of the leakage term. At a sufficient rate, the static term becomes negligible compared to the constant dynamic term hence giving a constant energy plateau. The third part implicates another phenomenon. In order for the output to span a full swing, the output rise and fall times t_{LH} and t_{HL} must be such that $t_{LH} + t_{HL} < T$ (for a data signal, if a clock signal is considered, the condition is $\max(t_{LH}, t_{HL}) < T$). If that condition is not met, the signal will suffer from attenuation. We plotted on fig. 2.2 on the next page the output pulse-amplitude-closure (PAC) defined as [Lee et al., 2000]

$$PAC = 1 - \frac{V_{out}^{peak-peak}}{V_{DD}} \quad (2.3)$$

We see that the pulse closure begins to increase at the frequencies corresponding to the end of our constant energy plateau. We can thus attribute the last energy decrease to the reduction of

the output swing, saving switching power. It can thus be beneficial to work in the third region. Note that the third region corresponds to the case when the gate is switching faster than the output node cut-off frequency. The node will thus have memory and previous bits will impact the node voltage (i.e. ISI).




Figure 2.2: Output PAC of a single $\times 2$ inverter charging a 40fF load

The driving inverter we considered can be modeled by an ideal NRZ source in series with an effective switching resistance R_n or R_p , depending whether the P or N network is on. Those effective switching resistances directly depend on the transistor's size, supply voltage and threshold voltage and can be expressed as

$$R_n = \frac{2V_{DD}}{\mu_n C_{ox} (V_{DD} - V_{THN})^2} = R'_n \cdot \frac{L}{W} \quad (2.4)$$

with μ_n the mobility of electrons, C_{ox} the oxide capacitance, V_{THN} the NMOS threshold voltage and L, W the physical dimensions of the device [Baker, 2010]. The typical equivalent resistance R'_n (resp. R'_p) of a LVTLP NMOS (resp. PMOS) at 0.8V device is 10k Ω (resp. 30k Ω). We can thus estimate that a LPLVT $\times 2$ inverter has an equivalent resistance $R_n = R_p \approx 3k\Omega$. Considering a 40fF load as in fig. 2.2, we can estimate that the inverter's 3-dB cut-off frequency is of about 1.3 GHz.

As mentioned above, the device threshold voltage has direct consequences on consumption and speed. Consider figure 2.3 on the next page, as expected, the cut-off frequency of GP devices is higher than that of LP devices. Besides, they present a much higher leakage consumption resulting in a low-efficiency when used in the first region (dominant leakage power). Finally, there is no difference between the libraries in the plateau-region (dominant dynamic power) as the threshold voltage does not impact dynamic power.




Figure 2.3: Energy per bit of a single $\times 2$ inverter charging a 40fF load

From this discussion, we extracted 4 levers on which we can play in order to control the power consumption:

Supply voltage Reducing the supply voltage will diminish the global consumption at the cost of a lower cut-off frequency.

Threshold voltage Through the choice of an adequate library, leakage consumption can be reduced also at the cost of a lower cut-off frequency.

Sizing The channel cut-off frequency can be raised by choosing larger output devices. This is done at the expense of a larger gate capacitance thus increasing the power consumption of the previous stage.

Frequency Working at higher rates is more costly in terms of power but advantageous in terms of energetical efficiency. Yet, there is an upper bound to the maximum achievable frequency due to the introduction of ISI.

2.1.2. TAPPED BUFFER

Consider the inverter chain depicted on fig. 2.4 on the facing page. Each inverter in the chain is larger by some factor F from the previous one. The idea behind a fixed tapering factor inverter chain is for each intermediate node to have the same time constant $\tau_n = R_n^{eq} C_{n+1} = R_N C_L$. [Jaeger and Linholm, 1975] proved that the total delay in the chain is minimized by choosing e , the natural base as tapering factor. We used e as tapering factor for the following reasoning

even though in our case, delay does not matter (as long as $t_{LH} + t_{HL} < T$).




Figure 2.4: Tapered chain of inverters [Frustaci et al., 2011]

Keeping e as our tapering factor, a chain designed to drive a capacitance C_{load} would be comprised of $\ln \frac{C_{load}}{C_{in}}$ stages and would have a total switching power of

$$P_{switch} = \alpha C_{in} \frac{V_{DD}^2}{2T_{ck}} \sum_{n=0}^{\ln \frac{C_{load}}{C_{in}}} e^n = \alpha \frac{V_{DD}^2}{2T_{ck}} \frac{C_{load} - C_{in}}{e - 1} \quad (2.5)$$

with C_{in} the input capacitance of the chain, T_{ck} the data period, V_{DD} the supply voltage and α the activity factor of the buffer. As in the case of a single buffer, power grows linearly with the load capacitance. The calculation for area yields similar results:

$$A_{tot} = A_0 \frac{1 - C_{load}/C_{in}}{1 - e} \quad (2.6)$$

A_0 being the area of the first stage. We represented the power consumption as a function of C_{load}/C_{in} on figure 2.5 on the next page. An input capacitance of 1fF was used (corresponding to the input capacitance reported in *ST Microelectronics* standard cell library [ST Microelectronics, 2008b] for a $\times 2$ inverter). The right-most point on the graph corresponds to the approximate capacitance of a 5cm microstrip (≈ 5.5 pF). According to equation (2.6), for a 5cm microstrip, the required area would be $3200\times$ larger than the the area of a $\times 2$ inverter, giving a total area of 0.005mm^2 . Tapered buffers are required for designs where ISI cannot be tolerated, they can theoretically be extended such that they could drive any capacitive load but that is at the cost of high static and dynamic power as well as area. We simulated the same chain using SPICE in order to get an idea of the total power consumption. As we can see on figure 2.6 on the following page, eventhough our working frequency is in the optimal plateau, the consumption of the chain is well above our specifications and other techniques to drive the line must be found. We should mention that variable tapering-factors techniques allowing different size ratios between stages have been developed to improve the power-delay product between stages [Choi and Lee, 1994]. More recently, [Frustaci et al., 2011] added a degree of freedom to the problem by allowing the threshold voltage to change between stages.




Figure 2.5: Tappered chain switching consumption for an alternating input of 1’s and 0’s with $C_{in} = 1\text{fF}$, $V_{DD} = 0.8\text{ V}$ and $f = 100\text{MHz}$




Figure 2.6: Energy per bit of a minimum-delay tapered chain charging a 5.5pF load

SECTION 2.2.

BASELINE SOLUTION

In this section, we first present the test setup used to simulate the design with *Eldo* (from *Mentor Graphics*). The architecture of the solution is presented, followed by an analysis of its performances.

2.2.1. TEST SETUP

The system was fully simulated on the testbench depicted on figure 2.7 on the next page. The testbench is composed of external loads and drivers used to simulate the designs with realistic boundary conditions. External loads/drivers, transmitter and receiver all received dedicated

power supplies. ESD protections were added to prevent antenna effects. Indeed, during manufacture, when etching the metal layers of the circuit, unwanted charges can build up on the metal layer. If a gate electrode is connected to the net, the device could break by reaching the insulating oxide breakdown point [Gabriel and Mcvittie, 1992]. ESD protections consists of diodes providing a non-destructive way to evacuate the built-up charge. By inserting them into our testbench, we make sure to take their parasitic capacitance into account eventhough it is negligible in front of the line capacitance (less than 0.02pF according to [ST Microelectronics, 2007]). The channel was implemented through *Eldo*'s lossy transmission line model described by its RLC distributed parameters [Mentor Graphics Corporation, 2005]. The following parameters were used

dC	11nF/m
dL	280nH/m
dR	160m Ω

Table 2.1: Channel RLC parameters

The line length is left as an adjustable parameter. The above values were computed with *Mantaro*'s online microstrip impedance calculator [Mantaro Product Development Services, 2015].




Figure 2.7: Testbench

Furthermore, this testbench was integrated in larger C/Python framework able to automatically generate testbenches for a given voltage, process and frequency. The framework is also able to detect transmission errors and can thus optimize the size of the output buffers in order to optimize BER and power.

2.2.2. INITIAL ARCHITECTURE

Figure 2.8 on the following page depicts the the architecture of the baseline solution. The transmitter consists of a Flip-Flop clocked on CK_{in}^{TX} and two output buffers implemented as a tapered chain of inverters with progressive sizing as discussed earlier. The receiver is a mirror copy of the transmitter and consists in two input buffers implemented as $\times 2$ inverters and a Flip-Flop clocked on CK_{in}^{RX} . Data are thus launched on the rising edge and captured on the

2.2. BASELINE SOLUTION

falling edge, which corresponds to the center of the data bit if disregarding timing errors. It was entirely implemented using LVTLP devices.




Figure 2.8: Transceiver architecture

The simulated BER performance of the solution considering perfect clock transmission is depicted on figure 2.9. When approaching the cut-off frequency, BER quickly degrades due to the apparition of ISI. In this case ISI starts when $T/\tau > 2.2$, which is consistent with $\max(t_{LH}, t_{HL}) < T$ as we can show that for symmetrical inverters $t_{LH} = t_{HL} \approx 2.2RC$ [Baker, 2010]. We are thus indeed constrained to design our driver/choose our rate to meet that condition. We also see that performances are not really dependent of the receiver process corners, even though PVT variations are what make inverters poor receive buffers.




Figure 2.9: MaTLAB simulation of the expected BER as a function of $f_{cut-off}/f_{signal}$

The power consumption of the solution is presented on figure 2.10 on the facing page. The inverter were sized such that if they had been chosen any smaller, errors would have been introduced. To achieve that optimal sizing, a bisection method was used. That means that we worked at the limit of the plateau region and that is the reason why the power consumption does not really vary with frequency. The bulk of the consumption is taken up by the drivers and, as expected, particularly the clock driver as it switches at least twice more often than the

data. Finally, we can acknowledge that the clock buffer is a bottleneck in terms of consumption and rate. Again, as the clock frequency is twice that of the data, clock transmission will fail first. Note that we don't report the power consumption of the system working at 0.6V for higher rates because the optimization algorithm could not find a suitable solution within the given size bounds (between $\times 2$ and $\times 500$). We also note that the optimal buffer size is much smaller than the expected buffer size. Indeed, equation eq. (2.6) anticipate an output size of 11000, while our output size is effectively 116. That is because we chose ϵ as our tapering factor, which is not optimal if optimizing power consumption. As mentioned earlier, the time constant seen by each node of the chain is (considering equal PMOS and NMOS strengths) $\tau = FR_1C_1$. The optimal F is chosen such that $T/\tau = 2.2$, that is $F = \frac{T}{2.2R_1C_1}$. Obviously, F is process and rate dependent.

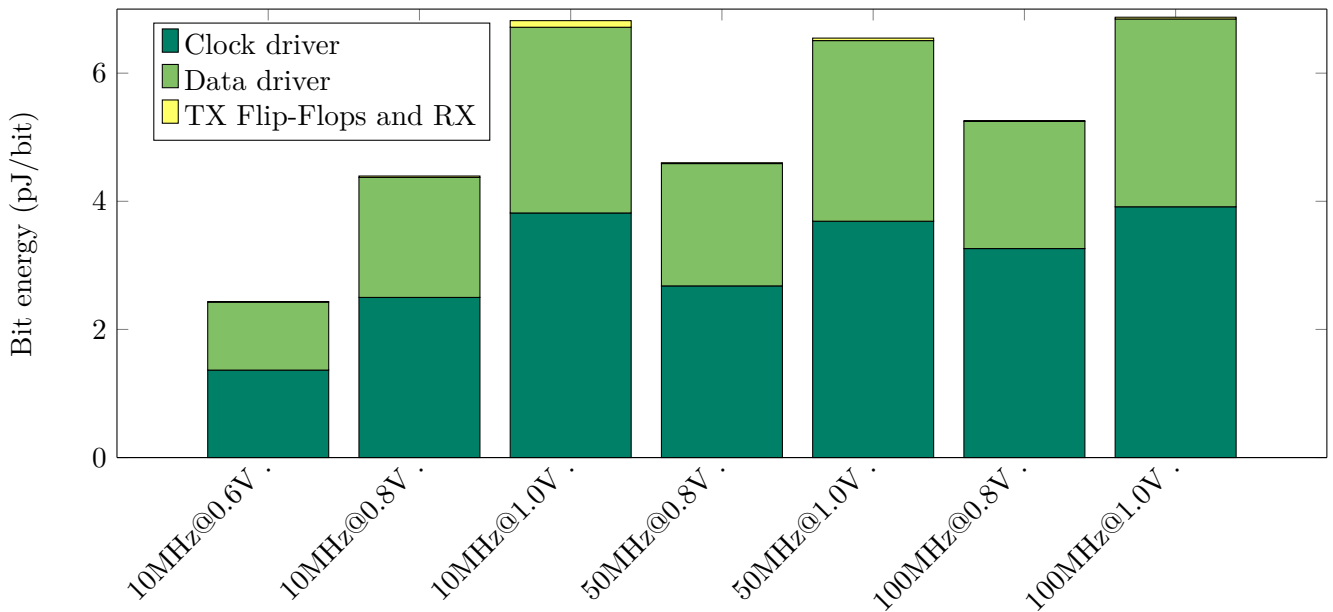


Figure 2.10: Power breakdown of the solution with optimally sized buffers

2.2.3. IMPACT OF JITTER

As there is no source of systematic offset in the system, the timing margin can be expressed as

$$t_{margin} = T_b/2 - t_{jd} - t_{jc} \quad (2.7)$$

For the timing margin to be positive jitter can amount up to 50% of the period which is not really constraining.

2.2.4. ABOUT INVERTERS USED AS INPUT STAGES

Inverters make poor receiver's input stages. In our system, the clock is transmitted along the data, if the driver is properly sized, the clock will suffer from attenuations at high speeds but its DC-level should not change. When using an inverter as input stage, PVT variations of the input offset will cause duty-cycle variations or cause the signal to go undetected (as the inverter

2.3. SUMMARY

switching point will be different from the clock DC level). Let's look at the switching point of an inverter, that is the input voltage for which the currents flowing in the NMOS and PMOS are identical [Baker, 2010]:

$$V_{SP} = \frac{\sqrt{\beta_n/\beta_p} V_{thn} + V_{DD} - V_{thp}}{1 + \sqrt{\beta_n/\beta_p}} \quad (2.8)$$

where $\beta_i = \mu_i C_{ox} W_i/L_i$. The switching point is directly proportional to the supply voltage and will thus greatly suffer from its variations. Furthermore, process corners will impact the switching point through V_{thn} and V_{thp} and finally, mismatch through $\sqrt{\beta_n/\beta_p}$. Temperature, on the other hand, will only affect the speed of the inverter as it will symmetrically impact both devices.

SECTION 2.3.

SUMMARY

In this chapter, we saw that several factors influence the power dissipation in a digital node: frequency, leakage currents, supply voltage and load capacitance. By working at a higher frequency, the static consumption can be distributed over multiple bits thus decreasing the energy required to drive a single bit on the capacitive node. We saw that, if we were willing to pay the cost in power and area, any load could be driven through the use of tapered inverter chains. Finally, reducing the output drive strength can save switching power but BER will quickly degrade due to the apparition of inter-symbol-interferences. An efficient solution would thus tolerate ISI as it would allow lighter drivers and lower output swing.

DESIGN OF THE PROPOSED SOLUTION

The first section of this chapter is dedicated to the description of the chosen architecture. The remaining sections describe the circuit operation and performances of each of the solution's building blocks. The last section is a review of the global performances of the system and a discussion of the gains relative to the baseline solution.

SECTION 3.1.

SOLUTION ARCHITECTURE

As discussed in the previous chapter, our solution must be able to drive datas on the channel at a low-swing and tolerate ISI. Furthermore, the size of the output driver should be limited in order to spare area and switching power.

3.1.1. WORKING WITH ISI

The optimal operation point of the baseline solution was located at the frequencies where ISI started to occur. Energy can be further spared by working beyond the system's cut-off frequency. If we are to allow ISI, proper reception mechanisms should be deployed. The clock deterministic nature can be exploited to easily recover its phase and frequency, while a simple equalization algorithm would enable the recovery of the datas.

RECOVERING THE CLOCK

Unlike the datas, the clock has a stationary mean. By extracting it and using $V_{clk} - \overline{V_{clk}}$ as the decision variable, we can determine whether the transmitted clock was high or low. As the system will be constrained to work beyond its cut-off frequency, the received clock will suffer a phase-shift of $\pi/2$ which must be corrected in order not to hinder the receiver's timing margin. Two approaches can be considered: static and dynamic phase correction. A dynamic phase correction is commonly implemented with a PLL. High-pass filtering could also be examined with the disadvantage that the filter-response should be tunable in order to always correspond to the system's cut-off frequency. Static phase correction presents the benefit of simplicity but will harm the timing margin if the transmission rate grows too close to the cut-off frequency. Using the clock's mean as a reference offers a major benefit over a static reference: process variations might result in non-symmetrical drivers pulling the clock mean high or low, even more, variation

3.1. SOLUTION ARCHITECTURE

on the reference voltage assumed to correspond to $\overline{V_{clk}^{out}}$ might introduce jitter and duty-cycle errors, using the mean as a reference ensures better robustness and resistance to PVT variations.

RECOVERING THE DATAS

Consider the expression of a simple low-pass filter

$$V_{out} - V_{in} = \tau \frac{dV_{out}}{dt} \quad (3.1)$$

When discretized it yields

$$\Delta V_o = V_{out}[n] - V_{out}[n-1] = (V_{in}[n] - V_{out}[n]) \frac{T}{\tau} \quad (3.2)$$

The voltage difference between two periods gives us a valuable insight about the voltage difference between V_{in} and V_{out} . Indeed, noting that for a limited swing

$$0 \leq V_{LO} \leq V_{out} \leq V_{HI} \leq V_{DD} \quad (3.3)$$

Considering NRZ-coding at the transmitter, we can deduce that

$$V_{in}[n] = \begin{cases} 0 & \text{if } \Delta V_o < 0 \\ V_{DD} & \text{if } \Delta V_o > 0 \\ V_{in}[n-1] & \text{if } \Delta V_o = 0 \end{cases} \quad (3.4)$$

By choosing ΔV_o as our decision variable, we can recover the input bits in the presence of ISI, as long as ΔV_o is greater than the comparator sensitivity. Yet, the drawback of this method is that, due to the system's memory, an error might cause further errors. Fortunately, this case only occurs for a sequence of consecutive identical bits and the probability of encountering such a block of size N decreases exponentially with N as $p(b)^N$, $p(b)$ being the source bits distribution. We call this structure "equalizer" but that denomination is not entirely correct. Indeed, our channel is in fact composed of two components: a linear low-pass filter and a non-linear saturator originating from the limited swing of the drivers. If the drivers had an infinite swing, we would not have to implement the third line of eq. (3.4) and our solution would have been a true linear equalizer but, as long as V_{out} does not reach the channel bounds, the structure behaves like a regular discrete time high-pass filter.

A MATLAB behavioral simulation of the proposed data receiver was performed and its results are visible on fig. 3.1 on the next page. Bit-error-rates were computed for different ratio R/f_c , R being the transmission rate and f_c the system's cut-off frequency. A 20mV threshold was assumed for the comparator. For low ratios, the system experience very poor BER, that is because the input signal is sampled on the rising edge and a slight timing error can result in huge sampling error (up to V_{swing}). When the ratios are too high, the voltage swing is severely reduced, falling below the comparator threshold. We can see that jitter does not really impact the performances of the system, indeed we will see in a future section that the system has a very high jitter tolerance (up to $1/R$).

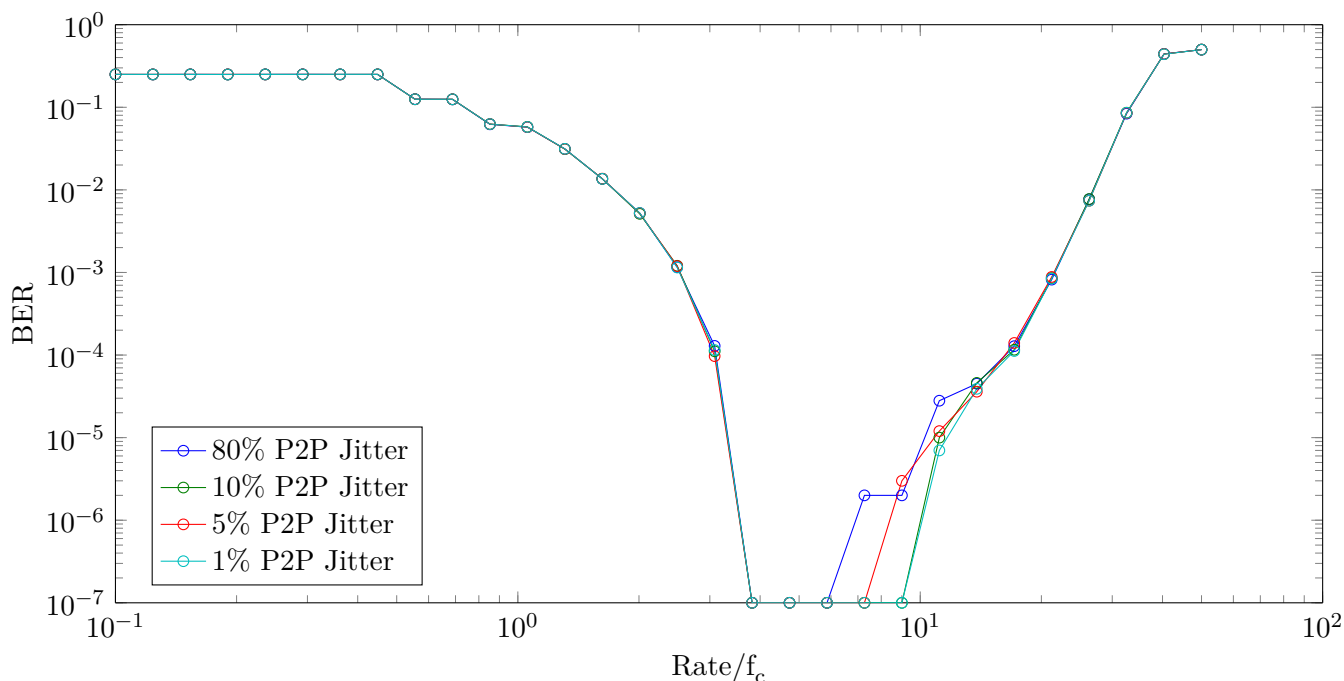


Figure 3.1: BER behavioral simulation for various peak-to-peak jitter amplitudes

3.1.2. GLOBAL ARCHITECTURE OF THE SOLUTION

Considering the previous points, we suggest the system depicted on figure 3.2 on the following page and 3.3 on page 37. In comparison to the baseline solution, the proposed system allows more control over its cut-off frequency through analog control voltages at the transmitter. Also, rather than injecting a significant amount of power into the channel, the swing and slew-rate of the drivers were reduced and an emphasis was made on the receiver ability to restore highly distorted signals.

RECEIVER

Figure 3.2 on the following page illustrates the receiver's architecture. The received signal is sampled at each clock rising edge and stored on a capacitor in a "Tic-Toc fashion" inspired from [Lee and Chandrakasan, 2007]. While a value is sampled on one of the three capacitors, the value stored on the other two capacitors are forwarded to the regenerative latch (RL) used to resolve ΔV_o to valid logic levels. As the RL must be pre-charged prior to any comparison, its output are only valid while CK is down, that is why a classical digital latch was inserted to retain the previously resolved value when CK is down. The sampling is controlled by a simple FSM consisting of resettable shift-register and the clock regeneration circuit consists of a simple RC network used to extract the mean and a comparator to resolve it back to valid logic levels. Note that the value are sampled on the clock rising edge, meaning that if the received wave is a perfectly square wave, timing errors are very likely to occur. That is why we ensured that the channel cut-off frequency is always higher than the signaling rate by using slew-rate limited

output drivers at the transmit side¹.

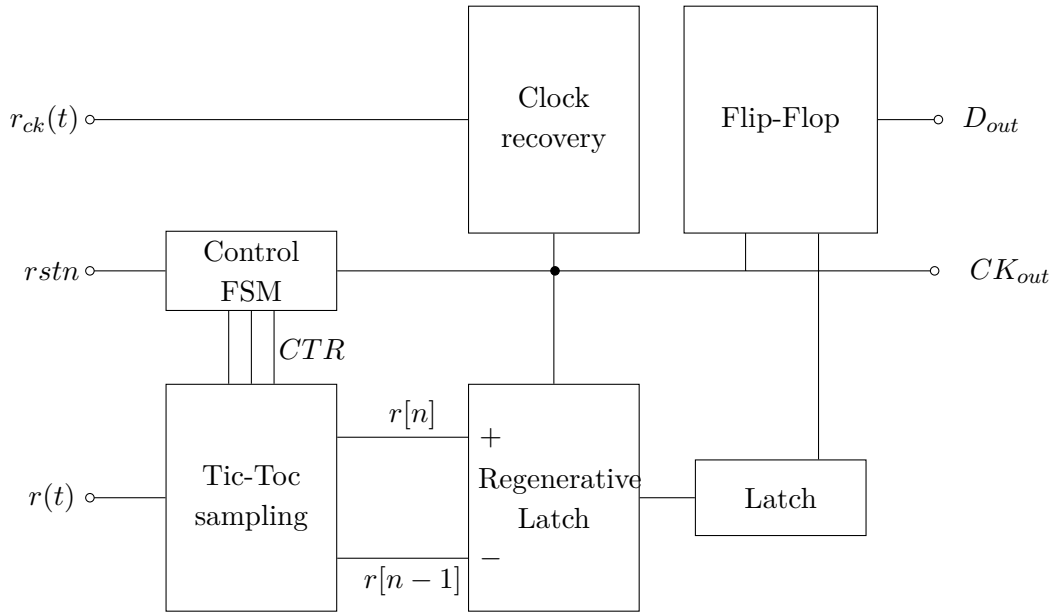


Figure 3.2: Proposed receiver block-diagram

TRANSMITTER

The transmitter consists of a clock and a data driver, both identically implemented. The output stage consists in pull-up/pull-down transistors. The gate of those transistor are controlled by a bootstrap circuit used to boost the output drive while maintaining a reasonable output size and thus avoiding the use of a tapered chain to drive the output stage. Furthermore, the driver are current-starved to ensure that the driver strength is such that the transmitted signal keeps smooth edges and that the transmission rate is higher than the system's cut-off frequency, yielding optimal power consumption and correct system operation. The swing-control logic senses the output amplitude in order to limit the output swing and thus save some output switching power while keeping the signal centered around $V_{DD}/2$. As mentionned above, the clock has to be delayed of a quarter of a cycle due to the nature of the reception. A static delay was chosen by feeding the transmitter a clock two times faster which is then divided and used to delay the resulting clock by the mean of a simple flip-flop.

¹In other words, the slew-rate is set such that the ratio R/f_c results in good BER performances as seen on fig. 3.1 on the previous page

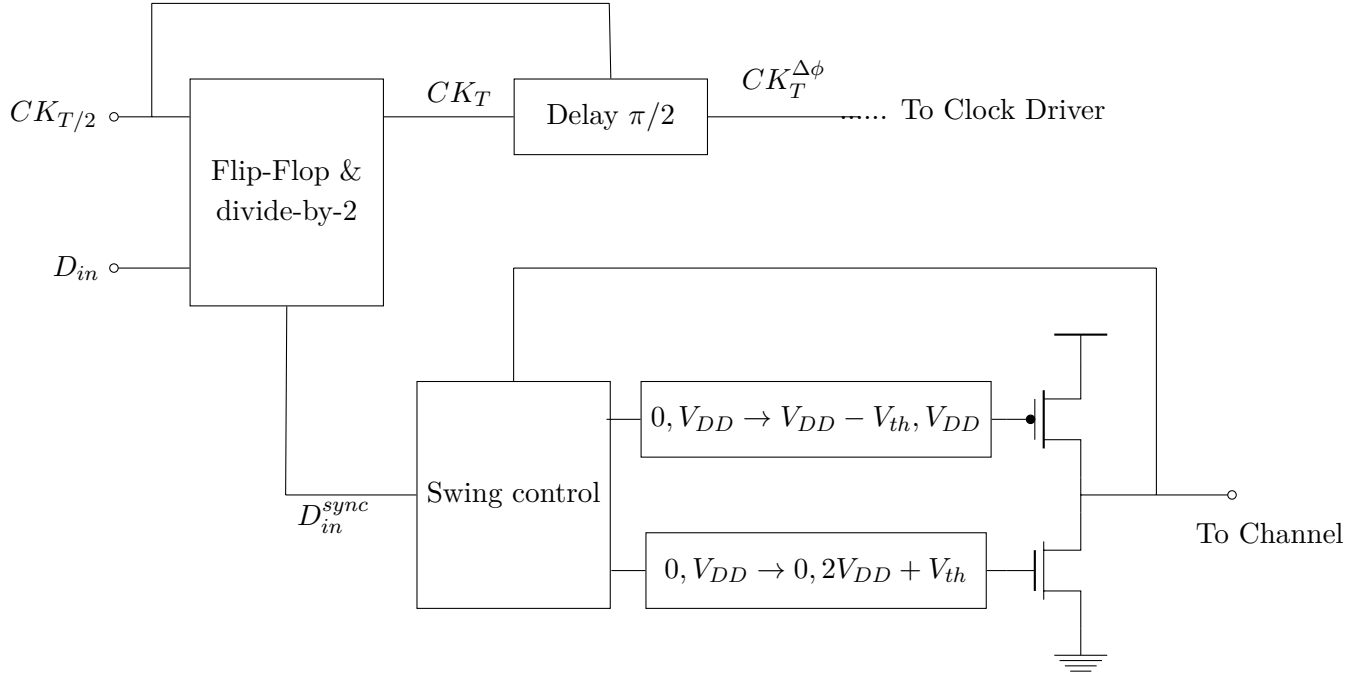


Figure 3.3: Proposed transmitter block-diagram

SECTION 3.2.

TIC-TOC SAMPLING

As stated earlier, the role of the Tic-Toc sampler is to provide the decision circuit with two consecutive received values in order to apply the decision strategy defined by equation (3.4). A block-diagram of the sampler can be seen on fig. 3.4 on the following page. Source followers (marked SF on the diagram) were used to isolate the sampling capacitors from the output loads and the output multiplexers' switching noise.

The control logic of the MUX/DEMUX is summarized in table 3.1. While the values of two capacitors are compared, the input is sampled on the third one. The control FSM was implemented by the mean of a simple 3-bits shift-register rotating between the three states. The initial value of the registers was hardcoded by using different flip-flops with either a set or reset input, allowing the use of a unique reset signal.

State	1	2	3
V_{in}	C_1	C_2	C_3
V_{out+}	C_3	C_1	C_2
V_{out-}	C_2	C_3	C_1

Table 3.1: MUX/DEMUX control logic




Figure 3.4: Tic-toc sequential sampling

3.2.1. SAMPLE & HOLD

Consider the transient operation of a NMOS S&H on figure 3.5 on the facing page. During the sampling phase, V_{out} perfectly tracks V_{in} , when the hold command is issued, an errors appears in the signal due to effects we discuss below. Looking carefully at the output signal when the input signal goes high, we see that the non infinite off-resistance of the switch also causes a small error on the output (amounting to a few mV).

When sizing a sample-and-hold circuit, we must account for its settling time, charge injection and clock feed-through. Consider the simple NMOS switch depicted on fig. 3.6 on the next page. When ϕ goes from high to low, the charge contained in the inverted channel is partly transferred onto the sampling capacitor, thus creating a slight voltage variation. The effect can be mitigated by choosing a larger sampling capacitor or a smaller device. This effect is referred to as "charge injection". Clock feed-through on the other hand, operates in the following fashion: when ϕ goes low, the V_{DD} voltage difference is seen at the output through the voltage divider composed of C_{gs} and C_{load} , resulting in a slight output voltage variation. The total sampling error resulting from those effects can be estimated as

$$\Delta V_{error} = \underbrace{\frac{C'_{ox} W L (V_{DD} - V_{in} - V_{th})}{2C_{load}}}_{\Delta V_{injection}} + \underbrace{\frac{C_{ov} V_{DD}}{C_{ov} + C_{load}}}_{\Delta V_{feedthrough}} \quad (3.5)$$




Figure 3.5: Transient response of the Sample-and-Hold

with $C'_{ox} = \frac{\epsilon_{ox}}{t_{ox}}$. It is thus proportional to the switch size and inversely proportional to the load capacitance. We can express a first order approximation of the settling time as

$$\tau_{S/H} \propto C_{load} R_{on}^{switch} \quad (3.6)$$




Figure 3.6: NMOS switch with parasitic capacitances

The trade-off between speed and precision is pretty straight-forward. Figure 3.7 on the following page depicts that trade-off. Each point corresponds to a circuit with different switch size and capacitance, the inverse relationship between speed and precision can distinctly be observed. One might wonder why PMOS switches perform so much better than NMOS switches. The answer can be found in the model definitions of the devices [ST Microelectronics, 2008a]. There we see that the NMOS devices have a much larger C_{gd} capacitance, resulting in a larger clock-feedthrough. A common technique to reduce the sampling error is the addition of a dummy switch. That is a switch controlled by $\bar{\phi}$, two times smaller and whose source and drain have been shorted. When ϕ goes low, the dummy switch opens and can absorb the parasitic charge. Unfortunately, in the case considered, we see that the addition of a dummy switch only provides

3.2. TIC-TOC SAMPLING

marginal improvement and we will thus not use any in order to spare area. Luckily, in our system, clock feed-through and charge injection will mostly be seen as common-mode noise and can thus be neglected. Of course, as the impact of clock feedthrough varies with the input amplitude, it won't be entirely compensated by the differential architecture. Considering the above discussion we choose the transistor size and capacitance as

$$\arg \min_{(W,C)} V_{error} \text{ with } (W, C) \in [(W', C') \mid \tau_{S/H} < 10 \text{ ns}] \quad (3.7)$$

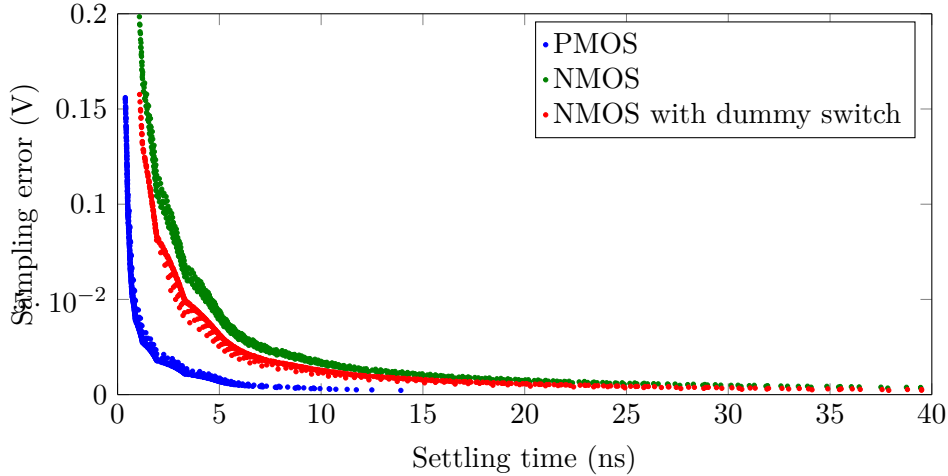


Figure 3.7: Simulated scatter plot of the settling time and sampling error obtained by sweeping over the load and size parameters

The sizing obtained are summarized in table 3.2. PMOS switches should be chosen for their better performances in terms of size, speed and precision but NMOS devices are selected to ensure that the receiver can accommodate for low input values. Finally, we should consider the impact of thermal noise on the sampled value. In the case of PMOS switches we expect a RMS noise amplitude at 300K

$$V_{noise,RMS} = \sqrt{\frac{kT}{C}} = 287\mu\text{V} \quad (3.8)$$

which can be deemed negligible next to the sampling error.

	PMOS	NMOS
W	390nm	1.1 μm
L	60nm	60nm
C	50fF	40fF
V_{error}	3mV	16mV
$t_{0.95}$	9.5ns	9.8ns

Table 3.2: Sizing and performances




Figure 3.8: Source-follower used to buffer the output

3.2.2. SOURCE FOLLOWER

The source-follower was implemented as fig. 3.8. The three relevant performance metrics are in our case: slew rate, static current and level-shift. As the circuit is small, it was sized by sweeping over M_1 , M_2 and V_{bias} and choosing the parameters in the following fashion:

$$\arg \min_{(W_1, W_2, V_{bias})} I \text{ with } (W_1, W_2, V_{bias}) \in [(W'_1, W'_2, V'_{bias}) \mid SR > 10\text{mV/ns and } \Delta V_{shift} < 0.3V] \quad (3.9)$$

We limit the output level-shift in order to ensure sufficient output swing and constraint the minimum slew-rate in order for the output to track the input with enough precision. The resulting parameters and performances are summarized in table (3.3). At 100Mhz, the static consumption will amount to 4fJ/b. We don't mind the large input transistor as its gate capacitance is negligible next to the 50fF capacitor at the output of the S&H.

W_1	$3.96\mu\text{m}$
W_2	$1.08\mu\text{m}$
L	60nm
V_{bias}	0.45V
SR	10mV/ns
ΔV_{shift}	233mV
I_D	$0.5\mu\text{A}$

Table 3.3: Sizing and performances

REGENERATIVE LATCH

Both outputs of the sampler are connected to the inputs of a regenerative latch inspired from [Cho and Gray, 1995]. The regenerative latch acts as a clocked amplifier and is able to take decisions at high speed thanks to its positive feedback network. As one of the requirement of the comparator was hysteresis, the latch has a built-in threshold which can be controlled by the mean of an external voltage. A schematic of the latch is visible on fig. 3.9 on the following page

and circuit operation can be seen on fig. 3.10 on the next page.

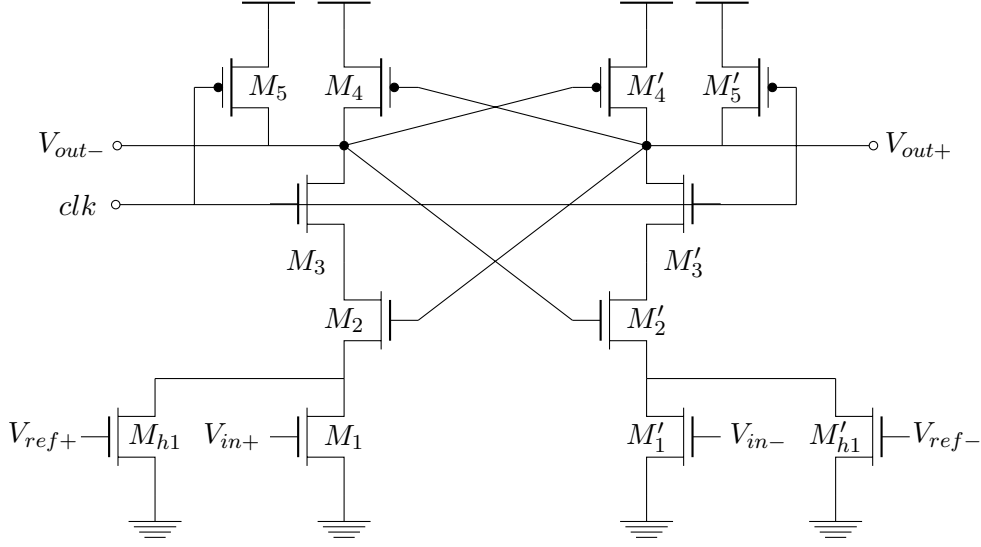


Figure 3.9: Regenerative latch

When clk is low, the output nodes are pre-charged to V_{DD} through M_5 and M'_5 and the sources of M_3 and M'_3 are pulled to ground via M_2 and M'_2 in order to suppress the latch's memory. During this phase, by closing M_3 and M'_3 , no current (save the switching currents) flow in the circuit, thus reducing static consumption [Baker, 2010]. When clk goes high, an imbalance is created between the two branches by M_1 and M'_1 , allowing the output nodes to converge to either V_{DD} or ground through a positive feedback loop implemented by M_2 , M'_2 , M_4 and M'_4 . Note that the input must be at least greater than V_{thn} for the input transistors to turn on. That will always be the case thanks to the level-shifting source-followers used in the previous stage. As the output of the circuit is only valid half a period, a standard digital latch as seen on fig 3.11 on the facing page is connected to the outputs. V_{out+} is connected to the S input of the latch and V_{out-} to the R input. During the pre-charge phase, S and R are both high, thus retaining the previously stored value. Transistors M_{h1} and M'_{h1} are used to add hysteresis to the comparator by providing a tunable threshold controlled by V_{ref+} and V_{ref-} . V_{ref+} is connected to the Q output of the latch and V_{ref-} to \bar{Q} .

When sizing this circuit, our main concern is to ensure that the decision can be resolved within half a clock cycle. The speed will mainly be determined by the capacitance of nodes V_{out+} and V_{out-} . With the help of [Allen, 2002], we can estimate the convergence time of the feedback network as

$$\tau_L \approx \max \left\{ 0.67 \sqrt{\frac{C_{ox} W_2 L_2^3}{2\mu_n I_1}}, 0.67 \sqrt{\frac{C_{ox} W_4 L_4^3}{2\mu_p I_1}} \right\} \quad (3.10)$$

Indeed, the latch speed is determined by the slowest of the two feedback loops. Hence, it is convenient to reduce the size of the feedback transistors as much as possible and increase the size of the input transistors (thus increasing the branch current) while keeping in mind the




Figure 3.10: Operation of the regenerative latch




Figure 3.11: Digital latch

effects of a greater load capacitance on the switches driving the inputs (reduced speed). We should also note that thanks to the previous level shifters pulling the bias point of M_1 higher, we benefit from a greater branch current I_1 . In practice though, we choose the same sizes for all NMOS devices. Indeed, if M_1 is significantly larger than the other devices, the source voltages of M_2 and M'_2 will be very small and the positive feedback could be triggered by noise or simply mismatch between the feedback transistors. Concerning the four lower transistors, they operate in the triode region and act as resistors. Considering that $V_{GS} \gg V_{DSAT}$ we can express the equivalent resistance of devices M_{h1} and M_1 as

$$R_{eq+}^{-1} = \mu_n C_{ox} \left(\frac{W_1}{L_1} (V_{in+} - V_{thn}) + \frac{W_{h1}}{L_{h1}} (V_{ref+} - V_{thn}) \right) \quad (3.11)$$

The output will converge towards V_{DD} for $R_{eq+} < R_{eq-}$, in other words, for symmetrical branches:

$$V_{in+} - V_{in-} > \frac{W_{h1}}{W_1} \underbrace{(V_{ref-} - V_{ref+})}_{=\pm V_{DD}} \quad (3.12)$$


The size ratio W_{h1}/W_1 therefore dictates the switching threshold. Considering the above comments, we choose the parameters of table 3.4 on the next page. Consequently, the expected threshold in nominal conditions is 50mV. The constraints on the threshold are that it must be greater than the noise floor of the design and smaller than the input voltage swing. We should also note that the threshold depends linearly on the supply voltage.

Device	W/L	Scale
M_{h1}	2	60 nm
Other PMOS	102.4	60 nm
Other NMOS	32	60 nm

Table 3.4: Transistors' sizes of the reg. latch

3.3.1. PERFORMANCES

Fig. 3.12 depicts the instantaneous power dissipation into the latch at 100MHz. As explained above, there are no DC currents apart when the latch is changing states, that is thanks to M_3 and M'_3 which ensure that no current can flow in the lower branches while the latch is charging. At 100MHz, the total consumption of the latch amounts to 11 fJ/b and at its maximum frequency of 2GHz, 19 fJ/b. The discussion we had on the baseline solution tended to say that at high frequencies, power efficiency should increase but, in the case of this circuit, the trend is not valid. The reason is that power is essentially dynamic in the circuit, thanks to the mechanisms explained above. Circuit operation at 2GHz can be seen on fig. 3.13 on the next page. The latch time constant is very close to $T/2$, faster operation would culminate in an attenuation of the outputs signals possibly resulting in unproper output levels. We also notice that the signals are much more sensible to switching noise. We can attribute the resulting voltage peaks to the very low impedances of sources used to drive the clock and data, indeed at high frequencies, the effects of capacitive coupling will be more important as our ideal sources can theoretically sink an infinite current. Therefore, as $C \frac{dV}{dt}$ gets larger, so does the parasitic switching currents.

Figure 3.12: Instantaneous power of the latch (*clk* is in red)

As stated earlier, the proper operation of this circuit relies on its symmetry. Indeed, a stronger branch will cause unwanted switching which translates as a wrong switching threshold. Consider the Monte-Carlo simulations depicted on fig. 3.14 on the facing page. The offset was computed as the minimum input difference to switch from a logic 0 to a logic 1. The dots on the x -axis correspond to the respective means of the distributions. The curves resemble regular gaussian bells which would have been stretched toward the right. That "stretching" is caused by our built-in threshold. We see that the means are around 40mV and are not heavily affected by process




Figure 3.13: High-speed operation of the regenerative latch

corners but that the curves have a large standard deviation. In the worst case (Fast-Slow), at least 43% of the circuits have an offset situated between 10 and 100mV. Fortunately, this offset could be easily corrected by using different supply voltages to generate V_{ref+} and V_{ref-} .




Figure 3.14: Offset distribution of the regenerative latch

3.4. DRIVERS

SECTION 3.4.

DRIVERS

The drivers consist of some logic used to limit the output swing which generates the output stage's control signals. Those signal are bootstrapped to provide increased drive strength and limit the output transistors' size. Furthermore, the output stage is current-starved in order to limit the output slew-rate and thus limit crosstalk between the clock and data driver. Additionally, the analog control voltage can be tuned to ensure that the system's cut-off frequency is inferior to the transmission rate. Bootstrapping the inputs and limiting the current can seem disconcerting as both method achieve opposite objectives but the end goal is to increase the range of channels over witch the system is functional. A diagram of the circuit can be seen on figure 3.15. The circuit works in the following fashion:

When $v_{out} < V_{high}$ and $D_{in} = 1$, the output is pulled towards V_{DD} .

When $v_{out} > V_{low}$ and $D_{in} = 0$, the output is pulled towards ground.

Otherwise, the driver is off.

Rather than using two comparators to determine whether the signal is below V_{low} , above V_{high} or between the two, an unique comparator was used and its voltage reference was switched between V_{high} and V_{low} depending on the value of D_{in} . This solution is efficient in terms of power and area but will introduce delay in some situations (the consequences are evaluated in section 3.6.1 on page 51).

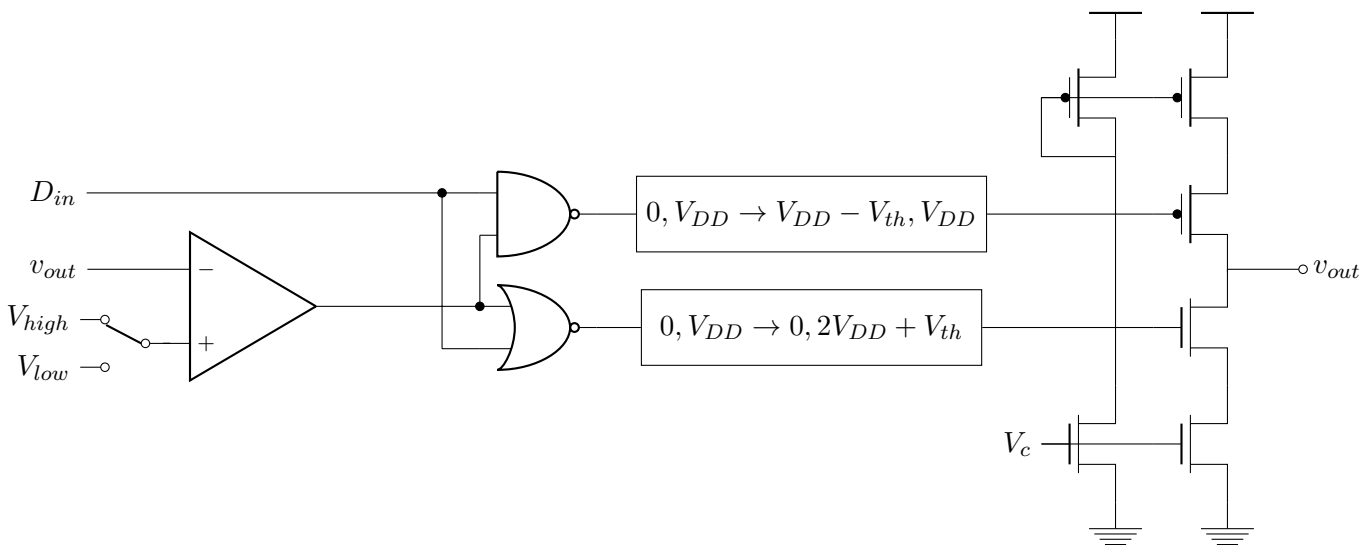


Figure 3.15: Output driver

3.4.1. SLEW-RATE LIMITED DRIVER

As discussed before, the output driver's slew rate is limited by four additional transistors. For proper reception, the output swing should be greater than the comparator built-in threshold.

By limiting the output slew-rate, we ensure that the amplitude error due to timing error won't be too great. For a given slew-rate SR and voltage swing V_{swing} , the amplitude sampling error resulting from timing error will be at most

$$V_{err} = \begin{cases} SR\Delta t & \text{if } \Delta t < V_{swing}/SR \\ V_{swing} & \text{otherwise} \end{cases} \quad (3.13)$$

This results from the fact that the receiver's sampling is clocked on the rising edge (as the data). Hence, in the presence of data with infinite slew-rate (.i.e. a perfect square-wave), any slight timing error would result in a paramount sampling error. It is thus preferable to keep the slew-rate low. Furthermore, limiting the slew-rate reduces switching noise and allow the system to perform its function at an optimal power cost on various channels. In order to avoid significant voltage drops on the starving transistors, they were chosen to be very large. On the other hand, to avoid significant static current in the left branch, small devices were chosen.

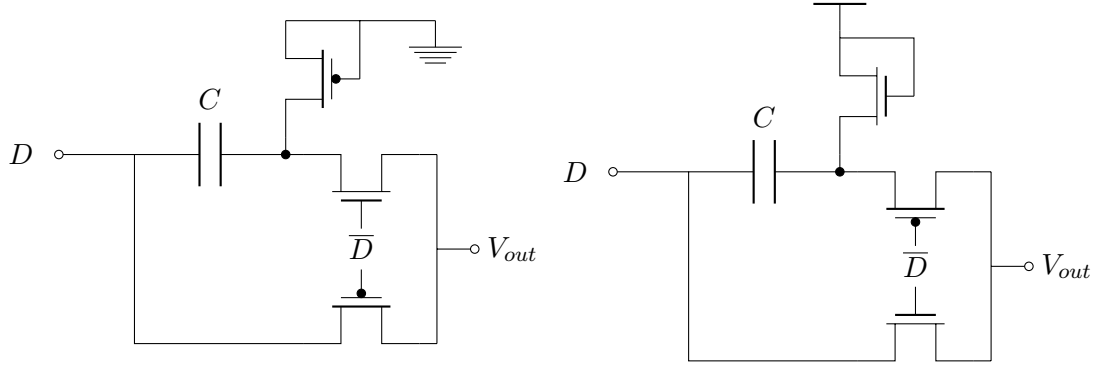
3.4.2. BOOTSTRAP CIRCUIT

Consider figure 3.17 on the next page showing the equivalent resistance of a NMOS device with varying gate voltages and $V_{DD} = 0.8V$. The highlighted points correspond to the nominal gate voltage and the bootstrapped voltages $2V_{DD} - V_{thn,lvt}$, $2V_{DD} - V_{thn,svt}$ and $2V_{DD} - V_{thn,hvt}$. Simple bootstrapping can thus at least double and nearly triple the drive strength for a given size. Figure 3.16 on the following page presents the proposed bootstrap circuits. Let's take a closer look to the positive bootstrap circuit operation. Suppose that D is low, then the node on the right of the capacitance is charged to $V_{DD} - V_{thn}$ through the diode-mounted device. At the same time the 0 is transmitted through the lower NMOS switch and $V_{out} = 0V$. Now, when D goes high, the central node is pulled up to $2V_{DD} - V_{thn}$ and the resulting signal is transmitted to the output through the PMOS switch. The negative bootstrap circuit operates in a very similar fashion: when D is initially high, the central node is dragged to V_{thp} . When D goes low, the node voltage is lowered to $V_{thp} - V_{DD}$. Both output drivers thus have an overdrive voltage of

$$V_{OV} = 2(V_{DD} - V_{th}) \quad (3.14)$$

3.4.3. COMPARATOR

We chose the architecture presented in [Chappell et al., 1988] (fig. 3.18 on page 49). This self-biased architecture is adapted to high-speed and high sensitivity operation and does not require any external voltage/current reference. M_3 , as well as providing self-biasing, increases the output resistance and total conductance of the design. The drawback of this design is that the large capacitance of the node connected to the gates of M_3 , M_1 and M'_1 introduces a low-frequency pole seen by v_{in+} only. Fortunately, in our application, v_{in+} will be connected to a DC reference and the effect of the pole can thus be neglected. When sizing this comparator, we must ensure that it has a sufficiently high cut-off frequency and that the gain is sufficient to ensure that the output will be within the inverter valid inputs. The gain will thus determine the comparator sensitivity. As the circuit is very small, we did an exhaustive search on parameters W_1 , W_2 and W_3 . The resulting parameters and performance can be found in table 3.5. Here



(a) Negative bootstrap circuit

(b) Positive bootstrap circuit

Figure 3.16: Bootstrap circuits implementations




Figure 3.17: Output driver equivalent resistance

we defined the sensitivity as the minimum positive input difference required for the output to reach $99\%V_{DD}$.

$(W/L)_1$	2
$(W/L)_2$	2
$(W/L)_3$	43
scale	60nm
V_{offset}	-5.7mV
Sensitivity	16.6mV
f_{3dB} first stage	130 MHz

Table 3.5: Sizing and performances

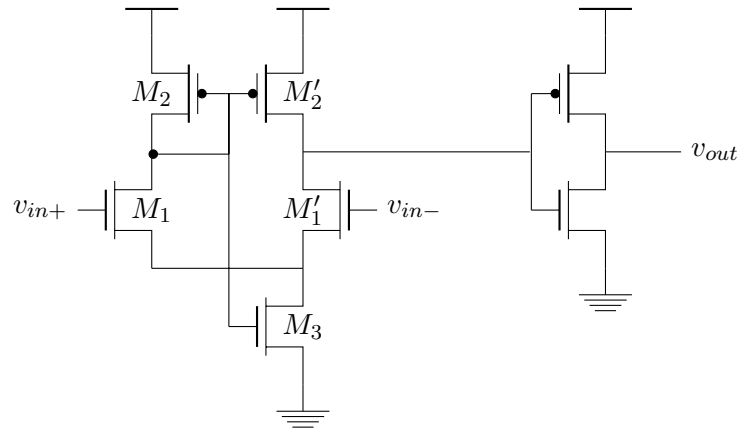


Figure 3.18: High-speed comparator

SECTION 3.5.

CLOCK REGENERATION

As described in the beginning of this chapter, clock regeneration is achieved by comparing the received clock to its mean. The circuit used is represented in figure 3.19. A simple RC network is used to extract the mean and the comparator presented in section 3.4.3 on page 47 was reused here. Its output is buffered by two inverters providing sharp edges and enough drive strength to propagate the clock to the rest of the receiver circuit.




Figure 3.19: Clock regeneration circuit

When choosing R and C , two elements should be considered: the filter's settling time and its gain-bandwidth. The parameters should be chosen such that the filter converge towards the mean fast enough while maintaining a stable oscillation-less output. The remaining oscillations will introduce duty-cycle variations in the system, resulting in jitter. The settling time will also impact the transceiver efficiency. Indeed, let's assume that the transceiver has the same power consumption during steady-state transmission than during the initial transient associated with the filter's settling time. If we wish to transmit a message of length T_m with a system having an initial transient settling time of T_s , the effective cost of a message bit will be

$$E_b^{effective} = \left(\frac{T_s}{T_m} + 1\right)E_b \quad (3.15)$$

Several solutions can be considered to overcome that problem. First, if the system has a suffi-

3.5. CLOCK REGENERATION

ciently high duty-cycle, a switch can be inserted between the resistor and the capacitor, holding the value of the mean when the system is not active. Of course, if the system duty-cycle is too low, the value will leak over time and should thus be refreshed periodically. Secondly, a strong switch could be inserted between the capacitor and a voltage reference corresponding to the expected mean, pulling the capacitor voltage towards that value before the beginning of the transmission. The drawback of this approach is that the switch used should have a significantly lower equivalent resistance than the resistor used in the filter so that the capacitor can be charged fast enough. The last solution would be to use higher-order filters. Figure 3.20 depicts the filter's attenuation at the target signal fundamental frequency against the settling-time of common filters' topologies. Second order Chebyshev and Butterworth filter present much more interesting characteristics than simple first order filters. Furthermore, they could be easily implemented with a Sallen & Key topology. In order to limit duty-cycle variations, we allow a transient equal to $12T$, giving only a 5% error on the mean. At 100MHz, that translates into $RC = 5.45e - 8s$. To meet that constraint, we choose $C = 1pF$ and $R = 54.5k\Omega$.




Figure 3.20: Filter attenuation versus settling-time

As mentioned before, the resulting clock will be phase-shifted of $\pi/2$, in order to correct that phase shift, an earlier phase-shift is applied on the transmit side. That is achieved by feeding the transmitter with a clock 2 times faster than the actual system clock. The system clock is generated by using a divide-by-2 flip-flop clock divider. The resulting clock is fed to another flip-flop triggered on the falling edge of the fast clock. This circuit is represented on figure 3.21 on the next page. As mentioned previously, when applying a static phase-shift, we must ensure

that the transmission rate is above the system's cut-off frequency (and consequently that the clock underwent a phase-shift of $-\pi/2$).

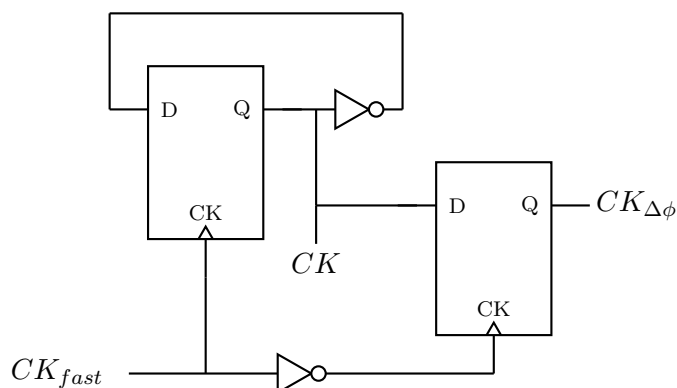


Figure 3.21: TX clock generation and phase-shift

SECTION 3.6.

OVERALL PERFORMANCES

In this section, we present the overall performances of the system by extracting the signals' eye diagrams at different points in the systems. Furthermore, we evaluate the amplitude and timing margins of the system before analyzing its power consumption and the gains over the baseline solution.

3.6.1. DATA EYE DIAGRAMS

Figure 3.22 on the following page depicts the receiver input eye diagram for $R = 50\text{Mbps}$ and $L = 1\text{cm}$ (with maximum slew-rate). In those conditions, the channel cut-off frequency is below the signal frequency but the limited swing driver prevent the apparition of significant ISI. As the eye was simulated with a simple transient simulation and no added noise, the jitter present on the eye consists only in the data-dependant jitter caused by the low-pass response of the channel. Notice how the crossing levels are not located in the center of the voltage swing. That is due to our pull-up and pull-down driving devices having different strength, resulting in duty-cycle errors. Figure 3.23 on the next page depicts the receiver input eye diagram for $R = 150\text{Mbps}$ and $L = 5\text{cm}$ (with maximum slew-rate). In those conditions, we can distinguish three levels: a low level at 0.36V , a high level at 0.44V and an intermediate level at 0.41V . The low and high levels correspond to the voltage bounds of the swing limited drivers. The intermediate level appears when a short pulse of width T_s is transmitted and is too short to saturate the driver. Some additional explanations must be given in order to understand the uncommon shape of the eye. The input data change at time -1 , 0 and 1 . When the input signal switches from 1 to 0 (resp. 0 to 1) while the received signal is at the high level (resp. low level), it starts decreasing (resp. increasing) immediately. On the other hand, when the input signal switches from high to low and the received signal is at the intermediate level, we observe a delay. That is because when at the high or low level, the received value is greater (resp. lower) than the two comparison thresholds of the comparators used for swing control. Thus, when the input signal switches

3.6. OVERALL PERFORMANCES

(and so the comparison threshold), the output of the comparator does not change. While at the intermediate level, the output of the comparator must change when its input changes abruptly, but, as shown earlier, its finite bandwidth introduce a delay. This is illustrated on fig 3.24 on the facing page. This issue could be overcome by using two separate comparators, one for each reference. We can estimate the maximum jitter at 150Mbps by computing the width of the trace at the crossing points between a rising and falling edge. It amounts to 1ns. The width of the horizontal levels is linked to the noise present in the system. In this case, it corresponds only to the noise due to ISI and reflections.

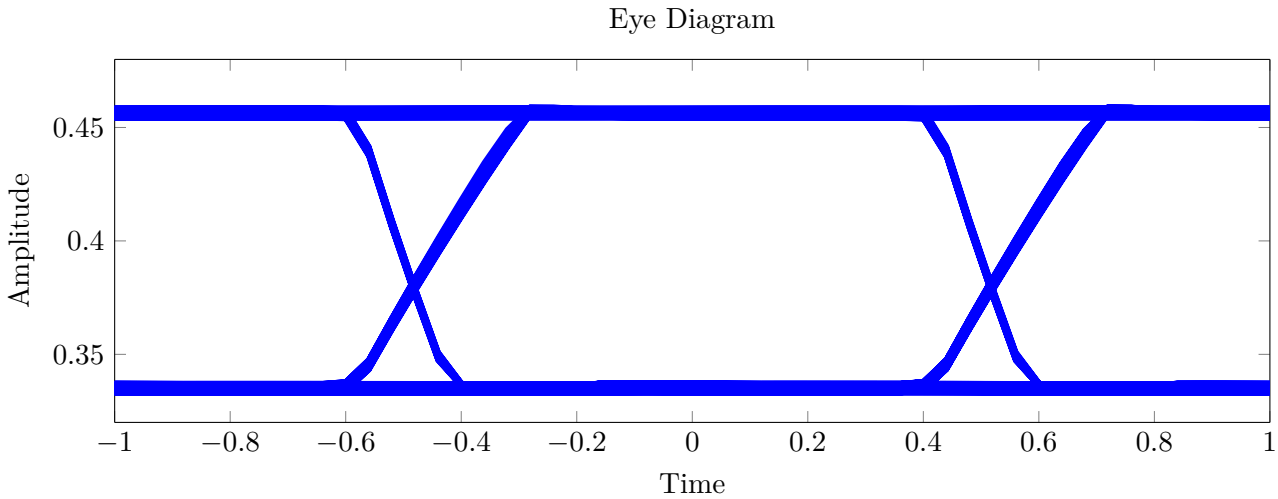


Figure 3.22: Receiver input eye diagram for $L = 1\text{cm}$ and $R = 50\text{Mbps}$ (maximum slew rate), time is normalized to the symbol period




Figure 3.23: Receiver input eye diagram for $L = 5\text{cm}$ and $R = 150\text{Mbps}$ (maximum slew rate), time is normalized to the symbol period

The eye diagram on figure 3.25 on page 54 gives us valuable information about the noise and amplitude margin of the system. For the system to perform its function correctly, the devia-




Figure 3.24: Upper graph: Red – swing reference, Blue – receiver input. Lower graph: Red – Comparator output, Blue – Input data. (A) receiver input above both high and low reference (B) receiver input between high and low references (C) comparator’s output does not change (D) comparator’s output changes

tion of the 0 level should be such that it is always smaller than the regenerative latch built-in threshold. Furthermore, The difference between the 0 level to any other level should be greater than the threshold. This is clearly the case for this first slow eye but consider the eye depicted on figure 3.26 on the following page. The most critical case in term of amplitude margin occurs when going from an intermediate level to a high level, between those levels, the eye opening is of only 19mV, lower than the threshold, resulting in negative amplitude margin. Fortunately, as we can only reach the intermediate level when going from 0 to 1, the threshold is negative (as a 1 was previously sent), resulting in a positive amplitude margin. Still, considering other transitions, the opening is of 28mV, resulting in near 0 margin for a threshold of 30mV. The system was still able to perform error-free, but security margin should be considered when choosing the swing bounds. Finally, the jitter on the clock amounts to 5% of the period.

We can obtain a more formal view of the amplitude and timing margin of the system by looking at fig. 3.27 on page 55. The plot was generated by generating input signals with different slew rates and delay and checking that they were correctly decoded. There we see that the minimum input swing is 38mV, greater than previously estimated and more consistent with the RL threshold. That is because we used the eye opening to compute the amplitude margin, that is, we considered the noise on the different levels to be uncorrelated. In practice, as the noise results from ISI, it will be correlated between successive levels, so, we can thus estimate the amplitude margin by measuring the distance between the lowest trace of a 0-level and the lowest possible trace of any other level. Proceeding this way, we measure 36mV. Concerning the clock delay, we obtain the expected result, the clock can be delayed by approximately $0.5T_s$ in either directions. Indeed consider a triangular wave, in the ideal case, the values will be sampled at the peak and at the base of the wave. In the worst case, $0.5T_s$ delay, the value sampled will

3.6. OVERALL PERFORMANCES

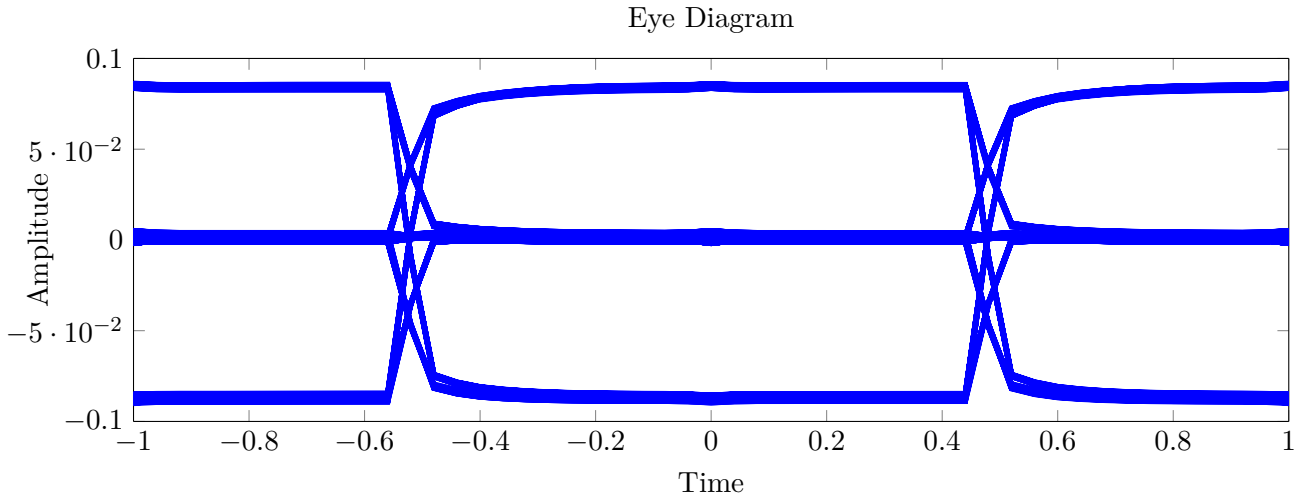


Figure 3.25: Comparator's input eye diagram (ΔV) for $L = 1\text{cm}$ and $R = 50\text{Mbps}$ (maximum slew rate), time is normalized to the symbol period




Figure 3.26: Comparator's input eye diagram (ΔV) for $L = 5\text{cm}$ and $R = 150\text{Mbps}$ (maximum slew rate), time is normalized to the symbol period

correspond to the middle of the rising and falling edge which are equal if the wave is symmetrical.

3.6.2. POWER CONSUMPTION

Figure 3.28 on page 56 depicts the consumption against the rate for $L = 5\text{cm}$. As expected, the consumption decreases with rate. This gain is mostly due to the even distributions of the static currents among multiple bits as the slew-rate was manually limited to ensure that we worked in the third operating region of the driver (see fig. 2.1 on page 24). At 150MHz , we achieve a power consumption as low as 0.45 pJ/b . Figure 3.29 on page 56 depicts the transmitter and receiver power breakdowns. As two times larger buffers were used for the clock and as its




Figure 3.27: Receiver allowable swing and clock delay for $L = 1\text{cm}$ and $R = 50\text{Mbps}$ (maximum slew rate), time is normalized to the symbol period

switching rate is at least two times larger than the data's switching rate, it consumes the bulk of the transmission power. On the receive end, the tic-toc sampler consumes most of the power. That is due to the charge and discharge of the sampling capacitor as well as the static currents running in the level-shifters. Clock regeneration takes a non-negligible part of the power due to the static consumption of the comparator and the large buffers used to forward the clock to the other blocks. Thanks to its architecture preventing large static currents, the latched comparator has a fairly low consumption compared to other elements in the system.

A comparison of the efficiency of the solution and the baseline solution is visible on fig. 3.30 on page 57. First, we note that we achieved a more adequate balance between TX and RX power compared to the initial solution for which the RX power was negligible behind the TX power. Furthermore, we reduced consumption by more than an order of magnitude.

3.6.3. MEASURED BER AND PVT ROBUSTNESS

In order to test the robustness of our system, we extracted the bit-error rates for different voltage, temperature and process corners. We made the voltage vary from 0.7 to 0.8 with no

3.6. OVERALL PERFORMANCES




Figure 3.28: Global consumption of the system. At 50Mbps, the slew-rate control voltage V_c was set to 0.35V, 0.4V at 100MHz and 0.8V at 150MHz




Figure 3.29: Transceiver power breakdown

difference amongst the different corners. As our computation of the BER had only a precision of 10^{-3} , we considered that the system "passed" if its measured BER was equal to 0. Table 3.6 summarizes the results obtained. In nominal conditions ($V_{DD} = 0.8, TT, 25^\circ$), we measured a BER inferior to 10^{-4} .




Figure 3.30: Power consumption of our solution against the consumption of the optimized baseline solution both at 100MHz

Temp./Proc.	TT	SS	FF	SF	FS
0°	✓	✗	✓	✓	✗
25°	✓	✓	✓	✓	✗
100°	✗	✓	✗	✗	✗

Table 3.6: Pass/fail test results for $V_{DD} : 0.7 \rightarrow 0.9$

3.6. OVERALL PERFORMANCES

CONCLUSION AND FUTURE PROSPECTS

The emergence of the IoT and possibly trillions of wireless sensor nodes requires the development of small and energy efficient systems. In those power constrained systems, classical ways of sending and receiving information off-chip such as SPI, I²C or USB have been discarded as they were not designed as power-efficient solutions. On the other hand, extensive research has already been conducted in the development of ultra-low-power high-speed serial transceivers able to exchange informations with external systems at an extremely low cost. In this context, we studied what was the most efficient way to send and recover data at an intermediate rate on a highly capacitive serial link such as the FR-4 microstrips posing as the interconnects of many printed circuit boards.

In the first chapter, we demonstrated the impairments resulting from long FR-4 channels. We saw that dispensing with impedance matching severely constrained the rate on long channels. Furthermore, the low-pass response of the channel results in significant ISI which must either be avoided by reducing the rate or corrected using relevant equalization strategies.

In the second chapter, we showed that a naive solution consisting of tapered inverters is able to successfully transmit information through virtually any channel but is impractical in terms of power consumption and area usage. Furthermore, the use of simple inverters as input stages would make the system ineffective against process, voltage and temperature variations. Besides, we saw that the average energetical cost of a transmitted bit decreased with the rate due to the invariability of static currents with increasing frequency. It is thus beneficial to work beyond the system's cut-off frequency, where inter-symbol-interference occurs. In this regime, dynamic switching consumption is dominant and thus reducing the supply voltage is also very profitable.

Our most significant contributions are located in the third chapter, where we present the final design of a low-power medium rate serial transceiver. With a simulated consumption as low as 0.45pJ/bit at 150Mbps on 5cm FR-4 links, it is well suited for current implementations of low-power vision sensor nodes having a throughput of several tens of Mbps, as these systems often have severe power constraints issuing from the tiny amounts of energy gathered by their energy harvesters. The proposed system was able to recover the transmitted clock and data through the use of simple yet efficient circuits. The clock regeneration circuit offers a simple alternative to more complex and costly solutions such as phase-locked loops, delay-locked loops and injection-locked oscillators. A simple equalization strategy adapted to the channel response was proven to perform well even in the presence of important ISI and very low slew-rates signals. Concerning the drivers, they are able to accommodate a wide range of line length

at a minimal power cost by the use of bootstrapped inputs and some slew-rate and swing control.

Even if the solution performs well in the tested conditions, four areas of improvement could still be distinguished: impedance matching, clock regeneration, calibration and mitigation of the system's settling time.

First, it is clear from the first chapter that the absence of impedance matching results in huge constraints on the rate/line length. The addition of impedance control to the drivers and the receiver input stage would make the solution more flexible and adapted to higher rates.

Next to the lack of impedance matching, the simple clock regeneration circuit could not sustain very high rates and presents a constraining settling time. Indeed, at 150Mbps, we reported a total jitter amounting to 1ns, making it impossible to consider Gbps communications. More complex and power-hungry timing recovery solutions should be implemented to properly track delays between clock and data.

Next, for the system to be practical, a calibration scheme is required. Three knobs can be easily tuned to ensure the system's correct operation: the slew-rate control voltage V_c , the upper driver voltage bound V_{high} and its counterpart V_{low} . If feedback is available, a simple BER-based scheme can be considered. Fortunately, we can expect the channel properties to stay constant over time and those parameters could in most cases be manually tuned at assembly time. We should also mention that the receiver sensitivity could be tuned by using different voltage references for the regenerative latch and thus accommodate for a larger range of signal's SNR's.

Lastly, the significant system's settling time (12 cycles) can be deterring for very short communications as no useful information is transmitted during the initial transient. Depending on the application's requirements, we could insert FIFO buffers between the data source and the transceiver to ensure that even with a low source throughput, data are sent in long streams over the channel, thus mitigating the cost of the transient.

BIBLIOGRAPHY

- [Allam and Elmasry, 2001] Allam, M. W. and Elmasry, M. I. (2001). Dynamic current mode logic (dycml): A new low-power high-performance logic style. *Solid-State Circuits, IEEE Journal of*, 36(3):550–558.
- [Allen, 2002] Allen, P. E. (2002). Discrete-time comparators. Lecture given at the Georgia Institute of Technology.
- [Analui et al., 2005] Analui, B., Buckwalter, J. F., and Hajimiri, A. (2005). Data-dependent jitter in serial communications. *Microwave Theory and Techniques, IEEE Transactions on*, 53(11):3388–3397.
- [Anritsu, 2010] Anritsu (2010). Understanding eye pattern measurements. Technical report, Anritsu.
- [Ashton, 2009] Ashton, K. (2009). That ‘internet of things’ thing. *RFID Journal*, 22(7):97–114.
- [Baker, 2010] Baker, R. J. (2010). *CMOS Circuit Design, Layout, and Simulation*. Wiley-IEEE Press, 3rd edition.
- [Bauer et al., 2013] Bauer, H., Veira, J., and Weig, F. (2013). Moore’s law: Repeal or renewal. *New York: McKinsey & Company*.
- [Bhansali and Roychowdhury, 2009] Bhansali, P. and Roychowdhury, J. (2009). Gen-adler: the generalized adler’s equation for injection locking analysis in oscillators. In *Proceedings of the 2009 Asia and South Pacific Design Automation Conference*, pages 522–527. IEEE Press.
- [Black, 1953] Black, H. S. (1953). *Modulation theory*. van Nostrand.
- [Bol et al., 2014] Bol, D., de Streel, G., Botman, F., Lusala, A. K., and Couniot, N. (2014). A 65-nm 0.5-v 17-pj/frame. pixel dps cmos image sensor for ultra-low-power socs achieving 40-db dynamic range. In *VLSI Circuits Digest of Technical Papers, 2014 Symposium on*, pages 1–2. IEEE.
- [Bol et al., 2013] Bol, D., De Vos, J., Botman, F., de Streel, G., Bernard, S., Flandre, D., and Legat, J.-D. (2013). Green socs for a sustainable internet-of-things. In *Faible Tension Faible Consommation (FTFC), 2013 IEEE*, pages 1–4. IEEE.
- [Bryzek, 2012] Bryzek, J. (2012). Emergence of a trillion mems sensor market.

- [Buckwalter et al., 2004] Buckwalter, J., Analui, B., and Hajimiri, A. (2004). Predicting data-dependent jitter. *Circuits and Systems II: Express Briefs, IEEE Transactions on*, 51(9):453–457.
- [Buckwalter and Hajimiri, 2006] Buckwalter, J. F. and Hajimiri, A. (2006). Analysis and equalization of data-dependent jitter. *Solid-State Circuits, IEEE Journal of*, 41(3):607–620.
- [Chappell et al., 1988] Chappell, B. A., Chappell, T. I., Schuster, S. E., Segmuller, H. M., Allan, J. W., Franch, R. L., and Restle, P. J. (1988). Fast cmos ecl receivers with 100-mv worst-case sensitivity. *Solid-State Circuits, IEEE Journal of*, 23(1):59–67.
- [Cho and Gray, 1995] Cho, T. B. and Gray, P. R. (1995). A 10 b, 20 msample/s, 35 mw pipeline a/d converter. *Solid-State Circuits, IEEE Journal of*, 30(3):166–172.
- [Choi et al., 2012] Choi, J., Park, S., Cho, J., and Yoon, E. (2012). A $1.36\mu\text{w}$ adaptive cmos image sensor with reconfigurable modes of operation from available energy/illumination for distributed wireless sensor network. In *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2012 IEEE International*, pages 112–114. IEEE.
- [Choi and Lee, 1994] Choi, J.-S. and Lee, K. (1994). Design of cmos tapered buffer for minimum power-delay product. *Solid-State Circuits, IEEE Journal of*, 29(9):1142–1145.
- [Cisco, 2015] Cisco (2015). The internet of everything cisco ioe value index study.
- [Diorio, 2004] Diorio, C. (2004). High-speed signaling. Lecture given at WashingtonUniversity.
- [Edson, 2015] Edson, B. (2015). Creating the internet of your things. *Microsoft Corporation*.
- [EIA, 2015] EIA, U. E. I. A. (2015). How many smart meters are installed in the united states, and who has them? <http://www.eia.gov/tools/faqs/faq.cfm?id=108&t=3>. [Online; accessed 20-May-2015].
- [Frustaci et al., 2011] Frustaci, F., Corsonello, P., and Alioto, M. (2011). Tapered-v th cmos buffer design for improved energy efficiency in deep nanometer technology. In *Circuits and Systems (ISCAS), 2011 IEEE International Symposium on*, pages 2075–2078. IEEE.
- [Fu, 2012] Fu, H. (2012). Equalization for high-speed serial interfaces in xilinx 7 series fpga transceivers. Technical report, Xilinx.
- [Gabriel and Mcvittie, 1992] Gabriel, C. T. and Mcvittie, J. P. (1992). How plasma etching damages thin gate oxides. *Solid State Technology*, 35(6):81–87.
- [Green and Singh, 2003] Green, M. M. and Singh, U. (2003). Design of cmos cml circuits for high-speed broadband communications. In *Circuits and Systems, 2003. ISCAS'03. Proceedings of the 2003 International Symposium on*, pages II–204.
- [Hajimiri et al., 1999] Hajimiri, A., Limotyakis, S., and Lee, T. H. (1999). Jitter and phase noise in ring oscillators. *Solid-State Circuits, IEEE Journal of*, 34(6):790–804.

- [Hu et al., 2012] Hu, K., Bai, R., Jiang, T., Ma, C., Ragab, A., Palermo, S., and Chiang, P. Y. (2012). 0.16-0.25 pj/bit, 8 gb/s near-threshold serial link receiver with super-harmonic injection-locking. *Solid-State Circuits, IEEE Journal of*, 47(8):1842–1853.
- [Hu et al., 2010] Hu, K., Jiang, T., Wang, J., O’Mahony, F., and Chiang, P. Y. (2010). A 0.6 mw/gb/s, 6.4–7.2 gb/s serial link receiver using local injection-locked ring oscillators in 90 nm cmos. *Solid-State Circuits, IEEE Journal of*, 45(4):899–908.
- [Jaeger and Linholm, 1975] Jaeger, R. C. and Linholm, L. (1975). Comments on "an optimized output stage for mos integrated circuits"[with reply]. *Solid-State Circuits, IEEE Journal of*, 10(3):185–186.
- [Kim et al., 2009] Kim, B., Liu, Y., Dickson, T. O., Bulzacchelli, J. F., and Friedman, D. J. (2009). A 10-gb/s compact low-power serial i/o with dfe-iir equalization in 65-nm cmos. *Solid-State Circuits, IEEE Journal of*, 44(12):3526–3538.
- [Kim et al., 2014] Kim, G., Lee, Y., Foo, Z., Pannuto, P., Kuo, Y.-S., Kempke, B., Ghaed, M. H., Bang, S., Lee, I., Kim, Y., et al. (2014). A millimeter-scale wireless imaging system with continuous motion detection and energy harvesting. In *VLSI Circuits Digest of Technical Papers, 2014 Symposium on*, pages 1–2. IEEE.
- [Kyeongho et al., 1995] Kyeongho, L., Sungjoon, K., Gijung, A., and Jeong, D.-K. (1995). A cmos serial link for fully duplexed data communication. *IEICE Transactions on Electronics*, 78(6):601–612.
- [Lee and Chandrakasan, 2007] Lee, F. S. and Chandrakasan, A. P. (2007). A 2.5 nj/bit 0.65 v pulsed uwb receiver in 90 nm cmos. *Solid-State Circuits, IEEE Journal of*, 42(12):2851–2859.
- [Lee et al., 2000] Lee, M.-J., Dally, W. J., and Chiang, P. (2000). Low-power area-efficient high-speed i/o circuit techniques. *Solid-State Circuits, IEEE Journal of*, 35(11):1591–1599.
- [Li et al., 2012] Li, S., Wang, H., Xu, T., and Zhou, G. (2012). Application study on internet of things in environment protection field. In *Informatics in Control, Automation and Robotics*, pages 99–106. Springer.
- [Mantaro Product Development Services, 2015] Mantaro Product Development Services (2015). *Microstrip Impedance Calculator*. Mantaro Product Development Services. http://www.mantaro.com/resources/impedance_calculator.htm.
- [Mentor Graphics Corporation, 2005] Mentor Graphics Corporation (2005). *Eldo User’s manual*. Mentor Graphics Corporation. Rev. 6.0.
- [Nyquist, 1928] Nyquist, H. (1928). Certain topics in telegraph transmission theory. *American Institute of Electrical Engineers, Transactions of the*, 47(2):617–644.
- [OnSemiconductors, 2014] OnSemiconductors (2014). Understanding data eye diagram methodology for analyzing high speed digital signals. Technical report, ON Semiconductors.
- [P. Sobieski, 2010] P. Sobieski, L. V. (2010). *Conception de Modems*. Université catholique de Louvain.

- [Palermo, 2010a] Palermo, S. (2010a). High-speed serial I/O design for channel limited and power-constrained systems. Texas A&M University.
- [Palermo, 2010b] Palermo, S. (2010b). Special topics in high-speed links circuits and systems. Lecture given at Texas A&M University.
- [Pannuto et al., 2015] Pannuto, P., Lee, Y., Kuo, Y.-S., Foo, Z., Kempke, B., Kim, G., Dreslinski Jr, R., Blaauw, D., and Dutta, P. (2015). Mbus: An ultra-low power interconnect bus for next generation nanopower systems. In *Proceedings of the 42nd International Symposium on Computer Architecture (ISCA'15)*.
- [Peffer, 2003] Peffer, M. (2003). The benefits of using linear equalization in backplane and cable applications. Technical report, Texas Instruments.
- [Poulton, 1998] Poulton, J. (1998). Signaling in high-performance memory systems. In *International Solid-State Circuits Conference*.
- [Ratasuk et al.,] Ratasuk, R., Prasad, A., Li, Z., Ghosh, A., and Uusitalo, M. A. Recent advancements in m2m communications in 4g networks and evolution towards 5g.
- [Schelkunoff, 1934] Schelkunoff, S. A. (1934). The electromagnetic theory of coaxial transmission lines and cylindrical shields. *Bell System Technical Journal*, 13(4):532–579.
- [Schrader et al., 2006] Schrader, J.-R., Klumperink, E. A., Visschers, J. L., and Nauta, M. (2006). Wireline equalization using pulse-width modulation. In *Custom Integrated Circuits Conference, 2006. CICC'06. IEEE*. IEEE.
- [Sehgal et al., 2015] Sehgal, A., Romascanu, D., and Ersue, M. (2015). Management of networks with constrained devices: Use cases. *Management*.
- [Severi et al., 2014] Severi, S., Sottile, F., Abreu, G., Pastrone, C., Spirito, M., and Berens, F. (2014). M2m technologies: Enablers for a pervasive internet of things. In *Networks and Communications (EuCNC), 2014 European Conference on*, pages 1–5. IEEE.
- [Sidiropoulos and Horowitz, 1997] Sidiropoulos, S. and Horowitz, M. A. (1997). A semidigital dual delay-locked loop. *Solid-State Circuits, IEEE Journal of*, 32(11):1683–1692.
- [Song et al., 2009] Song, H., Kim, S., and Jeong, D.-K. (2009). A reduced-swing voltage-mode driver for low-power multi-gb/s transmitters. *Journal of Semiconductor Technology and Science*, 9(2):104–109.
- [Song et al., 2013] Song, Y.-H., Bai, R., Hu, K., Yang, H.-W., Chiang, P. Y., and Palermo, S. (2013). A 0.47–0.66 pj/bit, 4.8–8 gb/s i/o transceiver in 65 nm cmos. *Solid-State Circuits, IEEE Journal of*, 48(5):1276–1289.
- [ST Microelectronics, 2007] ST Microelectronics (2007). *IO65LPHVT_ANA_50A_7M4X0Y2Z library*. ST Microelectronics. Rev. 24.
- [ST Microelectronics, 2008a] ST Microelectronics (2008a). *CMOS065 technology LVT_LP MOS transistor models*. ST Microelectronics. Rev. 1.3d.

- [ST Microelectronics, 2008b] ST Microelectronics (2008b). *CORE65LPLVT1.00V Standard cell library*. ST Microelectronics. Rev. 5.1.
- [Swan, 2012] Swan, M. (2012). Sensor mania! the internet of things, wearable computing, objective metrics, and the quantified self 2.0. *Journal of Sensor and Actuator Networks*, 1(3):217–253.
- [Wong et al., 2004] Wong, K.-L., Hatamkhani, H., Mansuri, M., and Yang, C.-K. (2004). A 27-mw 3.6-gb/s i/o transceiver. *Solid-State Circuits, IEEE Journal of*, 39(4):602–612.
- [Yang, 1998] Yang, C.-K. K. (1998). Design of high-speed serial links in cmos systems. Technical report, Stanford University.
- [Yole, 2013] Yole, D. (2013). Mems front-end manufacturing trends. *Research and Markets*, 3:96–114.